

# Practical Parallel Processing for Today's Rendering Challenges

Alan Chalmers   Tim Davis   Toshi Kato   Erik Reinhard



Course #40

SIGGRAPH 2001

Los Angeles  
12-17 August 2001



## About the authors



Alan Chalmers is a senior lecturer in the Department of Computer Science at the University of Bristol, UK. He has published over 70 papers in journals and international conferences on parallel photo-realistic graphics. He was the co-chairman of the recent IEEE Parallel Visualization and Graphics Symposium and is chairman of the Eurographics workshop series on Parallel Graphics and Visualisation. Recently he and Professor F. W. Jansen were the guest editors of the Journal of Parallel Computing for its special edition on Parallel Graphics and Visualisation. He is also currently vice-chair of ACM SIGGRAPH. His research interests include the application of parallel photo-realistic graphics to archaeological site visualisation in order to provide a flexible tool for investigating site reconstruction and utilization.



Timothy Davis is an assistant professor in the Computer Science Department at Clemson University where he works within the newly established MFAC (Master of Fine Arts in Computing) group, which trains students to produce special effects for entertainment and commercial projects. Most recently, he has published papers at the 1999 IEEE Parallel Visualization and Graphics Symposium and the 1998 Eurographics Workshop on Parallel Graphics and Visualisation. His research involves exploiting spatio-temporal coherence in a parallel environment to reduce rendering times for ray-traced animations. He received his Ph.D. from North Carolina State University and has worked in technical positions for the Environmental Protection Agency and NASA Goddard Space Flight Center.



For sixteen years Toshi Kato has consistently devoted his time to the development of high-end production-level renderers. Currently he is the project leader of an ongoing project "Kilauea," a parallel ray tracer developed at Square USA R&D department. The development of "Kilauea" involves various fields including system architecture design, implementation of a message passing layer, and a parallel shading and ray tracing engine. From 1993 to 1998 he developed an in-house scanline-based renderer at Rhythm & Hues in LA and participated in the production of movies such as Mouse Hunt (1998), Kazaam (1996) and Babe (1995). He also developed a special stereoscopic renderer for IMAX's Omnimax Solido dome.



Erik Reinhard is a researcher at the University of Utah in the fields of parallel ray tracing and visual perception. He received a 'TWAIO' diploma in parallel computer graphics from Delft University of Technology in 1996 and a PhD degree from the University of Bristol in 2000. His current research interests include algorithms and data structures for real-time ray tracing. This includes mechanisms to trade-off visual quality for interactivity and to add animation capabilities to real-time ray tracing. He has published more than 20 papers on parallel rendering.

# Structure of the Course

Parallel processing offers the potential of rendering high-quality images and animations in reasonable times. This course begins by reviewing the basic issues involved in rendering within a parallel or distributed computing environment. Specifically, various methods are presented for dividing the original rendering problem into subtasks and distributing them efficiently to independent processors. Careful consideration must be taken to balance the processing load across the processors, as well as reduce communication between these subtasks for faster processing.

The course continues by examining the strengths and weaknesses of multiprocessor machines and networked render farms for graphics rendering. Case studies of working applications including “real time raytracing”, and streamlining the creation of full CG movies (“Final Fantasy”) will be presented to demonstrate in detail practical ways of dealing with the issues involved.

Introduction (10 minutes) (Davis)

- The need for speed in satisfying the demand for high-quality graphics
- Rendering and parallel processing: a holy union
- The exciting possibilities of parallel processing in a world of advanced 3D graphics cards

Parallel/Distributed Rendering Issues (45 minutes) (Chalmers)

- Task subdivision
- Load balancing
- Communication
- Task migration
- Data Management

Classification of Parallel Rendering Systems (25 minutes) (Davis)

By rendering technique

- Polygon rendering
  - Sort first, sort middle , sort last
- Photo-realistic methods
  - Image space subdivision, object space subdivision, object subdivision

By hardware:

- Multiprocessor machines
  - Pros and cons
- Distributed Computing
  - Introduction to render farms
  - Pros and cons

Practical Applications (25 minutes) (Davis)

- Distributed computing and spatial/temporal coherence

- Animations

Practical Applications continued

Getting the most from your machine (40 minutes) (Reinhard)

- Real time raytracing
  - basic operations
  - animation and interactivity
  - adding complexity

Parallel rendering and the quest for realism (45 minutes) (Kato)

- The “Kilauea” massively parallel ray tracer and “Final Fantasy”
  - System design
  - Memory management for thread environments
  - Parallel shading calculations, space traversal, photo maps
  - Debugging and stability

Summary (10 minutes) (Chalmers)

Discussion and questions (10 minutes) (All)

## Structure of the notes

The notes contain important background information as well as detailed descriptions of the Parallel Processing techniques described. The notes are arranged as follows:

Section I introduces the concepts which are necessary to understand the difficulties confronting any parallel implementation. The section goes on in chapter 2 to describe the issues associated with scheduling tasks on a parallel system. A number of ways of decomposing a problem on a parallel machine are considered and the advantages and disadvantages of each approach are highlighted. Chapter 3 of Section I concentrates on managing large data requirements which are distributed across the parallel environment. Issues of consistency and latency are considered.

Section II considers the classification of parallel rendering systems according to the method of task subdivision and/or by the hardware used.

Section III considers Interactive Ray Tracing in depth including hardware considerations, animation and reuse techniques. Appendices are provided on the SGI Origin 2000.

Finally, Section IV provides details on the “Kilauea” massively parallel raytracer.

**Section I**  
**Parallel/Distributed Rendering Issues**

**Alan Chalmers**

# Contents

## Section I: Parallel/Distributed Rendering Issues

<b>I.1 Introduction</b>	<b>3</b>
1.1 Concepts . . . . .	4
1.1.1 Dependencies . . . . .	4
1.1.2 Scalability . . . . .	5
1.1.3 Control . . . . .	7
1.2 Classification of Parallel Systems . . . . .	9
1.2.1 Flynn's taxonomy . . . . .	9
1.2.2 Parallel versus Distributed systems . . . . .	12
1.3 The Relationship of Tasks and Data . . . . .	13
1.3.1 Inherent difficulties . . . . .	14
1.3.2 Tasks . . . . .	14
1.3.3 Data . . . . .	14
1.4 Evaluating Parallel Implementations . . . . .	15
1.4.1 Realisation Penalties . . . . .	15
1.4.2 Performance Metrics . . . . .	17
1.4.3 Efficiency . . . . .	21
<b>I.2 Task Scheduling</b>	<b>25</b>
2.1 Problem Decomposition . . . . .	25
2.1.1 Algorithmic decomposition . . . . .	26
2.1.2 Domain decomposition . . . . .	26
2.1.3 Abstract definition of a task . . . . .	27
2.1.4 System architecture . . . . .	27
2.2 Computational Models . . . . .	28
2.2.1 Data driven model . . . . .	29
2.2.2 Demand driven model . . . . .	33
2.2.3 Hybrid computational model . . . . .	36
2.3 Task Management . . . . .	36
2.3.1 Task definition and granularity . . . . .	36
2.3.2 Task distribution and control . . . . .	37
2.3.3 Algorithmic dependencies . . . . .	38
2.4 Task Scheduling Strategies . . . . .	41
2.4.1 Data driven task management strategies . . . . .	41
2.4.2 Demand driven task management strategies . . . . .	41
2.4.3 Task manager process . . . . .	45
2.4.4 Distributed task management . . . . .	47
2.4.5 Preferred bias task allocation . . . . .	48

<b>I.3 Data Management</b>	<b>50</b>
3.1 World Model of the Data: No Data Management Required . . . . .	50
3.2 Virtual Shared Memory . . . . .	50
3.2.1 Implementing virtual shared memory . . . . .	51
3.3 The Data Manager . . . . .	52
3.3.1 The local data cache . . . . .	52
3.3.2 Requesting data items . . . . .	54
3.3.3 Locating data items . . . . .	54
3.4 Consistency . . . . .	58
3.4.1 Keeping the data items consistent . . . . .	58
3.4.2 Weak consistency: repair consistency on request . . . . .	60
3.4.3 Repair consistency on synchronisation: Release consistency . . . . .	61
3.5 Minimising the Impact of Remote Data Requests . . . . .	61
3.5.1 Prefetching . . . . .	61
3.5.2 Multi-threading . . . . .	62
3.5.3 Profiling . . . . .	64
3.6 Data Management for Multi-Stage Problems . . . . .	65

# Chapter 1

## Introduction

*Parallel processing is like a dog's walking on its hind legs. It is not done well, but you are surprised to find it done at all.*

[Steve Fiddes (University of Bristol) with apologies to Samuel Johnson]

*Realistic computer graphics is an area of research which develops algorithms and methods to render images of artificial models or worlds as realistically as possible. Such algorithms are known for their unpredictable data accesses and their high computational complexity. Rendering a single high quality image may take several hours, or even days. Parallel processing offers the potential for solving such complex problems in reasonable times.*

*However, there are a number of fundamental issues: task scheduling, data management and caching techniques, which must be addressed if parallel processing is to achieve the desired performance when computing realistic images. These are applicable for all three rendering techniques presented in this tutorial: ray tracing, radiosity and particle tracing.*

*This chapter introduces the concepts of parallel processing, describes its development and considers the difficulties associated with solving problems in parallel.*

Parallel processing is an integral part of everyday life. The concept is so ingrained in our existence that we benefit from it without realising. When faced with a taxing problem, we involve others to solve it more easily. This co-operation of more than one worker to facilitate the solution of a particular problem may be termed parallel processing. The goal of parallel processing is thus to solve a given problem more rapidly, or to enable the solution of a problem that would otherwise be impracticable by a single worker.

The principles of parallel processing are, however, not new, as evidence suggests that the computational devices used over 2000 years ago by the Greeks recognised and exploited such concepts. In the Nineteenth Century, Babbage used parallel processing in order to improve the performance of his Analytical Engine [48]. Indeed, the first general purpose electronic digital computer, the ENIAC, was conceived as a highly parallel and decentralised machine with twenty-five independent computing units, co-operating towards the solution of a single problem [27].

However, the early computer developers rapidly identified two obstacles restricting the widespread acceptance of parallel machines: the complexity of construction; and, the seemingly high programming effort required [10]. As a result of these early set-backs, the developmental thrust shifted to computers with a single computing unit, to the detriment of parallel designs. Additionally, the availability of sequential machines resulted in the development of algorithms and techniques optimised for these particular architectures.

The evolution of serial computers may be finally reaching its zenith due to the limitations imposed on the design by its physical implementation and inherent bottlenecks [5]. As users continue to demand improved performance, computer designers have been looking increasingly at parallel approaches to overcome these limitations. All modern computer architectures incorporate a degree of parallelism. Improved hardware design and manufacture coupled with a growing understanding of how to tackle the difficulties of parallel programming has re-established parallel processing at the forefront of computer technology.

## 1.1 Concepts

Parallel processing is the solution of a single problem by dividing it into a number of sub-problems, each of which may be solved by a separate worker. Co-operation will always be necessary between workers during problem solution, even if this is a simple agreement on the division of labour. These ideas can be illustrated by a simple analogy of tackling the problem of emptying a swimming pool using buckets. This job may be sub-divided into the repeated *task* of removing one bucket of water.

A single person will complete all the tasks, and complete the job, in a certain time. This process may be speeded-up by utilising additional workers. Ideally, two people should be able to empty the pool in half the time. Extending this argument, a large number of workers should be able to complete the job in a small fraction of the original time. However, practically there are physical limitations preventing this hypothetical situation.

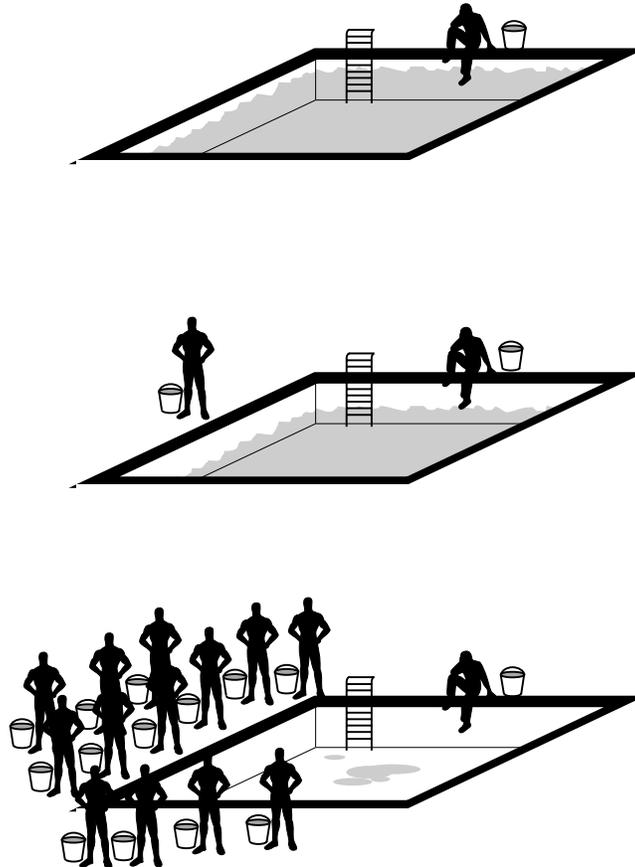


Figure 1.1: Emptying a pool by means of a bucket

The physical realisation of this solution necessitates a basic level of co-operation between workers. This manifests itself due to the contention for access to the pool, and the need to avoid collision. The time required to achieve this co-operation involves inter-worker communication which detracts from the overall solution time, and as such may be termed an *overhead*.

### 1.1.1 Dependencies

Another factor preventing an ideal parallel solution are termed: dependencies. Consider the problem of constructing a house. In simple terms, building the roof can only commence after the walls have been completed. Similarly, the walls can only be erected once the foundations are laid. The roof is thus dependent upon the walls, which are in turn dependent on the foundations. These dependencies divide the whole

problem into a number of distinct stages. The parallel solution of each stage must be completed before the subsequent stage can start.

The dependencies within a problem may be so severe that it is not amenable to parallel processing. A strictly sequential problem consists of a number of stages, each comprising a single task, and each dependent upon the previous stage. For example, in figure 1.2, building a tower of toy blocks requires a strictly sequential order of task completion. The situation is the antithesis of dependency-free problems, such as placing blocks in a row on the floor. In this case, the order of task completion is unimportant, but the need for co-operation will still exist.

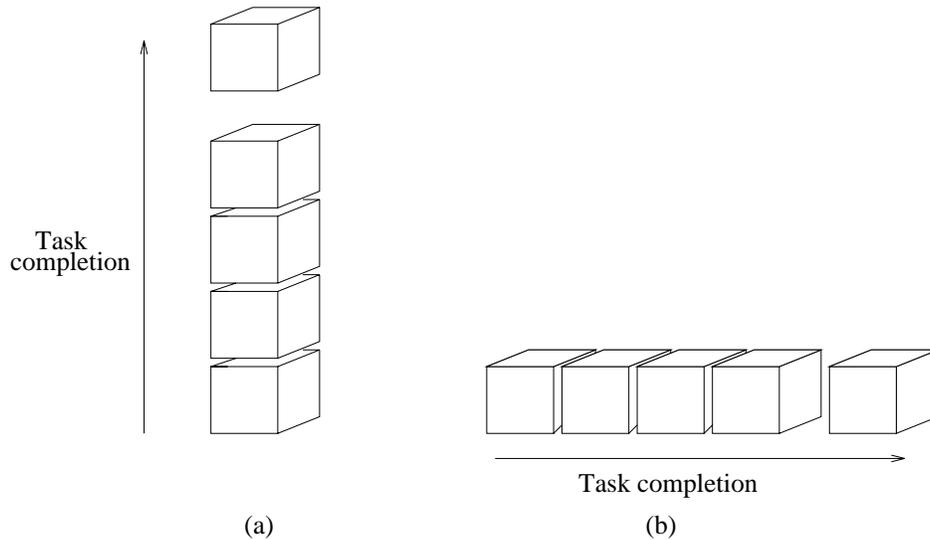


Figure 1.2: Building with blocks: (a) Strictly sequential (b) dependency-free

Pipelining is the classic methodology for minimising the effects of dependencies. This technique can only be exploited when a process, consisting of a number of distinct stages, needs to be repeated several times. An automotive assembly line is an example of an efficient pipeline. In a simplistic form, the construction of a car may consist of four linearly dependent stages as shown in figure 1.3: chassis fabrication; body assembly; wheel fitting; and, windscreen installation. An initial lump of metal is introduced into the pipeline then, as the partially completed car passes each stage, a new section is added until finally the finished car is available outside the factory.

Consider an implementation of this process consisting of four workers, each performing their task in one time unit. Having completed the task, the worker passes the partially completed car on to the next stage. This worker is now free to repeat its task on a new component fed from the previous stage. The completion of the first car occurs after four time units, but each subsequent car is completed every time unit.

The completion of a car is, of course, sensitive to the time taken by each worker. If one worker were to take longer than one time unit to complete its task then the worker after this difficult task would stand idle awaiting the next component, whilst those before the worker with the difficult task would be unable to move their component on to the next stage of the pipeline. The other workers would thus also be unable to do any further work until the difficult task was completed. Should there be any interruption in the input to the pipeline then the pipeline would once more have to be “refilled” before it could operate at maximum efficiency.

### 1.1.2 Scalability

Every problem contains an upper bound on the number of workers which can be meaningfully employed in its solution. Additional workers beyond this number will not improve solution time, and can indeed be detrimental. This upper bound provides an idea as to how suitable a problem is to parallel implementation: a measure of its *scalability*.

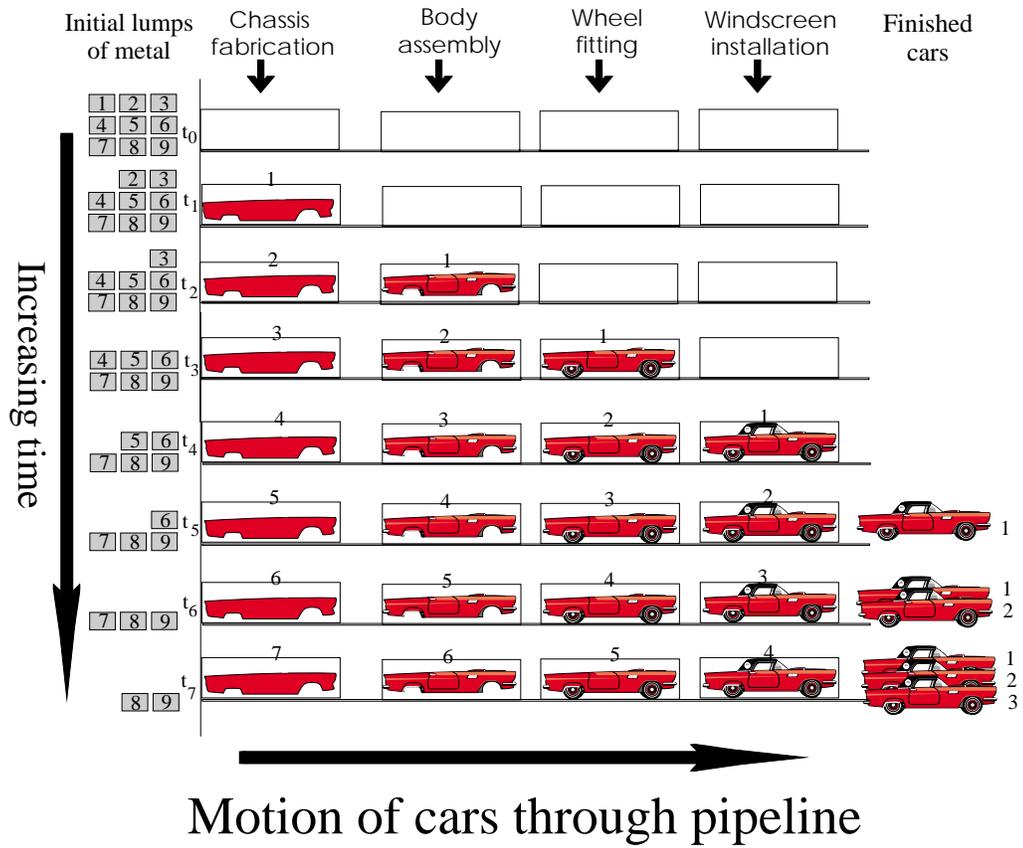


Figure 1.3: Pipeline assembly of a car

A given problem may only be divided into a finite number of sub-problems, corresponding to the smallest tasks. The availability of more workers than there are tasks, will not improve solution time. The problem of clearing a room of 100 chairs may be divided into 100 tasks consisting of removing a single chair. A maximum of 100 workers can be allocated one of these tasks and hence perform useful work.

The optimum solution time for clearing the room may not in fact occur when employing 100 workers due to certain aspects of the problem limiting effective worker utilisation. This phenomenon can be illustrated by adding a constraint to the problem, in the form of a single doorway providing egress from the room. A *bottleneck* will occur as large numbers of workers attempt to move their chairs through the door simultaneously, as shown in figure 1.4.

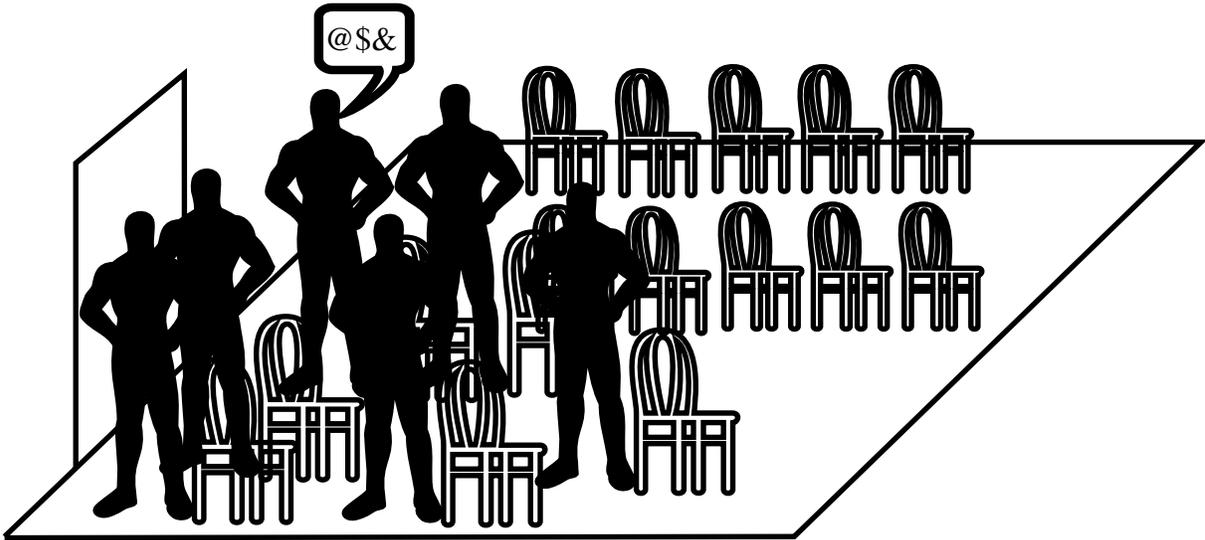


Figure 1.4: Bottleneck caused by doorway

The delays caused by this bottleneck may be so great that the time taken to empty the room of chairs by this large number of workers may in fact be *longer* than the original time taken by the single worker. In this case, reducing the number of workers can alleviate the bottleneck and thus reduce solution time.

### 1.1.3 Control

All parallel solutions of a problem require some form of control. This may be as simple as the control needed to determine what will constitute a task and to ascertain when the problem has been solved satisfactorily. More complex problems may require control at several stages of their solution. For example, solution time could be improved when clearing the room if a controller was placed at the door to schedule its usage. This control would ensure that no time was wasted by two (or more) workers attempting to exit simultaneously and then having to “reverse” to allow a single worker through. An alternative to this explicit *centralised* control would be some form of *distributed* control. Here the workers themselves could have a way of preventing simultaneous access, for example, if two (or more) workers reach the door at the same time then the biggest worker will always go first while the others wait.

Figure 1.5(a) shows the sequential approach to solving a problem. Computation is applied to the problem domain to produce the desired results. The controlled parallel approach shown in figure 1.5(b) achieves a parallel implementation of the same problem via three steps. In step 1, the problem domain is divided into a number of sub-problems, in this case four. Parallel processing is introduced in step 2 to enable each of the sub-problems to be computed in parallel to produce sub-results. In step 3, these results must now be collated to achieve the desired final results. Control is necessary in steps 1 and 3 to divide the problem amongst the workers and then to collect and collate the results that the workers have independently produced.

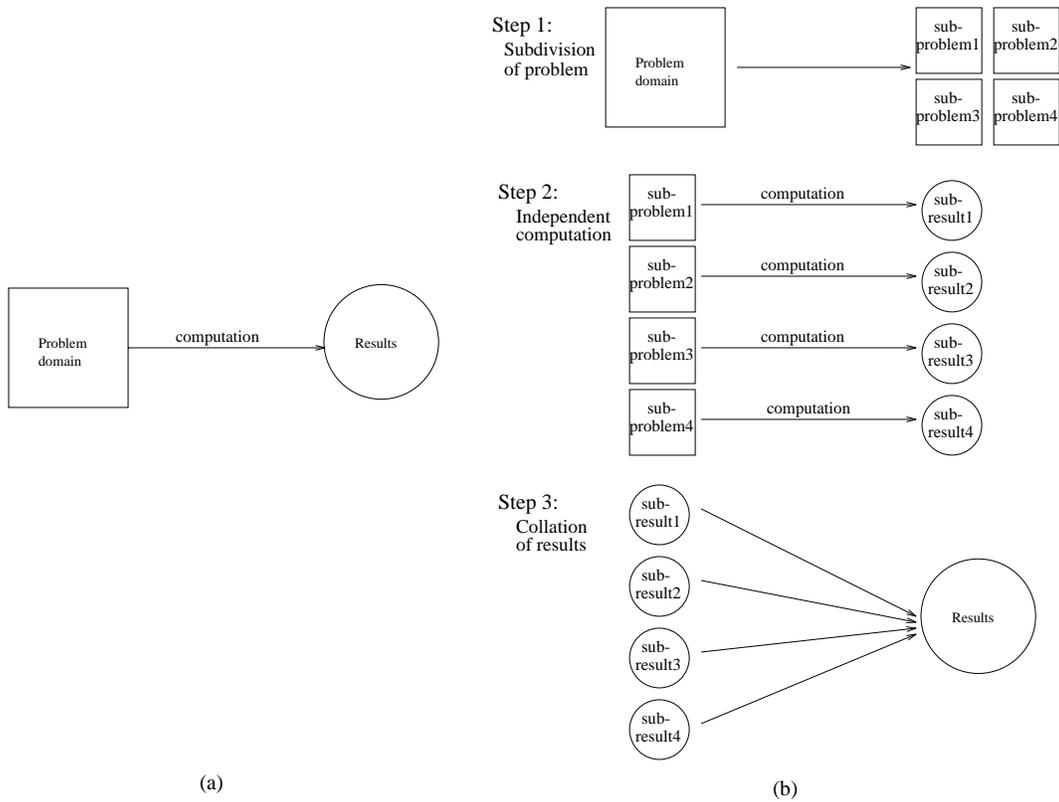


Figure 1.5: Control required in (a) a Sequential versus (b) a parallel implementation

## 1.2 Classification of Parallel Systems

A traditional sequential computer conforms to the *von Neumann* model. Shown in figure 1.6, this model comprises a processor, an associated memory, an input/output interface and various busses connecting these devices. The processor in the von Neumann model is the single computational unit responsible for the functions of fetching, decoding and executing a program's instructions. Parallel processing may be added to this architecture through pipelining using multiple functional units within a single computational unit or by replicating entire computational units (which may contain pipelining). With pipelining, each functional unit repeatedly performs the same operation on data which is received from the preceding functional unit. So in the simplest case, a pipeline for a computational unit could consist of three functional units, one to fetch the instructions from memory, one to decode these instructions and one to execute the decoded instructions. As we saw with the automobile assemblage example, a pipeline is only as effective as its slowest component. Any delay in the pipeline has repercussions for the whole system.

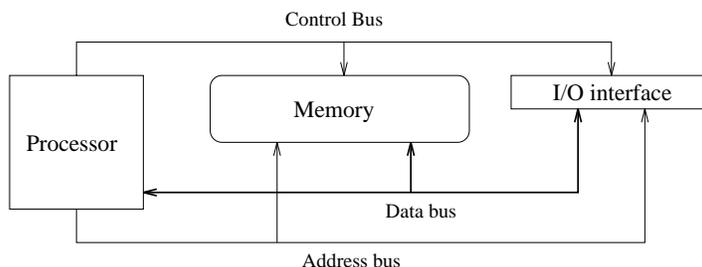


Figure 1.6: Von Neumann model architecture

*Vector processing* was introduced to provide efficient execution of program loops on large array data structures. By providing multiple registers as special vector registers to be used alongside the central processing unit, a vector processor is able to perform the *same operation on all elements* of a vector simultaneously. This simultaneous execution on every element of large arrays can produce significant performance improvements over conventional *scalar processing*. However, often problems need to be reformulate to benefit from this form of parallelism. A large number of scientific problems, such as weather forecasting, nuclear research and seismic data analysis, are well suited to vector processing.

Replication of the entire computational unit, the *processor*, allows individual tasks to be executed on different processors. Tasks are thus sometimes referred to as *virtual processors* which are allocated a physical processor on which to run. The completion of each task contributes to the solution of the problem.

Tasks which are executing on distinct processors at any point in time are said to be running in *parallel*. It may also be possible to execute several tasks on a single processor. Over a period of time the impression is given that they are running in parallel, when in fact at any point in time only *one* task has control of the processor. In this case we say that the tasks are being performed *concurrently*, that is their execution is being shared by the same processor. The difference between parallel tasks and concurrent tasks is shown in figure 1.7.

The workers which perform the computational work and co-operate to facilitate the solution of a problem on a parallel computer are known as *processing elements* and are often abbreviated as *PEs*. A processing element consists of a processor, one or more tasks, and the software to enable the co-operation with other processing elements. A parallel system comprises of *more than one* processing element.

### 1.2.1 Flynn's taxonomy

The wide diversity of computer architectures that have been proposed, and in a large number of cases realised, since the 1940's has led to the desire to classify the designs to facilitate evaluation and comparison. Classification requires a means of identifying distinctive architectural or behavioural features of a machine.

In 1972 Flynn proposed a classification of processors according to a macroscopic view of their principal interaction patterns relating to instruction and data *streams* [21]. The term stream was used by Flynn to

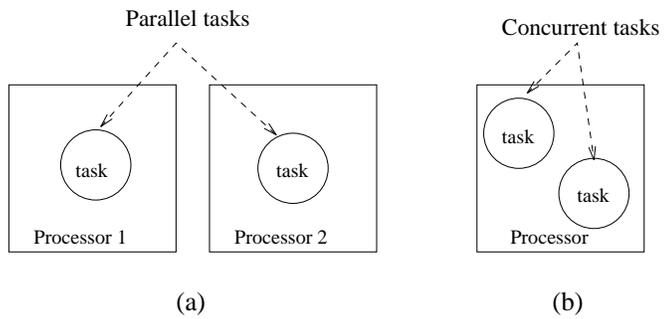


Figure 1.7: (a) Parallel tasks (b) Concurrent tasks

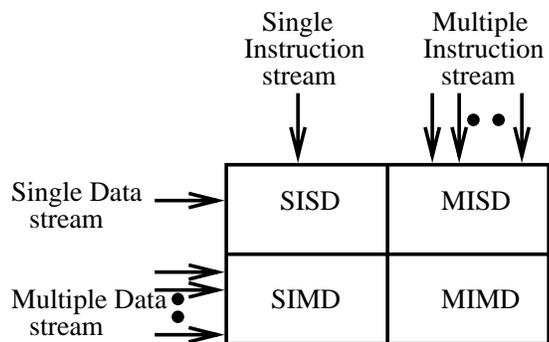


Figure 1.8: Flynn's taxonomy for processors

refer to the sequence of instructions to be executed, or data to be operated on, by the processor. What has become known as *Flynn's taxonomy* thus categorises architectures into the four areas shown in figure 1.8.

Since its inception, Flynn's taxonomy has been criticised as being too broad and has thus been enlarged by several other authors, for example, Shore in 1973 [55], Treleaven, Brownbridge and Hopkins in 1982 [58], Basu in 1984 [7], and perhaps one of the most detailed classifications was given by Hockney and Jesshope in 1988 [31].

Real architectures are, of course, much more complex than Flynn suggested. For example, an architecture may exhibit properties from more than one of his classes. However, if we are not too worried about the minute details of any individual machine then Flynn's taxonomy serves to separate fundamentally different architectures into four broad categories. The classification scheme is simple (which is one of the main reasons for its popularity) and thus useful to show an overview of the concepts of multiprocessor computers.

**SISD: Single Instruction Single Data** embraces the conventional sequential, or von Neumann, processor.

The single processing element executes instructions sequentially on a single data stream. The operations are thus ordered in time and may be easily traced from start to finish. Modern adaptations of this uniprocessor use some form of pipelining technique to improve performance and, as demonstrated by the Cray supercomputers, minimise the length of the component interconnections to reduce signal propagation times [54].

**SIMD: Single Instruction Multiple Data** machines apply a single instruction to a group of data items simultaneously.

A master instruction is thus acting over a vector of related operands. A number of processors, therefore, obey the same instruction in the same cycle and may be said to be executing in strict *lock-step*. Facilities exist to exclude particular processors from participating in a given instruction cycle. Vector processors, for example the Cyber 205, Fujitsu FACOM VP-200 and NEC SX1, and array processors, such as the DAP [53], Goodyear MPP (Massively Parallel Processor) [8], or the Connection Machine CM-1 [30], may be grouped in this category.

**MISD: Multiple Instruction Single Data** Although part of Flynn's taxonomy, no architecture falls obviously into the MISD category.

One the closest architecture to this concept is a pipelined computer. Another is systolic array architectures which derives their from the medical term "systole" used to describe the rhythmic contraction of chambers of the heart. Data arrives from different directions at regular intervals to be combined at the "cells" of the array. The Intel iWarp system was designed to support systolic computation [4]. Systolic arrays are well suited to specially designed algorithms rather than general purpose computing [40, 41].

**MIMD: Multiple Instruction Multiple Data** The processors within the MIMD classification autonomously

obey their own instruction sequence and apply these instructions to their own data. The processors are, therefore, no longer bound to the synchronous method of the SIMD processors and may choose to operate *asynchronously*. By providing these processors with the ability to communicate with each other, they may interact and therefore, co-operate in the solution of a single problem. This interaction has led to MIMD systems sometimes being classified as *tightly coupled* if the degree of interaction is high, or *loosely coupled* if the degree of interaction is low.

Two methods are available to facilitate this interprocessor communication. *Shared memory* systems allow the processors to communicate by reading and writing to a common address space. Controls are necessary to prevent processors updating the same portion of the shared memory simultaneously. Examples of such shared memory systems are the Sequent Balance [57] and the Alliant FX/8 [17].

In *distributed memory* systems, on the other hand, processors address only their private memory and communicate by passing messages along some form of communication path. Examples of MIMD processors from which such distributed memory systems can be built are the Intel i860 [50], the Immos transputer [33] and Analog Devices SHARC processor.

The conceptual difference between shared memory and distributed memory systems of MIMD processors is shown in figure 1.9. The interconnection method for the shared memory system, figure 1.9(a), allows all the processors to be connected to the shared memory. If two, or more, processors wish to access the same portion of this shared memory at the same time then some arbitration mechanism must be used to ensure only one processor accesses that memory portion at a time. This problem of memory contention may restrict the number of processors that can be interconnected

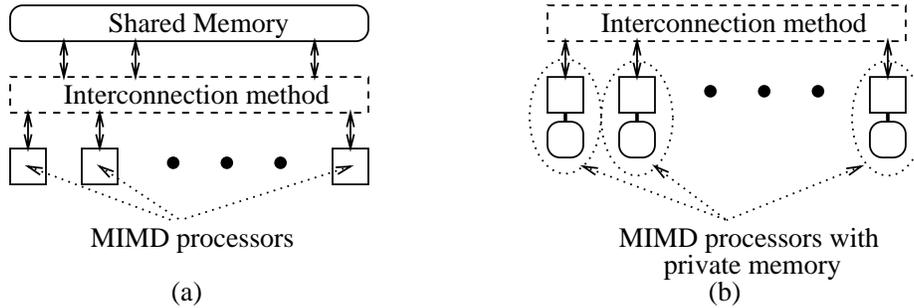


Figure 1.9: Systems of MIMD processors (a) shared memory (b) distributed memory

using the shared memory model. The interconnection method of the distributed memory system, figure 1.9(b), connects the processors in some fashion and if one, or more, processors wish to access another processor’s private memory, it, or they, can only do so by sending a message to the appropriate processor along this interconnection network. There is thus no memory contention as such. However, the density of the messages that result in distributed memory systems may still limit the number of processors that may be interconnected, although this number is generally larger than that of the shared memory systems.

Busses have been used successfully as an interconnection structure to connect low numbers of processors together. However, if more than one processor wishes to send a message on the bus at the same time, an arbiter must decide which message gets access to the bus first. As the number of processors increases, so the contention for use of the bus grows. Thus a bus is inappropriate for large multiprocessor systems. An alternative to the bus is to connect processors via dedicated links to form large networks. This removes the bus-contention problem by spreading the communication load across many independent links.

### 1.2.2 Parallel versus Distributed systems

Distributed memory MIMD systems consist of autonomous processors together with their own memory which co-operate in the solution of a single complex problem. Such systems may consist of a number of interconnected, dedicated processor and memory nodes, or interconnected “stand-alone” workstations. To distinguish between these two, the former configuration is referred to as a (dedicated) *parallel* system, while the latter is known as a *distributed* system, as shown in figure 1.10.

The main distinguishing features of these two systems are typically the computation-to-communication ratio and the cost. Parallel systems make use of fast, “purpose-built” (and thus expensive) communication infrastructures, while distributed systems rely on existing network facilities such as ethernet, which are significantly slower and susceptible to other non-related traffic.

The advantage of distributed systems is that they may consist of a cluster of existing workstations which can be used by many (sequential) users when not employed in a parallel capacity. A number of valuable tools have been developed to enable these workstations to act in parallel, such as Parallel Virtual Machine (PVM), and Message Passing Interface (MPI). These provide an easy framework for coupling heterogeneous computers including, workstations, mainframes and even parallel systems.

However, while some of the properties of a distributed computing system may be different from those of a parallel system, many of the underlying concepts are equivalent. For example, both systems achieve co-operation between computational units by passing messages, and each computational unit has its own distinct memory. Thus, the ideas presented in this tutorial should prove equally useful to the reader faced with implementing his or her realistic rendering problem on either system.

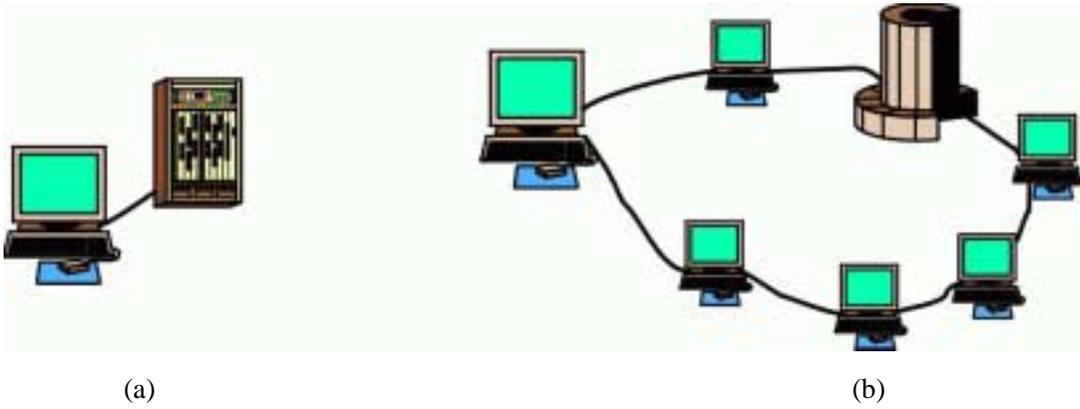


Figure 1.10: (a) Parallel system (b) Distributed system

### 1.3 The Relationship of Tasks and Data

The implementation of any problem on a computer comprises two components:

- the algorithm chosen to solve the problem; and,
- the domain of the problem which encompasses all the data requirements for that problem.

The algorithm interacts with the domain to produce the result for the problem, as shown diagrammatically in figure 1.11.

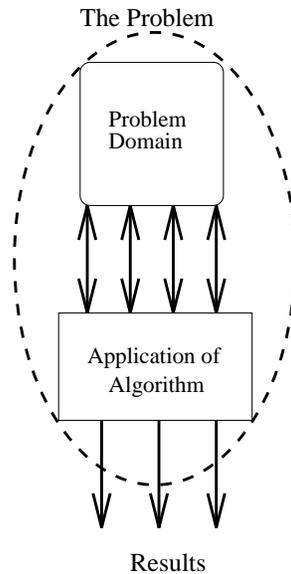


Figure 1.11: The components of a problem

A sequential implementation of the problem means that the entire algorithm and domain reside on a single processor. To achieve a parallel implementation it is necessary to divide the problem's components in some manner amongst the parallel processors. Now no longer resident on a single processor, the components will have to interact within the multiprocessor system in order to obtain the same result. This co-operation requirement introduces a number of novel difficulties into any parallel implementation which are not present in the sequential version of the same problem.

### 1.3.1 Inherent difficulties

User confidence in any computer implementation of a problem is bolstered by the successful termination of the computation and the fact that the results meet design specifications. The reliability of modern computer architectures and languages is such that any failure of a sequential implementation to complete successfully will point automatically to deficiencies in either the algorithm used or data supplied to the program. In addition to these possibilities of failure, a parallel implementation may also be affected by a number of other factors which arise from the manner of the implementation:

**Deadlock:** An active parallel processor is said to be deadlocked if it is waiting indefinitely for an event which will never occur. A simple example of deadlock is when two processors, using synchronised communication, attempt to send a message to each other at the same time. Each process will then wait for the other process to perform the corresponding input operation which will never occur.

**Data consistency:** In a parallel implementation, the problem's data may be distributed across several processors. Care has to be taken to ensure:

- if multiple copies of the same data item exists then the value of this item is kept consistent;
- mutual exclusion is maintained to avoid several processors accessing a shared resource simultaneously; and,
- the data items are fetched from remote locations efficiently in order to avoid processor idle time.

While there is meaningful computation to be performed, a sequential computer is able to devote 100% of its time for this purpose. In a parallel system it may happen that some of the processors become idle, not because there is no more work to be done, but because current circumstances prevent those processors being able to perform any computation.

Parallel processing introduces communication overheads. The effect of these overheads is to introduce latency into the multiprocessor system. Unless some way is found to minimise communication delays, the percentage of time that a processor can spend on useful computation may be significantly affected. So, as well as the factors affecting the successful termination of the parallel implementation, one of the fundamental considerations also facing parallel programmers is the *computation to communication* ratio.

### 1.3.2 Tasks

Subdividing a single problem amongst many processors introduces the notion of a task. In its most general sense, a task is a unit of computation which is assigned to a processor within the parallel system. In any parallel implementation a decision has to be taken as to what exactly constitutes a task. The *task granularity* of a problem is a measure of the amount of computational effort associated with any task. The choice of granularity has a direct bearing on the computation to communication ratio. Selection of too large a granularity may prevent the solution of the problem on a large parallel system, while too fine a granularity may result in significant processor idle time while the system attempts to keep processors supplied with fresh tasks.

On completion of a sequential implementation of a problem, any statistics that may have been gathered during the course of the computation, may now be displayed in a straightforward manner. Furthermore, the computer is in a state ready to commence the next sequential program. In a multiprocessor system, the statistics would have been gathered at each processor, so after the solution of the problem the programmer is still faced with the task of collecting and collating these statistics. To ensure that the multiprocessor system is in the correct state for the next parallel program, the programmer must also ensure that all the processors have *terminated gracefully*.

### 1.3.3 Data

The problem domains of many rendering applications are very large. The size of these domains are typically far more than can be accommodated within the local memory of any processing element (or indeed in the memory of many sequential computers). Yet it is precisely these complex problems that we wish to solve using parallel processing.

Consider a multiprocessor system consisting of sixty-four processing elements each with 4 MBytes of local memory. If we were to insist that the entire problem domain were to reside at each processing element then we would be restricted to solving problems with a maximum domain of 4 MBytes. The total memory within the system is  $64 \times 4 = 256$  MBytes. So, if we were to consider the memory of the multiprocessor system as a whole, then we could contemplate solving problems with domains of up to 256 MBytes in size; a far more attractive proposition. (If the problem domain was even larger than this, then we could also consider the secondary storage devices as part of the combined memory and that should be sufficient for most problems.)

There is a price to pay in treating the combined memory as a single unit. Data management strategies will be necessary to translate between the conceptual single memory unit and the physical distributed implementation. The aims of these strategies will be to keep track of the data items so that an item will always be available at a processing element when required by the task being performed. The distributed nature of the data items will thus be invisible to the application processes performing the computation. However, any delay between the application process requesting an item and this request being satisfied will result in idle time. As we will see, it is the responsibility of data management to avoid this idle time.

## 1.4 Evaluating Parallel Implementations

The chief reason for opting for a parallel implementation should be: *to obtain answers faster*. The time that the parallel implementation takes to compute results is perhaps the most natural way of determining the benefits of the approach that has been taken. If the parallel solution takes *longer* than any sequential implementation then the decision to use parallel processing needs to be re-examined. Other measurements, such as speed-up and efficiency, may also provide useful insight on the maximum scalability of the implementation.

Of course, there are many issues that need to be considered when comparing parallel and sequential implementations of the same problem, for example:

- Was the same processor used in each case?
- If not, what is the price of the sequential machine compared with that of the multiprocessor system?
- Was the algorithm chosen already optimised for sequential use, that is, did the data dependencies present preclude an efficient parallel implementation?

### 1.4.1 Realisation Penalties

If we assume that the same processor was used in both the sequential and parallel implementation, then we should expect, that the time to solve the problem decreases as more processing elements are added. The best we can reasonably hope for is that two processing elements will solve the problem twice as quickly, three processing elements three times faster, and  $n$  processing elements,  $n$  times faster. If  $n$  is sufficiently large then by this process, we should expect our large scale parallel implementation to produce the answer in a tiny fraction of the sequential computation, as shown by the “optimum time” curve in the graph in figure 1.12.

However, in reality we are unlikely to achieve these optimised times as the number of processors is increased. A more realistic scenario is that shown by the curve “actual times” in figure 1.12. This curve shows an initial decrease in time taken to solve the example problem on the parallel system up to a certain number of processing elements. Beyond this point, adding more processors actually leads to an *increase* in computation time.

Failure to achieve the optimum solution time means that the parallel solution has suffered some form of *realisation* penalty. A realisation penalty can arise from two sources:

- an *algorithmic* penalty; and,
- an *implementation* penalty.

The algorithmic penalty stems from the very nature of the algorithm selected for parallel processing. The more inherently sequential the algorithm, the less likely the algorithm will be a good candidate for parallel processing.

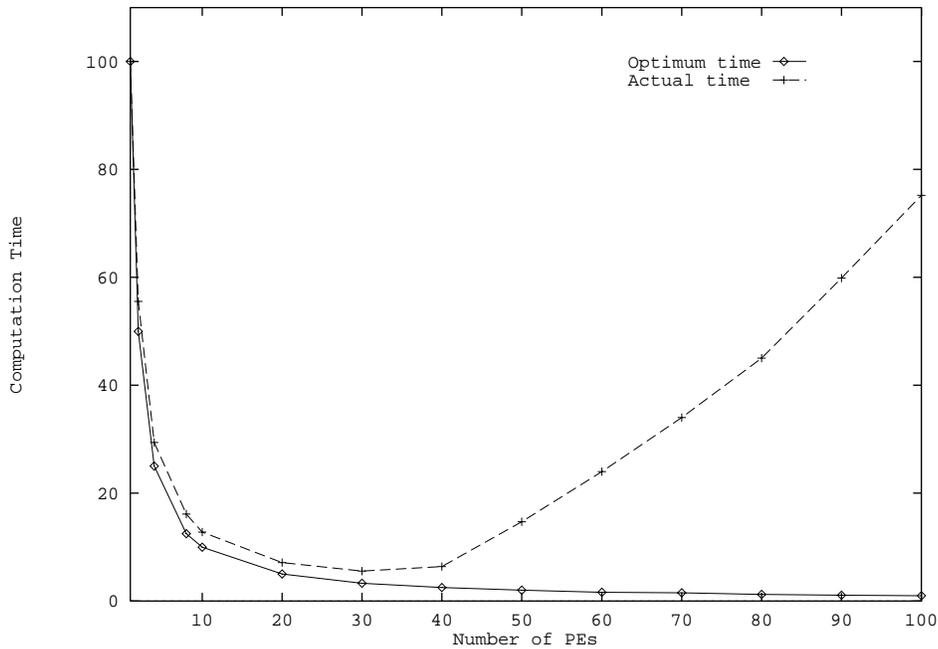


Figure 1.12: Optimum and actual parallel implementation times

**Aside:** It has also been shown, albeit not conclusively, that the more experience the writer of the parallel algorithm has in sequential algorithms, the less parallelism that algorithm is likely to exhibit [13].

This sequential nature of an algorithm and its implicit data dependencies will translate, in the domain decomposition approach, to a requirement to *synchronise* the processing elements at certain points in the algorithm. This can result in processing elements standing idle awaiting messages from other processing elements. A further algorithmic penalty may also come about from the need to reconstruct sequentially the results generated by the individual processors into an overall result for the computation.

Solving the same problem twice as fast on two processing elements implies that those two processing elements must spend 100% of their time on computation. We know that a parallel implementation requires some form of communication. The time a processing element is forced to spend on communication will naturally impinge on the time a processor has for computation. Any time that a processor cannot spend doing useful computation is an implementation penalty. Implementation penalties are thus caused by:

- **the need to communicate**

As mentioned above, in a multiprocessor system, processing elements need to communicate. This communication may not only be that which is necessary for a processing element's own actions, but in some architectures, a processing element may also have to act as a intermediate for other processing elements' communication.

- **idle time**

Idle time is any period of time when an application process is available to perform some useful computation, but is unable to do so because either there is no work locally available, or its current task is suspended awaiting a synchronisation signal, or a data item which has yet to arrive.

It is the job of the local task manager to ensure that an application process is kept supplied with work. The computation to communication ratio within the system will determine how much time a task manager has to fetch a task before the current one is completed. A *load imbalance* is said to exist if some processing elements still have tasks to complete, while the others do not.

While synchronisation points are introduced by the algorithm, the management of data items for a processing element is the job for the local data manager. The domain decomposition approach means that the problem domain is divided amongst the processing elements in some fashion. If an application process requires a data item that is not available locally, then this must be fetched from some other processing element within the system. If the processing element is unable to perform other useful computation while this fetch is being performed, for example by means of multi-threading as discussed in section 3.5.2, then the processing element is said to be idle

- **concurrent communication, data management and task management activity**

Implementing each of a processing element's activities as a separate concurrent process on the same processor, means that the physical processor has to be shared. When another process other than the application process is scheduled then the processing element is not performing useful computation even though its current activity is necessary for the parallel implementation.

The fundamental goal of the system software is to minimise the implementation penalty. While this penalty can never be removed, intelligent communication, data management and task scheduling strategies can avoid idle time and significantly reduce the impact of the need to communicate.

## 1.4.2 Performance Metrics

Solution time provides a simple way of evaluating a parallel implementation. However, if we wish to investigate the relative merits of our implementation then further insight can be gained by additional metrics. A range of metrics will allow us to compare aspects of different implementations and perhaps provide clues as to how overall system performance may be improved.

### Speed-up

A useful measure of any multiprocessor implementation of a problem is *speed-up*. This relates the time taken to solve the problem on a single processor machine to the time taken to solve the same problem using the parallel implementation. We will define the speed-up of a multiprocessor system in terms of the elapsed time that is taken to complete a given problem, as follows:

$$\text{Speed-up} = \frac{\text{elapsed time of a uniprocessor}}{\text{elapsed time of the multiprocessors}} \quad (1.1)$$

The term *linear speed-up* is used when the solution time on an  $n$  processor system is  $n$  times faster than the solution time on the uniprocessor. This linear speed-up is thus equivalent to the optimum time shown in section 1.4.1. The optimum and actual computation times in figure 1.12 are represented as a graph of linear and actual speed-ups in figure 1.13. Note that the actual speed-up curve increases until a certain point and then subsequently decreases. Beyond this point we say that the parallel implementation has suffered a *speed-down*.

The third curve in figure 1.13 represents so-called *super-linear speed-up*. In this example, the implementation on 20 processors has achieved a computation time which is approximately 32 times faster than the uniprocessor solution. It has been argued, see [19], that it is not possible to achieve a speed-up greater than the number of processors used. While in practice it certainly is possible to achieve super-linear speed-up, such implementation may have exploited "unfair" circumstances to obtain such timings. For example, most modern processors have a limited amount of cache memory with an access time significantly faster compared with a standard memory access. Two processors would have double the amount of this cache memory. Given we are investigating a fixed size problem, this means that a larger proportion of the problem domain is in the cache in the parallel implementation than in the sequential implementation. It is not unreasonable, therefore, to imagine a situation where the two processor solution time is more than twice as fast than the uniprocessor time.

Although super-linear speed-up is desirable, in this tutorial we will assume a "fair" comparison between uniprocessor and multiprocessor implementations. The results that are presented in the case studies thus make no attempt to exploit any hardware advantages offered by the increasing number of processors. This will enable the performance improvements offered by the proposed system software extensions to be highlighted without being masked by any variations in underlying hardware. In practice, of course, it would be

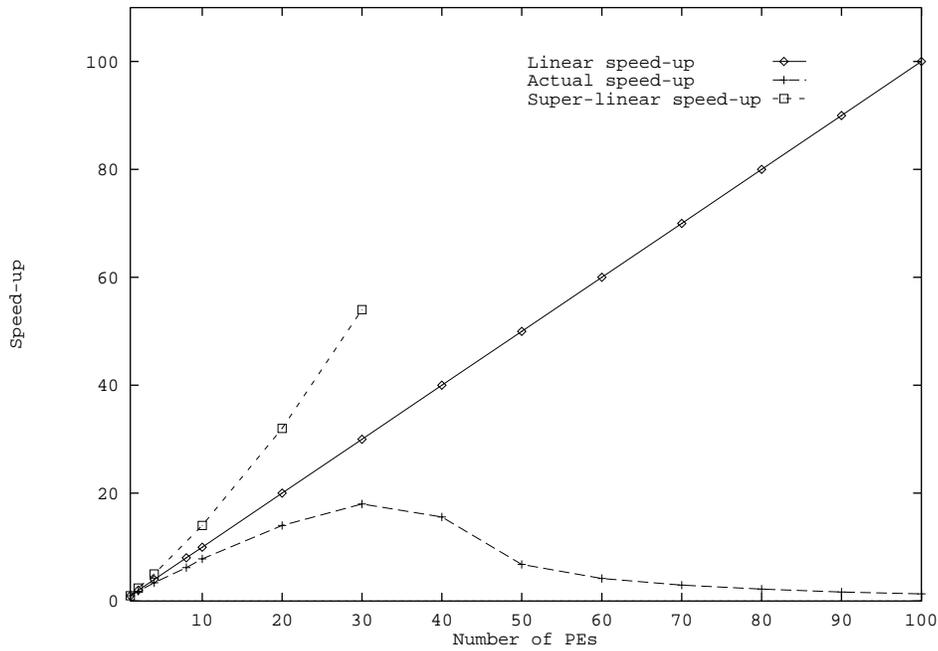


Figure 1.13: Linear and actual speed-ups

foolish to ignore these benefits and readers are encouraged to “squeeze every last ounce of performance” out of their parallel implementation.

Two possibilities exist for determining the “elapsed time of a uniprocessor”. This could be the time obtained when executing:

1. an optimised sequential algorithm on a single processor,  $T_s$ ; or,
2. the parallel implementation on *one* processing element,  $T_1$ .

The time taken to solve the problem on  $n$  processing elements we will term  $T_n$ . The difference between how the two sequential times are obtained is shown in figure 1.14. There are advantages in acquiring both these sequential times. Comparing the parallel to the optimised sequential implementation highlights any algorithmic efficiencies that had to be sacrificed to achieve the parallel version. In addition, none of the parallel implementation penalties are hidden by this comparison and thus the speed-up is not exaggerated. One of these penalties is the time taken simply to supply the data to the processing element and collect the results.

The comparison of the single processing element with the multiple processing element implementation shows how well the problem is “coping” with an increasing number of processing elements. Speed-up calculated as  $\frac{T_1}{T_n}$ , therefore, provides the indication as to the *scalability* of the parallel implementation. Unless otherwise stated, we will use this alternative for speed-up in the case studies in this book as it better emphasizes the performance improvements brought about by the system software we shall be introducing.

As we can see from the curve for “actual speed-up” in figure 1.13, the speed-up obtained for that problem increased to a maximum value and then subsequently decreased as more processing elements were added. In 1967 Amdahl presented what has become known as “Amdahl’s law” [3]. This “law” attempts to give a maximum bound for speed-up from the nature of the algorithm chosen for the parallel implementation. We are given an algorithm in which the proportion of time that needs to be spent on the purely sequential parts is  $s$ , and the proportion of time that might be done in parallel is  $p$ , by definition. The total time for the algorithm on a single processor is  $s + p = 1$  (where the 1 is for algebraic simplicity), and the maximum speed-up that can be achieved on  $n$  processors is:

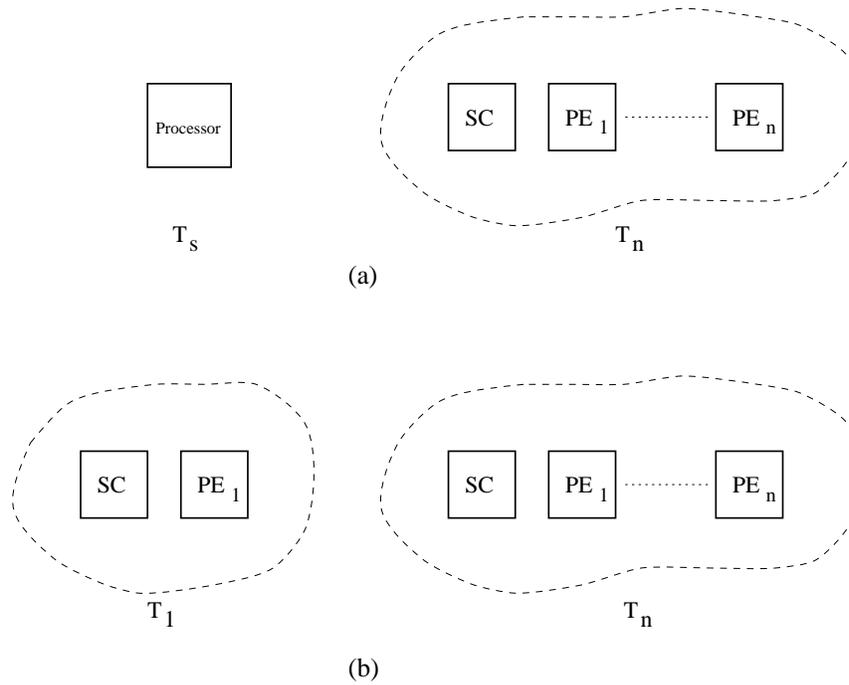


Figure 1.14: Systems used to obtain  $T_n$  and (a)  $T_s$  (b)  $T_1$

$$\begin{aligned}
 \text{maximum speed-up} &= \frac{(s + p)}{s + \frac{p}{n}} \\
 &= \frac{1}{s + \frac{p}{n}}
 \end{aligned} \tag{1.2}$$

Figure 1.15 shows the maximum speed-up predicted by Amdahl's law for a sequential portion of an algorithm requiring 0.1%, 0.5%, 1% and 10% of the total algorithm time, that is  $s = 0.001$ ,  $0.005$ ,  $0.01$  and  $0.1$  respectively. For 1000 processors the maximum speed-up that can be achieved for a sequential portion of only 1% is less than 91. This rather depressing forecast put a serious damper on the possibilities of massive parallel implementations of algorithms and led Gustafson in 1988 to issue a counter claim [26]. Gustafson stated that a problem size is virtually never independent of the number of processors, as it appears in equation (1.2), but rather:

*... in practice, the problem size scales with the number of processors.*

Gustafson thus derives a maximum speed-up of:

$$\begin{aligned}
 \text{maximum speed-up} &= \frac{(s + (p \times n))}{s + p} \\
 &= n + (1 - n) \times s
 \end{aligned} \tag{1.3}$$

This maximum speed-up according to Gustafson is also shown in figure 1.15. As the curve shows, the maximum achievable speed-up is nearly linear when the problem size is increased as more processing elements are added. Despite this optimistic forecast, Gustafson's premise is not applicable in a large number of cases. Most scientists and engineers have a particular problem they want to solve in as short a time as possible. Typically, the application already has a specified size for the problem domain. For example, in parallel radiosity we will be considering the diffuse lighting within a particular environment subdivided into

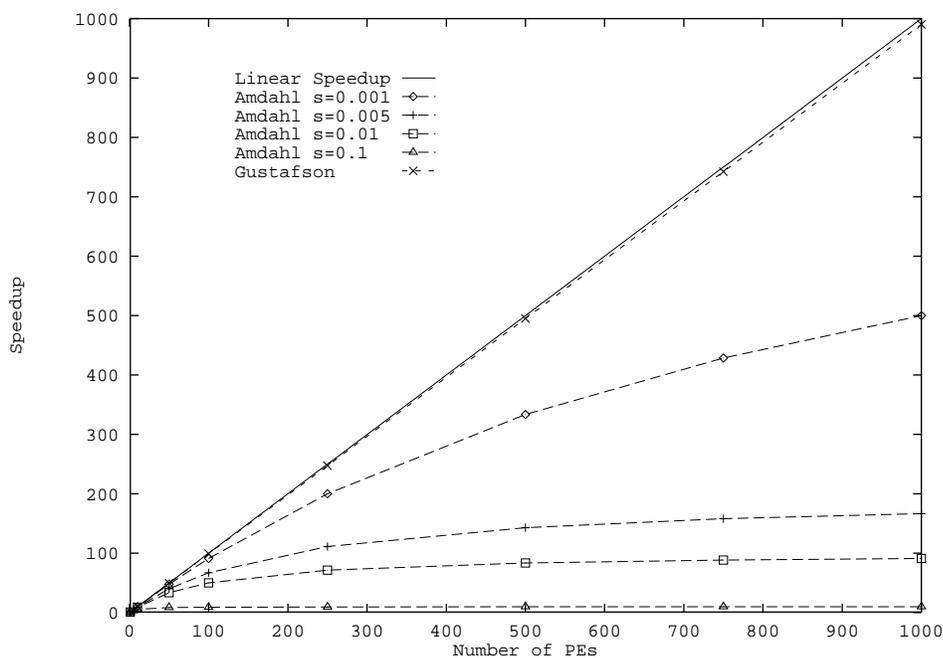


Figure 1.15: Example maximum speed-up from Amdahl and Gustafson's laws

a necessary number of patches. In this example it would be inappropriate for us to follow Gustafson's advice and increase the problem size as more processing elements were added to their parallel implementation, because to do so would mean either:

- the physical size of the three dimensional objects within the environment would have to be increased, which is of course not possible; or,
- the size of the patches used to approximate the surface would have to be reduced, thereby increasing the number of patches and thus the size of the problem domain.

This latter case is also not an option, because the computational method is sensitive to the size of the patches relative to their distances apart. Artificially significantly decreasing the size of the patches may introduce numerical instabilities into the method. Furthermore, artificially increasing the size of the problem domain may improve speed-up, but it *will not* improve the time taken to solve the problem.

For fixed sized problems it appears that we are left with Amdahl's gloomy prediction of the maximum speed-up that is possible for our parallel implementation. However, all is not lost, as Amdahl's assumption that an algorithm can be separated into a component which has to be executed sequentially and part which can be performed in parallel, may not be totally appropriate for the domain decomposition approach. Remember, in this model we are retaining the complete sequential algorithm and exploiting the parallelism that exists in the problem domain. So, in this case, an equivalent to Amdahl's law would imply that the data can be divided into two parts, that which must be dealt with in a strictly sequential manner and that which can be executed in parallel. Any data dependencies will certainly imply some form of sequential ordering when dealing with the data, however, for a large number of problems such data dependencies may not exist. It may also be possible to reduce the effect of dependencies by clever scheduling.

The achievable speed-up for a problem using the domain decomposition approach is, however, bounded by the number of tasks that make up the problem. Solving a problem comprising a maximum of twenty tasks on more than twenty processors makes no sense. In practice, of course, any parallel implementation suffers from realisation penalties which increase as more processing elements are added. The actual speed-up obtained will thus be less than the maximum possible speed-up.

### 1.4.3 Efficiency

A relative efficiency based on the performance of the problem on one processor, can be a useful measure as to what percentage of a processor's time is being spent in useful computation. This, therefore, determines what the system overheads are. The relative efficiency we will measure as:

$$\text{Efficiency} = \frac{\text{speed-up} \times 100}{\text{number of processors}} \quad (1.4)$$

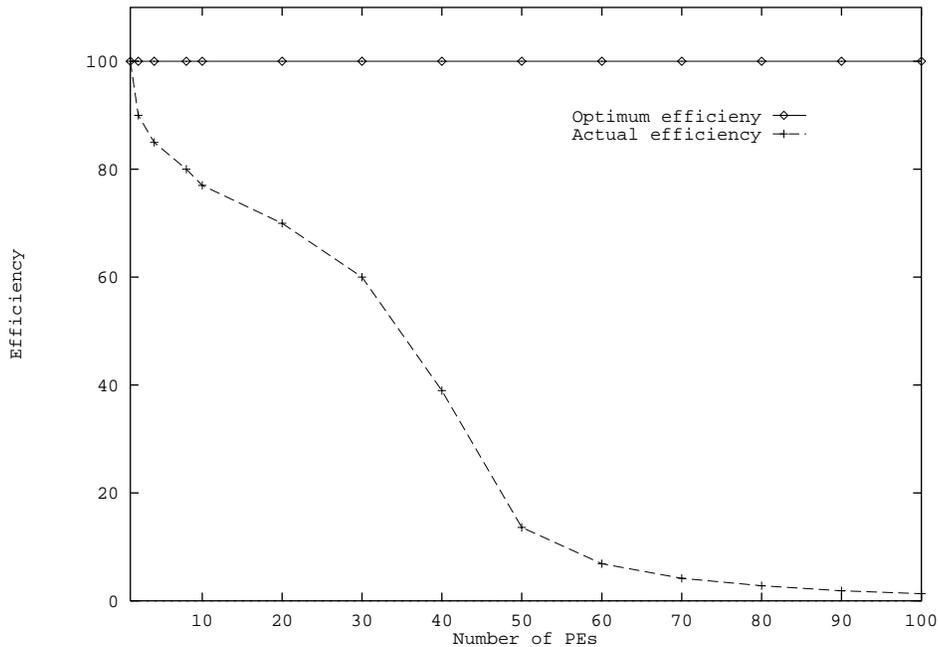


Figure 1.16: Optimum and actual processing element efficiency

Figure 1.16 shows the optimum and actual computation times given in figure 1.12 represented as processing element efficiency. The graph shows that optimum computation time, and therefore linear speed-up, equates to an efficiency of 100% for each processing element. This again shows that to achieve this level of efficiency every processing element must spend 100% of its time performing useful computation. Any implementation penalty would be immediately reflected by a decrease in efficiency. This is clearly shown in the curve for the actual computation times. Here the efficiency of each processing element decreases steadily as more are added until by the time 100 processing elements are incorporated, the realisation penalties are so high that each processing element is only able to devote just over 1% of its time to useful computation.

#### Optimum number of processing elements

Faced with implementing a fixed size problem on a parallel system, it may be useful to know the optimum number of processing elements on which this particular problem should be implemented in order to achieve the best possible performance. We term this optimum number  $n_{opt}$ . We shall judge the *maximum performance* for a particular problem with a fixed problem domain size, as the shortest possible time required to produce the desired results for a certain parallel implementation. This optimum number of processing elements may be derived directly from the “computation time” graph. In figure 1.12 the minimum actual computation time occurred when the problem was implemented on 30 processing elements. As figure 1.17 shows, this optimum number of processing elements is also the point on the horizontal axis in figure 1.13 at which the maximum speed-up was obtained.

The optimum number of processing elements is also the upper bound for the scalability of the problem for that parallel implementation. To improve the scalability of the problem it is necessary to re-examine the

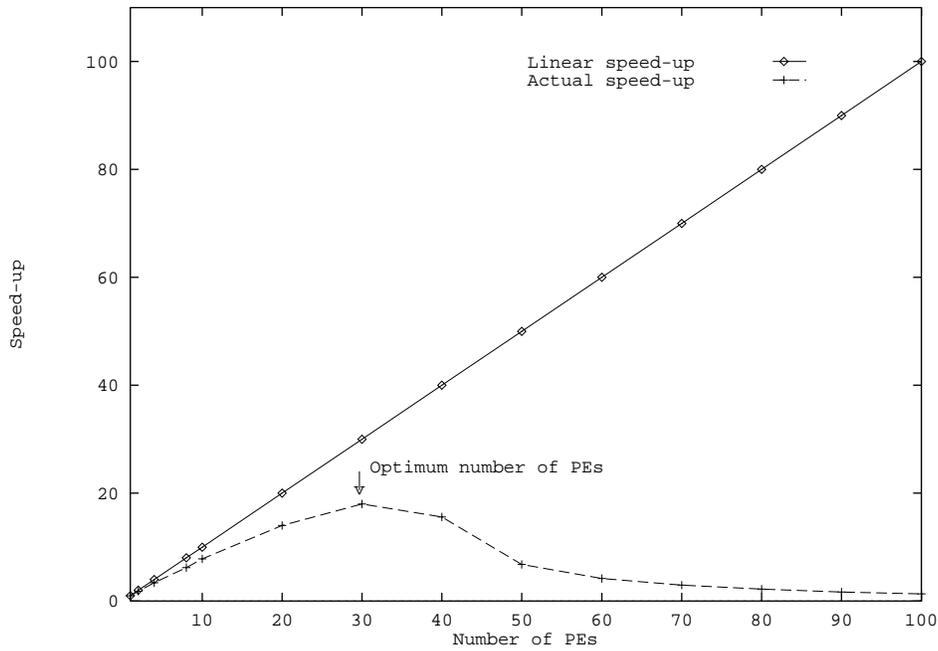


Figure 1.17: Optimum number of processing elements related to speed-up

decisions concerning the algorithm chosen and the make-up of the system software that has been adopted for supporting the parallel implementation. As we will see in the subsequent chapters, the correct choice of system software can have a significant effect on the performance of a parallel implementation.

Figure 1.18 shows the speed-up graphs for different system software decisions for the *same* problem. The goal of a parallel implementation may be restated as:

*“to ensure that the optimum number of processing elements for your problem is greater than the number of processing elements physically available to solve the problem!”*

### Other metrics

Computation time, speed-up and efficiency provide insight into how successful a parallel implementation of a problem has been. As figure 1.18 shows, different implementations of the same algorithm on the same multiprocessor system may produce very different performances. A multitude of other metrics have been proposed over the years as a means of comparing the relative merits of different architectures and to provide a way of assessing their suitability as the chosen multiprocessor machine.

The performance of a computer is frequently measured as the rate of some number of events per second. Within a multi-user environment the elapsed time to solve a problem will comprise the user’s CPU time plus the system’s CPU time. Assuming that the computer’s clock is running at a constant rate, the user’s CPU performance may be measured as:

$$\text{CPU time} = \frac{\text{CPU clock cycles for a program}}{\text{clock rate (eg. 100MHz)}}$$

The average clock cycles per instruction (CPI) may be calculated as:

$$\text{CPI} = \frac{\text{CPU clock cycles for a program}}{\text{Instruction count}}$$

We can also compute the CPU time from the time a program took to run:

$$\text{CPU time} = \frac{\text{seconds}}{\text{program}}$$

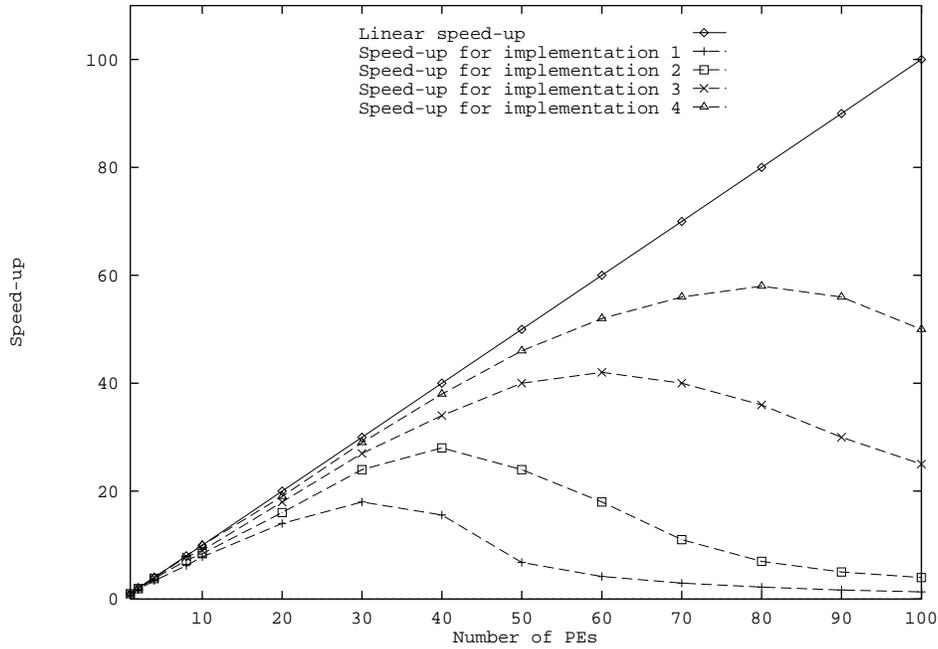


Figure 1.18: Speed-up graphs for different system software for the same problem

$$= \frac{\text{seconds}}{\text{clock cycle}} \times \frac{\text{clock cycles}}{\text{instructions}} \times \frac{\text{instructions}}{\text{program}}$$

Such a performance metric is dependent on:

**Clock rate:** this is determined by the hardware technology and the organisation of the architecture;

**CPI:** a function of the system organisation and the instruction set architecture; and,

**Instruction count:** this is affected by the instruction set architecture and the compiler technology utilised.

One of the most frequently used performance metrics is the *MIPS* rating of a computer, that is how many Million Instructions Per Second the computer is capable of performing:

$$\text{MIPS} = \frac{\text{instruction count}}{\text{execution time} \times 10^6} = \frac{\text{clock rate}}{\text{CPI} \times 10^6}$$

However, the MIPS value is dependent on the instruction set used and thus any comparison between computers with different instruction sets is not valid. The MIPS value may even vary between programs running on the same computer. Furthermore, a program which makes use of hardware floating point routines may take *less time* to complete than a similar program which uses a software floating point implementation, but the first program will have a *lower* MIPS rating than the second [28]. These anomalies have led to MIPS sometimes being referred to as “*Meaningless Indication of Processor Speed*”.

Similar to MIPS is the “Mega-FLOPS” (MFLOPS) rating for computers, where MFLOPS represents Million Floating point Operations per Second:

$$\text{MFLOPS} = \frac{\text{no. of floating point operations in a program}}{\text{execution time} \times 10^6}$$

MFLOPS is not universally applicable, for example a word processor utilising no floating point operations would register no MFLOPS rating. However, the same program executing on different machines should be comparable, because, although the computers may execute a different number of instructions,

they should perform the same number of operations, provided the set of floating point operations is consistent across both architectures. The MFLOPS value will vary for programs running on the same computer which have different mixtures of integer and floating point instructions as well as a different blend of “fast” and “slow” floating point instructions. For example, the *add* instruction often executes in less time than a *divide* instruction.

A MFLOPS rating for a single program can not, therefore, be generalised to provide a single performance metric for a computer. A suite of benchmark programs, such as the LINPACK or Livermore Loops routines, have been developed to allow a more meaningful method of comparison between machines. When examining the relative performance of computers using such benchmarks it is important to discover the *sustained* MFLOPS performance as a more accurate indication of the machines’ potential rather than merely the *peak* MFLOPS rating, a figure that “*can be guaranteed never to be exceeded*”.

Other metrics for comparing computers include:

**Dhrystone:** A CPU intensive benchmark used to measure the integer performance especially as it pertains to system programming.

**Whetstone:** A synthetic benchmark without any vectorisable code for evaluating floating point performance.

**TPS:** Transactions Per Second measure for applications, such as airline reservation systems, which require on-line database transactions.

**KLIPS:** Kilo Logic Inferences Per Second is used to measure the relative inference performance of artificial intelligence machines

Tables showing the comparison of the results of these metrics for a number of architectures can be found in several books, for example [32, 37].

Cost is seldom an issue that can be ignored when purchasing a high performance computer. The desirability of a particular computer or even the number of processors within a system may be offset by the extraordinarily high costs associated with many high performance architectures. This prompted an early “law” by Grosch that the speed of a computer is proportional to its cost [25, 24]. Fortunately, although this is no longer completely true, multiprocessor machines are nevertheless typically more expensive than their general purpose counterparts. The parallel computer eventually purchased should provide acceptable computation times for an affordable price, that is maximise: “*the bangs per buck*” (performance per unit price).

# Chapter 2

## Task Scheduling

*The efficient solution of a problem on a parallel system requires the computational performance of the processing elements to be fully utilised. Any processing element that is not busy performing useful computations is degrading overall system performance. Task scheduling strategies may be used to minimise these potential performance limitations.*

### 2.1 Problem Decomposition

A problem may be solved on a parallel system by either exploiting the parallelism inherent in the algorithm, known as *algorithmic decomposition*, or by making use of the fact that the algorithm can be applied to different parts of the problem domain in parallel, which is termed *domain decomposition*. These two decomposition methods can be further categorised as shown in figure 2.1.

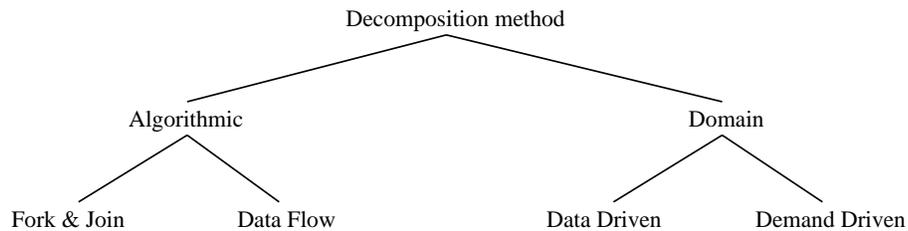


Figure 2.1: Methods of decomposing a problem to exploit parallelism

Over the years, an abundance of algorithms have been developed to solve a multitude of problems on sequential machines. A great deal of time and effort has been invested in the production of these sequential algorithms. Users are thus loathed to undertake the development of novel parallel algorithms, and yet still demand the performance that multiprocessor machines have to offer.

Algorithmic decomposition approaches to this dilemma have led to the development of compilers, such as those for High Performance Fortran, which attempt to parallelise automatically these existing algorithms. Not only do these compilers have to identify the parallelism hidden in the algorithm, but they also need to decide upon an effective strategy to place the identified segments of code within the multiprocessor system so that they can interact efficiently. This has proved to be an extremely hard goal to accomplish.

The domain decomposition approach, on the other hand, requires little or no modification to the existing sequential algorithm. There is thus no need for sophisticated compiler technology to analyse the algorithm. However, there will be a need for a parallel framework in the form of system software to support the division of the problem domain amongst the parallel processors.

### 2.1.1 Algorithmic decomposition

In algorithmic decomposition the algorithm itself is analysed to identify which of its features are capable of being executed in parallel. The finest granularity of parallelism is achievable at the operation level. Known as *dataflow*, at this level of parallelism the data “flows” between individual operands which are being executed in parallel [1]. An advantage of this type of decomposition is that little data space is required per processor [29], however, the communication overheads may be very large due to the very poor computation to communication ratio.

Fork & join parallelism, on the other hand, allocates portions of the algorithm to separate processors as the computation proceeds. These portions are typically several statements or complete procedures. The difference between the two algorithmic forms of decomposition is shown for a simple case in figure 2.2.

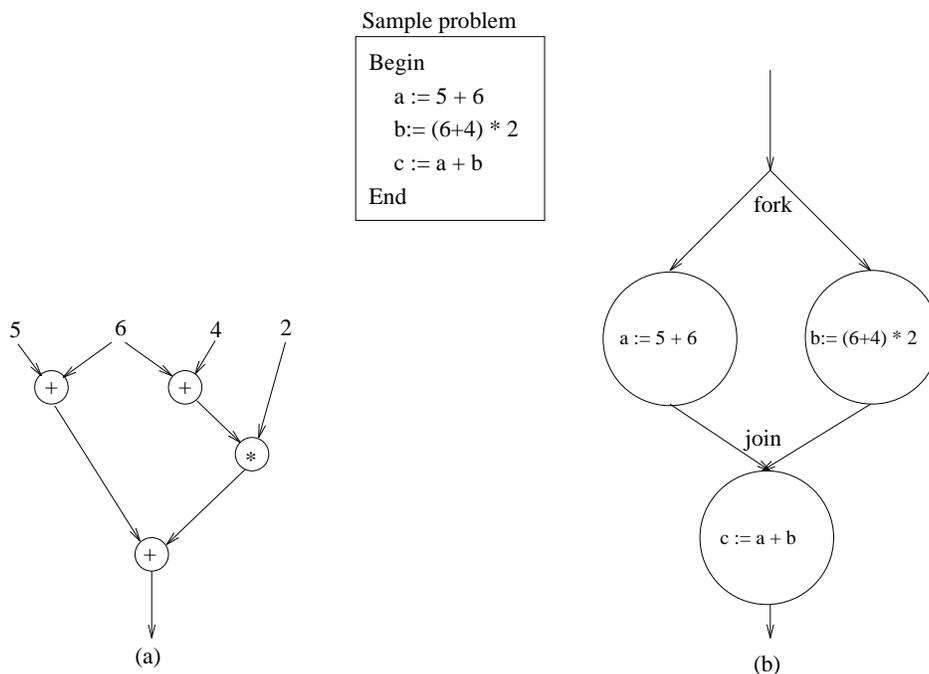


Figure 2.2: Algorithmic decomposition: (a) dataflow (b) fork & join

### 2.1.2 Domain decomposition

Instead of determining the parallelism inherent in the algorithm, domain decomposition examines the problem domain to ascertain the parallelism that may be exploited by solving the algorithm on distinct data items in parallel. Each parallel processor in this approach will, therefore, have a complete copy of the algorithm and it is the problem domain that is divided amongst the processors. Domain decomposition can be accomplished using either a data driven or demand driven approach.

As we shall see, given this framework, the domain decomposition approach is applicable to a wide range of problems. Adoption of this approach to solve a particular problem in parallel, consists of two steps:

1. **Choosing the appropriate sequential algorithm.**

Many algorithms have been honed over a number of years to a high level of perfection for implementation on sequential machines. The data dependencies that these highly sequential algorithms exhibit may substantially inhibit their use in a parallel system. In this case alternative sequential algorithms which are more suitable to the domain decomposition approach will need to be considered.

2. **Analysis of the problem in order to extract the criteria necessary to determine the optimum system software.**

The system software provides the framework in which the sequential algorithm can execute. This system software takes care of ensuring each processor is kept busy, the data is correctly managed, and any communication within the parallel system is performed rapidly. To provide maximum efficiency, the system software needs to be tailored to the requirements of the problem. There is thus no *general purpose* parallel solution using the domain decomposition approach, but, as we shall see, a straightforward analysis of any problem's parallel requirements, will determine the correct construction of the system software and lead to an efficient parallel implementation.

Before commencing the detailed description of how we intend to tackle the solution of realistic rendering problems in parallel, it might be useful to clarify some of the terminology we shall be using.

### 2.1.3 Abstract definition of a task

The domain decomposition model solves a single problem in parallel by having multiple processors apply the same sequential algorithm to different data items from the problem domain in parallel. The lowest unit of computation within the parallel system is thus the application of the algorithm to one data item within the problem domain.

The data required to solve this unit of computation consists of two parts:

1. the *principal data items* (or PDIs) on which the algorithm is to be applied; and
2. *additional data items* (or ADIs) that may be needed to complete this computation on the PDIs.

For example, in ray tracing, we are computing the value at each pixel of our image plane. Thus these pixels would form our PDIs, while all the data describing the scene would constitute the ADIs. The problem domain is thus the pixels *plus* the scene description.

The application of the algorithm to a specified principal data item may be regarded as performing a single *task*. The task forms the elemental unit of computation within the parallel implementation. This is shown diagrammatically in figure 2.3.

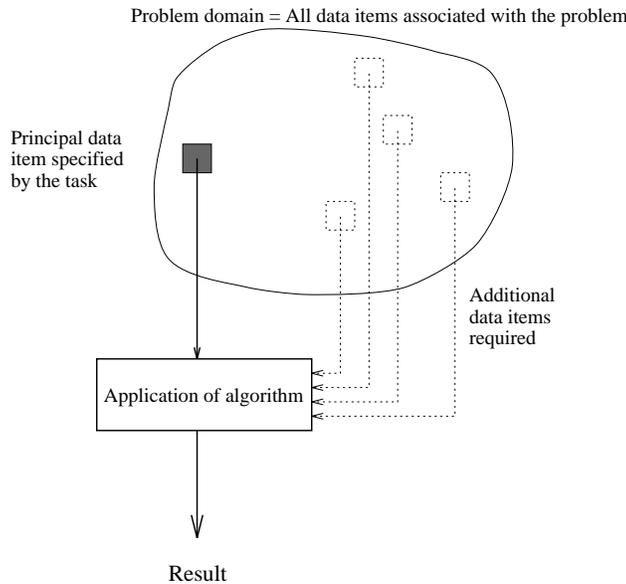


Figure 2.3: A task: the processing of a *principal* data item

### 2.1.4 System architecture

This tutorial is concentrating on implementing realistic rendering techniques on distributed memory systems (either a dedicated parallel machine or a distributed system of workstations). These processors may be

connected together in some manner to form a *configuration*. A *process* is a segment of code that runs concurrently with other processes on a single processor. Several processes will be needed at each processor to implement the desired application and provide the necessary system software support. A *processing element* consists of a single processor together with these application and system processes and is thus the building block of the *multiprocessor system*. (We shall sometimes use the abbreviation *PE* for processing element in the figures and code segments.) When discussing configurations of processing elements, we shall use the term *links* to mean the communication paths between processes.

### Structure of the system controller

To provide a useful parallel processing platform, a multiprocessor system must have access to input/output facilities. Most systems achieve this by designating at least one processing element as the *system controller* (SC) with the responsibilities of providing this input/output interface, as shown in figure 2.4. If the need for input/output facilities becomes a serious bottleneck then more than one system controller may be required. Other processing elements perform the actual computation associated with the problem.

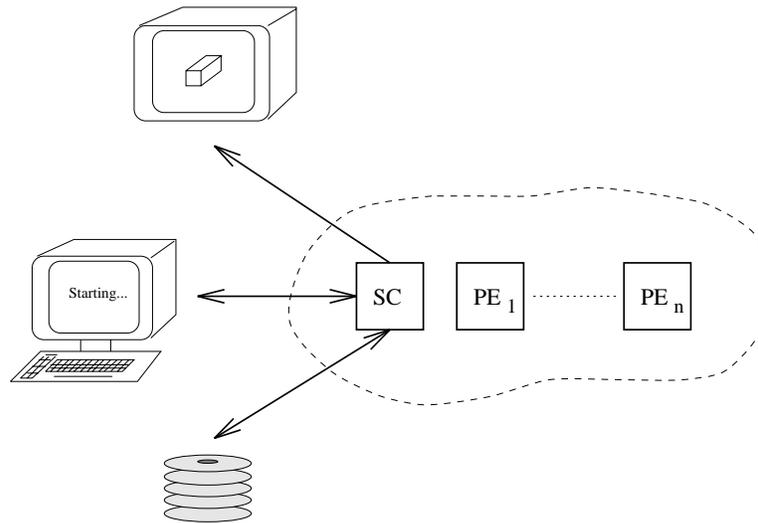


Figure 2.4: The system controller as part of a parallel system

In addition to providing the input/output facilities, the system controller may also be used to collect and collate results computed by the processing elements. In this case the system controller is in the useful position of being able to determine when the computation is complete and gracefully terminate the concurrent processes at every processing element.

## 2.2 Computational Models

The computational model chosen to solve a particular problem determines the manner in which work is distributed across the processors of the multiprocessor system. In our quest for an efficient parallel implementation we must maximise the proportion of time the processors spend performing necessary computation. Any imbalance may result in processors standing idle while others struggle to complete their allocated work, thus limiting potential performance. Load balancing techniques aim to provide an even division of computational effort to all processors.

The solution of a problem using the domain decomposition model involves each processing element applying the specified algorithm to a set of principal data items. The computational model ensures that every principal data item is acted upon and determines how the tasks are allocated amongst the processing elements. A choice of computation model exists for each problem. To achieve maximum system performance, the model chosen must see that the total work load is distributed evenly amongst the processing

elements. This balances the overheads associated with communicating principal data items to processing elements with the need to avoid processing element idle time. A simplified ray tracing example illustrate the differences between the computational models.

A sequential solution to this problem may be achieved by dividing the image plane into twenty-four distinct regions, with each region constituting a single principal data item, as shown in figure 2.5, and then applying the ray tracing algorithm at each of these regions in turn. There are thus twenty-four tasks to be performed for this problem where each task is to compute the pixel value at one area of the image plane. To understand the computational models, it is not necessary to know the details of the algorithm suffice to say that each principal data item represents an area of the image plane on which the algorithm can be applied to determine the value for that position. We will assume that no additional data items are required to complete any task.

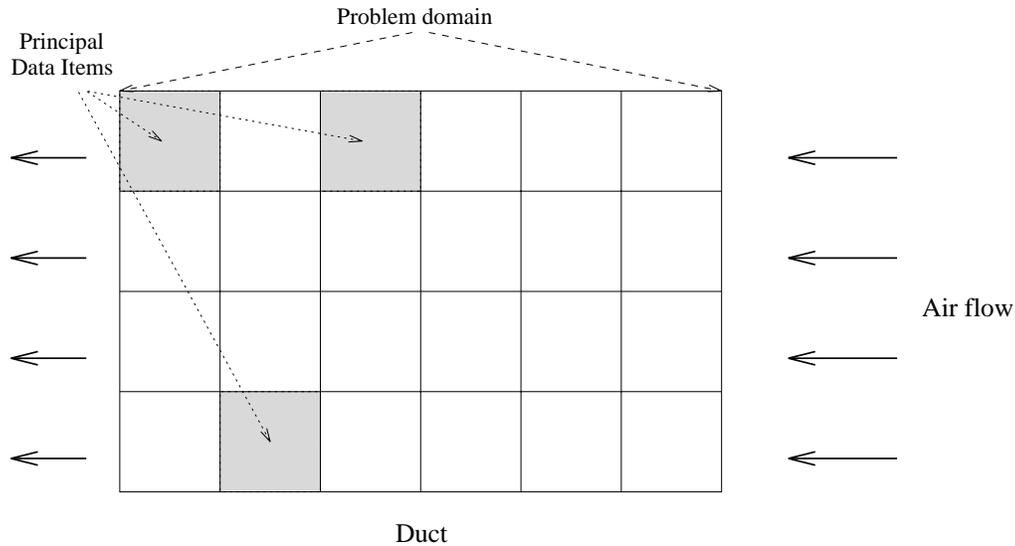


Figure 2.5: Principal data items for calculating the pixels in the image plane

### 2.2.1 Data driven model

The data driven model allocates all the principal data items to specific processing elements before computation commences. Each processing element thus knows *a priori* the principal data items to which they are required to apply the algorithm. Providing there is sufficient memory to hold the allocated set at each processing element, then, apart from the initial distribution, there is no further communication of principal data items. If there is insufficient local memory, then the extra items must be *fetched* as soon as memory space allows. This fetching of remote data items will be discussed further when data management is examined in Chapter 3.

#### Balanced data driven

In balanced data driven systems (also known as geometric decompositions), an equal number of principal data items is allocated to each processing element. This portion is determined simply by dividing the total number of principal data items by the number of processing elements:

$$\text{portion at each PE} = \frac{\text{number of principal data items}}{\text{number of PEs}}$$

If the number of principal data items is not an exact multiple of the number of processing elements, then

$$(\text{number of principal data items}) \text{ MOD } (\text{number of PEs})$$

will each have one extra principal data item, and thus perform one extra task. The required start task and the number of tasks is communicated by the system controller to each processing element and these can then apply the required algorithm to their allotted principal data items. This is similar to the way in which problems are solved on arrays of SIMD processors.

In this example, consider the simple ray tracing calculation for an empty scene. The principal data items (the pixels) may be allocated equally to three processing elements, labelled  $PE_1$ ,  $PE_2$  and  $PE_3$ , as shown in figure 2.6. In this case, each processing element is allotted eight principal data items.

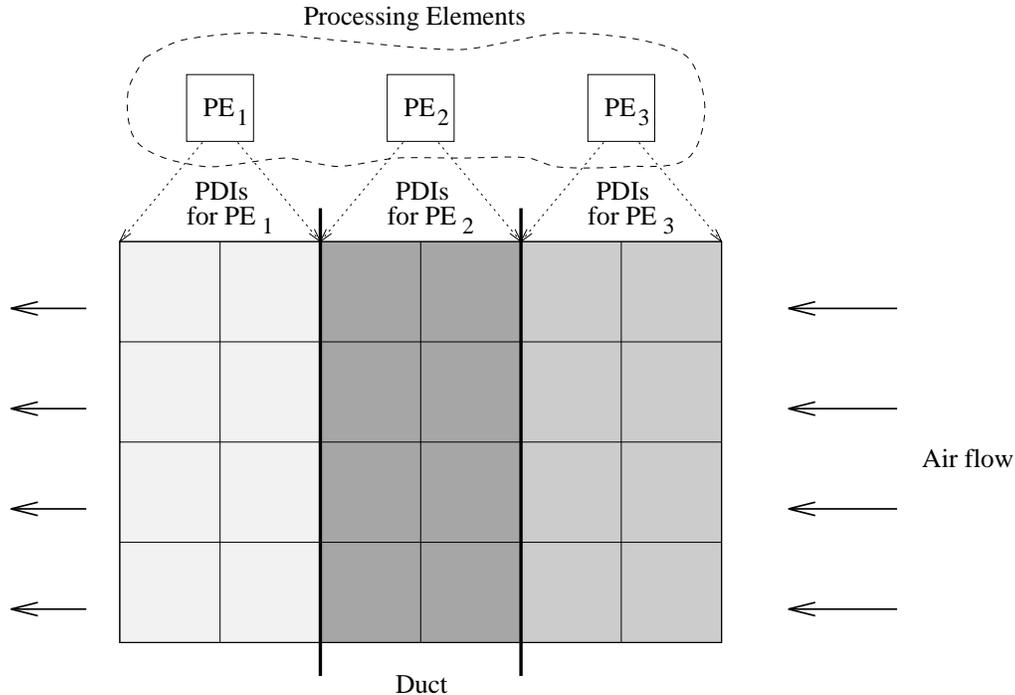


Figure 2.6: Equal allocation of data items to processing elements

As no further principal data item allocation takes place after the initial distribution, a balanced work load is only achieved for the balanced data driven computational model if the computational effort associated with each portion of principal data items is *identical*. If not, some processing elements will have finished their portions while others still have work to do. With the balanced data driven model the division of principal data items amongst processing elements is geometric in nature, that is each processing element simply may be allocated an equal number of principal data items irrespective of their position within the problem domain. Thus, to ensure a balanced work load, this model should only be used if the computational effort associated with each principal data item is the same, and preferably where the number of principal data items is an exact multiple of the number of processing elements. This implies *a priori* knowledge, but given this, the balanced data driven approach is the simplest of the computational models to implement.

Using figure 2.6, if the computation of each pixel 1 *time unit* to complete, then the sequential solution of this problem would take 24 *time units*. The parallel implementation of this problem using the three processing elements each allocated eight tasks should take approximately 8 *time units*, a third of the time required by the sequential implementation. Note, however, that the parallel solution will not be exactly one third of the sequential time as this would ignore the time required to communicate the portions from the system controller to the processing elements. This also ignores time required to receive the results back from the processing elements and for the system controller to collate the solution. A balanced data driven version of this problem on the three processing elements would more accurately take:

$$\text{Solution time} = \text{initial distribution} + \lceil \frac{24}{3} \rceil + \text{result collation}$$

Assuming low communication times, this model gives the solution in approximately one third of the time

of the sequential solution, close to the maximum possible linear speed-up. Solution of the same problem on five processing elements would give:

$$\text{Solution time} = \text{initial distribution} + \lceil \frac{24}{5} \rceil + \text{result collation}$$

This will be solved in even longer than the expected 4.8 *time units* as, in this case, one processing element is allocated 4 principal data items while the other four have to be apportioned 5. As computation draws to a close, one processing element will be idle while the four others complete their extra work. The solution time will thus be slightly more than 5 *time units*.

### Unbalanced data driven

Differences in the computational effort associated with the principal data items will increase the probability of substantial processing element idle time if the simplistic balanced data driven approach is adopted. If the individual computation efforts differ, and are known *a priori*, then this can be exploited to achieve optimum load balancing.

The unbalanced data driven computational model allocates principal data items to processing elements based on their computational requirements. Rather than simply apportioning an equal number of tasks to each processing element, the principal data items are allocated to ensure that each processing element will complete its portion at approximately *the same time*.

For example, the complexity introduced into the ray tracing calculations by placing object into the scene, as shown in figure 2.7, will cause an increased computational effort required to solve the portions allocated to  $PE_1$  and  $PE_2$  in the balanced data driven model. This will result in these two processing elements still being busy with their computations long after the other processing element,  $PE_3$ , has completed its less computationally complex portion.

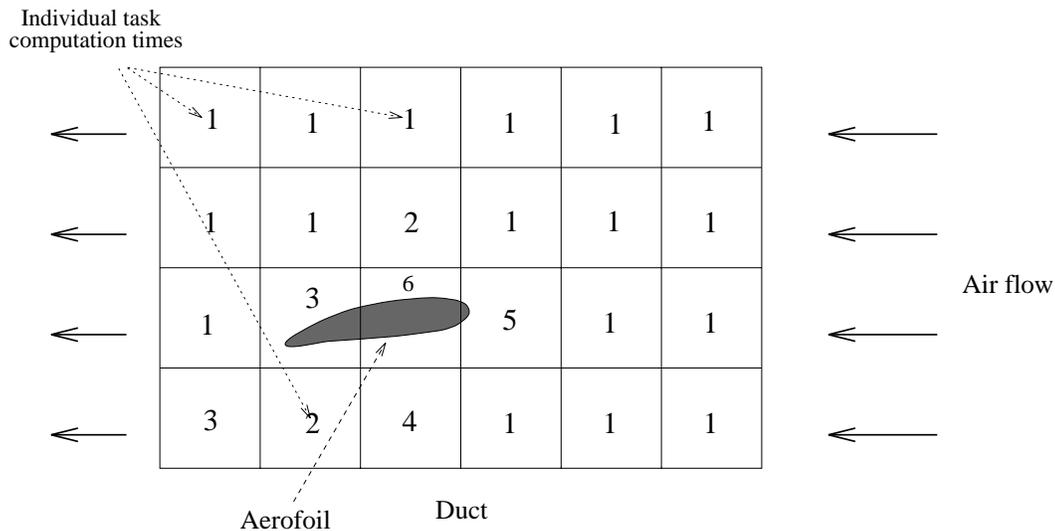


Figure 2.7: Unequal computational effort due to presence of objects in the scene

Should *a priori* knowledge be available regarding the computational effort associated with each principal data item then they may be allocated *unequally* amongst the processing elements, as shown in figure 2.8. The computational effort now required to process each of these unequal portions will be approximately the same, minimising any processing element idle time.

The sequential time required to solve the ray tracing with objects in the scene is now 42 *time units*. To balance the work load amongst the three processing elements, each processing element should compute for 14 *time units*. Allocation of the portions to each processing element in the unbalanced data driven model involves a preprocessing step to determine precisely the best way to subdivide the principal data items. The optimum compute time for each processing element can be obtained by simply dividing the

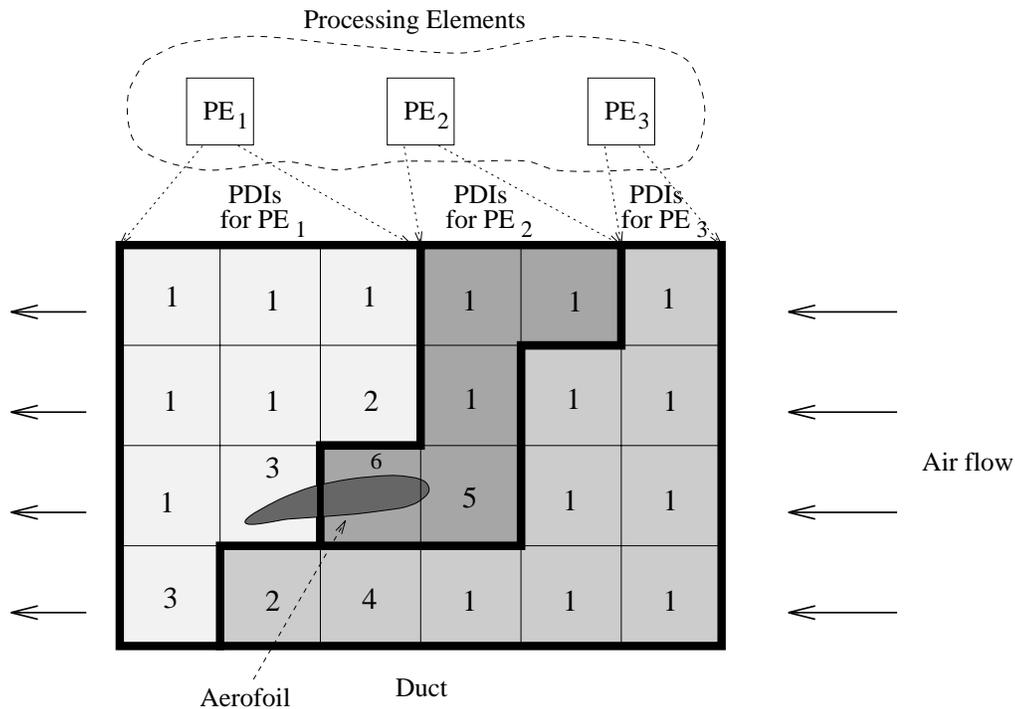


Figure 2.8: Unequal allocation of data items to processing elements to assist with load balancing

total computation time by the number of processing elements. If possible, no processing element should be allocated principal data items whose combined computation time exceeds this optimum amount. Sorting the principal data items in descending computation times can facilitate the subdivision.

The total solution time for a problem using the unbalanced data driven model is thus:

$$\begin{aligned} \text{Solution time} = & \text{preprocessing} + \text{distribution} \\ & + \text{longest portion time} + \text{result collation} \end{aligned}$$

So comparing the naive balanced distribution from section 2.2.1

$$\begin{aligned} \text{Balanced solution time} = & \text{distribution} + 21 + \\ & \text{result collation} \end{aligned}$$

$$\begin{aligned} \text{Unbalanced solution time} = & \\ & \text{preprocessing} + \text{distribution} + 14 + \text{result collation} \end{aligned}$$

The preprocessing stage is a simple sort requiring far less time than the ray tracing calculations. Thus, in this example, the unbalanced data driven model would be significantly faster than the balanced model due to the large variations in task computational complexity.

The necessity for the preprocessing stage means that this model will take more time to use than the balanced data driven approach should the tasks have the same computation requirement. However, if there are variations in computational complexity and they are known, then the unbalanced data driven model is the most efficient way of implementing the problem in parallel.

### 2.2.2 Demand driven model

The data driven computational models are dependent on the computational requirements of the principal data items being known, or at least being predictable, before actual computation starts. Only with this knowledge can these data items be allocated in the correct manner to ensure an even load balance. Should the computational effort of the principal data items be unknown or unpredictable, then serious load balancing problems can occur if the data driven models are used. In this situation the demand driven computational model should be adopted to allocate work to processing elements evenly and thus optimise system performance.

In the demand driven computational model, work is allocated to processing elements *dynamically* as they become idle, with processing elements no longer bound to any particular portion of the principal data items. Having produced the result from one principal data item, the processing elements demand the next principal data item from some work supplier process. This is shown diagrammatically in figure 2.9 for the simple ray tracing calculation.

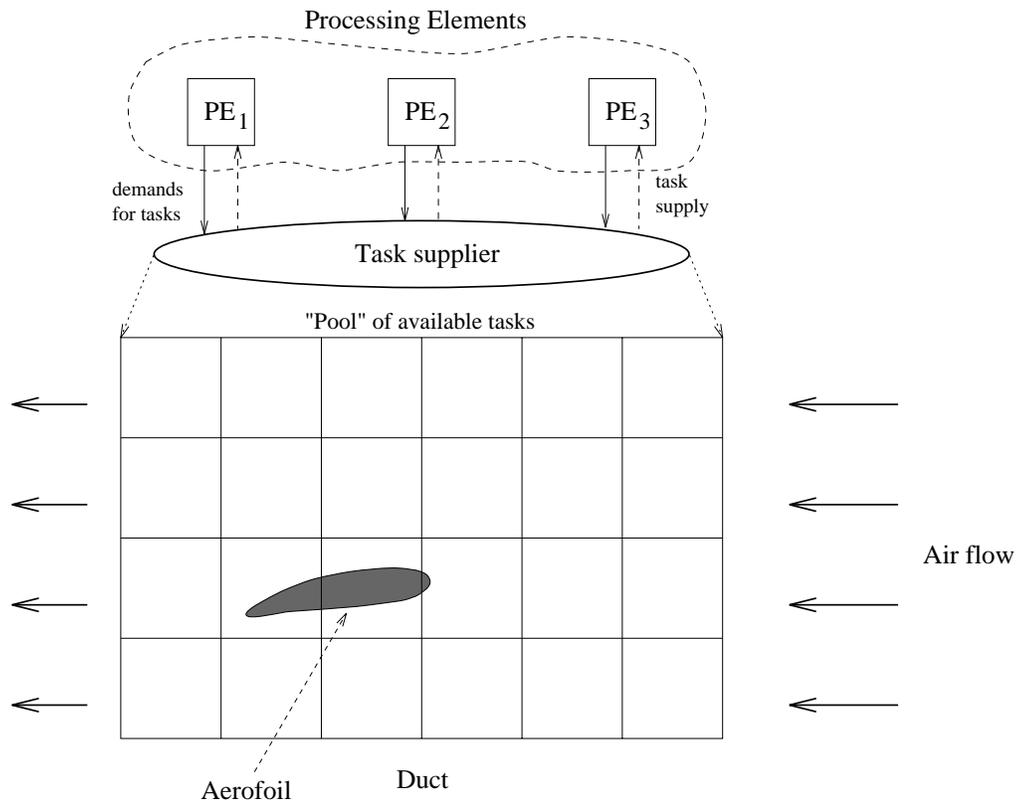


Figure 2.9: A demand driven model for a simple ray tracing calculation

Unlike the data driven models, there is no initial communication of work to the processing elements, however, there is now the need to send requests for individual principal data items to the supplier and for the supplier to communicate with the processing elements in order to satisfy these requests. To avoid unnecessary communication it may be possible to combine the return of the results from one computation with the request for the next principal data item.

The optimum time for solving a problem using this simple demand driven model is thus:

$$\text{Solution time} = \frac{2 \times \text{total communication time} + \text{total computation time for all PDIs}}{\text{number of PEs}}$$

This optimum computation time,  $\frac{\text{total computation time for all PDIs}}{\text{number of PEs}}$ , will only be possible if the work can be allocated so that all processing elements complete the last of their tasks at exactly the same time. If this is not so then some processing elements will still be busy with their final task while the others have completed. It may also be possible to reduce the communication overheads of the demand driven model by overlapping the communication with the computation in some manner. This possibility will be discussed later in section 2.3.

On receipt of a request, if there is still work to be done, the work supplier responds with the next available task for processing. If there are no more tasks which need to be computed then the work supplier may safely ignore the request. The problem will be solved when all principal data items have been requested and all the results of the computations on these items have been returned and collated. The dynamic allocation of work by the demand driven model will ensure that while some processing elements are busy with more computationally demanding principal data items, other processing elements are available to compute the less complex parts of the problem.

Using the computational times for the presence of objects in the scene as shown in figure 2.8, figure 2.10 shows how the principal data items may be allocated by the task supplier to the processing elements using a simple serial allocation scheme. Note that the processing elements do not complete the same number of tasks. So, for example, while processing elements 2 and 3 are busy completing the computationally complex work associated with principal data items 15 and 16, processing element 1 can compute the less computationally taxing tasks of principal data items 17 and 18.

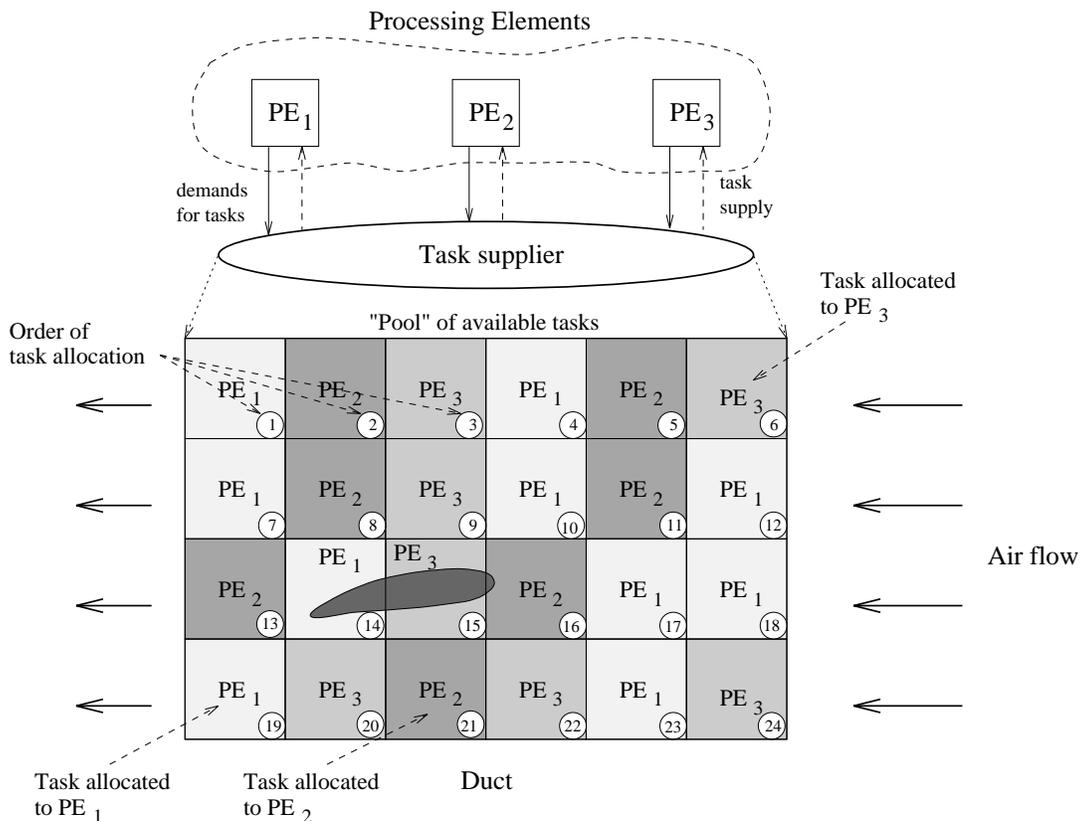


Figure 2.10: Allocation of principal data items using a demand driven model

The demand driven computational model facilitates dynamic load balancing when there is no prior knowledge as to the complexity of the different parts of the problem domain. Optimum load balancing is still dependent on all the processing elements completing the last of the work at the same time. An unbalanced solution may still result if a processing element is allocated a complex part of the domain towards the end of the solution. This processing element may then still be busy well after the other processing elements

have completed computation on the remainder of the principal data items and are now idle as there is no further work to do. To reduce the likelihood of this situation it is important that the computationally complex portions of the domain, the so called *hot spots*, are allocated to processing elements early on in the solution process. Although there is no *a priori* knowledge as to the exact computational effort associated with any principal data item (if there were, an unbalanced data driven approach would have been adopted), nevertheless, any insight as to possible hot spot areas should be exploited. The task supplier would thus assign principal data items from these areas first.

In the ray tracing example, while the exact computational requirement associated with the principal data items in proximity of the objects in the scene may be unknown, it is highly likely that the solution of the principal items in that area will more complex than those elsewhere. In this problem, these principal data items should be allocated first.

If no insight is possible then a simple serial allocation, as shown in figure 2.10, or spiral allocation, as shown in figure 2.11 or even a random allocation of principal data items will have to suffice. While a random allocation offers perhaps a higher probability of avoiding late allocation of principal data items from hot spots, additional effort is required when choosing the next principal data item to allocate to ensure that no principal data item is allocated more than once.

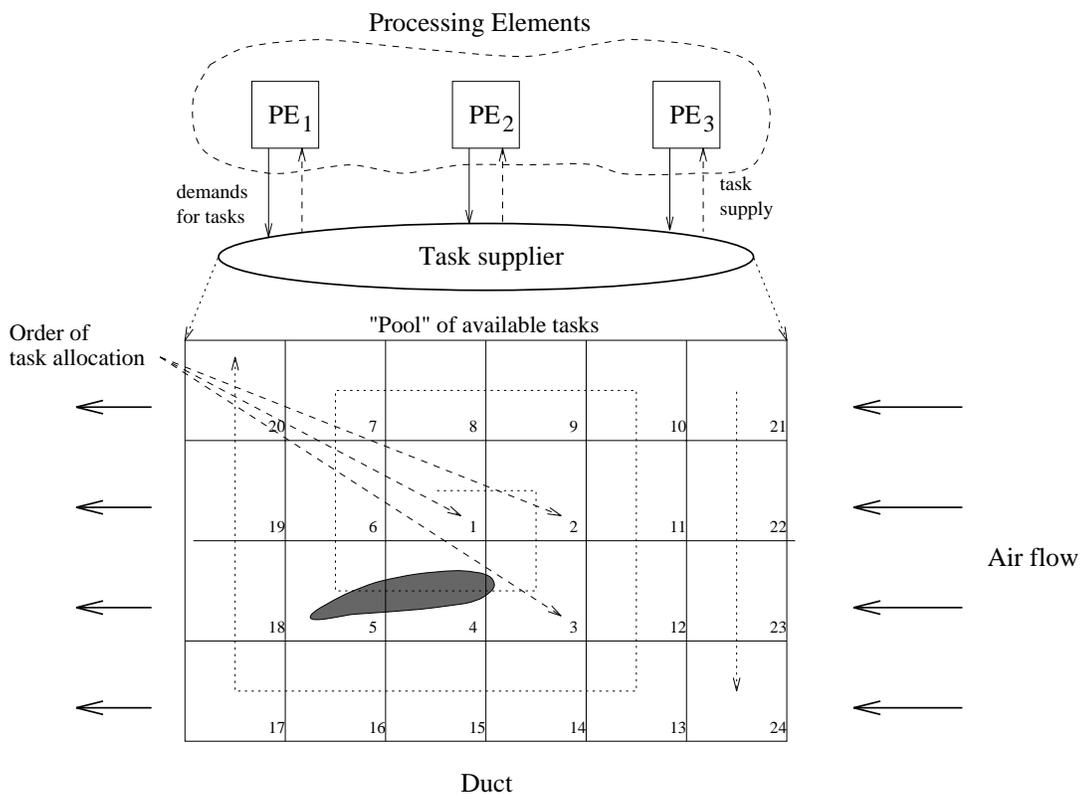


Figure 2.11: Allocation of principal data items in a spiral manner

As with all aspects of parallel processing, extra levels of sophistication can be added in order to exploit any information that becomes available as the parallel solution proceeds. Identifying possible hot spots in the problem domain may be possible from the computation time associated with each principal data item as these become known. If this time is returned along with the result for that principal data item, the work supplier can build a dynamic profile of the computational requirements associated with areas of the domain. This information can be used to adapt the allocation scheme to send principal data items from the possible hot spot regions. There is, of course, a trade off here between the possible benefits to load balancing in the early allocation of principal data items from hot spots, and the overhead that is introduced by the need to:

- time each computation at the processing elements;

- return this time to the work supplier;
- develop the time profile at the work supplier; and,
- adapt the allocation strategy to take into account this profile.

The benefits gained by such an adaptive scheme are difficult to predict as they are dependent on the problem being considered and the efficiency of the scheme implementation. The advice in these matters is always: “*implement a simple scheme initially and then add extra sophistication should resultant low system performance justify it.*”

### 2.2.3 Hybrid computational model

For most problems, the correct choice of computational model will either be one of the data driven strategies or the demand driven approach. However, for a number of problems, a hybrid computational model, exhibiting properties of both data and demand driven models, can be adopted to achieve improved efficiency. The class of problem that can benefit from the hybrid model is one in which an initial set of principal data items of known computational complexity may spawn an unknown quantity of further work.

In this case, the total number of principal data items required to solve the problem is *unknown* at the start of the computation, however, there are at least a known number of principal data items that must be processed first. If the computational complexity associated with these initial principal data items is unknown then a demand driven model will suffice for the whole problem, but if the computational complexity is known then one of the data driven models, with their lower communication overheads, should at least be used for these initial principal data items. Use of the hybrid model thus requires the computational model to be switched from data driven to demand driven mode as required.

## 2.3 Task Management

Task management encompasses the following functions:

- the definition of a task;
- controlling the allocation of tasks;
- distribution of the tasks to the processing elements; and,
- collation of the results, especially in the case of a problem with multiple stages.

### 2.3.1 Task definition and granularity

An *atomic element* may be thought of as a problem’s lowest computational element within the sequential algorithm adopted to solve the problem. As introduced in section 2.1.2, in the domain decomposition model a single task is the application of this sequential algorithm to a principal data item to produce a result for the sub-parts of the problem domain. The task is thus the smallest element of computation for the problem within the parallel system. The *task granularity* (or *grain size*) of a problem is the number of atomic elements, which are included in one task. Generally, the task granularity remains constant for all tasks, but in some cases it may be desirable to alter dynamically this granularity as the computation proceeds. A task which includes only one atomic element is said to have the *finest granularity*, while a task which contains many is *coarser grained*, or has a *coarser granularity*. The actual definition of what constitutes a principal data item is determined by the granularity of the tasks.

A parallel system solves a problem by its constituent processing elements executing tasks in parallel. A *task packet* is used to inform a processing element which task, or tasks, to perform. This task packet may simply indicate which tasks require processing by that processing element, thus forming the lowest level of distributed work. The packet may include additional information, such as additional data items, which the tasks require in order to be completed.

To illustrate the differences in this terminology, consider again the simple ray tracing problem. The atomic element of a sequential solution of this problem could be to perform a single ray-object intersection

test. The principal data item is the pixel being computed and the additional data item required will be object being considered. A sequential solution of this problem would be for a single processing element to consider each ray-object intersection in turn. The help of several processing elements could substantially improve the time taken to perform the ray tracing.

The finest task granularity for the parallel implementation of this problem is for each task to complete one atomic element, that is perform one ray-object intersection. For practical considerations, it is perhaps more appropriate that each task should instead be to trace the complete path of a single ray. The granularity of each task is now the number of ray-object intersections required to trace this single ray and each pixel is a principal data item. A sensible task packet to distribute the work to the processing elements would include details about one or more pixels together with the necessary scene data (if possible, see Chapter 3).

To summarise our choices for this problem:

**atomic element:** to perform one ray-object intersection;

**task:** to trace the complete path of one ray (may consists of a number of atomic elements);

**PDI:** the pixel location for which we are computing the colour;

**ADI:** the scene data; and,

**task packet:** one or more rays to be computed.

Choosing the task granularity for the parallel implementation of a problem is not straightforward. Although it may be fairly easy to identify the atomic element for the sequential version of the problem, such a fine grain may not be appropriate when using many processing elements. Although the atomic element for ray tracing was specified as computing a single ray-object intersection in the above example, the task granularity for the parallel solution was chosen as computing the complete colour contribution at a particular pixel. If one atomic element had been used as the task granularity then additional problems would have introduced for the parallel solution, namely, the need for processors to exchange partial results. This difficulty would have been exacerbated if, instead, the atomic element had been chosen as tracing a ray into a voxel and considering whether it does in fact intersect with an object there. Indeed, apart from the higher communication overhead this would have introduced, the issue of dependencies would also have to be checked to ensure, for example, that a ray was not checked against an object more than once.

As well as introducing additional communication and dependency overheads, the incorrect choice of granularity may also increase computational complexity variations and hinder efficient load balancing. The choice of granularity is seldom easy, however, a number of parameters of the parallel system can provide an indication as to the desirable granularity. The computation to communication ratio of the architecture will suggest whether additional communication is acceptable to avoid dependency or load balancing problems. As a general rule, where possible, data dependencies should be avoided in the choice of granularity as these imply unnecessary synchronisation points within the parallel solution which can have a significant effect on overall system performance.

### 2.3.2 Task distribution and control

The task management strategy controls the distribution of packets throughout the system. Upon receipt, a processing element performs the tasks specified by a packet. The composition of the task packet is thus an important issue that must be decided before distribution of the tasks can begin. To complete a task a processing element needs a copy of the algorithm, the principal data item(s), and any additional data items that the algorithm may require for that principal data item. The domain decomposition paradigm provides each processing element with a copy of the algorithm, and so the responsibility of the task packet is to provide the other information.

The principal data items form part of the problem domain. If there is sufficient memory, it may be possible to store the entire problem domain as well as the algorithm at each processing element. In this case, the inclusion of the principal data item as part of the task packet is unnecessary. A better method would be simply to include the identification of the principal data item within the task packet. Typically, the identification of a principal data item is considerably smaller, in terms of actual storage capacity, than the item itself. The communication overheads associated with sending this smaller packet will be significantly less than sending the principal data item with the packet. On receipt of the packet the processing element

could use the identification simply to fetch the principal data item from its local storage. The identification of the principal data item is, of course, also essential to enable the results of the entire parallel computation to be collated.

If the additional data items required by the task are known then they, or if possible, their identities, may also be included in the task packet. In this case the task packet would form an integral unit of computation which could be directly handled by a processing element. However, in reality, it may not be possible to store the whole problem domain at every processing element. Similarly, numerous additional data items may be required which would make their inclusion in the task packet impossible. Furthermore, for a large number of problems, the additional data items which are required for a particular principal data item may not be known in advance and will only become apparent as the computation proceeds.

A task packet should contain as a minimum either the identity, or the identity and actual principal data items of the task. The inability to include the other required information in the packet means that the parallel system will have to resort to some form of *data management*. This topic is described fully in Chapter 3.

### 2.3.3 Algorithmic dependencies

The algorithm of the problem may specify an order in which the work must be undertaken. This implies that certain tasks must be completed before others can commence. These dependencies must be preserved in the parallel implementation. In the worst case, algorithmic dependencies can prevent an efficient parallel implementation, as shown with the tower of toy blocks in figure 1.2. Amdahl's law, described in section 1.4, shows the implications to the algorithmic decomposition model of parallel processing of the presence of even a small percentage of purely sequential code. In the domain decomposition approach, algorithmic dependencies may introduce two phenomena which will have to be tackled:

- *synchronisation points* which have the effect of dividing the parallel implementation into a number of distinct stages; and,
- *data dependencies* which will require careful data management to ensure a consistent view of the data to all processing elements.

#### Multi-stage algorithms

Many problems can be solved by a single stage of computation, utilising known principal data items to produce the desired results. However, the dependencies inherent in other algorithms may divide computation into a number of distinct stages. The *partial results* produced by one stage become the principal data items for the following stage of the algorithm, as shown in figure 2.12. For example, many scientific problems involve the construction of a set of simultaneous equations, a distinct stage, and the subsequent solution of these equations for the unknowns. The partial results, in this case elements of the simultaneous equations, become the principal data for the tasks of the next stage.

Even a single stage of a problem may contain a number of distinct substages which must first be completed before the next substage can proceed. An example of this is the use of an iterative solver, such as the Jacobi method [22, 35], to solve a set of simultaneous equations. An iterative method starts with an approximate solution and uses it in a recurrence formula to provide another approximate solution. By repeatedly applying this process a sequence of solutions is obtained which, under suitable conditions, converges towards the exact solution.

Consider the problem of solving a set of six equations for six unknowns,  $\mathbf{Ax} = \mathbf{b}$ . The Jacobi method will solve this set of equations by calculating, at each iteration, a new approximation from the values of the previous iteration. So the value for the  $x_i$ 's at the  $n^{th}$  iteration are calculated as:

$$\begin{aligned} x_1^n &= \frac{b_i - a_{12}x_2^{n-1} - \dots - a_{16}x_6^{n-1}}{a_{11}} \\ x_2^n &= \frac{b_i - a_{21}x_1^{n-1} - \dots - a_{26}x_6^{n-1}}{a_{22}} \\ &\vdots \\ x_6^n &= \frac{b_i - a_{61}x_1^{n-1} - \dots - a_{65}x_5^{n-1}}{a_{66}} \end{aligned}$$

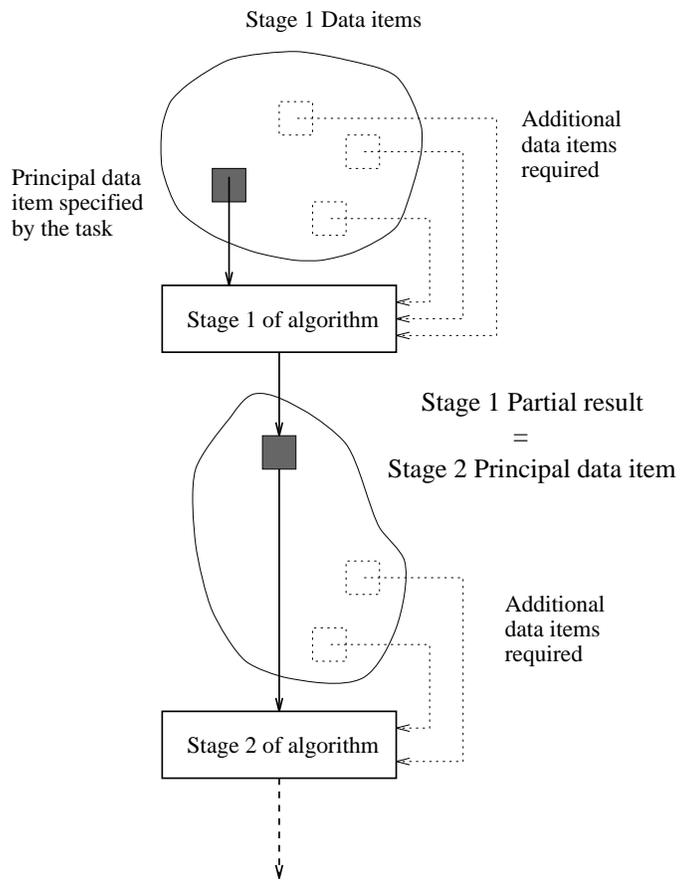


Figure 2.12: The introduction of partial results due to algorithmic dependencies

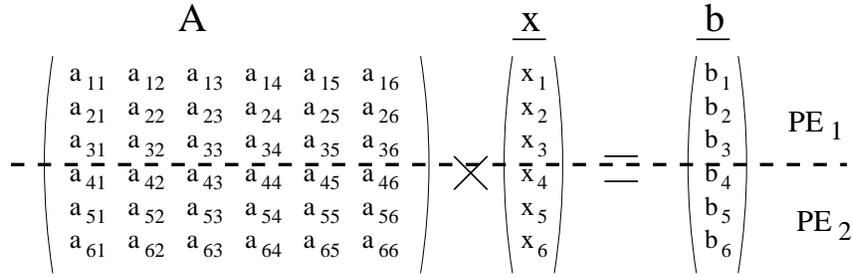


Figure 2.13: Solving an iterative matrix solution method on two processing elements

A parallel solution to this problem on two processing elements could allocate three rows to be solved to each processing element as shown in figure 2.13. Now  $PE_1$  can solve the  $n^{th}$  iteration values  $x_1^n$ ,  $x_2^n$  and  $x_3^n$  in parallel with  $PE_2$  computing the values of  $x_4^n$ ,  $x_5^n$  and  $x_6^n$ . However, neither processing element can proceed onto the  $(n+1)^{st}$  iteration until both have finished the  $n^{th}$  iteration and exchanged their new approximations for the  $x_i^n$ 's. Each iteration is, therefore, a substage which must be completed before the next substage can commence. This point is illustrated by the following code segment from  $PE_1$ :

```

PROCEDURE Jacobi() (* Executing on PE 1 *)
Begin
  Estimate x[1] ... x[6]
  n := 0 (* Iteration number *)
  WHILE solution_not_converged DO
    Begin
      n := n + 1
      Calculate new x[1], x[2] & x[3] using old x[1] ... x[6]
      PARALLEL
        SEND new x[1], x[2] & x[3] TO PE_2
        RECEIVE new x[4], x[5] & x[6] FROM PE_2
      End
    End
  End (* Jacobi *)

```

### Data dependencies

The concept of dependencies was introduced in section 1.1.1 when we were unable to construct a tower of blocks in parallel as this required a strictly sequential order of task completion. In the domain decomposition model, data dependencies exist when a task may not be performed on some principal data item until another task has been completed. There is thus an implicit ordering on the way in which the task packets may be allocated to the processing elements. This ordering will prevent certain tasks being allocated, even if there are processing elements idle, until the tasks on which they are dependent have completed.

A linear dependency exists between each of the iterations of the Jacobi method discussed above. However, no dependency exists for the calculation of each  $x_i^n$ , for all  $i$ , as all the values they require,  $x_j^{n-1}$ ,  $\forall j \neq i$ , will already have been exchanged and thus be available at every processing element.

The Gauss-Seidel iterative method has long been preferred in the sequential computing community as an alternative to Jacobi. The Gauss-Seidel method makes use of new approximations for the  $x_i$  as soon as they are available rather than waiting for the next iteration. Provided the methods converge, Gauss-Seidel will converge more rapidly than the Jacobi method. So, in the example of six unknowns above, in the  $n^{th}$  value of  $x_1^n$  would still be calculated as:

$$x_1^n = \frac{b_1 - a_{12}x_2^{n-1} - \dots - a_{16}x_6^{n-1}}{a_{11}},$$

but the  $x_2^n$  value would now be calculated by:

$$x_2^n = \frac{b_2 - a_{21}x_1^n - a_{23}x_3^{n-1} - \dots - a_{26}x_6^{n-1}}{a_{22}}$$

Although well suited to sequential programming, the strong linear dependency that has been introduced, makes the Gauss-Seidel method poorly suited for parallel implementation. Now within each iteration no value of  $x_i^n$  can be calculated until all the values for  $x_j^n$ ,  $j < i$  are available; a strict sequential ordering of the tasks. The less severe data dependencies within the Jacobi method thus make it a more suitable candidate for parallel processing than the Gauss-Seidel method which is more efficient on a sequential machine.

It is possible to implement a hybrid of these two methods in parallel, the so-called ‘‘Block Gauss-Seidel - Global Jacobi’’ method. A processing element which is computing several rows of the equations, may use the Gauss-Seidel method for these rows as they will be computed sequentially within the processing element. Any values for  $x_i^n$  not computed locally will assume the values of the previous iteration, as in the Jacobi method. All new approximations will be exchanged at each iteration. So, in the example,  $PE_2$  would calculate the values of  $x_4^n$ ,  $x_5^n$  and  $x_6^n$  as follows:

$$x_4^n = \frac{b_i - a_{11}x_1^{n-1} - a_{12}x_2^{n-1} - a_{13}x_3^{n-1} - a_{15}x_5^{n-1} - a_{16}x_6^{n-1}}{a_{44}}$$

$$x_5^n = \frac{b_i - a_{11}x_1^{n-1} - a_{12}x_2^{n-1} - a_{13}x_3^{n-1} - a_{14}x_4^n - a_{16}x_6^{n-1}}{a_{55}}$$

$$x_6^n = \frac{b_i - a_{11}x_1^{n-1} - a_{12}x_2^{n-1} - a_{13}x_3^{n-1} - a_{14}x_4^n - a_{15}x_5^n}{a_{66}}$$

## 2.4 Task Scheduling Strategies

### 2.4.1 Data driven task management strategies

In a data driven approach, the system controller determines the allocation of tasks prior to computation proceeding. With the unbalanced strategy, this may entail an initial sorting stage based on the known computational complexity, as described in section 2.2.1. A single task-packet detailing the tasks to be performed is sent to each processing element. The application processes may return the results upon completion of their allocated portion, or return individual results as each task is performed, as shown in this code segment:

```
PROCESS Application_Process()
  Begin
    RECEIVE task_packet FROM SC via R
    FOR i = start_task_id TO finish_task_id DO
      Begin
        result[i] := Perform_Algorithm(task[i])
        SEND result[i] TO SC via R
      End
    End (* Application_Process *)
```

In a data driven model of computation a processing element may initially be supplied with as many of its allocated principal data items as its local memory will allow. Should there be insufficient storage capacity a simple data management strategy may be necessary to prefetch the missing principal data items as computation proceeds and local storage allows. This is discussed further when considering the management of data in Chapter 3.

### 2.4.2 Demand driven task management strategies

Task management within the demand driven computational model is explicit. The work supplier process, which forms part of the system controller, is responsible for placing the tasks into packets and sending these packets to requesting processing elements. To facilitate this process, the system controller maintains a *pool* of already constituted task packets. On receipt of a request, the work supplier simply dispatches the next available task packet from this task pool, as can be seen in figure 2.14.

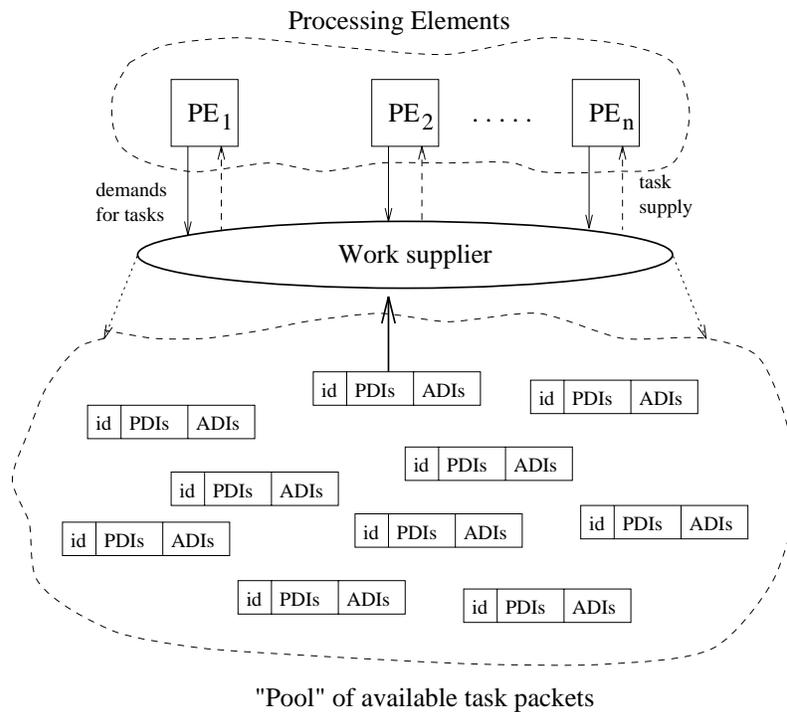


Figure 2.14: Supplying task packets from a task pool at the system controller

The advantage of a task pool is that the packets can be inserted into it in advance, or concurrently as the solution proceeds, according to the allocation strategy adopted. This is especially useful for problems that create work dynamically, such as those using the hybrid approach as described in section 2.2.3. Another advantage of the task pool is that if a hot spot in the problem domain is identified, then the ordering within the task pool can be changed dynamically to reflect this and thus ensure that potentially computationally complex tasks are allocated first.

More than one task pool may be used to reflect different levels of task priority. High priority tasks contained in the appropriate task pool will always be sent to a requesting processing element first. Only once this high priority task pool is (temporarily) empty will tasks from lower priority pools be sent. The multiple pool approach ensures that high priority tasks are not ignored as other tasks are allocated.

In the demand driven computational model, the processing elements demand the next task as soon as they have completed their current task. This demand is translated into sending a request to the work supplier, and the demand is only satisfied when the work supplier has delivered the next task. There is thus a definite delay period from the time the request is issued until the next task is received. During this period the processing element will be computationally idle. To avoid this idle time, it may be useful to include a buffer at each processing element capable of holding at least one task packet. This buffer may be considered as the processing element's own private task pool. Now, rather than waiting for a request to be satisfied from the remote system controller, the processing element may proceed with the computation on the task packet already present locally. When the remote request has been satisfied and a new task packet delivered, this can be stored in the buffer waiting for the processing element to complete the current task.

Whilst avoiding delays in fetching tasks from a remote task pool, the use of a buffer at each processing element may have serious implications for load balancing, especially towards the end of the problem solution. We will examine these issues in more detail after we have considered the realisation of task management for a simple demand driven system - the processor farm.

### A first approach: The processor farm

Simple demand driven models of computation have been implemented and used for a wide range of applications. One realisation of such a model, often referred to in the literature, is that implemented by May and Shepherd [47]. This simple demand driven model, which they term a *processor farm*, has been used for solving problems with high computation to communication ratios. The model proposes a single system controller and one or more processing elements connected in a linear configuration, or chain. The structure of a processing element in this model is shown in figure 2.15.

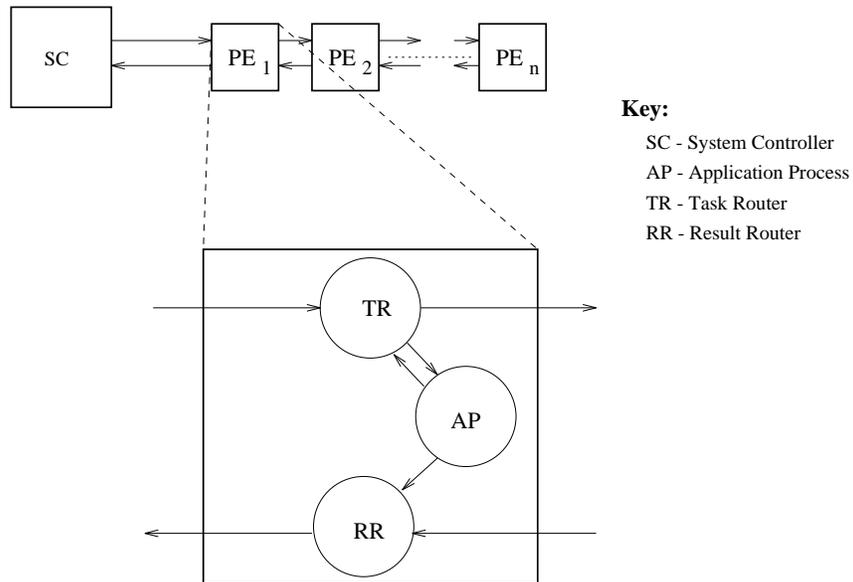


Figure 2.15: A processing element for the processor farm model

The application process performs the desired computation, while the communication within the system is dealt with by two router processes, the Task Router (TR) and the Result Router (RR). As their names suggest, the task router is responsible for distributing the tasks to the application process, while the result router returns the results from the completed tasks back to the system controller. The system controller contains the initial pool of tasks to be performed and collates the results. Such a communication strategy is simple to implement and largely problem independent.

To reduce possible processing element idle time, each task router process contains a single buffer in which to store a task so that a new task can be passed to the application process as soon as it becomes idle. When a task has been completed the results are sent to the system controller. On receipt of a result, the system controller releases a new task into the system. This synchronised releasing of tasks ensures that there are never more tasks in the system than there is space available.

On receipt of a new task, the task router process either:

1. passes the task directly to the application process if it is waiting for a task; or
2. places the task into its buffer if the buffer is empty; or, otherwise
3. passes the task onto the next processing element in the chain.

The processor farm is initialised by loading sufficient tasks into the system so that the buffer at each task router is full and each application process has a task with which to commence processing. Figure 2.16 shows the manner in which task requests are satisfied within a simple two processing element configured in a chain.

The simplicity of this realisation of a demand driven model has contributed largely to its popularity. Note that because of the balance maintained within the system, the only instance at which the last processing element is different from any other processing element in the chain is to ensure the `closedown_command`

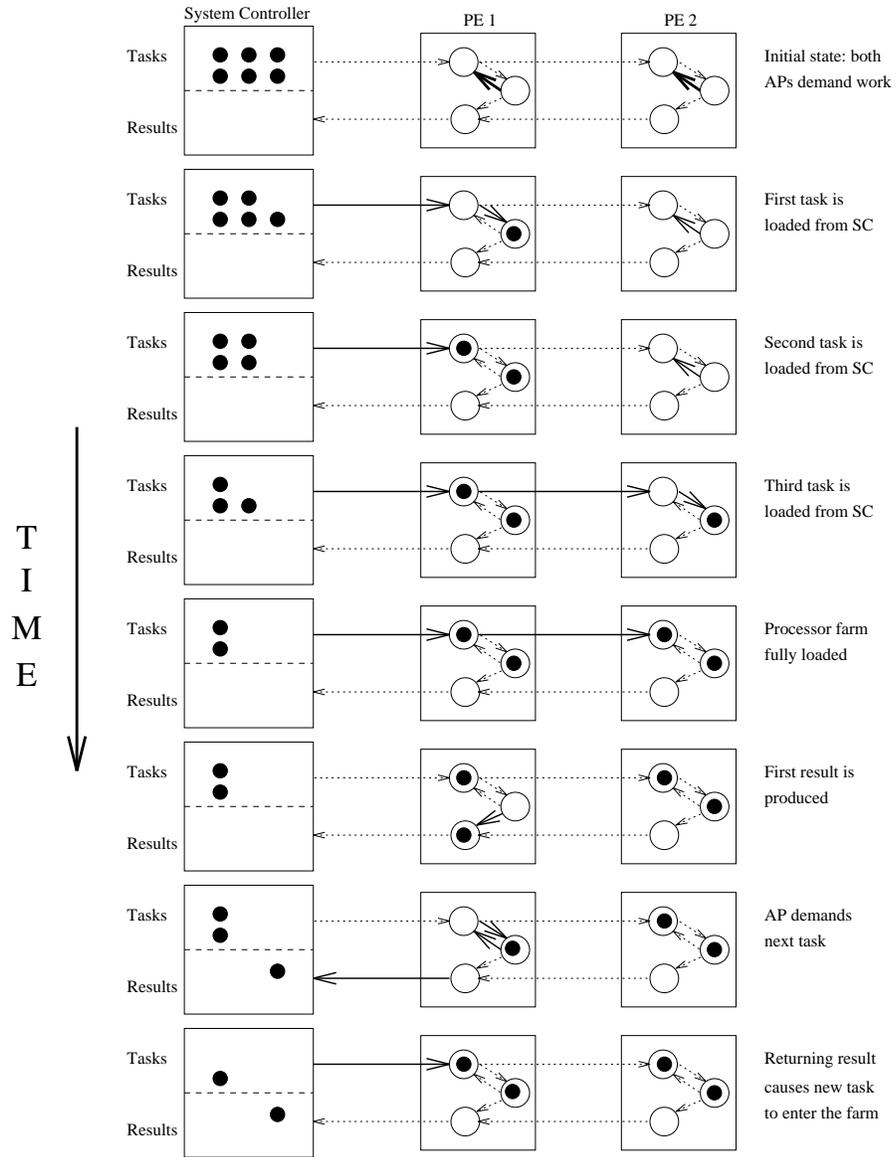


Figure 2.16: Task movement within a two PE processor farm

does not get passed any further. However, such a model does have disadvantages which may limit its use for more complex problems.

The computation to communication ratio of the desired application is critical in order to ensure an adequate performance of a processor farm. If this ratio is too low then significant processing element idle time will occur. This idle time occurs because the computation time for the application process to complete its current task and the task buffered at the task router may be lower than the combined communication time required for the results to reach the system controller plus the time for the new tasks released into the system to reach the processing element. This problem may be partially alleviated by the inclusion of several buffers at each task router instead of just one. However, without *a priori* knowledge as to the computation to communication ratio of the application, it may be impossible to determine precisely what the optimum number of buffers should be. This analysis is particularly difficult if the computational complexity of the tasks vary; precisely the type of problem demand driven models are more apt at solving. The problem independence of the system will also be compromised by the use of any *a priori* knowledge.

If the number of buffers chosen is too small, then the possibility of application process idle time will not be avoided. Provision of too many buffers will certainly remove any immediate application process idle time, but will re-introduce the predicament as the processing draws to a close. This occurs once the system controller has no further tasks to introduce into the system and now processing must only continue until all tasks still buffered at the processing elements have been completed. Obviously, significant idle time may occur as some processing elements struggle to complete their large number of buffered tasks.

The computation to communication ratio of the processor farm is severely exacerbated by the choice of the chain topology. The distance between the furthest processing element in the chain and the system controller grows linearly as more processing elements are added. This means that the combined communication time to return a result and receive a new task also increases. Furthermore, this communication time will also be adversely affect by the message traffic of all the intermediate processing elements which are closer to the system controller.

### 2.4.3 Task manager process

The aim of task management within a parallel system is to ensure the efficient supply of tasks to the processing elements. A Task Manager process (TM) is introduced at each processing element to assist in maintaining a continuous supply of tasks to the application process. The application process no longer deals with task requests directly, but rather indirectly using the facilities of the task manager. The task manager process assumes the responsibility for ensuring that every request for additional tasks from the application process will be satisfied immediately. The task manager attempts to achieve this by maintaining a local task pool.

In the processor farm, the task router process contains a single buffered task in order to satisfy the next local task request. As long as this buffer is full, task supply is immediate as far as the application process is concerned. The buffer is refilled by a new task from the system controller triggered on receipt of a result. The task router acts in a *passive* manner, awaiting replenishment by a new task within the farm. However, if the buffer is empty when the application process requests a task then this process must remain idle until a new task arrives. This idle time is wasted computation time and so to improve system performance the passive task router should be replaced by a “intelligent” task manager process more capable of ensuring new tasks are always available locally.

The task management strategies implemented by the task manager and outlined in the following sections are *active*, dynamically requesting and acquiring tasks during computation. The task manager thus assumes the responsibility of ensuring local availability of tasks. This means that an application process should *always* have its request for a task satisfied immediately by the task manager unless:

- at the start of the problem the application processes make a request before the initial tasks have been provided by the system controller;
- there are no more tasks which need to be solved for a particular stage of the parallel implementation; or,
- the task manager’s replenishment strategy has failed in some way.

## A local task pool

To avoid any processing element idle time, it is essential that the task manager has at least one task available locally at the moment the application process issues a task request. This desirable situation was achieved in the processor farm by the provision of a single buffer at each task router. As we saw, the single buffer approach is vulnerable to the computation to communication ratio within the system. Adding more buffers to the task router led to the possibility of serious load imbalances towards the end of the computation.

The task manager process maintains a local task pool of tasks awaiting computation by the application process. This pool is similar to the task pool at the system controller, as shown in figure 2.14. However, not only will this local pool be much smaller than the system controller's task pool, but also it may be desirable to introduce some form of "status" to the number of available tasks at any point in time.

Satisfying a task request will free some space in the local task pool. A simple replenishment strategy would be for the task manager immediately to request a new task packet from the system controller. This request has obvious communication implications for the system. If the current message densities within the system are high and as long as there are still tasks available in the local task pool, this request will place an unnecessary additional burden on the already overloaded communication network.

As an active process, it is quite possible for the task manager to delay its replenishment request until message densities have diminished. However, this delay must not be so large that subsequent application process demands will deplete the local task pool before any new tasks can be fetched causing processor idle time to occur. There are a number of indicators which the task manager can use to determine a suitable delay. Firstly, this delay is only necessary if current message densities are high. Such information should be available for the router. Given a need for delay, the number of tasks in the task pool, the approximate computation time each of these tasks requires, and the probable communication latency in replenishing the tasks should all contribute to determining the request delay.

In a demand driven system, the computational complexity variations of the tasks are not known. However, the task manager will be aware of how long previous tasks have taken to compute (the time between application process requests). Assuming some form of preferred biased allocation of tasks in which tasks from similar regions of the problem domain are allocated to the same processing element, as discussed in section 2.4.5, the task manager will be able to build up a profile of task completion time which can be used to predict approximate completion times for tasks in the task pool. The times required to satisfy previous replenishment requests will provide the task manager with an idea of likely future communication responses. These values are, of course, mere approximations, but they can be used to assist in determining reasonable tolerance levels for the issuing of replenishment requests.

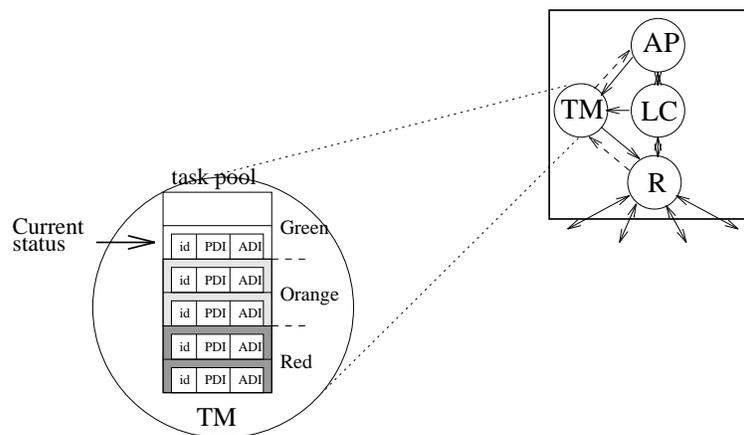


Figure 2.17: Status of task manager's task pool

The task manager's task pool is divided into three regions: *green*, *orange* and *red*. The number of tasks available in the pool will determine the current status level, as shown in figure 2.17. When faced with the need to replenish the task pool the decision can be taken based on the current status of the pool:

**green:** Only issue the replenishment request if current message traffic density is low;

**orange:** Issue the replenishment request unless the message density is very high; and,

**red:** Always issue the replenishment request.

The boundaries of these regions may be altered dynamically as the task manager acquires more information. At the start of the computation the task pool will be all red. The computation to communication ratio is critical in determining the boundaries of the regions of the task pool. The better this ratio, that is when computation times are high relative to the time taken to replenish a task packet, the smaller the red region of the task pool need be. This will provide the task manager with greater flexibility and the opportunity to contribute to minimising communication densities.

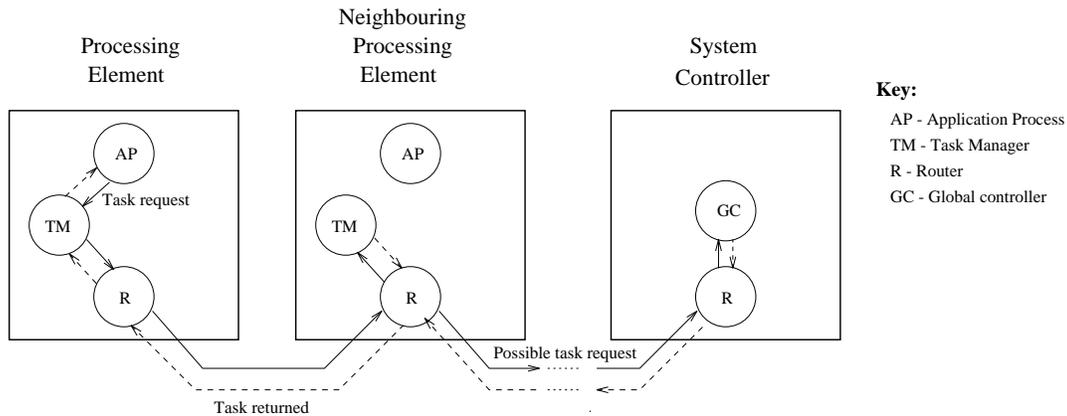


Figure 2.18: Task request propagating towards the system controller

## 2.4.4 Distributed task management

One handicap of the centralised task pool system is that all replenishment task requests from the task managers must reach the system controller before the new tasks can be allocated. The associated communication delay in satisfying these requests can be significant. The communication problems can be exacerbated by the bottleneck arising near the system controller. Distributed task management allows task requests to be handled at a number of locations remote from the system controller. Although all the tasks originate from the system controller, requests from processing elements no longer have to reach there in order to be satisfied.

The closest location for a task manager to replenish a task packet is from the task pool located at the task manager of one of its nearest neighbours. In this case, a replenishment request no longer proceeds directly to the system controller, but simply via the appropriate routers to the neighbouring task manager. If this neighbouring task manager is able to satisfy the replenishment request then it does so from its task pool. This task manager may now decide to in turn replenish its task pool, depending on its current status and so it will also request another task from one of its neighbouring task managers, but obviously not the same neighbour to which it has just supplied the task. One sensible strategy is to propagate these requests in a “chain like” fashion in the direction towards the main task supplier at the system controller, as shown in figure 2.18.

This distributed task management strategy is referred to as a *producer-consumer* model. The application process is the initial consumer and its local task manager the producer. If a replenishment request is issued then this task manager becomes the consumer and the neighbouring task manager the producer, and so on. The task supplier process of the system controller is the overall producer for the system. If no further tasks exist at the system controller then the last requesting task manager may change the direction of the search. This situation may occur towards the end of a stage of processing and facilitates load balancing of any tasks remaining in task manager buffers. As well as reducing the communication distances for task replenishment, an additional advantage of this “chain reaction” strategy is that the number of request

messages in the system is reduced. This will play a major rôle helping maintain a lower overall message density within the system.

If a task manager is unable to satisfy a replenishment request as its task pool is empty, then to avoid “starvation” at the requesting processing element, this task manager must ensure that the request is passed on to another processing element.

A number of variants of the producer-consumer model are also possible:

- Instead of following a path towards the system controller, the “chain reaction” could follow a predetermined Hamiltonian path (the system controller could be one of the processors on this path).

**Aside:** A Hamiltonian path is a circuit starting and finishing at one processing element. This circuit passes through each processor in the network once only.

Such a path would ensure that a processing element would be assured of replenishing a task if there was one available and there would be no need to keep track of the progress of the “chain reaction” to ensure no task manager was queried more than once per chain.

- In the course of its through-routing activities a router may handle a task packet destined for a distant task manager. If that router’s local task manager has an outstanding “red request” for a task then it is possible for the router to *poach* the “en route task” by diverting it, so satisfying its local task manager immediately. Care must be taken to ensure that the task manager for whom the task was intended is informed that the task has been poached, so it may issue another request. In general, tasks should only be poached from “red replenishment” if to do so would avoid local application process idle time.

## 2.4.5 Preferred bias task allocation

The preferred bias method of task management is a way of allocating tasks to processing elements which combines the simplicity of the balanced data driven model with the flexibility of the demand driven approach. To reiterate the difference in these two computational models as they pertain to task management:

- Tasks are allocated to processing elements in a predetermined manner in the balanced data driven approach.
- In the demand driven model, tasks are allocated to processing elements on demand. The requesting processing element will be assigned the next available task packet from the task pool, and thus no processing element is bound to any area of the problem domain.

Provided no data dependencies exist, the order of task completion is unimportant. Once all tasks have been computed, the problem is solved. In the preferred bias method the problem domain is divided into equal regions with each region being assigned to a particular processing element, as is done in the balanced data driven approach. However, in this method, these regions are purely *conceptual* in nature. A demand driven model of computation is still used, but the tasks are not now allocated in an arbitrary fashion to the processing elements. Rather, a task is dispatched to a processing element from its conceptual portion. Once all tasks from a processing element’s conceptual portion have been completed, only then will that processing element be allocated its next task from the portion of another processing element which has yet to complete its conceptual portion of tasks. Generally this task should be allocated from the portion of the processing element that has completed the least number of tasks. So, for example, from figure 2.19, on completion of the tasks in its own conceptual region,  $PE_3$  may get allocated task number 22 from  $PE_2$ ’s conceptual region. Preferred bias allocation is sometimes also referred to as *conceptual task allocation*.

The implications of preferred bias allocation are substantial. The demand driven model’s ability to deal with variations in computational complexity is retained, but now the system controller and the processing elements themselves know to whom a task that they have been allocated conceptually belongs. As we will see in section 3.6, this can greatly facilitate the even distribution of partial results at the end of any stage of a multi-stage problem.

The exploitation of data coherence is a vital ploy in reducing idle time due to remote data fetches. Preferred bias allocation of tasks can ensure that tasks from the same region of the problem are allocated to the same processing element. This can greatly improve the cache hit ratio at that processing element.

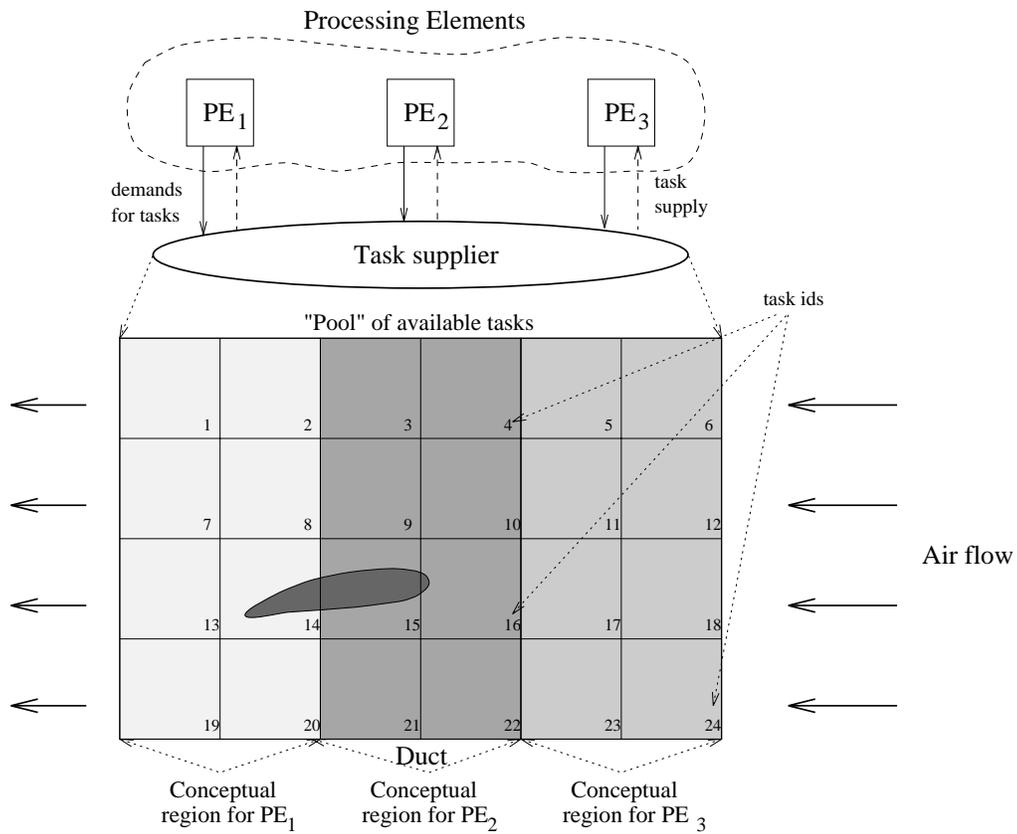


Figure 2.19: Partial result storage balancing by means of conceptual regions

## Chapter 3

# Data Management

The data requirements of many problems may be far larger than can be accommodated at any individual processing element. Rather than restricting ourselves to only solving those problems that fit completely within every processing element's local memory, we can make use of the combined memory of all processing elements. The large problem domain can now be distributed across the system and even secondary storage devices if necessary. For this class of application some form of data management will be necessary to ensure that data items are available at the processing elements when required by the computations.

Virtual shared memory regards the whole problem domain as a single unit in which the data items may be individually referenced. This is precisely how the domain could be treated if the problem was implemented on a shared memory multiprocessor system. However, on a distributed memory system, the problem domain is distributed across the system and hence the term **virtual**. Virtual shared memory systems may be implemented at different levels, such as in hardware or at the operating system level. In this chapter we will see how the introduction of a data manager process at each processing element can provide an elegant virtual shared memory at the system software level of our parallel implementation.

### 3.1 World Model of the Data: No Data Management Required

Not all problems possess very large data domains. If the size of the domain is such that it may be accommodated at every processing element then we say that the processing elements have a "world model" of the data. A world model may also exist if all the tasks allocated to a processing element only ever require a *subset* of the problem domain and this subset can be accommodated completely. In the world model, all principal and additional data items required by an application process will always be available locally at each processing element and thus there is no need for any data item to be fetched from another remote location within the system. If there is no requirement to fetch data items from remote locations as the solution of the problem proceeds then there is no need for any form of data management.

The processor farm described in section 2.4.2 is an example of a parallel implementation which assumes a world model. In this approach, tasks are allocated to processing elements in an arbitrary fashion and thus there is no restriction on which tasks may be computed by which processing element. No provision is made for data management and thus to perform any task, the entire domain must reside at each processing element.

Data items do not always have to be present at the processing element from the start of computation to avoid any form of data management. As discussed in section 2.3.1, both principal and additional data items may be included within a task packet. Provided no further data items are required to complete the tasks specified in the task packet then no data management is required and this situation may also be said to be demonstrating a world data model.

### 3.2 Virtual Shared Memory

Virtual shared memory provides all processors with the concept of a single memory space. Unlike a traditional shared memory model, this physical memory is distributed amongst the processing elements. Thus, a virtual shared memory environment can be thought of providing each processing element with a *virtual*

*world model* of the problem domain. So, as far as the application process is concerned, there is no difference between requesting a data item that happens to be local, or remote; only the speed of access can be (very) different.

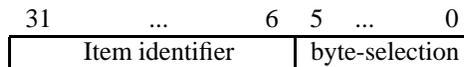
Virtual shared memory can be implemented at any level in the computer hierarchy. Implementations at the hardware level provide a transparent interface to software developers, but requires a specialised machine, such as the DASH system [43]. There have also been implementations at the operating system and compiler level. However, as we shall see, in the absence of dedicated hardware, virtual shared memory can also be easily provided at the system software level. At this level, a great deal of flexibility is available to provide specialised support to minimise any implementation penalties when undertaking the solution of problems with very large data requirements on multiprocessor systems. Figure 3.1 gives four levels at which virtual shared memory (VSM) can be supported, and examples of systems that implement VSM at that particular level.

<i>Higher level</i>	System Software	Provided by the Data Manager process
	Compiler	High Performance Fortran[36], ORCA[6]
	Operating System	Coherent Paging[45]
<i>Lower level</i>	Hardware	DDM [61], DASH [43], KSR-1 [38]

Figure 3.1: The levels where virtual shared memory can be implemented.

### 3.2.1 Implementing virtual shared memory

At the *hardware level* virtual shared memory intercepts all memory traffic from the processor, and decides which memory accesses are serviced locally, and which memory accesses need to go off-processor. This means that everything above the hardware level (machine code, operating system, etc.) sees a virtual shared memory with which it may interact in exactly the same manner as a physically shared memory. Providing this, so called, transparency to the higher levels, means that the size of data is not determined by the hardware level. However, in hardware, a data item becomes a fixed consecutive number of bytes, typically around 16-256. By choosing the size to be a power of 2, and by aligning data items in the memory, the physical memory address can become the concatenation of the “item-identifier” and the “byte selection”. This strategy is easier to implement in hardware.



In this example, the most significant bits of a memory address locates the data item, and the lower bits address a byte within the item. The choice of using 6 bits as the byte selection in this example is arbitrary.

If a data structure of some higher level language containing two integers of four bytes each happened to be allocated from, say, address ...1100 111100 to ...1101 000100, then item ...1100 will contain the first integer, and item ...1101 will contain the other one. This means that two logically related integers of data are located in two physically separate items (although they could fit in a single data item).

Considered another way, if two unrelated variables, say  $x$  and  $y$  are allocated at addresses ...1100 110000 and ...1100 110100, then they reside in the same data item. If they are heavily used on separate processors, this can cause inefficiencies when the machine tries to maintain sequentially consistent copies of  $x$  and  $y$  on both processors. The machine cannot put  $x$  on one processor and  $y$  on the other, because it does not recognise  $x$  and  $y$  as different entities; the machine observes it as a single item that is shared between two processors. If sequential consistency has to be maintained the machine must update every write to  $x$  and  $y$  on both processors, even though the variables are not shared at all. This phenomenon is known as *false sharing*.

Virtual shared memory implemented at the *operating system level* also use a fixed size for data items, but these are typically much larger than at the hardware level. By making an item as large as a page of the operating system (around 1-4 KByte), data can be managed at the page level. This is cheaper, but slower than a hardware implementation.

When the *compiler* supports virtual shared memory, a data item can be made exactly as large as any user data structure. In contrast with virtual shared memory implementations at the hardware or operating

system level, compiler based implementations can keep logically connected variables together and distribute others. The detection of logically related variables is in the general case very hard, which means that applications written in existing languages such as C, Modula-2 or Fortran cannot be compiled in this way. However, compilers for specially designed languages can provide some assistance. For example, in High Performance Fortran the programmer indicates how arrays should be divided and then the compiler provides the appropriate commands to support data transport and data consistency.

Implementing virtual shared memory at the *system software* level provides the greatest flexibility to the programmer. However, this requires explicit development of system features to support the manipulation of the distributed data item. A *data manager* process is introduced at each processing element especially to undertake this job.

### 3.3 The Data Manager

Virtual shared memory is provided at the system software level by a data manager process at each processing element. The aim of data management within the parallel system is to ensure the efficient supply of data items to the processing elements. The data manager process manages data items just as the task manager was responsible for maintaining a continuous supply of tasks. Note that the data items being referred to here are the principal and additional data items as specified by the problem domain and not every variable or constant the application process may invoke for the completion of a task.

The application process now no longer deals with the principal and additional data items directly, but rather indirectly using the facilities of the data manager. The application process achieves this by issuing a data request to the data manager process every time a data item is required. The data manager process assumes the responsibility for ensuring that every request for a data item from the application process will be satisfied. The data manager attempts to satisfy these requests by maintaining a local data cache.

The data management strategies implemented by the data manager and outlined in the following sections are *active*, dynamically requesting and acquiring data items during computation. This means that an application process should *always* have its request for a data item satisfied immediately by the data manager unless:

- at the start of the problem the application processes make requests before any initial data items have been provided by the system controller;
- the data manager's data fetch strategy has failed in some way.

#### 3.3.1 The local data cache

The concept of *data sharing* may be used to cope with very large data requirements [14, 23]. Data sharing implements virtual shared memory by allocating every data item in the problem domain an unique identifier. This allows a required item to be "located" from somewhere within the system, or from secondary storage if necessary. The size of problem that can now be tackled is, therefore, no longer dictated by the size of the local memory at each processing element, but rather only by the limitations of the combined memory plus the secondary storage.

The principal data item required by an application process is specified by the task it is currently performing. Any additional data item requirements are determined by the task *and* by the algorithm chosen to solve the problem. These additional data items may be known *a priori* by the nature of the problem, or they may only become apparent as the computation of the task proceeds.

To avoid any processing element idle time, it is essential that the data manager has the required data item available locally at the moment the application process issues a request for it. In an attempt to achieve this, the data manager maintains a local cache of data items as shown in figure 3.2. The size of this cache, and thus the number of data items it can contain, is determined by the size of a processing element's local memory.

Each data item in the system is a packet containing the unique identifier, shown in figure 3.2 as *id*, together with the actual data which makes up the item. The data items may be permanently located at a specific processing element, or they may be free to migrate within the system to where they are required. When a data manager requires a particular data item which is not already available locally, this data item must be fetched from some remote location and placed into the local cache. This must occur before the application

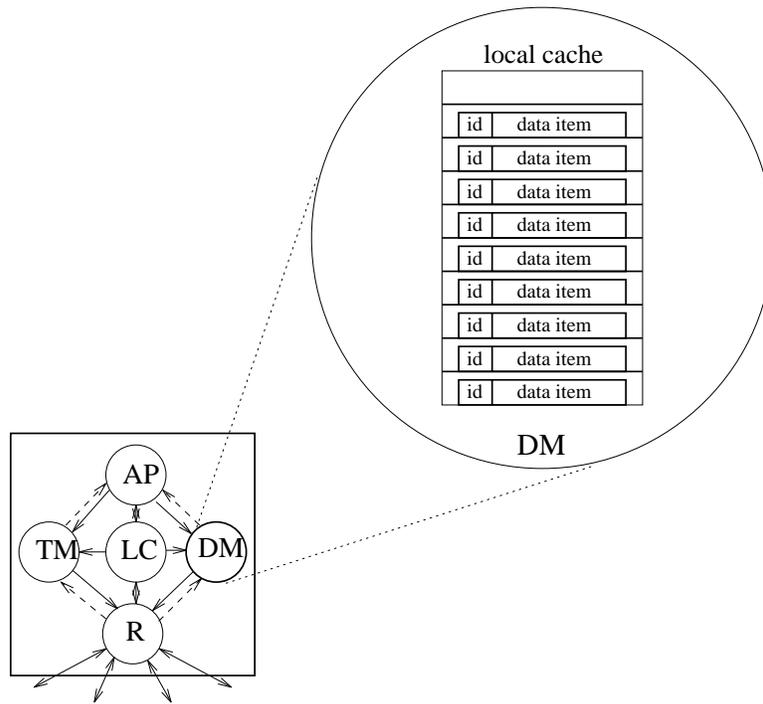


Figure 3.2: The local cache at the data manager

process can access the data item. The virtual shared memory of the system is thus the combination of the local caches at all the processing elements plus the secondary storage which is under the control of the file manager at the system controller.

In certain circumstances, as will be seen in the following sections, rather than removing the data item from the local cache in which it was found, it may be sufficient simply to take a copy of the data item and return this to the local cache. This is certainly the case when the data items within the problem domain are *read-only*, that is the values of the data items are not altered during the course of the parallel solution of the problem (and indeed the same would be true of the sequential implementation). This means that it is possible for copies of the same data item to be present in a number of local caches. Note that it is no advantage to have more than one copy of any data item in one local cache.

There is a limited amount of space in any local cache. When the cache is full and another data item is acquired from a remote location, then one of the existing data items in the local cache must be replaced by this new data item. Care must be taken to ensure that no data item is inadvertently completely removed from the system by being replaced in all local caches. If this does happen then, assuming the data item is read-only, a copy of the entire problem domain will reside on secondary storage, from where the data items were initially loaded into the local caches of the parallel system. This means that should a data item being destroyed within the system, another copy can be retrieved from the file manager (FM) of the system controller.

If the data items are *read-write* then their values may be altered as the computation progresses. In this case, the data managers have to beware of consistency issues when procuring a data item. The implications of consistency will be discussed in section 3.4.

As we will now see, the strategies adopted in the parallel implementation for acquiring data items and storing them in the local caches can have a significant effect on minimising the implementation penalties and thus improving overall system performance. The onus is on the data manager process to ensure these strategies are carried out efficiently.

### 3.3.2 Requesting data items

The algorithm being executed at the application process will determine the next data item required. If the data items were all held by the application process, requesting the data item would be implemented within the application process as an “assignment statement”. For example a request for data item  $i$  would simply be written as  $x := \text{data\_item}[i]$ . When all the data items are held instead by the data manager process, this “assignment statement” must be replaced by a request from the application process to the data manager for the data item followed by the sending of a copy of the data item from the local cache of the data manager to the waiting application process, as shown in figure 3.3.

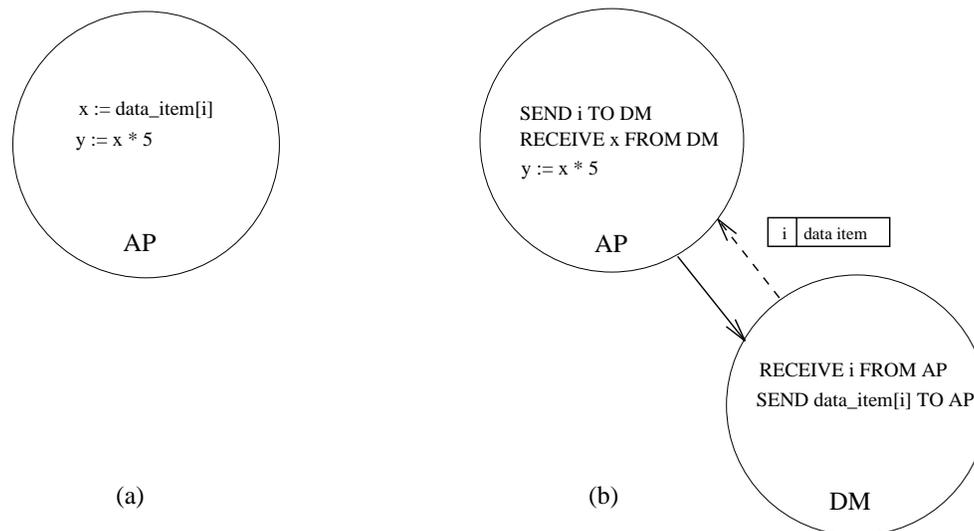


Figure 3.3: Accessing a data item (a) with, and (b) without a data manager

The data item’s unique identifier enables the data manager to extract the appropriate item from its local cache. If a data item requested by the application process is available, it is immediately transferred, as shown in figure 3.4(a). The only slight delay in the computation of the application process will occur by the need to schedule the concurrent data manager and for this process to send the data item from its local cache. However, if the data item is not available locally then the data manager must “locate” this item from elsewhere in the system. This will entail sending a message via the router to find the data item in another processing element’s local cache, or from the file manager of the system controller. Having been found, the appropriate item is returned to the requesting data manager’s own local cache and then finally a copy of the item is transferred to the application process.

If the communicated request from the application process is asynchronous and this process is able to continue with its task while awaiting the data item then no idle time occurs. However, if the communication with the data manager is synchronous, or if the data item is essential for the continuation of the task then idle time will persist until the data item can be fetched from the remote location and a copy given to the application process, as shown in figure 3.4(b). Unless otherwise stated, we will assume for the rest of this chapter that an application process is unable to continue with its current task until its data item request has been satisfied by the data manager.

### 3.3.3 Locating data items

When confronted with having to acquire a remote data item, two possibilities exist for the data manager. Either it knows exactly the location of the data item within the system, or this location is unknown and some form of search will have to be instigated.

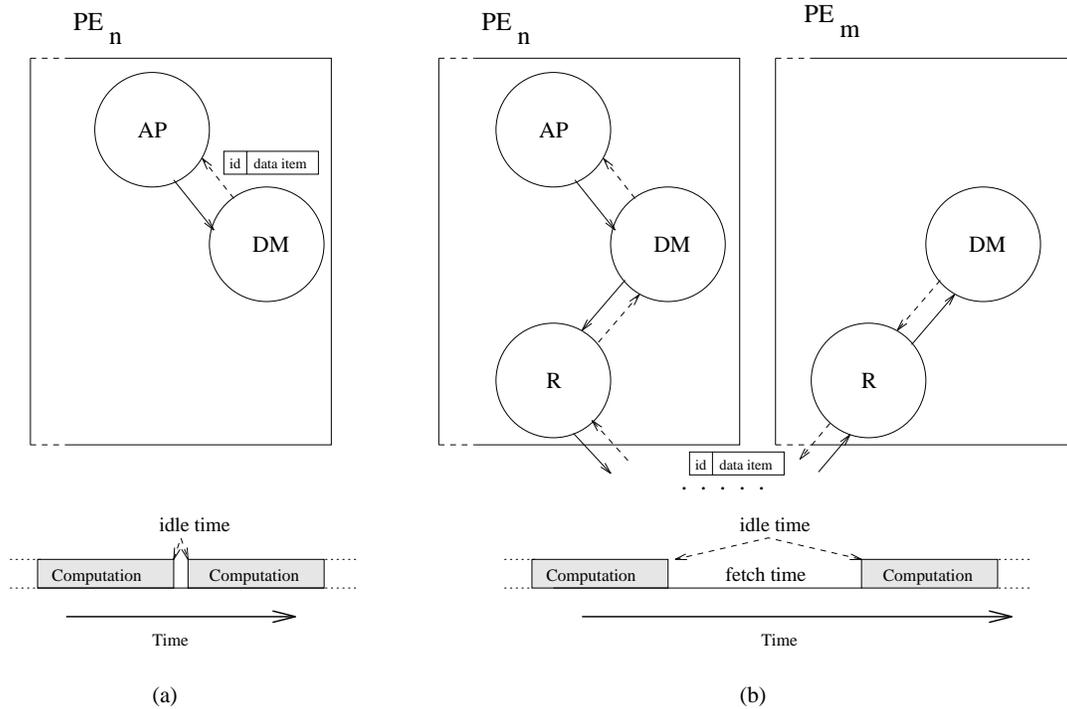


Figure 3.4: AP idle time due to: (a) Data item found locally (b) Remote data item fetch

### Resident sets

Knowing the precise location of the requested data item within the system enables the data manager to instruct the router to send the request for the data item, directly to the appropriate processing element.

One of the simplest strategies for allocating data items to each processing element's local cache is to divide all the data items of the problem domain evenly amongst the processing elements before the computation commences. Providing there is sufficient local memory and assuming there are  $n$  processing elements, this means that each processing element would be allocated  $\frac{1}{n}$  of the total number of data items. If there isn't enough memory at each processing element for even this fraction of the total problem domain then as many as possible could be allocated to the local caches and the remainder of the data items would be held at the file manager of the system controller. Such a simplistic scheme has its advantages. Provided these data items remain at their predetermined local cache for the duration of the computation, then the processing element from which any data item may be found can be computed directly from the identity of the data item.

For example, assume there are twelve data items, given the unique identification numbers  $1, \dots, 12$ , and three processing elements,  $PE_1$ ,  $PE_2$ , and  $PE_3$ . A predetermined allocation strategy may allocate permanently data items  $1, \dots, 4$  to  $PE_1$ , data items  $5, \dots, 8$  to  $PE_2$  and  $9, \dots, 12$  to  $PE_3$ . Should  $PE_2$  wish to acquire a copy of data item 10, it may do so directly from the processing element known to have that data item, in this case  $PE_3$ .

It is essential for this simple predetermined allocation strategy that the data items are not overwritten or moved from the local cache to which they are assigned initially. However, it may be necessary for a processing element to also acquire copies of other data items as the computation proceeds, as we saw with  $PE_2$  above. This implies that the local cache should be partitioned into two distinct regions:

- a region containing data items which may never be replaced, known as the *resident set*; and,
- a region for data items which may be replaced during the parallel computation.

The size of the resident set should be sufficient to accommodate all the pre-allocated data items, as

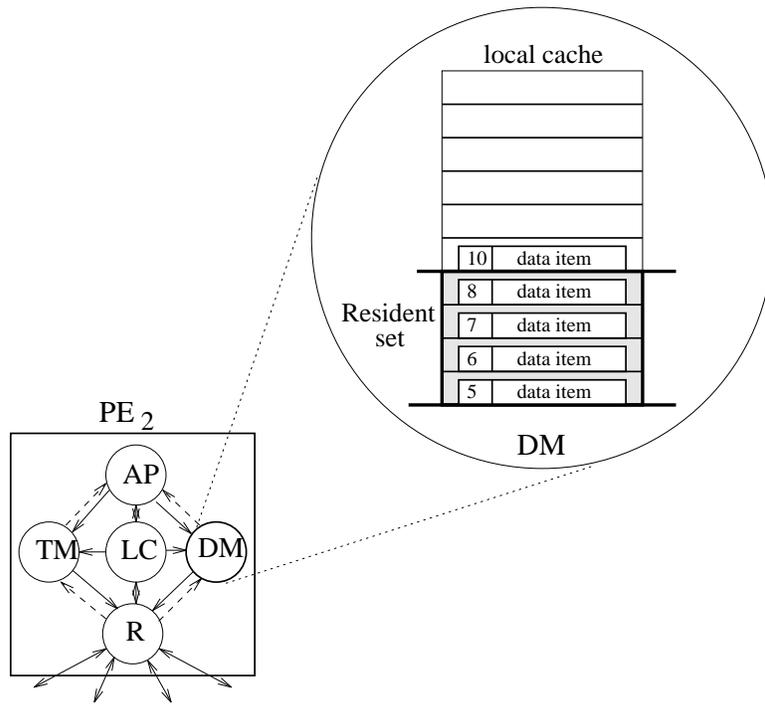


Figure 3.5: Resident set of the local cache

shown for  $PE_2$  from the above example in figure 3.5. The remaining portion of the local cache will be as large as allowed by the local memory of the processing element. Note that this portion needs to have sufficient space to hold a minimum of *one* data item as this is the maximum that the application process can require at any specific point during a task's computation. To complete a task an application process may require many data items. Each of these data items may in turn replace the previously acquired one in the single available space in the local cache.

The balanced data driven model of computation is well suited to a simple pre-determined even data item allocation scheme. In this model the system controller knows prior to the computation commencing precisely which tasks are to be assigned to which processing elements. The same number of tasks is assigned to each processing element and thus the principal data items for each of these tasks may be pre-allocated evenly amongst the local caches of the appropriate processing elements. Similar knowledge is available to the system controller for the unbalanced data driven model, but in this case the number of tasks allocated to each processing element is not the same and so different numbers of principal data items will be loaded into each resident set. Note that the algorithm used to solve the problem may be such that, even if a data driven model is used and thus the principal data items are known in advance, the additional data items may not be known *a priori*. In this case, these additional data items will have to be fetched into the local caches by the data managers as the computation proceeds and the data requirements become known.

More sophisticated pre-allocation strategies, for example some form of hashing function, are possible to provide resident sets at each processing element. It is also not necessary for each data item to be resident at only one processing element. Should space permit, the same data item may be resident at several local caches.

The pre-allocation of resident sets allows the location of a data item to be determined from its unique identifier. A pre-allocated resident set may occupy a significant portion of a local cache and leave little space for other data items which have not been pre-allocated. The shortage of space would require these other data items to be replaced constantly as the computation proceeds. It is quite possible that one data item may be needed often by the same application process either for the same task or for several tasks. If this data item is not in the resident set for that processing element, then there is the danger that the data item will be replaced during the periods that it is not required and thus will have to be re-fetched when it is

required once more. Furthermore, despite being pre-allocated, the data items of a resident set may in fact never be required by the processing element to which they were allocated. In the example given earlier,  $PE_2$  has a resident set containing data items 5, . . . , 8. Unless there is *a priori* knowledge about the data requirements of the tasks, there is no guarantee that  $PE_2$  will ever require any of these data items from its resident set. In this case, a portion of  $PE_2$ 's valuable local cache is being used to store data items which are never required, thus reducing the available storage for data items which are needed. Those processing elements that do require data items 5, . . . , 8 are going to have to fetch them from  $PE_2$ . Not only will the fetches of these data items imply communication delays for the requesting data managers, but also, the need for  $PE_2$ 's data manager to service these requests will imply concurrent activity by its data manager which will detract from the computation of the application process.

The solution to this dilemma is not to pre-allocate resident sets, but to build up such a set as computation proceeds and information is gained by each data manager as to the data items most frequently used by its processing element. Profiling can also assist in establishing these resident sets, as explained in section 3.5.3. The price to pay for this flexibility is that it may no longer be possible for a data manager to determine precisely where a particular data item may be found within the system.

### Searching for data at unknown locations

Acquiring a specific data item from an unknown location will necessitate the data manager requesting the router process to "search" the system for this item. The naive approach would be for the router to send the request to the data manager process of each processing element in turn. If the requested data manager has the necessary data item it will return a copy and then there is no need for the router to request any further processing elements. If the requested data manager does not have the data item then it must send back a `not_found` message to the router, whereupon the next processing element may be tried. The advantages of this *one-to-one* scheme is that as soon as the required data item is found, no further requests need be issued and only one copy of the data item will ever be returned. However, the communication implications of such a scheme for a large parallel system are substantial. If by some quirk of fate (or Murphy's law), the last processing element to be asked is the one which has the necessary data item, then one request will have resulted in  $2 \times (\text{number of PEs} - 1)$  messages, a quite unacceptable number for large systems. Furthermore, the delay before the data item is finally found will be large, resulting in long application process idle time.

An alternative to this communication intensive *one-to-one* approach, is for the router process to issue a global broadcast of the request; a *one-to-many* method. A bus used to connect the processing elements is particularly suited to such a communication strategy, although, as discussed in section 1.2.1, a bus is not an appropriate interconnection method for large multiprocessor systems. The broadcast strategy may also be used efficiently on a more suitable interconnection method for large systems, such as interconnections between individual processors. In this case, the router issues the request to its directly-connected neighbouring processing elements. If the data managers at these processing elements have the required data item then it is returned, if not then these neighbouring processing elements in turn propagate the request to their neighbours (excluding the one from which they received the message). In this way, the requests propagates through the system like ripples on a pond. The message density inherent in this approach is significantly less than the *one-to-one* approach, however one disadvantage is that if the requested data item is replicated at several local caches, then several copies of the same data item will be returned to the requesting data manager, when only one is required.

For very large multiprocessor systems, even this *one-to-many* approach to discovering the unknown location of a data item may be too costly in terms of communication latency and its contribution to message density within the system. A compromise of the direct access capabilities of the pre-allocated resident set approach and the flexibility of the dynamic composition of the local caches is the notion of a *directory* of data item locations.

In this approach, it is not necessary to maintain a particular data item at a fixed processing element. We can introduce the notion of a *home*-processing element that knows where that data item is, while the data item is currently located at the *owner*-processing element. The home-processing element is fixed and its address may be determined from the identifier of the data item. The home processing element knows which processing element is currently owning the data item. Should this data item be subsequently moved and the one at the owner-process removed, then either the home-processing element must be informed as to the new location of the data item or the previous owner-processing element must now maintain a pointer to this

new location. The first scheme has the advantage that a request message may be forwarded directly from the home-processing element to the current owner, while the second strategy may be necessary, at least for a while after the data item has been moved from an owner, to cope with any requests forwarded by the home-processing element before it has received the latest location update.

Finally, it is also possible to do away with the notion of a home-processing element, by adding a hierarchy of directories. Each directory on a processing element “knows” which data items are present on the processing element. If the required data item is not present, a directory higher up in the hierarchy might know if it is somewhere nearby. If that directory does not know, yet another directory might if it is further away. This is much like the organisation of libraries: you first check the local library for a book, if they do not have it you ask the central library, and so on until you finally query the national library. With this organisation there is always a directory that knows the whereabouts of the data item, but it is very likely that the location of the data item will be found long before asking the highest directory. (The Data Diffusion Machine [61] and KSR-1 [38] used a similar strategy implemented in hardware).

## 3.4 Consistency

Copies of *read-only* data items may exist in numerous local caches within the system without any need to “keep track” of where all the copies are. However, if copies of *read-write* data items exist then, in a virtual shared memory system, there is the danger that the data items may become *inconsistent*. The example in figure 3.6 illustrates this problem of inconsistency. Suppose that we have two processing elements  $PE_1$  and  $PE_2$ , and a data item  $\gamma$  with a value 0, that is located at processing element  $PE_1$ . Processing Element  $PE_2$  needs  $\gamma$ , so it requests and gets a copy of  $\gamma$ . The data manager on processing element  $PE_2$  decides to keep this copy for possible future reference. When the application at processing element  $PE_1$  updates the value of  $\gamma$ , for example by overwriting it with the value 1, processing element  $PE_2$  will have a *stale* copy of  $\gamma$ . This situation is called *inconsistent*: if the application running at processing element  $PE_1$  requests  $\gamma$  it will get the new value (1), while the application at processing element  $PE_2$  will still read the old value of  $\gamma$  (0). This situation will exist until the data manager at processing element  $PE_2$  decides to evict  $\gamma$  from its local memory.

The programming model of a physical shared memory system maintains only one copy of any data item; the copy in the shared memory. Because there is only one copy, the data items cannot become inconsistent. Hence, naive *virtual* shared memory differs from *physical* shared memory in that virtual shared memory can become inconsistent.

To maintain consistency all copies of the data items will have to be “tracked down” at certain times during the parallel computation. Once again one-to-one or many-to-one methods could be used to determine the unknown locations of copies of the data items. If the directory approach is used then it will be necessary to maintain a complete “linked list” through all copies of any data item, where each copy knows where the next copy is, or it knows that there are no more copies. A consistency operation is performed on this list by sending a message to the first copy on the list, which then ripples through the list. These operations thus take a time linear in the number of copies. This is expensive if there are many copies, so it can be more efficient to use a tree structure (where the operation needs logarithmic time). (A combination of a software and hardware tree directory of this form is used in the LimitLESS directory [12].)

There are several ways to deal with this inconsistency problem. We will discuss three options: data items are kept consistent at all times (known as sequential consistency); the actual problem somehow copes with the inconsistencies (known as weak consistency); and finally, inconsistent data items are allowed to live for a well defined period (the particular scheme discussed here is known as release consistency).

### 3.4.1 Keeping the data items consistent

The first option is that the data manager will keep the data items consistent at all times. To see how the data items can be kept consistent, observe first that there are two conditions that must be met before a data item can become inconsistent. Firstly, the data item must be duplicated; as long as there is only a single copy of the data item, it cannot be inconsistent. Secondly, some processing element must update one of the copies, without updating the other copies. This observation leads to two protocols that the data manager can observe to enforce consistency, while still allowing copies to be made:

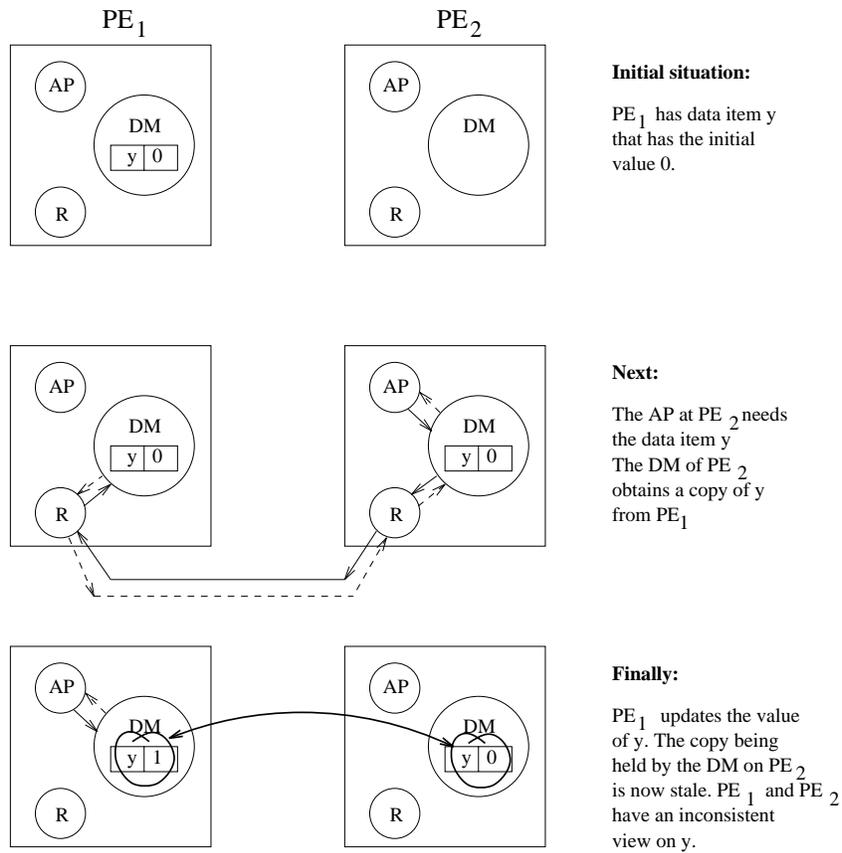


Figure 3.6: An example how an inconsistency arises. There are two processing elements,  $PE_1$  and  $PE_2$  and a data item  $y$ .  $PE_2$  keeps a copy of  $y$ , while  $y$  is updated at  $PE_1$ .

1. Ensure that there is not more than a single copy of the data item when it is updated. This means that before a write all but one of the copies must be deleted. This solution is known as an *invalidating protocol*.
2. Ensure that all copies of the data item are replaced when it is updated. This solution is known as an *updating protocol*.

It is relatively straightforward to check that the invalidating option will always work: all copies are always identical, because a write only occurs when there is only a single copy. In the example, the copy of  $y$  at processing element  $PE_2$  will be destroyed before  $y$  is updated on processing element  $PE_1$ .

For the updating protocol to be correct, the protocol must ensure that all copies are replaced “at the same time”. Suppose that this is not the case: in the example the value on processing element  $PE_1$  might be updated, while processing element  $PE_2$  still has an old value for  $y$ . If the data managers running on processing elements  $PE_1$  and  $PE_2$  communicate, they can find out about this inconsistency. In order for the update protocol to work, the updating data manager must either ensure that no other data manager is accessing the data item while it is being updated, or that it is impossible for any communication (or other update) to overtake this update.

It is not easy to decide in general whether an invalidating or an updating protocol is better. Below are two examples that show that invalidating and updating protocols both have advantages and disadvantages. In both cases we assume that the problem is running on a large number of processing elements, and that there is a single shared data item that is initially replicated over all processing elements.

1. A task, or tasks, being performed by an application process at one processing element might require that the data item be updated at this data manager many times, without any of the other processing elements using it. An updating protocol will update all copies on all processing elements during every update, even though the copies are not being used on any of the other processing elements.

An invalidating protocol is more efficient, because it will invalidate all outstanding copies once, whereupon the application process can continue updating the data item without extra communication.

2. Suppose that instead of ignoring the data item, all other processing elements do need the updated value. An invalidating protocol will invalidate all copies and update the data item, whereupon all other processing elements have to fetch the value again. This fetch is on demand, which means that they will have to wait on the data item.

An updating protocol does a better job since it distributes the new value, avoiding the need for the other processing elements to wait for it.

There is a case for (and against) both protocols. It is for this reason that these two protocols are sometimes combined. This gives a protocol that, for example, invalidates all copies that have not been used since the last update, and updates the copies that were used since the last update. Although these hybrid protocols are potentially more efficient, they are unfortunately often more complex than a pure invalidating or updating protocol.

### 3.4.2 Weak consistency: repair consistency on request

The option to maintain sequential consistency is an expensive one. In general, an application process is allowed to proceed with its computation only after the invalidate or update has been completed. In the example of the invalidating protocol, all outstanding copies must have been erased and the local copy must have been updated before the application process can proceed. This idle time may be an unacceptable overhead. One of the ways to reduce this overhead is to forget about maintaining consistency automatically. Instead, the local cache will stay inconsistent until the application process orders the data manager to repair the inconsistency.

There are two important advantages of weak consistency. Firstly, the local cache is made consistent at certain points in the task execution only, reducing the overhead. Secondly, local caches can be made consistent in parallel. Recall for example, the updating protocol of the previous section. In a weakly consistent system we can envisage that every write to a data item is asynchronously broadcasted to all remote copies. Asynchronously means that the processing element performing the write continues whether the update has been completed or not. Only when a consistency-command is executed must the application process wait

until all outstanding updates are completed. In the same way, a weakly consistent invalidating protocol can invalidate remote copies in parallel. These optimisations lead to further performance improvement. The disadvantage of weak consistency is the need for the explicit commands within the algorithm at each application process so that when a task is being executed, at the appropriate point, the data manager can be instructed to make the local cache consistent.

### 3.4.3 Repair consistency on synchronisation: Release consistency

A weak consistency model as sketched above requires the programmer of the algorithm to ensure consistency at any moment in time. Release consistency is based on the observation that algorithms do not go from one phase to the other without first synchronising. So it suffices to make the local caches consistent during the synchronisation operation. This means that immediately after each synchronisation the local caches are guaranteed to be consistent. This is in general slightly more often than strictly necessary, but it is far less often than would be the case when using sequential consistency. More importantly, the application process itself does not have to make the local caches consistent anymore, it is done “invisibly”.

Note that although invisible, consistency is only restored during an explicit synchronisation operation; release consistency behaves still very differently from sequential consistency. As an example, an application process at  $PE_1$  can poll a data item in a loop, waiting for the data item to be changed by the application process at  $PE_2$ . Under sequential consistency any update to the data item will be propagated, and cause the application process at  $PE_1$  to exit the loop. Under release consistency updates do not need to be propagated until a synchronisation point, and because it does not recognise that the polling loop is actually a synchronisation point the application process at  $PE_1$  might be looping forever.

## 3.5 Minimising the Impact of Remote Data Requests

Failure to find a required data item locally means that the data manager has to acquire this data item from elsewhere within the system. The time to fetch this data item and, therefore, the application process idle time, can be significant. This *latency* is difficult to predict and may not be repeatable due to other factors, such as current message densities within the system. The overall aim of data management is to maximise effective processing element computation by minimising the occurrence and effects of remote data fetches. A number of techniques may be used to reduce this latency by:

**Hiding the Latency:** - overlapping the communication with the computation, by:

**Prefetching** - anticipating data items that will be required

**Multi-threading** - keeping the processing element busy with other useful computation during the remote fetch

**Minimising the Latency:** - reducing the time associated with a remote fetch by:

**Caching & profiling** - exploiting any coherence that may exist in the problem domain

### 3.5.1 Prefetching

If it is known at the start of the computation which data items will be required by each task then these data items can be *prefetched* by the data manager so that they are available locally when required. The data manager thus issues the requests for the data items *before* they are actually required and in this way overlaps the communication required for the remote fetches with the ongoing computation of the application process. This is in contrast with the simple fetch-upon-demand strategy where the data manager only issues the external request for a data item at the moment it is requested by the application process and it is not found in the local cache.

By treating its local cache as a “circular buffer” the data manager can be loading prefetched data items into one end of the buffer while the application process is requesting the data items from the other end, as shown in figure 3.7. The “speed” at which the data manager can prefetch the data items will be determined by the size of the local cache and the rate at which the application process is “using” the data items.

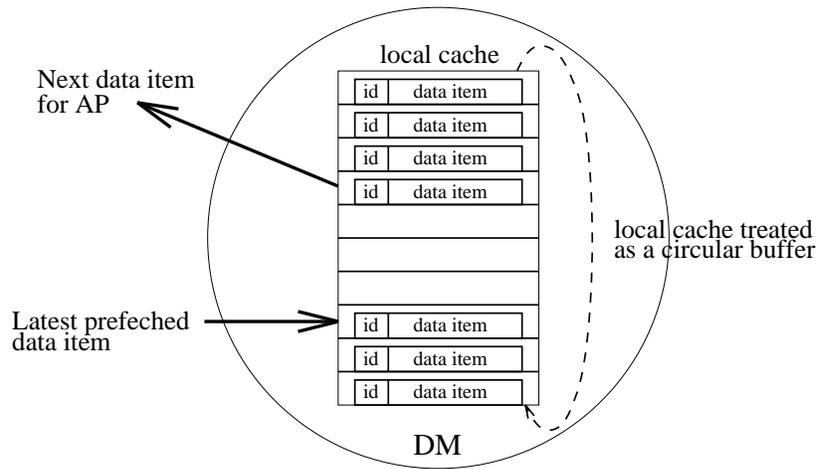


Figure 3.7: Storing the prefetched data items in the local cache

This knowledge about the data items may be known *a priori* by the nature of problem. For example, in a parallel solution of the hemi-cube radiosity method, the data manager knows that each task, that is the computation of a single row of the matrix of form factors, requires all the environment's patch data. The order in which these data items are considered is unimportant, as long as all data items are considered. The data manager can thus continually prefetch those data items which have yet to be considered by the current task. Note that in this problem, because all the data items are required by every task and the order is unimportant (we are assuming that the local cache is not sufficiently big to hold all these data items), those data items which remain in the local cache at the end of one task are also required by the subsequent task. Thus, at the start of the next task, the first data item in the local cache can be forwarded to the application process and prefetching can commence once more as soon as this has happened.

The choice of computation model adopted can also provide the information required by the data manager in order to prefetch. The principal data items for both the balanced and unbalanced data driven models will be known by the system controller before the computation commences. Giving this information to the data manager will enable it to prefetch these data items. A prefetch strategy can also be used for principal data items within the preferred bias task allocation strategy for the demand driven computation model, as described in section 2.4.5. Knowledge of its processing element's conceptual region can be exploited by the data manager to prefetch the principal data items within this region of the problem domain.

### 3.5.2 Multi-threading

Any failure by the data manager to have the requested data item available locally for the application process will result in idle time unless the processing element can be kept busy doing some other useful computation.

One possibility is for the application process to save the current state of a task and commence a new task whenever a requested data item is not available locally. When the requested data item is finally forthcoming either this new task could be suspended and the original task resumed, or processing of the new task could be continued until it is completed. This new task may be suspended awaiting a data fetch and so the original task may be resumed. Saving the state of a task may require a large amount of memory and indeed, several states may need to be saved before one requested data item finally arrives. Should the nature of the problem allow these stored tasks to in turn be considered as task packets, then this method has the additional advantage that these task packets could potentially be completed by another processing element in the course of load balancing, as explained in the section 2.4.4 on distributed task management.

Another possible option is multi-threading. In this method there is not only one, but several application processes on each processing element controlled by an application process controller (APC), as shown in figure 3.8. Each application process is known as a separate *thread* of computation. Now, although one thread may be suspended awaiting a remote data item, the other threads may still be able to continue. It

may not be feasible to determine just how many of these application processes will be necessary to avoid the case where all of them are suspended awaiting data. However, if there are sufficient *threads* (and of course sufficient tasks) then the processing element should always be performing useful computation. Note that multi-threading is similar to the Bulk Synchronous Parallel paradigm [60].

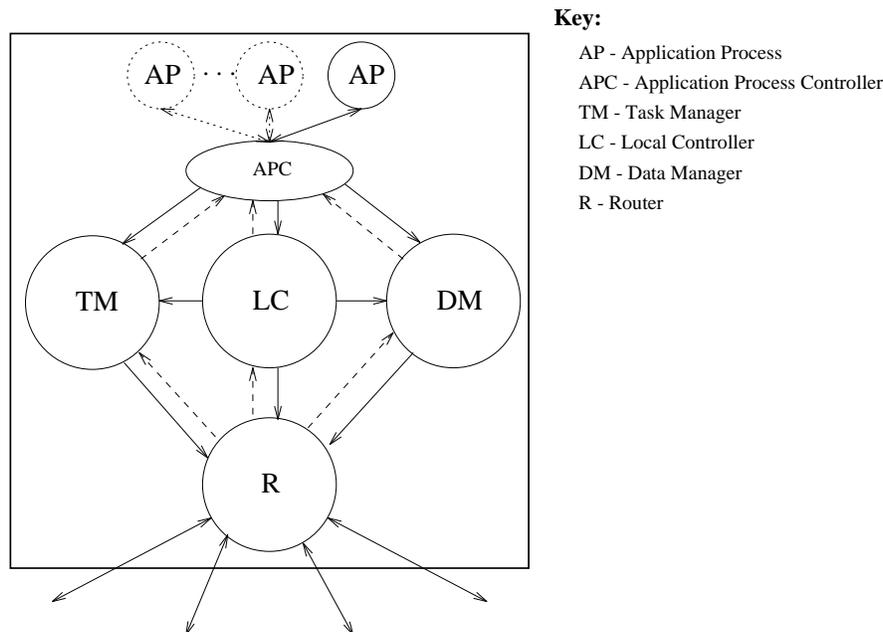


Figure 3.8: Several application processes per processing element

One disadvantage of this approach is the overhead incurred by the additional context switching between all the application processes and the application process controller, as well as the other system software processes: the router, task manager and the data manager, that are all resident on the same processor. A variation of multiple active threads is to have several application processes existing on each processing element, but to only have *one* of them active at any time and have the application process controller manage the scheduling of these processes explicitly from information provided by the data manager. When an application processes' data item request cannot be satisfied locally, that process will remain descheduled until the data item is forthcoming. The data manager is thus in a position to inform the application process controller to activate another application process, and only reactivate the original application process once the required data has been obtained. Note the application process controller schedules a new application process by sending it a task to perform. Having made its initial demand for a task to the application process controller (and not the task manager as discussed in section 2.4.2) an application process will remain descheduled until explicitly rescheduled by the application process controller.

Both forms of multi-threading have other limitations. The first of these is the extra memory requirements each thread places on the processing elements local memory. The more memory that each thread will require, for local constants and variables etc, the less memory there will be available for the local cache and thus fewer data items will be able to be kept locally by the data manager. A "catch 22" (or is that "cache 22") situation now arises as fewer local data items implies more remote data fetches and thus the possible need for yet more threads to hide this increase in latency. The second difficulty of a large number of threads running on the same processing element is the unacceptably heavy overhead that may be placed on the data manager when maintaining the local cache. For example, a dilemma may exist as to whether a recently fetched data item for one thread should be overwritten before it has been used if its "slot" in the local cache is required by the currently active thread.

Figure 3.9 shows results for a multi-threaded application. The graph shows the time in seconds to solve a complex parallel ray tracing problem with large data requirements using more than one application process per processing element. As can be seen, increasing the number of application processes per processing

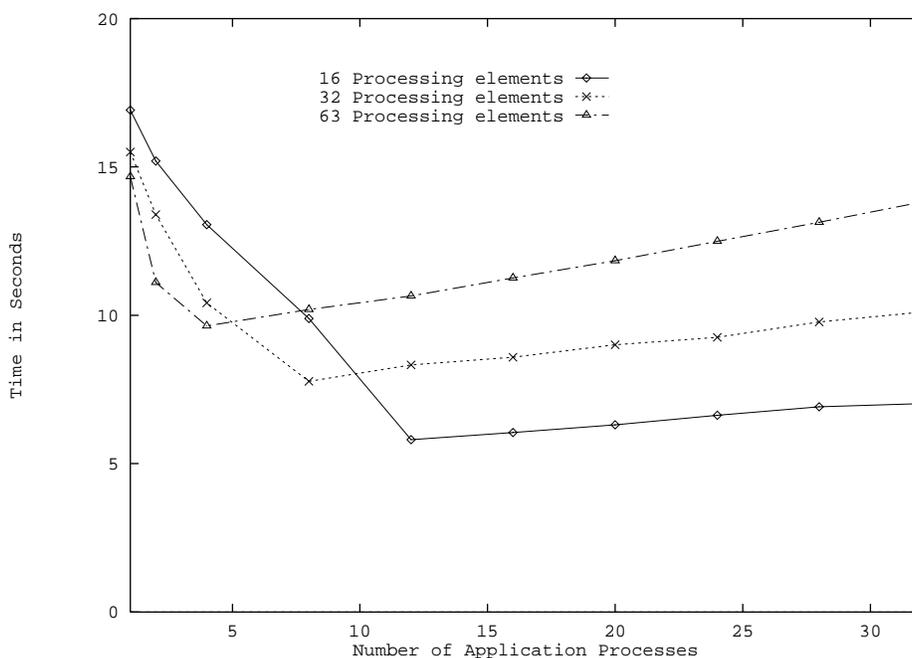


Figure 3.9: Problem solution time in seconds

element produces a performance improvement until a certain number of threads have been added. Beyond this point, the overheads of having the additional threads are greater than the benefit gained, and thus the times to solve the problem once more increase. The number of threads at which the overheads outweigh the benefits gained is lower for larger numbers of processing elements. This is because the more application processes there are per processing element, the larger the message output from each processing element will be (assuming an average number of remote fetches per thread). As the average distances the remote data fetches have to travel in larger systems is greater, the impact of increasing numbers of messages on the overall system density is more significant and thus the request latency will be higher. Adding more threads now no longer helps overcome communication delays, but in fact, the increasing number of messages actually exacerbates the communication difficulties. Ways must be found of dynamically scheduling the optimum number of application processes at each processing element depending on the current system message densities.

Despite these shortcomings, multi-threading does work well, especially for low numbers of threads and is a useful technique for avoiding idle time in the face of unpredictable data item requirements. Remember that multiple threads are only needed at a processing element if a prefetch strategy is not possible and the data item required by one thread was not available locally. If ways can be found to try and guess which data items are likely to be required next then, if the data manager is right at least some of the time, the number of remote fetches-on-demand will be reduced. Caching and profiling assist the data manager with these predictions.

### 3.5.3 Profiling

Although primarily a task management technique, profiling is used explicitly to assist with data management, and so is discussed here. At the start of the solution of many problems, no knowledge exists as to the data requirements of any of the tasks. (If this knowledge did exist then a prefetching strategy would be applicable). Monitoring the solution of a single task provides a list of all the data items required by that task. If the same monitoring action is carried for all tasks then at the completion of the problem, a complete “picture” of the data requirements of all tasks would be known. Profiling attempts to predict the

data requirements of future tasks from the list of data requirements of completed tasks.

Any spatial coherence in the problem domain will provide the profiling technique with a good estimate of the future data requirements of those tasks from a similar region of the problem domain. The data manager can now use this profiling information to *prefetch* those data items which are *likely* to be used by subsequent tasks being performed at that processing element. If the data manager is always correct with its prediction then profiling provides an equivalent situation to prefetching in which the application process is never delayed awaiting a remote fetch. Note in this case there is no need for multi-threading.

A simple example of spatial coherence of the problem domain is in shown in figure 3.10. This figure is derived from figure 2.3 which showed how the principal data item (PDI) and additional data items (ADIs) made up a task. In figure 3.10 we can see that task *i* and task *j* come from the same region of the problem domain and spatial coherence of the problem domain has meant that these two tasks have three additional data items in common. Task *k*, on the other hand, is from a different region of the problem domain, requires only one additional data item which is not common to either task *i* or task *j*.

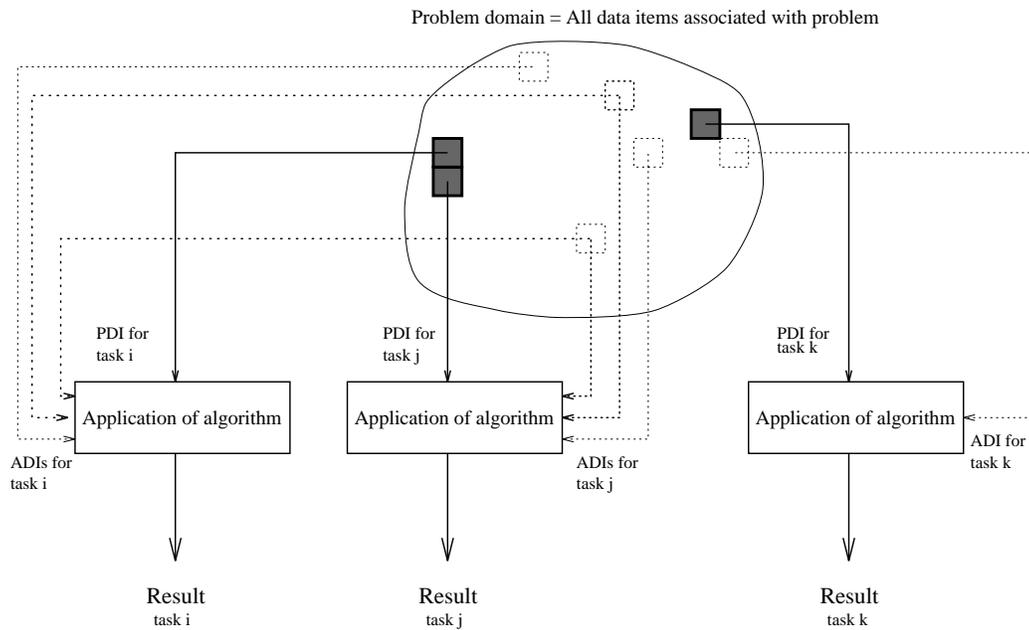


Figure 3.10: Common additional data items due to spatial coherence of the problem domain

Thus, the more successful the predictions are from the profiling information, the higher will be the cache-hit ratios. From figure 3.10 on page 65 we can see that if the completion of task *i* was used to profile the data item requirements for task *j* then, thanks to the spatial coherence of task *i* to task *j* in the problem domain, the data manager would have a 66% success rate for the additional data items for task *j*. However, a similar prediction for the additional data items for task *k* would have a 0% success rate and result in a 100% cache-miss, that is all the additional data items for task *k* would have to be fetched-on-demand.

### 3.6 Data Management for Multi-Stage Problems

In section 2.3.3 we discussed the algorithmic and data dependencies that can arise in problems which exhibit more than one distinct stage. In such problems, the results from one stage become the principal data items for the subsequent stage as was shown in figure 2.12. So, in addition to ensuring the application processes are kept supplied with data items during one stage, the data manager also needs to be aware as to how the partial results from one stage of the computation are stored at each processing element in anticipation of the following stage.

This balancing of partial result storage could be achieved statically by all the results of a stage being returned to the system controller. At the end of that current stage the system controller is in a position to

distribute this data evenly as the principal and additional data items for the next stage of the problem. The communication of these potentially large data packets twice, once during the previous stage to the system controller and again from the system controller to specific processing elements, obviously may impose an enormous communication overhead. A better static distribution strategy might be to leave the results in place at the processing elements for the duration of the stage and then have them distributed from the processing elements in a manner prescribed by the system controller. Note that in such a scheme the local cache of each processing element must be able to hold not only the principal and additional data items for the current stage, but also have space in which to store these partial results in anticipation of the forthcoming stage. It is important that these partial results are kept separate so that they are not inadvertently overwritten by data items during the current stage.

In a demand driven model of computation the uneven computational complexity may result in a few processing elements completing many more tasks than others. This produces a flaw in the second static storage strategy. The individual processing elements may simply not have sufficient space in their local cache to store more than their fair share of the partial results until the end of the stage.

Two dynamic methods of balancing this partial result data may also be considered. Adoption of the preferred bias task management strategy, as discussed in section 2.4.5, can greatly facilitate the correct distribution of any partial results. Any results produced by one processing element from another's conceptual portion, due to task load balancing, may be sent directly to this other processing element. The initial conceptual allocation of tasks ensures that the destination processing element will have sufficient storage for the partial result.

If this conceptual allocation is not possible, or not desirable, then balancing the partial results dynamically requires each processing element to be kept informed of the progress of all other processing elements. This may be achieved by each processing element broadcasting a short message on completion of every task to all other processing elements. To ensure that this information is as up to date as possible, it is advisable that these messages have a special high priority so that they may be handled immediately by the router processes, by-passing the normal queue of messages. Once a data manager's local cache reaches its capacity the results from the next task are sent in the direction of the processing element that is known to have completed the least number of tasks and, therefore, the one which will have the most available space. To further reduce the possible time that this data packet may exist in the system, any processing element on its path which has storage capacity available may absorb the packet and thus not route it further.

# Bibliography

- [1] W. B. Ackerman. Data flow languages. In N. Gehani and A. D. McGettrick, editors, *Concurrent Programming*, chapter 3, pages 163–181. Addison-Wesley, 1988.
- [2] G. S. Almasi and A. Gottlieb. *Highly Parallel Computing*. Benjamin/Cummings, Redwood City, California, 2<sup>nd</sup> edition, 1994.
- [3] G. M. Amdahl. Validity of the single-processor approach to achieving large scale computing capabilities. In *AFIPS*, volume 30, Atlantic City, Apr. 1967. AFIPS Press, Reston, Va.
- [4] M. Annaratone et al. Warp architecture and implementation. In *13<sup>th</sup> Annual International Symposium on Computer Architecture*, pages 346–356, Tokyo, June 1986.
- [5] J. Backus. Can programming be liberated from the von Neumann style functional style and its algebra of programs. *Communications of the ACM*, 21(8):613–641, 1978.
- [6] H. Bal. *Programming Distributed Systems*. Silicon Press, Summit, New Jersey, 1990.
- [7] A. Basu. A classification of parallel processing systems. In *ICCD*, 1984.
- [8] K. Batcher. Design of a massively parallel processor. *IEEE Transactions on Computers*, 29(9):836–840, Sept. 1980.
- [9] M. Ben-Ari. *Principles of Concurrent and Distributed Programming*. Addison-Wesley, Wokingham, England, 1990.
- [10] A. W. Burks. Programming and structural changes in parallel computers. In W. Händler, editor, *Conpar*, pages 1–24, Berlin, 1981. Springer.
- [11] N. Carriero and D. Gelernter. *How to Write Parallel Programs*. MIT Press, Cambridge, Massachusetts, 1990.
- [12] D. Chaiken, J. Kubiawicz, and A. Agarwal. LimitLESS directories: A scalable cache coherence scheme. In *Proceedings of the 4<sup>th</sup> International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS-IV*, pages 224–234, Apr. 1991.
- [13] A. G. Chalmers. Occam - the language for educating future parallel programmers? *Microprocessing and Microprogramming*, 24:757–760, 1988.
- [14] A. G. Chalmers and D. J. Paddon. Communication efficient MIMD configurations. In *4<sup>th</sup> SIAM Conference on Parallel Processing for Scientific Computing*, Chicago, 1989.
- [15] P. Chaudhuri. *Parallel Algorithms: Design and analysis*. Prentice-Hall, Australia, 1992.
- [16] B. Codenotti and M. Leonici. *Introduction to parallel processing*. Addison-Wesley, Wokingham, England, 1993.
- [17] A. L. DeCegama. *The Technology of Parallel Processing: Parallel Processing Architectures and VLSI Design*. Prentice-Hall International Inc., 1989.
- [18] D. L. Eager, J. Zahorjan, and E. D. Lazowska. Speedup versus efficiency in parallel systems. *IEEE Transactions on Computers*, 38(3):408–423, Mar. 1989.

- [19] V. Faber, O. M. Lubeck, and A. B. White Jr. Super-linear speedup of an efficient sequential algorithm is not possible. *Parallel Computing*, 3:259–260, 1986.
- [20] H. P. Flatt and K. Kennedy. Performance of parallel processors. *Parallel Computing*, 12:1–20, 1989.
- [21] M. J. Flynn. Some computer organisations and their effectiveness. *IEEE Transactions on Computers*, 21(9):948–960, 1972.
- [22] C. F. Gerald and P. O. Wheatley. *Applied numerical analysis*. World Student Series. Addison-Wesley, Reading, MA, 5<sup>th</sup> edition, 1994.
- [23] S. A. Green and D. J. Paddon. A non-shared memory multiprocessor architecture for large database problems. In M. Cosnard, M. H. Barton, and M. Vanneschi, editors, *Proceedings of the IFIP WG 10.3 Working Conference on Parallel Processing*, Pisa, 1988.
- [24] H. A. Grosch. High speed arithmetic: The digital computer as a research tool. *Journal of the Optical Society of America*, 43(4):306–310, Apr. 1953.
- [25] H. A. Grosch. Grosch’s law revisited. *Computerworld*, 8(16):24, Apr. 1975.
- [26] J. L. Gustafson. Re-evaluating Amdahl’s law. *Communications of the ACM*, 31(5):532–533, May 1988.
- [27] D. R. Hartree. The ENIAC, an electronic computing machine. *Nature*, 158:500–506, 1946.
- [28] J. L. Hennessy and D. A. Patterson. *Computer Architecture: A quantitative approach*. Morgan Kaufmann, San Mateo, CA, 1990.
- [29] T. Hey. Scientific applications. In G. Harp, editor, *Transputer Applications*, chapter 8, pages 170–203. Pitman Publishing, 1989.
- [30] D. W. Hillis. *The Connection Machine*. The MIT Press, 1985.
- [31] R. W. Hockney and C. R. Jesshope. *Parallel Computers 2: Architecture, Programming and Algorithms*, chapter 1, pages 60–81. Adam Hilger, 1988.
- [32] R. W. Hockney and C. R. Jesshope. *Parallel Computers 2: Architecture, Programming and Algorithms*. Adam Hilger, Bristol, 1988.
- [33] M. Homewood, M. D. May, D. Shepherd, and R. Shepherd. The IMS T800 transputer. *IEEE Micro*, pages 10–26, 1987.
- [34] R. M. Hord. *Parallel Supercomputing in MIMD Architectures*. CRC Press, Boca Raton, 1993.
- [35] R. J. Hosking, D. C. Joyce, and J. C. Turner. *First steps in numerical analysis*. Hodder and Stoughton, London, 1978.
- [36] HPF Forum. High Performance Fortran language specification. *Scientific Programming*, 2(1), June 1993.
- [37] K. Hwang. *Advanced Computer Architecture*. McGraw-Hill Series in Computer Science. McGraw-Hill, New York, 1993.
- [38] KSR. *KSR Technical Summary*. Kendall Square Research, Waltham, MA, 1992.
- [39] V. Kumar, A. Grama, A. Gupta, and G. Karyps. *Introduction to Parallel Computing*. Benjamin/Cummings, Redwood City, California, 1994.
- [40] H. T. Kung. *VLSI array processors*. Prentice-Hall, Englewood Cliffs, NJ, 1988.
- [41] H. T. Kung and C. E. Leiserson. Systolic arrays (for VLSI). In Duff and Stewart, editors, *Sparse Matrix proceedings*, Philadelphia, 1978. SIAM.
- [42] C. Lazou. *Supercomputers and Their Use*. Clarendon Press, Oxford, revised edition, 1988.

- [43] D. Lenoski, J. Laudon, T. Joe, D. Nakahira, L. Stevens, A. Gupta, and J. Hennessy. The DASH prototype: Logic overhead and performance. *IEEE Transactions on Parallel and Distributed Systems*, 4(1):41–61, Jan. 1993.
- [44] T. Lewis and H. El-Rewini. *Introduction to parallel computing*. Prentice-Hall, 1992.
- [45] K. Li. Ivy: A shared virtual memory system for parallel computing. *Proceedings of the 1988 International Conference on Parallel Processing*, 2:94–101, Aug. 1988.
- [46] G. J. Lipovski and M. Malek. *Parallel Computing: Theory and comparisons*. John Wiley, New York, 1987.
- [47] M. D. May and R. Shepherd. Communicating process computers. Inmos technical note 22, Inmos Ltd., Bristol, 1987.
- [48] L. F. Menabrea and A. Augusta(translator). Sketch of the Analytical Engine invented by Charles Babbage. In P. Morrison and E. Morrison, editors, *Charles Babbage and his Calculating Engines*. Dover Publications, 1961.
- [49] D. Nussbaum and A. Argarwal. Scalability of parallel machines. *Communications of the ACM*, 34(3):56–61, Mar. 1991.
- [50] B. Purvis. Programming the Intel i860. *Parallelogram International*, pages 6–9, Oct. 1990.
- [51] M. J. Quinn. *Parallel Computing: Theory and practice*. McGraw-Hill, New York, 1994.
- [52] V. Rajaraman. *Elements of parallel computing*. Prentice-Hall of India, New Dehli, 1990.
- [53] S. F. Reddaway. DAP - a Distributed Array Processor. In *1<sup>st</sup> Annual Symposium on Computer Architecture*, 1973.
- [54] R. M. Russel. The CRAY-1 computer system. *Communications of the ACM*, 21:63–72, 1978.
- [55] J. E. Shore. Second thoughts on parallel processing. *Comput. Elect. Eng.*, 1:95–109, 1973.
- [56] R. J. Swam, S. H. Fuller, and D. P. Siewiorek. ‘Cm\*—A Modular, Multi-Microprocessor’. In *Proc. AFIPS 1977 Fall Joint Computer Conference 46*, pages 637–644, 1977.
- [57] S. Thakkar, P. Gifford, and G. Fiellamd. The Balance multiprocessor system. *IEEE Micro*, 8(1):57–69, Feb. 1988.
- [58] P. C. Treleaven, D. R. Brownbridge, and R. P. Hopkins. Data driven and demand-driven computer architecture. *Communications of the ACM*, 14(1):95–143, Mar. 1982.
- [59] A. Trew and G. Wilson, editors. *Past, Present and Parallel: A survey of available parallel computer systems*. Springer-Verlag, London, 1991.
- [60] L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, Aug. 1990.
- [61] D. H. D. Warren and S. Haridi. The Data Diffusion Machine—a scalable shared virtual memory multiprocessor. In *Proceedings of the 1988 International Conference on Fifth Generation Computer Systems*, pages 943–952, Tokyo, Japan, Dec. 1988.
- [62] R. P. Weicker. Dhrystone: A synthetic systems programming benchmark. *Communications of the ACM*, 27(10):1013–1030, 1984.

**Section II**  
**Classification of Parallel Rendering**  
**Systems**

**Tim Davis**

# Section II Classification of Parallel Rendering Systems

In the preceding chapters, we have reviewed the fundamental concepts of parallel processing and given some indication of how it might be effectively used in graphics rendering. Since many types of parallel rendering have been investigated [Gree91] [Whit92], classifying the various schemes is important to characterize the behavior of each. A parallel rendering system can be classified according to the method of task subdivision and/or by the hardware used to implement the scheme. Often, the choice of one influences the other.

Classifying by task subdivision refers to the method in which the original rendering task is broken into smaller pieces to be processed in parallel. Obviously, such subdivision strongly depends on the type of rendering employed. A task for rendering polygons will offer a different set of subdivision opportunities than a ray tracing task. Also included in these decisions is the type of load balancing technique to employ.

Ultimately the rendering scheme is implemented within some sort of parallel environment. The system may run on parallel hardware (e.g., a general multiprocessor or specialized hardware) or in a distributed computing environment (a group of individual machines working together to solve a single problem). The advantages and disadvantages associated with each environment are discussed below.

## II.1 Classification by Task Subdivision

In this section, we will look at two different types of rendering (polygon-based rendering and ray tracing) and various methods for subdividing the original task into subtasks for parallel processing. Although many subdivision techniques exist for each, we will focus on the schemes most widely used. For each technique, recall that our goal is to subdivide the original task in such a way as to maximize parallelism, while not creating excessive overhead.

### II.1.1 Polygon Rendering

For polygon rendering, we often deal with a very large number of primitives (e.g., triangles)

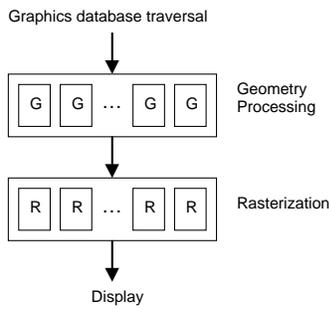


Figure 1: the graphics pipeline

which can often be processed in a parallel manner. To handle this type of rendering, a graphics pipeline is usually employed (see figure X-1). Stages in this pipeline include geometry processing and rasterization.

Geometry processing comprises transformation, clipping, lighting, and other tasks associated with a primitive. A straightforward method for parallelizing geometry processing is to assign each processor a subset of primitives (or objects) in the scene to render. In rasterization, scan-conversion, shading, and visibility determination are performed. To parallelize this processing, each processor could perform the pixel calculations for a small part of the final image.

One way to view the processing of primitives is as a problem of sorting primitives to the

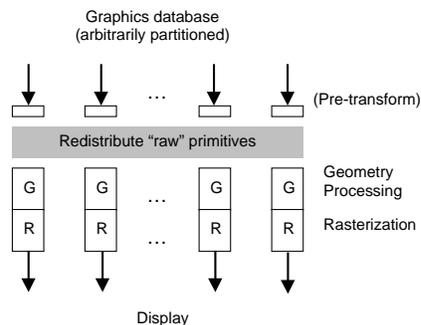


Figure 2: sort-first polygon rendering

screen since a graphics primitive can fall anywhere on or off the screen [Moln94]. For a parallel system, we need to distribute data across processors to keep the load balanced. Actually, this sort can occur anywhere in the rendering pipeline:

- during geometry processing (sort-first)
- between geometry processing and rasterization (sort-middle)
- during rasterization (sort-last)

The structure of the parallel rendering system is determined by the location of this sort. The following discussion follows that in [Moln94].

### II.1.1.1 Sort-First

The main idea behind sort-first is to distribute primitives early (during geometry processing) in the rendering pipeline (see figure X-2). The screen is divided into regions of equal size (see figure X-3), and each processor (or renderer) is assigned a region. Each processor is responsible for all the pixel calculations that are associated with its screen region.

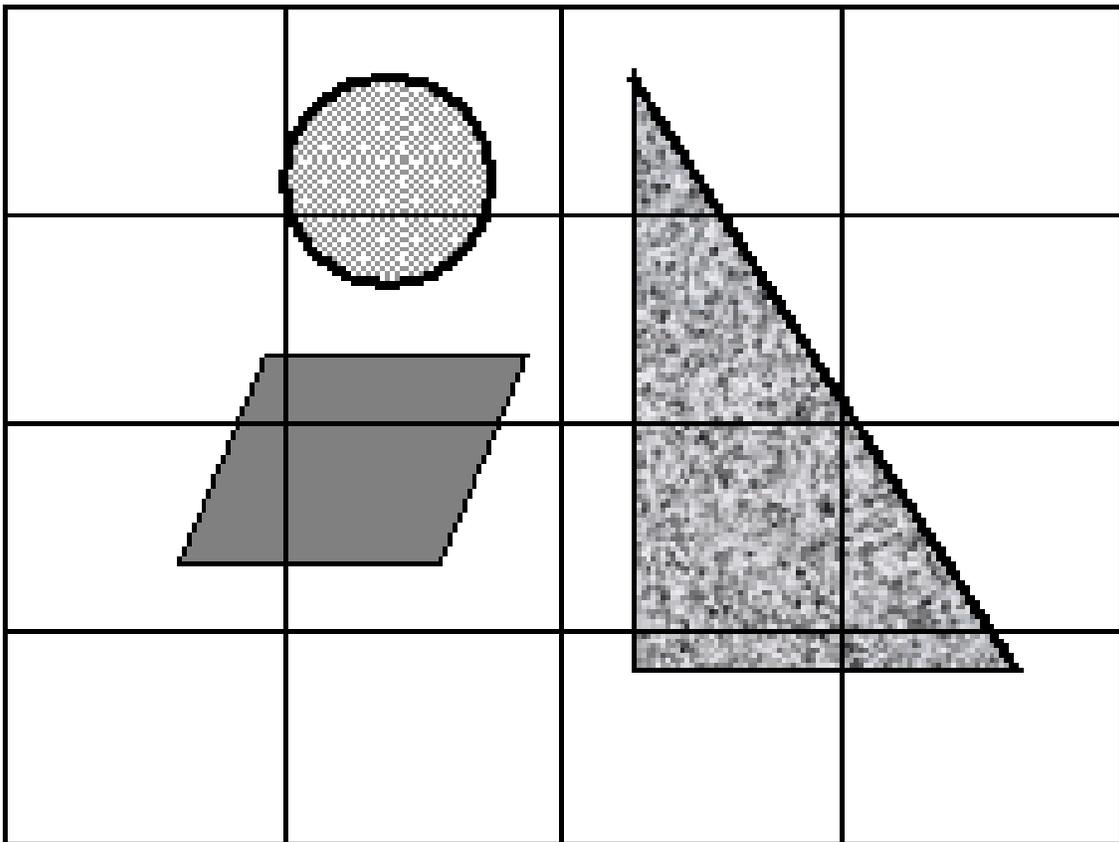


Figure 3: image subdivision

In an actual implementation, primitives would initially be assigned to processors in an arbitrary way. Each renderer then performs enough transformation processing to determine the screen region into which the primitive falls. If this region belongs to another processor, the primitive is sent over the interconnection network to that processor for rendering. After each primitive has been placed with the proper renderer, all of the processors can work in parallel to complete the final image.

With this method, each processor implements the entire rendering pipeline for its portion of the screen. Communication costs can be kept comparatively low compared with other methods if features such as frame coherence are properly exploited. For rendering a single frame, however, almost all the primitives will have to be redistributed after the initial random assignment. Some duplication of effort may occur if a primitive falls into more than one region, or if the results of the original geometry processing are not sent with the primitive when it is transmitted to the appropriate renderer. Also, the system is susceptible to load imbalance since primitives may be concentrated into particular regions or may simply take longer to render. Both of these situations will cause the affected processor to consume more time in processing its screen region. Very few, if any, sort-first renderers have been built.

### II.1.1.2 Sort-Middle

In a sort-middle renderer, primitives are sorted and redistributed in the middle of the pipeline: between geometry processing and rasterization (see figure X-4). By this point, the screen coordinates of the primitives have been determined through transformation processing, but the primitives have not yet been rasterized. This point is a natural breaking position in the rendering pipeline.

In an actual implementation, primitives are arbitrarily assigned to processors as before. The

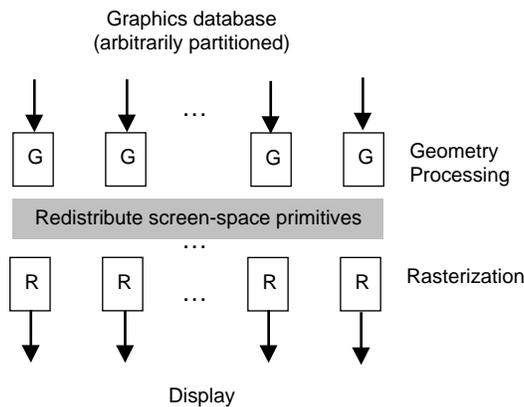


Figure 4: sort-middle polygon rendering

geometry processors perform transformation, lighting, and other processing on the primitives originally assigned to them, and then classify the primitives according to screen region. Screen-space primitives are then sent to the appropriate rasterizing processors, which have been assigned unique regions of the screen, for the remaining processing.

The sort-middle strategy is general and straightforward and has been implemented in both hardware (including Pixel-Planes 5 [Fuch89] and the SGI Reality Engine [Akel93]) and software [Whit94] [Ells94]. Like sort-first, however, this method is susceptible to load imbalance due to the uneven distribution of primitives across screen space. Communication times can be higher under certain conditions. Also, primitives that overlap regions may require some additional processing.

### **II.1.1.3 Sort-Last**

Under the sort-last strategy, sorting is deferred until the end of the rendering pipeline (see figure X-5). Primitives are rasterized into pixels, samples, or pixel fragments, which are then transmitted to the appropriate processor for visibility determination.

In practice, primitives are initially distributed to processors in an arbitrary manner, as in the other methods. Each renderer then performs the operations necessary to compute the pixel values for its primitives, regardless of where those pixels may reside on the screen. These values are then sent to the appropriate processors according to screen location. At this point, the rasterizing processors perform visibility calculations and composite the pixels for final display.

As in sort-first, each processor implements the entire graphics pipeline for its primitives. While the overall technique is less prone to load imbalance, the pixel traffic in the final sort can be very high. Numerous rendering systems using the sort-last method have been constructed in various forms, including [Evan92] and [Kubo93].

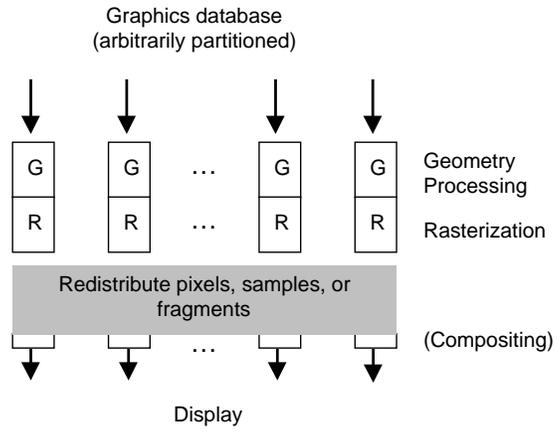


Figure 5: sort-last polygon rendering

## **II.1.2 Ray Tracing**

Ray tracing is a powerful rendering technique that can produce high-quality graphics images; however, this quality comes at a price of intensive calculation and long rendering times. Even relatively simple ray-traced animations can prohibitively expensive to render on a single processor. For longer, more complex animations, the rendering time can be intractable. Fortunately, ray tracing is a prime candidate for parallelization since its processing is readily amenable to subdivision. Specifically, ray tracing inherently contains a large amount of parallelism due to the independent nature of its pixel calculations [Whit80]; therefore, most ray tracing rendering algorithms lend themselves to parallelization in screen space.

Other partitioning schemes are employed in ray tracing as well. Instead of dividing the image space, the object space can be split into smaller regions, or the objects themselves may be assigned to individual processors. These techniques are discussed more fully below.

### **II.1.2.1 Image Space Partitioning**

Using this scheme, the viewing plane is divided into regions, each of which is completely rendered by an individual processor (see figure X-3). That is, for each pixel in a region, the processor assigned to that region computes its entire ray tree. While this technique is conceptually straightforward, the entire database of scene objects must be accessible to every processor.

The benefits of this approach are simplicity and low interprocessor communication as compared with other partitioning methods; the largest drawback is its limitation to multiprocessor architectures with significant local processor memory. Another potential problem is load imbalance, since image detail may be concentrated in certain regions of the screen. To combat this situation, the load balancing algorithm may further subdivide complex regions to provide idle processors with additional tasks.

### **II.1.2.2 Object Space Partitioning**

Here the 3-dimensional space where the scene objects reside is divided into subvolumes, or voxels. Voxels may not be equally sized in order to achieve better load balancing. In the initialization phase of ray tracing, each voxel is parceled out to a particular processor. When rays are cast during rendering, they are passed from processor to processor as they travel through the object space. Each processor, therefore, needs only the scene information associated with its assigned voxels.

While this technique may not suffer from frequent load imbalance, it does incur costs in other ways. First, as new rays are shot, they must be tracked through voxel space; this processing is not required for other schemes. Additionally, since potentially millions of rays are fired for each image, communication could become excessive as rays enter and exit regions of object space during rendering.

### **II.1.2.3 Object Partitioning**

This partitioning scheme parallelizes the rendering task by assigning each object to an individual processor. As in object space partitioning, rays are passed as messages between processors, which in turn test the ray for intersection with the objects they are assigned. Object partitioning also shares some of the benefits and detriments of object space partitioning. Specifically, the load may be fairly well-balanced, but the communication costs may be high due to the large amount of ray message traffic.

### **II.1.2.4 Load Balancing Scheme**

Finally, parallel rendering schemes can be classified according to their load balancing method. Of course, the primary goal of any load balancing scheme is to distribute the work among processors as evenly as possible and thus exploit the highest degree of parallelism available in the application. Many different types of load balancing schemes exist, but each falls into one of two categories:

- *Static Load Balancing.* In this scheme, partitioning is performed up front and processors are assigned subtasks for the entire duration of the rendering process. In this way, overhead is minimized later in the rendering; however, a good deal of care must be taken to ensure that the load will be balanced. Otherwise, the algorithm will suffer from poor performance.
- *Dynamic Load Balancing.* With this scheme, some processing assignments are determined at the start, but later assignments are demand-driven. That is, when a processor determines that it needs more work to do, it will request a new assignment. In this way, processors will never be idle as long as more work is left to do. The key here is to distribute the load as evenly as possible without incurring excessive overhead.

### II.1.2.5 Comments

Hybrid schemes have also been proposed, combining image space partitioning with object space partitioning [Bado94] and image space partitioning with object partitioning [Kim96]. When choosing a partitioning scheme, the architecture of the parallel machine should be considered. For instance, an image space partitioning algorithm will perform better on a MIMD machine than on a SIMD machine. In general, tradeoffs exist between the type of partitioning algorithm used and the architecture chosen.

## II.2 Classification by Hardware

As previously stated, for some computationally intensive rendering tasks, parallel processing provides the only practical means to a solution. One way to perform parallel rendering is to use a single multiprocessor machine, such as a Thinking Machines CM-5, Intel Paragon, Cray T3E, or specialized parallel processor. In these machines, enormous computing power is provided by up to tens of thousands of processing elements able to access many gigabytes of memory and to work in concert through a high-speed interconnection network. Multiprocessors are the most powerful computers in the world and play an active role in solving Grand Challenge problems, such as weather prediction, fluid dynamics, and drug design [Hwan93].

An alternative to using traditional multiprocessor systems for parallel processing is to employ a network of workstations acting as a single machine. This approach, termed distributed or cluster computing, is conceptually similar to a multiprocessor, but each processing element consists of an independent machine connected to a network usually much slower than a multiprocessor interconnection network. While this network can be of any type (e.g., Ethernet, ATM) or topology, the computers connected to it are generally UNIX-based machines which support some type of distributed programming environment, such as Parallel Virtual Machine (PVM) [Geis94] or Message Passing Interface (MPI) [Grop94]. Many types of applications can benefit from distributed computing, including computation-intensive graphics tasks, such as ray tracing [Sung96].

This section focuses on past work that has been documented using traditional multiprocessors and clusters of machines to accomplish graphics rendering tasks, particularly in the area of ray tracing. Included in this discussion is relevant background concerning PVM and MPI, as well as motivation for using these systems in a clustered environment.

## II.2.1 Parallel Hardware

Since rendering consumes such a large amount of computing resources and time, a good deal of effort has gone into exploring parallel solutions on multiprocessor machines. Some of the schemes proposed are designed to run on general-purpose parallel machines, such as the CM-5, while others rely on specialized hardware built especially for ray tracing rendering. A brief survey of these techniques appears in the following sections. Although current research continues in the design and implementation of parallel rendering systems, a flurry of activity in this area occurred in the late 1980s and early 1990s, as reflected in many of the references.

### II.2.1.1 General-Purpose Multiprocessors

In [Plun85], a vectorized ray tracer is proposed for the CDC Cyber 205. In a given execution cycle, rays awaiting processing are distributed to individual processors and ray-object calculations are performed object by object in a lock-step SIMD fashion.

Similarly, [Crow88] implements a SIMD ray tracing algorithm, but for the Connection Machine (CM-2). Image subdivision is used with one pixel being assigned to each of the 16K processors to produce a 128x128 image, with ray-object intersections performed on an object by object basis. The algorithm proposed by [Schr92] also runs on a CM-2, but uses an object-space subdivision coupled with processor remapping capabilities to achieve dynamic load balancing.

Rounding out the SIMD field, [Goel96] describes a ray casting method developed on the MasPar MP-1 for volume rendering, another computation-intensive graphics application used for viewing complex structures in medical imaging and other forms of scientific visualization. To handle the large amount of data and processing involved, machines are assigned portions of the volume to render, which are composited to produce a final image. This system allows users to rotate a volume, magnify areas of interest, and perform other viewing operations.

In the MIMD category, [Reis97] employs an IBM SP/2 running an image-space partitioning scheme with dynamically adjustable boundaries to render frames of an animation progressively. In this form of rendering, termed progressive rendering, an image is initially rendered quickly at low resolution and progressively refined when little or no user interaction takes place. Progressive rendering is useful in interactive environments where frame generation rate is important. The goal of [Keat95] also involves progressive rendering, although their renderer makes use of object-space partitioning on the Kendall Square Research KSR1 machine.

Several research efforts have focused on the Intel iPSC machines as the architectural environment for implementing a parallel ray tracer. Interestingly, whether the partitioning scheme is image-based [Isle91] [Silv94], object space-based [Prio88] [Prio89] or a hybrid of the two [Akti94] [Bado94], the load balancing scheme is almost always of a static nature ([Isle91]

also tests a dynamic scheme). This choice results from a concern that dynamic load balancing schemes produce a large number of messages, which in turn, may dramatically affect the performance of a distributed machine [Prio89].

In the area of transputer-based machines, [Gree90] uses an image subdivision technique combined with memory and cache local to each processor to deal with the many required accesses to the scene description database. Here, the granularity of parallelism is controlled through the size of the image subregion, which also relates directly to the effectiveness of the dynamic load balancing scheme. To render ray-traced animations, [Maur93] use a static object-space partitioning scheme on a system of 36 transputers. Progressive ray tracing and volume rendering on transputers is addressed by [Sous90] and [Pito93], respectively.

### **II.2.1.2 Specialized Multiprocessors**

Probably one of the most noteworthy examples of a specialized multiprocessor for polygon rendering is the series of Pixel-Planes machines developed at UNC-Chapel Hill. The Pixel-Planes 4 machine [Eyle88] is a SIMD machine with three basic components: a host workstation, a graphics processor, and a frame buffer. Each of the customized processors is responsible for a column of display pixels. For its time, it provided good performance; however, the system used processors with slow clock speeds and did not provide effective load balancing.

Pixel-Planes 5 [Fuch89] provided some improvements over Pixel-Planes 4 by incorporating faster processors and employing a more flexible MIMD architecture. The system implemented a sort-middle algorithm, with each processor in charge of a particular region of the screen. To handle the communication, a ring architecture capable of handling eight messages simultaneously is employed. Ultimately, the ring network imposes a limit on scalability.

The PixelFlow machine [Moln92] was developed to overcome the limitations of the previous architectures through parallel image composition. Each individual processor works to create a full-screen image using only the primitives assigned to it. All of these images are collected and composited to form the final display.

For ray tracing, [Lin91] employs a specialized SIMD machine to perform stochastic ray tracing. The stochastic method adds extra processing to the ray tracing algorithm to handle antialiasing, an important aspect of any renderer. To overcome some of the inefficiencies found in other SIMD approaches, a combination of image space partitioning and object space partitioning is used. That is, a block of pixels is rendered by casting rays and using scene coherence to restrict the parts of object space which must be tested.

In [Gaud88] a special-purpose MIMD architecture using image space subdivision and a static load distribution is described. To overcome the problem of having the entire object database resident at each processor, a central broadcast processor issues data packets describing

the object database cyclically. Here, the processors make requests for various pieces of the database, and only those parts are broadcast in a given cycle. Using a somewhat different approach, [Shen95] uses object space partitioning on clusters of processors, but each processor operates in a pipelined fashion, a scheme previously explored in the LINKS-1 architecture [Nish83].

One of the few multiprocessor architectures which allocates work based on object subdivision combined with image subdivision is proposed by [Kim96]. Each processor handles ray-object intersection tests with its assigned objects, which are spread across the object space. If the load becomes unbalanced, objects are dynamically transferred to other processors.

Other specialized multiprocessor machines of note are the Pixel Machine [Potm89] (useful for several types of rendering including ray tracing and the RayCasting Engine [Meno94] (specifically built for CSG modeling).

### **II.2.1.3 Distributed Computing Environments**

Parallel rendering using distributed computing environments continues to grow in popularity, especially in the fields of entertainment and scientific visualization. Below are a few interesting examples.

Perhaps the most popular example is the Disney film, *Toy Story*, which used a network of 117 Sun workstations and the Pixar Renderman system to produce the animation [Henn96]. To generate its 144,000 individual frames, *Toy Story* required about 43 years of CPU time. If not for the many machines participating in the computation, the movie's production could not be realized. For some tasks, such as applying surface textures, one machine was chosen as a server for the rest. For other tasks, such as final rendering, the machines were basically used independently to render individual frames.

Another entertainment application used a network of 40 Amiga machines to render special effects for the television series *SeaQuest* [Worl93]. Although the delivered product contained only two to three minutes of computer graphics per episode (3,600 to 5,400 frames of animation), the rendering activity was so time-intensive that the team struggled to deliver the graphics within its weekly deadline.

For the average user, some popular commercial animation packages (e.g., Alias/Wavefront, Maya, and 3D Studio) employ coarse-grain parallelism to allow rendering of individual frames of an animation across a network of machines. This technique can mean the difference between an animation being ready in hours or in days. POV-Ray has also been ported to run in a clustered environment; however, the parallelization scheme works on single images only.

Other computation-intensive graphics problems have taken advantage of the processing

power of distributed computing, specifically volume rendering and virtual reality. Although real-time interaction in these systems is constrained by the relatively slow network connecting the machines, significant speedups have been reported using a network of IBM RS/6000 machines for volume rendering [Gier93] [Ma93], and a network of Sun Sparcstations and HP workstations for virtual environments [Pan96].

Distributed computing is also being applied to computer vision algorithms [Judd94], a field closely related to graphics. Here, researchers use PVM on a cluster of 25 Sun Sparcstations for an edge-detection algorithm. One remarkable result of their experiments is that they achieved superlinear speedup on the cluster over the sequential version. This result is due to the large aggregate memory of the clustered machines, which reduced the amount of paging as compared to the single processor.

Not nearly as much research has been conducted concerning ray tracing in distributed computing environments as in traditional multiprocessor machines. Perhaps this fact is due to the relatively recent introduction of PVM and MPI. Regardless of the reason, more advanced parallel ray tracing algorithms combined with a distributed computing environment remain a largely unexplored area. Several related projects are summarized below.

For single images, [Jeva89] uses a dynamic load balancing technique with spatial partitioning and a novel warp synchronization method. At the other end of the spectrum, [Ris94] applies a static load balancing scheme using object partitioning on a network composed of both sequential workstations and parallel computers. Surprisingly few systems use an image partitioning scheme in a distributed computing environment, even though it represents the technique with the highest potential for speedup [Clea86] and overcomes the problems of limited local memory that exist in traditional multiprocessors.

For animations, [DeMa92] describes the DESIRE (Distributed Environment System for Integrated Rendering) system, which incorporates a coarse-level dynamic load balancing scheme that distributes individual frames of an animation to networked workstations. The goal of [Stob88] is similar, except that the system is designed to run without affecting the regular users of the workstations. By stealing idle cycles from 22-34 workstations, a ray-traced animation lasting five minutes (7550 frames) was rendered in two months, although the overall task was estimated at 32 CPU-months.

The work presented in [Cros95] uses a relatively small (three-machine) distributed environment for ray tracing animations in virtual reality applications. Here, each of the machines has a special task assigned to it according to its processing specialty. In order to achieve close to interactive rates, the system, which takes advantage of progressive refinement, is composed of fairly powerful individual processors connected by an ATM network.

## II.2.2 Discussion of Architectural Environments

For parallel processing tasks, the fastest systems will generally be the specialized multiprocessor machines, since they are built with a specific task in mind. Next will be general-purpose multiprocessor machines. Although distributed environments may provide the same number of processors as a multiprocessor machine, computations will be performed more quickly on multiprocessors due to their high-speed interconnection networks. Even so, several factors have motivated a trend toward distributed computing.

First, and perhaps most importantly, not many organizations can afford a parallel machine, which can easily cost millions of dollars [Geis94]. Many sites, however, already have some type of network of computers. Second, multiprocessors often employ specialized or exotic hardware and software resources that significantly increase the complexity, and hence the cost, of the machine; conversely, great expense is rarely incurred to perform distributed computing because the network and the machines are usually already in place. Surprisingly, distributed computing has proven to be so cost-effective that networks of standard workstations have been purchased specifically to run parallel applications that were previously executed on more expensive supercomputers [Grop94].

Due to the fact that networks of workstations are loosely coupled, distributed computing environments allow the network to grow in stages and take advantage of the latest network technology. As CPUs evolve to faster speeds, workstations can be swapped out for the latest model. Such flexibility in network and processor choice is not usually available on a multiprocessor. Another consideration is system software: operating system interface, editors, compilers, debuggers, etc. A benefit of workstation platforms is that they remain relatively stable over time, allowing programmers to work in familiar environments. To use multiprocessor systems, developers may have to climb a steep learning curve.

Additionally, in a distributed computing networked environment, the interconnected computers often consist of a wide variety of architectures and capabilities. This heterogeneity leads to a rich variety of machine combinations and computing possibilities, which can be tailored to specific applications to reduce overall execution time. On the other hand, a multiprocessor machine does not spend processing time converting data between various machine types, as a distributed computing environment might.

Finally, while utilization and efficiency are extremely important in the multiprocessor world, users on a network of traditional machines rarely consider these issues. The results are under-utilized computers which spend much of their time idle. With distributed computing, some of those idle cycles can be put to good use without impacting the primary users of the machines.

### II.2.3 Message-Passing Software for Distributed Computing Environments

To realize distributed computing, computers in a network must support some type of distributed programming environment that allows users to write parallel applications for networked machines. This programming environment should provide a common interface for developers to pass messages easily across various network types and between machines of differing architectures. Although many additional features are usually included, a distributed programming environment need only provide a minimum set of capabilities to be useful [Grop94]:

- First, some method must exist to start up and initialize the parallel processes on all participating machines. This procedure may be as simple as specifying each machine and an associated command in a static file, or spawning the processes directly within the program of the master process. Here, the master process refers to a user-initiated process responsible for delegating work and compositing results; conversely, slave processes perform only the work assigned to them and report results back to the master.
- Once start-up is complete, a process should be able to identify itself, as well as other processes running on the local machine or remote machines participating in the work. Such identification is useful for specifying the source and destination of transmitted messages.
- Since a distributed computing environment often consists of machines with widely varying architectures, message transmission must account for differing data formats so that all computers on the network understand the data exchanged between them. This capability is often built into the programming library, which first transforms the data into a common format that can be easily decoded on the receiver's side. For this reason, among others, a version of the distributed programming environment must exist for every type of machine architecture participating in the computation.
- Finally, once an application is complete, some way of terminating all the processes must be available.

Many distributed programming environments have received attention in the last five years, including p4 [Butl94], Express [Flow94], Linda [Carr94], and TCGMSG [Harr91]; however, by far the most popular systems are PVM and MPI. PVM was developed at Emory University and Oak Ridge National Laboratory and was first released in 1991. The MPI standard, an international effort, was introduced in 1993.

Both the PVM and MPI message-passing environments, freely available on the world-wide web, provide common interfaces for communication on both multiprocessors and networks of workstations. Both run on many different machines, allowing networked computers of diverse

architectures to emulate a distributed-memory multiprocessor. Each of the machines in the network may be a single-processor or multiprocessor system.

A running process in either PVM or MPI can view a network of computers as a single, virtual machine, ignoring architectural details, or as a set of specialized processors with unique computational abilities. The master process runs on a single computer from which other tasks participating in the computation are initiated. Processes, or tasks, roughly correspond to UNIX processes and operate independently and sequentially, performing both communication and computation. Multiple tasks can run on multiple machines, on a single machine, or a combination of the two.

PVM and MPI are not programming languages; rather, they provide libraries specifying the names, parameters, and results of Fortran and C routines used in message passing. Any program making use of these routines can be compiled with standard compilers by linking in the PVM or MPI library. Note that the developer controls the parallelism in the program by writing master and slave tasks and explicitly specifying the high-level message passing protocol between them. Both PVM and MPI support functional parallelism, in which each task is assigned one function of a larger process; data parallelism, in which identical tasks solve the same problem but for small subsets of the data; or a combination of either approach.

# References

- [Acke93] K. Ackeley, "RealityEngine Graphics," *Proceedings of SIGGRAPH 93 in ACM Computer Graphics*, Aug. 1993, pp. 109-116.
- [Akti94] M. Aktihanoglu, B. Ozguc, and C. Aykanat, "MARS: A Tool-based Modeling, Animation, and Parallel Rendering System," *The Visual Computer*, Vol. 11, No. 1, 1994, pp. 1-14.
- [Bado94] D. Badouel, K. Bouatouch, and T. Priol, "Distributing Data and Control for Ray Tracing in Parallel," *IEEE Computer Graphics and Applications*, Vol. 14, No. 4, 1994, pp. 69-77.
- [Butl94] R. Butler and E. Lusk, "Monitors, Messages, and Clusters: the p4 Programming System," *Parallel Computing*, Vol. 20, No. 4, 1994, pp. 547-564.
- [Carr94] N. Carriero, D. Gelernter, T. Mattson, and A. Sherman, "The Linda Alternative to Message-Passing Systems," *Parallel Computing*, Vol. 20, No. 4, 1994, pp. 633-655.
- [Clea86] J. G. Cleary, G. Wyvill, and B. M Birtwistley, "Multiprocessor Ray Tracing," *Computer Graphics Forum*, Vol. 5, No. 1, 1986, pp. 3-12.
- [Cros95] R. A. Cross, "Interactive Realism for Visualization Using Ray Tracing," *Proceedings of the 1995 IEEE Visualization Conference*, Atlanta, GA, 1995, pp. 19-26.
- [Crow88] F. C. Crow, G. Demos, J. Hardy, J. McLaughlin, and K. Sims, "3D image Synthesis on the Connection Machine," *Proceedings of the International Conference on Parallel Processing for Computer Vision and Display in Parallel Processing for Computer Vision and Display*, P. M. Dew, T. R. Heywood, and R. A. Earnshaw (Eds.), Addison-Wesley, 1988, pp. 254-269.
- [DeMa92] J. M. De Martino and R. Kohling, "Production Rendering on a Local Area Network," *Computers and Graphics*, Vol. 16, No. 3, 1992, pp. 317-324.

- [Ells94] D. Ellsworth, "A New Algorithm for Interactive Graphics on Multicomputers," *IEEE Computer Graphics and Applications*, Vol. 14, No. 4, July 1994, pp. 33-40.
- [Evan92] Evans and Sutherland Computer Corporation, *Freedom Series Technical Report*, Salt Lake City, Utah, October, 1992.
- [Eyle88] J. Eyles, J. Austin, H. Fuchs, T. Greer, and J. Poulton, "Pixel-Planes 4: A Summary," *Advances in Computer Graphics Hardware II*, Springer-Verlag, Berlin, 1988, pp. 183-207.
- [Flow94] J. Flower and A. Kolawa, "Express Is Not Just a Message Passing System-Current and Future Directions in Express," *Parallel Computing*, Vol. 20, No. 4, 1994, pp. 597-614.
- [Fuch89] H. Fuchs, J. Poulton, J. Eyles, T. Greer, J. Goldfeather, D. Ellsworth, S. Molnar, G. Turk, B. Tebbs, and L. Israel, "Pixel-Planes 5: A Heterogeneous Multiprocessor Graphics System Using Processor-Enhanced Memories," *ACM Computer Graphics*, Vol. 23, No. 3, 1989, pp. 79-88.
- [Gaud88] S. Gaudet, R. Hobson, P. Chilka, and T. Calvert, "Multiprocessor Experiments for High-Speed Ray Tracing," *ACM Transactions on Graphics*, Vol. 7, No. 3, 1988, pp. 151-179.
- [Geis94] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam, *PVM: Parallel Virtual Machine – A Users' Guide and Tutorial for Network Parallel Computing*, MIT Press, Cambridge, MA, 1994.
- [Gier93] C. Giertsen and J. Petersen, "Parallel Volume Rendering on a Network of Workstations," *IEEE Computer Graphics and Applications*, Vol. 13, No. 6, 1993, pp. 16-23.
- [Goel96] V. Goel and A. Mukherjee, "An Optimal Parallel Algorithm for Volume Ray Casting," *The Visual Computer*, Vol. 12, No. 1, 1996, pp. 26-39.
- [Gree90] S. A. Green and D. J. Paddon, "A Highly Flexible Multiprocessor Solution for Ray Tracing," *The Visual Computer*, Vol. 6, No. 2, 1990, pp. 62-73.
- [Gree91] S. A. Green, *Parallel Processing for Computer Graphics*, MIT Press, Cambridge, MA, 1991.
- [Grop94] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI – Portable Parallel Programming with the Message-Passing Interface*, MIT Press, Cambridge,

MA, 1994.

- [Harr91] R. J. Harrison, "Portable Tools and Applications for Parallel Computers," *International Journal of Quantum Chemistry*, Vol. 40, 1991, pp. 847-863.
- [Henn96] M. Henne, H. Hickel, E. Johnson, and S. Konishi, "The Making of *Toy Story*," *IEEE COMPCON '96 Digest of Papers*, IEEE Computer Society Press, Los Alamitos, CA, 1996, 463-468.
- [Hwan93] K. Hwang, *Advanced Computer Architecture – Parallelism, Scalability, Programmability*, McGraw-Hill, New York, 1993.
- [Isle91] V. Isler and B. Ozguc, "Fast Ray Tracing 3D Models," *Computers and Graphics*, Vol. 15, No. 2, 1991, pp. 205-216.
- [Jeva89] D. A. Jevans, "Optimistic Multi-Processor Ray Tracing," *New Advances in Computer Graphics (Proceedings of Computer Graphics International '89)*, R. P. Earnshaw and B. M. Wyvill (Eds.), Springer-Verlag, Berlin, 1989, pp. 507-522.
- [Keat95] M. J. Keates and R. J. Hubbard, "Interactive Ray Tracing on a Virtual Shared-Memory Parallel Computer," *Computer Graphics Forum*, Vol. 14, No. 4, 1995, pp. 189-202.
- [Kim96] H-J. Kim and C-M Kyung, "A New Parallel Ray-Tracing System Based on Object Decomposition," *The Visual Computer*, Vol. 12, No. 5, 1996, pp. 244-253.
- [Kubo93] Kubota Pacific Computer, *Denali Technical Overview*, version 1.0, Santa Clara, CA, March 1993.
- [Lin91] T. T. Y. Lin and M. Slater, "Stochastic Ray Tracing Using SIMD Processor Arrays," *The Visual Computer*, Vol. 7, No. 4, 1991, pp. 187-199.
- [Ma93] K-L. Ma and J. S. Painter, "Parallel Volume Visualization on Workstations," *Computers and Graphics*, Vol. 17, No. 1, 1993, pp. 31-37.
- [Maur93] H. Maurel, Y. Duthen, and R. Caubet, "A 4D Ray Tracing," *Proceedings of Eurographics '91 in Computer Graphics Forum*, Vol. 12, No. 3, 1993, pp. 285-294.
- [Meno94] J. Menon, R. J. Marisa, and J. Zagajac, "More Powerful Solid Modeling through Ray Representations," *IEEE Computer Graphics and*

*Applications*, Vol. 14, No. 3, 1994, pp. 22-35.

- [Moln92] S. Molnar, J. Eyles, and J. Poulton, "PixelFlow: High-Speed Rendering Using Image Composition," *Proceedings of SIGGRAPH 92 in ACM Computer Graphics*, Vol. 26, No. 2, pp. 231-240.
- [Moln94] S. Molnar, M. Cox, D. Ellsworth, and H. Fuchs, "A Sorting Classification of Parallel Rendering," *IEEE Computer Graphics and Applications*, Vol. 14, No. 4, July 1994, pp. 23-32.
- [Nish83] H. Nishimura, H. Ohno, T. Kawata, I. Shirakawa, and K. Omura, "LINKS-1: A Parallel Pipelined Multimicrocomputer System for Image Creation," *Proceedings of the 10<sup>th</sup> Symposium on Computer Architecture (SIGARCH)*, 1983, pp. 387-394.
- [Pan96] Z. Pan, J. Shi, and M. Zhang, "Distributed Graphics Support for Virtual Environments," *Computers and Graphics*, Vol. 20, No. 2, 1996, pp. 191-197.
- [Pito93] P. Pitot, "The Voxar Project," *IEEE Computer Graphics and Applications*, Vol. 13, No. 1, 1993, pp. 27-33.
- [Plun85] D. J. Plunkett and M. J. Bailey, "The Vectorization of a Ray-Tracing Algorithm for Improved Execution Speed," *IEEE Computer Graphics and Applications*, Vol. 8, No. 5, 1985, pp. 52-60.
- [Potm89] M. Potmesil and E. M. Hoffert, "The Pixel Machine: A Parallel Image Computer," *Proceedings of SIGGRAPH '89 in ACM Computer Graphics*, Vol. 23, No. 3, 1989, pp. 69-78.
- [Prio88] T. Priol and k. Bouatouch, "Experimenting with a Parallel Ray-Tracing Algorithm on a Hypercube Machine," *Proceedings of Eurographics '88*, D. A. Duce and P. Jancene (Eds.), North-Holland, 1988, pp. 243-259.
- [Prio89] T. Prioland K. Bouatouch, "Static Load Balancing for Parallel Ray Tracing on a MMD Hypercube," *The Visual Computer*, Vol. 5, Nos. 1 and 2, 1989, pp. 109-119.
- [Reis97] A. Reisman, C. Gotsman, and A. Schuster, "Parallel Progressive Rendering of Animation Sequences at Interactive Rates on Distributed-Memory Machines," *Proceedings of the 1997 IEEE Parallel Rendering Symposium*, Phoenix, AZ, 1997, pp. 39-47.
- [Ris94] P. Ris and D. Arques, "Parallel Ray Tracing Based upon a Multilevel Topological Knowledge Acquisition of the Scene," *Proceedings of*

*Eurographics '94 in Computer Graphics Forum*, Vol. 13, No. 3, 1994, pp. 221-232.

- [Schr92] P. Schroder and S. M Drucker, "A Data Parallel Algorithm for Raytracing of Heterogeneous Databases," *Proceedings of Computer Graphics Interface '92*, Vancouver, British Columbia, 1992, pp. 167-175.
- [Shen95] L-S. Shen, E. F. Deprettere, and p. Dewilde, "A Parallel Image-Rendering Algorithm and Architecture Based on Ray Tracing and Radiosity Shading," *Computers and Graphics*, Vol. 19, No. 2, 1995, pp. 281-296.
- [Silv94] C. T. Silva and A. E. Kaufman, "Ray Parallel Performance Measures for Volume Ray Casting," *Proceedings of the 1994 IEEE Visualization Conference*, Washington, D. C., 1994, pp. 196-203.
- [Sous90] A. A. Sousa, A. M. C. Costa, and F. N. Ferreira, "Interactive Ray-Tracing for Image Production with Increasing Realism," *Proceedings of Eurographics '90*, North-Holland, 1990, pp. 449-457.
- [Stob88] A. Stober, A. Schmitt, B. Neidecker, H. Muller, T. Maus, and W. Leister, "'Occursus Cum Novo' – Tools for Efficient Photo-Realistic Computer Animation," *Proceedings of Eurographics '88*, D. A. Duce and P. Jancene (Eds.), North-Holland, Amsterdam, 1988, pp. 31-41.
- [Sung96] K. Sung, J. L. J. Shiuan, and A. L. Ananda, "Ray Tracing in a Distributed Environment," *Computers and Graphics*, Vol. 20, No. 1, 1996, pp. 41-49.
- [Whit92] S. Whitman, *Multiprocessor Methods for Computer Graphics Rendering*, Jones and Bartlett Publishers, Boston, 1992.
- [Whit94] S. Whitman, "A Task Adaptive Parallel Graphics Renderer," *IEEE Computer Graphics and Applications*, Vol. 14, No. 4, July 1994, pp. 41-48.
- [Whit80] T. Whitted, "An Improved Illumination Model for Shaded Display," *Communications of the ACM*, Vol. 23, No. 6, 1980, pp. 343-349.
- [Worl93] L. World, "Low-End System Animates the Depths in *SeaQuest*," *IEEE Computer Graphics and Applications*, Vol. 13, No. 6, 1993, p. 93.

**Section III**  
**Interactive Ray Tracing**

**Erik Reinhard**

# Contents

<b>III Interactive ray tracing</b>	<b>3</b>
III.1 Introduction	4
III.1.1 Hardware considerations	5
III.2 Interactive ray tracing	7
III.2.1 Organization of the algorithm	7
III.2.2 Frameless rendering	7
III.2.3 Tracing a single ray	7
III.2.4 Low level optimization	8
III.2.5 Profiling	9
III.2.6 Memory and CPU placement	9
III.3 Animation and interactive ray tracing	11
III.3.1 Algorithm	11
III.3.2 Results	15
III.3.3 Discussion	19
III.4 Sample reuse techniques	20
III.4.1 Render cache algorithm	21
III.4.2 Parallel render cache	22
III.4.3 Implementation details	24
III.4.4 Measuring scalability	25
III.4.5 Results	26
III.4.6 Discussion	31
III.5 Summary and discussion	32
IIIA SGI Origin 2000	35
IIIA.1 The MIPS R10000 processor	35
IIIA.2 Origin 2000 layout	35
IIIB Profiling on the SGI Origin 2000	38
IIIB.1 Performance analysis using perfex	39
IIIB.2 Absolute counts of one or two events	39
IIIB.3 Statistical counts of all events	40
IIIB.4 Analytic output with the -y option	41
IIIB.5 Using SpeedShop	44
IIIB.6 PC sampling profiling	44
IIIB.7 Using prof	46
IIIB.8 Ideal time profiling	48

III.B.9	Operation counts . . . . .	49
III.B.10	Gprof . . . . .	50
III.B.11	Ustertime profiling . . . . .	52
III.C	Dynamic Acceleration Structures for Interactive Ray Tracing . . . . .	54
III.C.1	Introduction . . . . .	54
III.C.2	Acceleration Structures for Ray Tracing . . . . .	55
III.C.3	Grids . . . . .	55
III.C.4	Hierarchical grids . . . . .	57
III.C.5	Evaluation . . . . .	58
III.C.6	Conclusions . . . . .	61

## III.1 Introduction

It is our belief that in the near future ray tracing will become the *de facto* standard for interactive rendering. For certain applications, multi-processor ray tracing already outperforms hardware based solutions. Whereas hardware rendering becomes ever more complex to keep up with today's demands, ray tracing is easily implemented in software, scales sub-linearly with scene-complexity and (nearly) linearly with the number of processors used. These are indeed favorable conditions.

Research has shown that using current high-end general purpose hardware, interactive ray tracing is possible for moderately complex scenes and scales well to reasonably large numbers of processors, albeit using fairly small image sizes (typically 512 by 512 pixels). One of the key features of such interactive ray tracing systems is that the inherent parallelism of ray tracing is exploited as much as possible. In other words, it is paramount that the complexity of the algorithm needs to be kept as low as possible. A simple algorithm is better than a complex one. Small data structures are better than large ones.

The choice of hardware is also quite important. High level choices such as interconnection network and memory topology can make or break an interactive ray tracing project. While shared-memory architectures are said to not scale beyond a certain number of processors due to the bus-architecture commonly employed, we believe that with current technology the latency and network throughput experienced in distributed memory systems just is not good enough for the purposes of interactive ray tracing. So this section of the course will be exclusively targeted at shared memory architectures. In fact, the interactive ray tracer described here is implemented on a Silicon Graphics Origin 2000 and results in the following chapters are shown using a Silicon Graphics Origin 3800. Both are shared memory machines with a (cache-coherent) non-uniform memory access (ccNUMA) architecture. Hardware selection and its implications are further discussed in section III.1.1

Another key feature is that at the code level the algorithm needs to be optimized as far as possible. The particular features of the hardware on which the algorithm is going to run, need to be exploited as much as possible. Low level optimization is an extremely important aspect of interactive ray tracing. Given the choice of hardware, in Chapter III.2 we show which issues should be considered to optimize the ray tracing process. We would also like to refer to Siggraph 2001 course 13 which contains a more elaborate discussion of these issues [24].

After optimizing the ray tracer, one may want to extend the feature set of the ray tracer. Walking through a scene at interactive rates is useful, but being able to interact with it is even more useful. While ray tracing is traditionally good at rendering static scenes, we show that it is possible to render animated scenes, or even interact with the scene in real-time. This requires some basic modifications to the ray tracing algorithm (in particular the spatial subdivisions need a tweak) that are easy to implement, incur a small performance penalty, but allow animation to take place interactively and give the user the ability to manoeuvre objects. Chapter III.3 shows the details.

A second wish users will have once they've implemented their interactive ray tracer, is that it would be great if it could deal with higher scene complexity and/or larger image sizes. As the ray tracer is already fully optimized, it is unlikely that further low-

level tweaks are going to significantly improve performance. A high-level optimization is required. One approach would be to try and implement techniques that allow results from previous frames to be reused. There are a number of techniques in existence that can be employed. Reusing previous results involves displaying pixels at a higher frame-rate than new pixels can be produced. Storing pixels in a 3D point cloud and reprojecting the points for each new frame is one such method. Several choices need to be made to optimize such reprojection techniques, including strategies to decide which pixels to render for the next frame, whether frames are going to be used in the first place or whether the algorithm is to operate in asynchronous mode (frameless rendering), and whether the point reprojection is also going to operate in parallel or not. Chapter III.4 presents an implementation and discusses its merits and weaknesses.

### **III.1.1 Hardware considerations**

As discussed briefly above, the choice of hardware is important. Ray tracing can be implemented to perform interactively, but the first thing that appears to be required is some form of parallelism. Without it, the options are extremely limited. Assuming that parallel hardware is available, the next question to ask is whether this hardware is suitable or not.

We limit the discussion about what hardware to use to general purpose hardware and leave dedicated hardware solutions to others. General purpose hardware can generally be grouped into two broad categories: distributed and shared memory architectures. Putting this distinction before any others is no coincidence: its impact on the performance of our interactive ray tracer will be profound. Whereas distributed memory architectures (either as parallel machines or as networks of workstations) can be cheap, they do exhibit a few disadvantages that are more or less important depending on the application. The first of these is the fact that fetching a data item from a remote processor's local memory can be a couple of orders of magnitude slower than a local memory access, especially on clusters of workstations. Rendering large scenes using distributed memory architectures would therefore require the scene database to be replicated. Even then, after each frame is rendered, the pixel data needs to be transferred to a framebuffer which is typically not distributed. It therefore requires communication of all pixel data to one processor which holds the frame buffer. Given the relatively high latency and low throughput which one can expect from distributed memory systems, this is unlikely to be a fruitful approach. One solution would be to research the use of distributed frame buffers, which may make distributed memory interactive ray tracing feasible.

As far as we are aware, currently there are no distributed memory systems with distributed framebuffers commercially available. This is likely to change in the future, though. When this happens, the above arguments will have to be re-evaluated.

As a result, a solution with the necessary latency and throughput figures, is currently available by using shared memory systems. Although presented to the user as containing just one contiguous block of memory, these machines usually employ caching mechanisms to further speed-up memory accesses. Hence, there is still a distinction between local and remote memory accesses, albeit that the time figures for these accesses are much better than for distributed memory machines.

Another advantage of such systems is that in theory, programming is much simpler than programming on a distributed memory architecture. On shared memory architectures, a sequential program may run unaltered. If the performance is unsatisfactory, the inner loops can be parallelized incrementally. There is no need to parallelize the full code in one go. After each incremental step, the program could be tested for correctness and performance. Also, within the shared memory paradigm, loops can be parallelized without knowledge about where the data resides that is accessed within the loop. These are facilities that make programming on a shared-memory architecture relatively straightforward.

In practice, however, ease of programming is strongly related to the level of optimization desired. As ray tracing is extremely computationally intensive, we cannot afford to waste computer cycles and we therefore have to pay careful attention to performance optimization. Unfortunately, this negatively affects ease of programming. On the hardware architecture described in Appendix IIIA, the memory is physically distributed, although it is presented to the programmer as one block of contiguous memory, having a single address space. This means that it is faster to access some parts of the memory than it is to access other parts of the memory. Thus, there is a performance gain to be had from anticipating where data is located. Although the parallelized program will work without such knowledge, its performance will not be optimal.

We have now stated our case for shared memory architectures. It is motivated by the nature of our application, which requires extremely fast communication, just to get the pixels to the screen. For other applications, the choice of hardware may be different and for coarse grain applications, it may be extremely cost effective to choose distributed memory machines. If, for example, the lighting simulation is to include diffuse inter-reflection, the cost of rendering a single frame may be prohibitively expensive. Using a distributed memory system is then an attractive alternative to turn a nearly intractable problem into a feasible one.

Whereas we've so far considered memory architectures in general, the interactive ray tracer described in this course is actually implemented on a particular shared memory system: a Silicon Graphics Origin 2000. In Appendix IIIA we describe this architecture in some detail, as knowledge of this hardware allows our code to be optimized further at a later stage. Note also that the results presented in the following chapters were obtained on a faster Silicon Graphics Origin 3800. From a programmer's point of view, the differences are small, although the algorithms run around 1.7 times faster than on the Origin 2000.

## III.2 Interactive ray tracing

In this section we focus on the basics of interactive ray tracing. The general approach is to keep the algorithm as simple as possible. In practice this means that a master-slave configuration is employed. The slaves produce pixels, which are displayed by a single display thread, which also doles out new tasks.

The display thread uses double buffering. While one set of pixels is being displayed, another pixel array is filled by the renderers. Once the new frame is complete, the display thread swaps the two pixel arrays and displays the new frame.

### III.2.1 Organization of the algorithm

For the renderers, a task consists of rendering a number of neighboring pixels. The task size is chosen such that the pixel data in a single task fits on a cache line. Because the pixels are close to each other in screen space, their associated rays are likely to intersect the same objects, thus improving cache performance by minimizing the amount of object data that needs to be fetched from memory.

New tasks are doled out starting from the top of the image. Tasks near the bottom of the image, which are computed later during the current frame, consist of fewer pixels than the ones near the top. This ensures a good load balance, while also keeping the total number of tasks executed per frame small. The latter is important, because there is inevitably some overhead associated with each task and fewer tasks may improve cache coherence.

### III.2.2 Frameless rendering

An alternative strategy is to abandon the concept of frames and switch to frameless rendering [4, 7, 18, 30], where pixels are displayed as soon as they are rendered. Here, a task consists of a single ray, which for each processor is chosen randomly. Although the latency between computing a pixel and displaying it is reduced, the time required to render the equivalent of a complete frame is increased. Thus, from a performance point of view, this is a less successful organization of the algorithm. Cache coherence is preserved less well and extra overhead is incurred. However, when scene complexity increases, frame-based rendering would slow down to the point where the frame rate would be too slow for practical purposes. Although frameless rendering is in fact slightly slower than frame-based rendering, the effect of randomizing the order in which pixels are traced, and splatting the pixels on screen as soon as results become available, gives the impression of smooth movement long after frame-based rendering ceased to be effective. In Chapter III.4 this technique is compared with other mechanisms to increase the production of pixels in an interactive ray tracing context. This technique is discussed further in Section III.2.4.

### III.2.3 Tracing a single ray

Tracing each ray is optimized in the usual way by employing a spatial subdivision. While the exact type of spatial subdivision is not fixed, in our system it is either a grid,

which may be nested to accommodate local scene complexity, or a hierarchical spatial subdivision such as a bounding volume hierarchy or an octree. More elaborate schemes are possible, but we have obtained good results with the above spatial subdivisions. We will therefore not go into great detail describing them, with the exception of the discussion in Chapter III.3 where object animation has a direct impact on the organization of any spatial subdivision.

While the above discusses some important design decisions, a simple screen space subdivision combined with a standard spatial subdivision does not explain why this algorithm is capable of achieving interactive rates. The performance gain for this type of rendering on SGI Origin 2000 and 3800 architectures lies in the low level optimization employed. This is the topic of the following sections.

### **III.2.4 Low level optimization**

In this section, we discuss simple ways to optimize data structures. Efficiency can be gained by optimizing data access patterns as well as by ensuring that data structures fit on a single cache line as much as possible.

Ray tracing in general can be quite efficient in terms of data access patterns. Neighboring primary rays have a reasonable probability of intersecting the same objects. Hence, if neighboring rays are traced one after the other, chances are that the objects fetched for the first ray, still reside in the local cache when tracing the second ray. Such data coherence can for example be exploited by tiling the screen into sufficiently large tiles and assign these tiles as tasks to processors. Although this improves cache coherence, such tiling approaches reduce the number of tasks available per frame. By making tiles too large, load balancing issues may appear. In the interactive ray tracer a trade-off is reached by assigning large tiles at the start of the frame and reducing tile sizes when the frame progresses. This ensures high cache efficiency throughout most of the computations associated with a frame, while at the same time the workload is well balanced since the processors are likely to finish their work for the current frame roughly at the same time.

In the case of frameless rendering, results are displayed as soon as they become available. This is in contrast with standard frame-based rendering where the results of the current frame are displayed as soon as the complete frame is finished. Frameless rendering therefore has the advantage that most of the results are displayed quicker than in frame-based rendering of the same scene. However, to achieve fluidity, task assignment will have to be randomized. Because of this, cache coherence is effectively destroyed and the time taken to render the equivalent of a complete frame would therefore take longer than it would take to compute a single frame using frame-based rendering. One advantage of frameless rendering is that for slightly more complex scenes, the animation or walk-through is perceived more fluid and responsive where in frame-based rendering the update rate of the display would become distractingly slow. Frameless rendering is compared to other optimization strategies in Chapter III.4.

Assuming frame-based rendering for now, the data structures employed can be optimized by noting that on the Origin 2000, a cache line is 128 bytes long. When a cache invalidation occurs, a whole cache line is swapped out, to be replaced by a new one. If two different data items reside on a single cache line, and one gets invalidated, by

necessity the other data item will also be removed from the cache. This may incur a performance penalty that can be fairly easily avoided.

Many data structures routinely employed for ray tracing can be cast in the form of an array of structs. If this is the case, then the size of each of these structs should be considered. In general, it is recommended that the size of these structures is as small as possible. One could ensure that cache lines are completely filled by making the size of these structures a power of two. Those data structures that are crucial to performance can be artificially increased in size to be multiples of 128 bytes by adding a character array with the required size to make the whole struct 128 bytes long. As a result, these padded data items can never be removed from the cache because another item on the same cache line is invalidated. Such padding is therefore an important weapon in the arsenal of the programmer, but should nonetheless be used with care. Before and after padding such data structures, one should profile the result to establish if the effect was beneficial or not.

Finally, source code in general can be optimised by minimising the number of branch instructions that occur. This is especially the case for branch instructions that are located within inner loops. With the advent of branch prediction facilities, branching is less detrimental to performance than it once was, but still should be considered an opportunity for optimization.

### **III.2.5 Profiling**

Other non-trivial reasons for suboptimal performance may be determined using profiling and can result from unoptimized code, unoptimized memory access patterns and perhaps even a lack of understanding of the underlying architecture. As it can be difficult to predict where unnecessary performance penalties are incurred, a thorough analysis of the implementation can be useful. Rather than giving a detailed analysis of the profiling steps undertaken to make the interactive ray tracer faster, we would like to refer to appendices IIIA and IIIB which describe the SGI Origin 2000 architecture in some detail and provide a tutorial on profiling on these machines.

### **III.2.6 Memory and CPU placement**

As it is realized that even in architectures that are presented to the programmer as shared memory configurations, the memory is in fact physically distributed, performance can be gained by matching memory with processors. Normally, memory is allocated using system calls, which means that the operating system decides which part of the address space to use for each particular memory allocation. Therefore, unhappy memory allocations may occur where the process requesting the memory is physically located far away from the memory.

If it is known in advance that a particular data structure will be predominantly accessed by a particular process, it may be advantageous to actively place the memory and the processors on fixed locations, rather than leaving this placement up to the operating system. For such an undertaking to be successful, knowledge of the machine's physical architecture is necessary. For the Silicon Graphics Origin 2000, the architecture is described in Appendix IIIA. Briefly summarizing: the machine consists of

nodes which are interconnected using hubs. Each node contains a block of memory and two processors (four on the Origin 3800). Each processor has a separate primary and secondary cache. Memory access is therefore fastest if the data item requested is located in a local cache. Second fastest are memory accesses within the node. Memory that is located with other nodes are slowest to access.

If an algorithm causes many cache misses and if profiling has revealed that this is not easily fixable, then mapping the data structures that cause the cache misses to specific nodes may improve performance. One way to achieve this is to use **mmap** to place memory on a node close to the requesting process and can be achieved using:

```
devzero_fd = open ("/dev/zero", O_RDWR);
local_memory = (localmemory_t *) mmap (0, sizeof (localmemory_t),
                                       PROT_READ | PROT_WRITE,
                                       MAP_PRIVATE | MAP_LOCAL,
                                       devzero_fd, 0);
```

For further information, we would like to refer to the *mmap* manual page. A less involved method for placing memory with specific nodes is described in the *dplace* manual page.

If the *mmap* mechanism is used for memory placement, it may be advantageous to also pin processes to specific CPU's. This can be achieved with the **sysmp** and **sproc** commands (see the *sysmp* and *sproc* manual pages). To spawn a new process on a specific processor, the following statements may be used:

```
sysmp (MP_MUSTRUN, cpu);
pid = sproc (render_process, PR_SALL);
```

### III.3 Animation and interactive ray tracing

A fully optimized ray tracer which allows interactive walk-throughs is attractive over other real-time rendering algorithms because it allows a large set of effects to be rendered which are more difficult or even impossible to obtain using graphics hardware. In addition, ray tracing scales sub-linearly in the number of objects due to the use of spatial subdivisions. It also scales sub-linearly in the number of pixels rendered, provided cache coherency can be exploited fully.

To make interactive ray tracing more attractive, we have looked into ways to enable objects to be manipulated in real-time ([22], reproduced with permission in Appendix IIIC). In the following we assume that animation paths are not known prior to the rendering, and so updates to the scene need to be achieved in real-time with as little overhead as possible. In addition it is important that the effect of time-varying scenes on the performance of the renderer is as small as possible.

Changing the coordinates of an object in real-time is not particularly difficult to achieve, so we will not address this issue in any detail. However, an object's change in location, size or rotation does imply that after the transformation, the object may occupy a different portion of space. In the absence of a spatial subdivision to speed up the intersection tests, this would not constitute a problem.

However, the current speed of the hardware, combined with the number of computations required to ray trace an image, does not allow us to do away with spatial subdivisions altogether. Additionally, spatial subdivisions are usually built as a pre-process to rendering. The cost of building a spatial subdivision is not negligible. Hence, spatial subdivisions are required to obtain interactive frame-rates, but at the same time they are not flexible enough to accommodate time-varying data.

In this section we describe a simple adaptation to both grid and octree spatial subdivisions which caters for a small number of animated objects. These objects can either be animated according to pre-defined motion splines or they can be picked up by the user and placed elsewhere in the scene. Animating all objects at the same time in a complex scene is not yet possible. It would require rebuilding the entire spatial subdivision for each frame and this is too costly to achieve using current technology. Focusing on just a small number of objects to be animated/manipulated allows the design of spatial subdivisions which can be incrementally updated after each frame.

In the following sub-sections the basic idea is explained (Section III.3.1) and results are shown (Section III.3.2). We would also like to refer to Appendix IIIC which includes a full publication regarding this subject. The results presented in this chapter are obtained using an SGI Origin 3800, while appendix IIIC contains older results using an SGI Origin 2000.

#### III.3.1 Algorithm

In this section modifications to grid and octree spatial subdivisions are discussed. The octree is a hierarchical extension to the grid. We assume the reader is familiar with these spatial subdivisions [1, 5, 8, 10, 12, 14, 16, 17, 27, 29].

Grid spatial subdivisions for static scenes, without any modifications, are already useful for animated scenes, as traversal costs are low and insertion and deletion of

objects is reasonably straightforward. Insertion and deletion are considered basic operations necessary for the animation of objects. The general approach is to remove an object from the spatial subdivision, modify its coordinates and re-insert the object into the acceleration structure. Insertion is usually accomplished by mapping the axis-aligned bounding box of an object to the voxels of the grid. The object is inserted into all voxels that overlap with this bounding box. Deletion can be achieved in a similar way.

However, when an object moves outside the extent of the spatial subdivision, the acceleration structure would normally have to be rebuilt. As this is too expensive to perform repeatedly, we propose to logically replicate the grid over space. If an object exceeds the bounds of the grid, the object wraps around before re-insertion. Ray traversal then also wraps around the grid when a boundary is reached. In order to provide a stopping criterion for ray traversal, a logical bounding box is maintained which contains all objects, including the ones that have crossed the original perimeter. As this scheme does not require grid re-computation whenever an object moves far away, the cost of maintaining the spatial subdivision will be substantially lower. On the other hand, because rays now may have to wrap around, more voxels may have to be traversed per ray, which will slightly increase ray traversal time.

During a pre-processing step, the grid is built as usual. We will call the bounding box of the entire scene at start-up the 'physical bounding box'. If during the animation an object moves outside the physical bounding box, either because it is placed by the user in a new location, or its programmed path takes it outside, the logical bounding box is extended to enclose all objects. Initially, the logical bounding box is equal to the physical bounding box. Insertion of an object which lies outside the physical bounding box is accomplished by wrapping the object around within the physical grid, as depicted in Figure III.1 (left).

As the logical bounding box may be larger than the physical bounding box, ray traversal now starts at the extended bounding box and ends if an intersection is found or if the ray leaves the logical bounding box. In the example in Figure III.1 (right), the ray pointing to the sphere starts within a logical voxel, voxel (0, -2), which is mapped to physical voxel (0, 2). The logical coordinates of the sphere are checked and found to be outside of the currently traversed voxel and thus no intersection test is necessary. The ray then progresses to physical voxel (1, 2). For the same reason, no intersection with the sphere is computed again. Traversal then continues until the sphere is intersected in logical voxel (4, 2), which maps to physical voxel (0, 2).

Objects that are outside the physical grid are tagged, so that in the above example, when the ray aimed at the triangle enters voxels (0, 2) and (1, 2), the sphere does not have to be intersected. Similarly, when the ray is outside the physical grid, objects that are within the physical grid need not be intersected. As most objects will initially lie within the physical bounds, and only a few objects typically move away from their original positions, this scheme speeds up traversal considerably for parts of the ray that are outside the physical bounding box.

When the logical bounding box becomes much larger than the physical bounding box, there is a tradeoff between traversal speed (which deteriorates for large logical bounding boxes) and the cost of rebuilding the grid. In our implementation, the grid is rebuilt when the length of the diagonals of the physical and logical bounding boxes

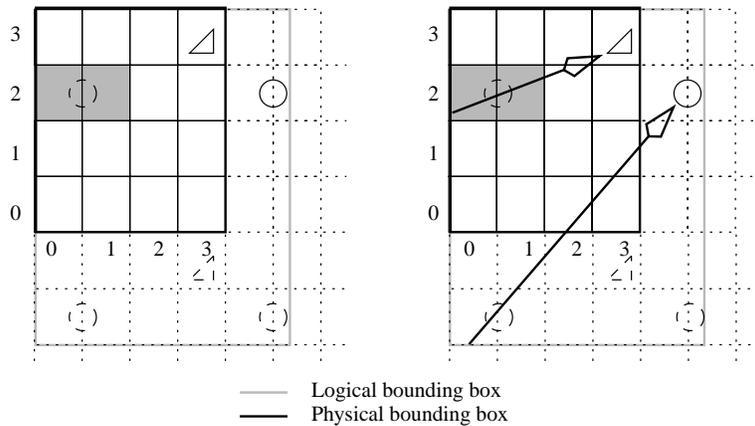


Figure III.1: Grid insertion (left). The sphere has moved outside the physical grid, now overlapping with voxels (4, 2) and (5, 2). Therefore, the object is inserted at the location of the shaded voxels. The logical bounding box is extended to include the newly moved object. Right: ray traversal through extended grid. The solid lines are the actual objects whereas the dashed lines indicate voxels which contain objects whose actual extents are not contained in that voxel.

differ by a factor of two. This heuristic aims to provide a trade-off between traversal speed and the frequency with which the spatial subdivision needs to be re-generated.

Hence, there is a hierarchy of operations that can be performed on grids. For small to moderate expansions of the scene, wrapping both rays and objects is relatively quick without incurring too high a traversal cost. For larger expansions, rebuilding the grid will become a more viable option.

This grid implementation shares the advantages of simplicity and cheap traversal with commonly used grid implementations. However, it adds the possibility of increasing the size of the scene without having to completely rebuild the grid every time there is a small change in scene extent.

The cost of deleting and inserting a single object is not constant and depends largely on the size of the object relative to the size of the scene. The size of an object relative to each voxel in a grid influences how many voxels will contain that object. This in turn negatively affects insertion and deletion times. Hence, it would make sense to find a spatial subdivision whereby the voxels can have different sizes. If this is accomplished, then insertion and deletion of objects can be made independent of their sizes and can therefore be executed in constant time. Such spatial subdivisions are not new and are known as hierarchical spatial subdivisions. Octrees, bintrees and hierarchical grids are all examples of hierarchical spatial subdivisions. However, normally such spatial subdivisions store all their objects in leaf nodes and would therefore still incur non-constant insertion and deletion costs. We extend the use of hierarchical grids in such a way that objects can also reside in intermediary nodes or even in the root node for objects that are nearly as big as the entire scene.

Because such a structure should also be able to deal with expanding scenes, our

efforts were directed towards constructing a hierarchy of grids (similar to Sung [27]), thereby extending the functionality of the grid structure presented in the previous section. Effectively, the proposed method constitutes a balanced octree.

Object insertion now proceeds similarly to grid insertion, except that the grid level needs to be determined before insertion. This is accomplished by comparing the size of the object in relation to the size of the scene. A simple heuristic is to determine the grid level from the diagonals of the two bounding boxes. Specifically, the length of the grid's diagonal is divided by the length of the object's diagonal, the result determining the grid level. Insertion and deletion progresses as explained above.

The gain of better control over insertion time is offset by a slightly more complicated traversal algorithm. Hierarchical grid traversal is effectively the same as grid traversal with the following modifications. Traversal always starts at a leaf node which may first be mapped to a physical leaf node as described earlier in this section. The ray is intersected with this voxel and all its parents until the root node is reached. This is necessary because objects at all levels in the hierarchy may occupy the same space as the currently traversed leaf node. If an intersection is found within the space of the leaf node, then traversal is finished. If not, the next leaf node is selected and the process is repeated.

This traversal scheme is wasteful because the same parent nodes may be repeatedly traversed for the same ray. To combat this problem, note that common ancestors of the current leaf node and the previously intersected leaf node, need not be traversed again (Figure III.2). If the ray direction is positive, the current voxel's number can be used to derive the number of levels to go up in the tree to find the common ancestor between the current and the previously visited voxel. For negative ray directions, the previously visited voxel's number is used instead. Finding the common ancestor is achieved using simple bit manipulation, as detailed in Figure III.3.

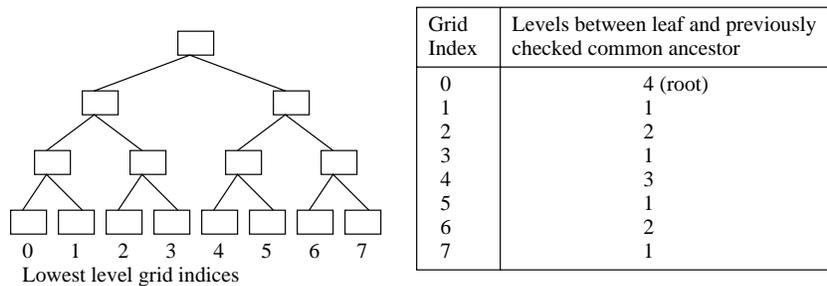


Figure III.2: Hierarchical grid traversal. Assuming that ray traversal starts at node 0 and goes in positive direction, then after each step, the common ancestor is found  $n$  levels above the leaf node as indicated in the table.

As the highest levels of the grid may not contain any objects, ascending all the way to the highest level in the grid is not always necessary. Ascending the tree for a particular leaf node can stop when the largest voxel containing objects is visited.

This hierarchical grid structure has the following features. The traversal is only marginally more complex than standard grid traversal. In addition, wrapping of objects

```

bitmask = (raydir_x > 0) ? x : x + 1
forall levels in hierarchical grid
{
  cell = hgrid[level][x>>level][y>>level][z>>level]
  forall objects in cell
    intersect(ray, object)
  if (bitmask & 1)
    return
  bitmask >>= 1
}

```

Figure III.3: Hierarchical grid traversal algorithm in C-like pseudo-code. The bitmask is set assuming that the last step was along the x-axis.

in the face of expanding scenes is still possible. If all objects are the same size, this algorithm effectively defaults to grid traversal. Insertion and deletion times are much better controlled than for the interactive grid<sup>1</sup>.

### III.3.2 Results

The grid and hierarchical grid spatial subdivisions were implemented using an interactive ray tracer [18], which runs on an SGI Origin 3800 with 32 processors and 16GB of main memory<sup>2</sup>. Each processor is an R12k running at 400Mhz and manages an 8MB secondary cache. We have chosen to use 30 processors for rendering and one extra thread to take care of user input, displaying the frames, and also for updating and rebuilding the spatial subdivision when necessary (one processor remained unused by our application to allow for system processes to run smoothly). The reason to include the scene update routines with the display thread is that querying the keyboard and displaying the images takes very little time. The remainder of the time to calculate a frame could therefore easily be spent animating objects. In addition, it is important that the scene updates are completed within the time to compute a new frame, as longer update times would either cause delays or result in jerky movement of objects. As the frame rate depends on both the scene complexity and the number of processors that participate in the calculation, the time to update the scene is dependent on both of these parameters.

For evaluation purposes, two test scenes were used. In each scene, a number of objects were animated using pre-programmed motion paths. The scenes as they are at start-up are depicted in Figure III.21 (top, Appendix IIIC). An example frame taken during the animation is given for each scene in Figure III.21 (bottom, Appendix IIIC). All images were rendered at a resolution of 512<sup>2</sup> pixels.

#### Traversal performance - static scenes

The performance penalty incurred by the new grid and hierarchical grid implementations are assessed by comparing these with a standard grid implementation. The stan-

<sup>1</sup>Note that this also obviates the need for mailbox systems to avoid redundant intersection tests.

<sup>2</sup>Note that the original work, presented in Appendix IIIC, reported results obtained on a slower Origin 2000.

standard grid data structure consists of a single array of object pointers. This design allows better cache efficiency on the SGI Origin series. Finally, we have also implemented a hierarchical grid with a higher branching factor. Instead of subdividing a voxel into eight children, here nodes are split into 64 children (4 along each axis).

From here on we will refer to the new grid implementation as ‘interactive grid’ to distinguish between the two grid traversal algorithms. As all these spatial subdivision methods have a user defined parameter to set the resolution (voxels along one axis and maximum number of grid levels, respectively), various settings are evaluated. The overall performance is given in Figure III.4 and is measured in frames per second.

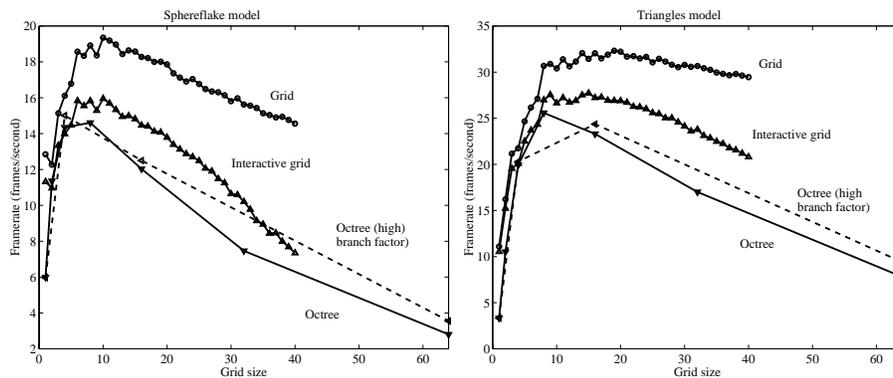


Figure III.4: Performance (in frames per second) for the grid, the interactive grid and the hierarchical grids for two static scenes.

The extra flexibility gained by both the interactive grid and hierarchical grid implementations results in a somewhat slower frame rate. This is according to expectation, as the traversal algorithm is a little more complex and the Origin’s cache structure cannot be exploited as well with either of the new grid structures. The graphs in Figure III.4 should be compared to our previous results given in Figure III.19 in Appendix IIIC. For the hierarchical grid with the higher branching factor, the observed frame rates are very similar to the hierarchical grid.

### Object update rate - dynamic scenes

The object update rates were slightly better for the sphereflake and triangle scenes, because the size differences between the objects matches this acceleration structure better than both the interactive grid and the hierarchical grid.

The non-zero cost of updating the scene effectively limits the number of objects that can be animated within the time-span of a single frame. However, for both scenes, this limit was not reached. For each of these tests, the hierarchical grid is more efficiently updated than the interactive grid, which confirms its usefulness.

The size difference between different objects should cause the update efficiency to be variable for the interactive grid, while remaining relatively constant for the hierarchical grid. In order to demonstrate this effect, both the ground plane and one of the

triangles in the triangle scene was interactively repositioned during rendering. Similarly, in the sphereflake scene one of the large spheres and one of the small spheres were interactively manipulated. The update rates for different size parameters for both the interactive grid and the hierarchical grid, are presented in Figure III.5. Comparing the grid size of 16 for the interactive grid with the size parameter of 4 for the interactive grid in this figure, shows that for similar numbers of voxels (at the deepest level of the hierarchical grid) along each axis, the update rate varies much more dependent on object size for the interactive grid than for the hierarchical grid. Hence, the hierarchical grid copes much better with objects of different sizes than the interactive grid. Dependent on the number of voxels in the grid, there is one to two orders of magnitude difference between inserting a large and a small object. For larger grid sizes, the update time for the ground plane of the triangles scene is roughly half a frame. This leads to visible artifacts when using the interactive grid, as during the update the processors that are rendering the next frame temporarily cannot intersect this object (it is simply taken out of the spatial subdivision). In practice, the hierarchical grid implementation does not show this disadvantage.

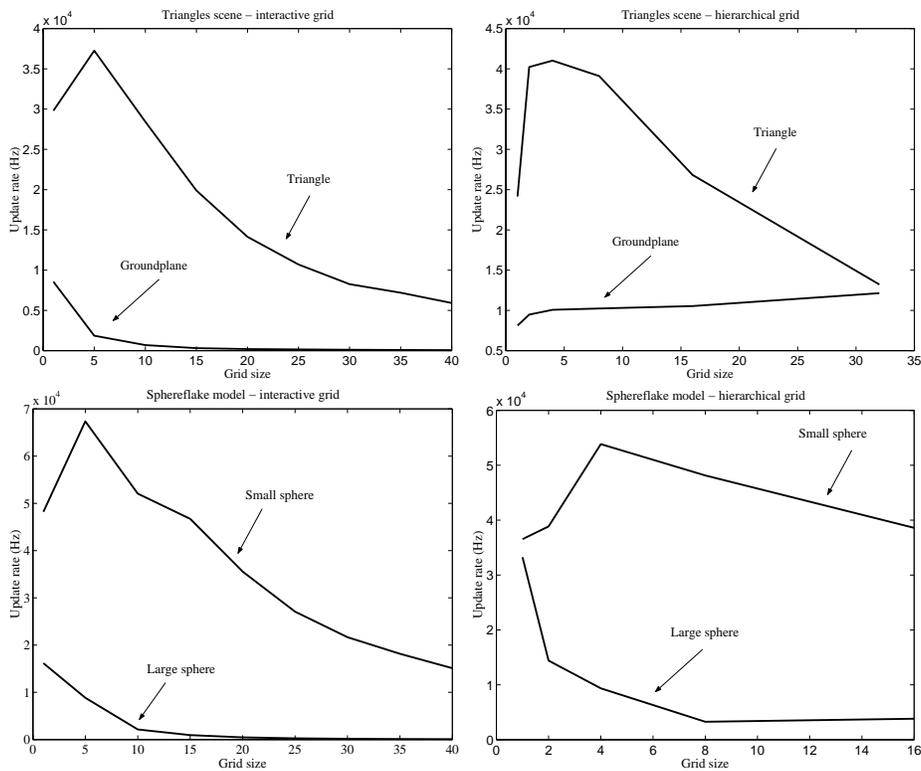


Figure III.5: Update rate as function of grid size for the interactive and hierarchical grids.. We compare the update rates for a small as well as a large object in both the triangles model (top) and the sphereflake model (bottom).

The time to rebuild a spatial subdivision from scratch is expected to be considerably higher than the cost of re-inserting a small number of objects. For the triangles scene, where 200 out of 201 objects were animated, the update rate was still a factor of two faster than the cost of completely rebuilding the spatial subdivision. This was true for both the interactive grid and the hierarchical grid. A factor of two was also found for the animation of 81 spheres in the sphereflake scene. When animating only 9 objects in this scene, the difference was a factor of 10 in favor of updating. We believe that the performance difference between rebuilding the acceleration structure and updating all objects is largely due to the cost of memory allocation, which occurs when rebuilding. The cost of rebuilding the spatial subdivision will become prohibitive when much larger scenes are rendered.

### Traversal cost - dynamic scenes

In the case of expanding scenes, the logical bounding box will become larger than the physical bounding box. The number of voxels that are traversed per ray will therefore on average increase. This is the case in the triangles scene. The variation over time of the frame rate is given in Figure III.6. In this example, the objects are first stationary. At some point the animation starts and the frame rate drops because the scene immediately starts expanding. For the sphereflake scene, the animated objects do not cause the scene to expand, and therefore no drop in framerate is observed.

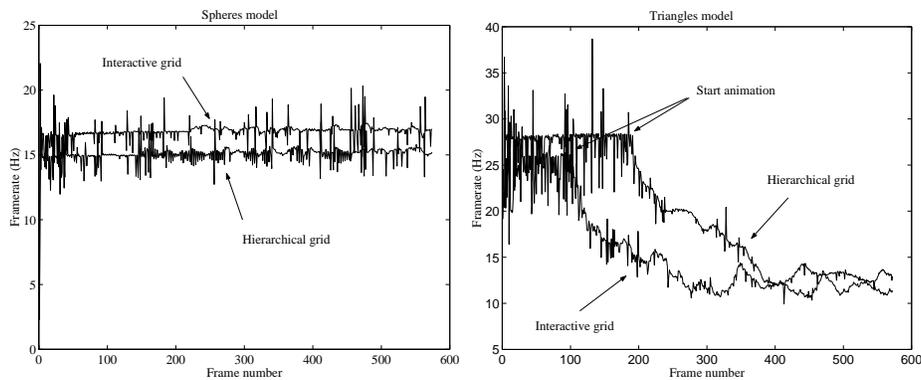


Figure III.6: Framerate as function of time for the triangles scene and the sphereflake scene. Note that the sphereflake scene does not expand over time and therefore starting the animation does not appreciably affect the framerate.

### Animating clusters of objects

For many applications it will be necessary to animate clusters of objects in a coherent manner. For example, if a teapot such as depicted in Figure III.7, needs to be repositioned, it would not make sense to individually move each of its 25,000 individual

polygons. Encapsulating the teapot within its own spatial subdivision will improve rendering time but will not improve insertion and deletion time, as after moving the teapot, all its polygons would still require updating. Here, the use of instancing provides a good solution as only the transformation matrix specifying where the teapot is positioned in space will need to be updated. For this example, updating the spatial subdivision as well as the transformation matrix can be performed around 2400 times per second. The benchmark for this scene resulted in a frame rate of 12.1 fps<sup>3</sup>.

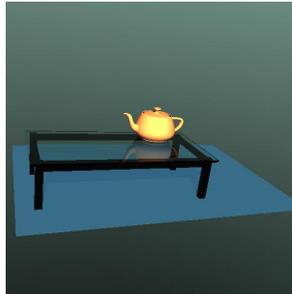


Figure III.7: Example of instancing. Moving the teapot requires a cheap update of a transformation matrix.

Finally, Figure III.22 shows that interactively updating scenes using drag and drop interaction is feasible.

### III.3.3 Discussion

When objects are interactively manipulated and animated within a ray tracing application, much of the work that is traditionally performed during a pre-processing step becomes a limiting factor. Especially spatial subdivisions which are normally built once before the computation starts, do not exhibit the flexibility that is required for animation. The insertion and deletion costs can be both unpredictable and variable. We have argued that for a small cost in traversal performance flexibility can be obtained and insertion and deletion of objects can be performed in a well controlled amount of time.

By logically extending the (hierarchical) grids into space, these spatial subdivisions deal with expanding scenes rather naturally. For modest expansions, this does not significantly alter the frame rate. When the scenes expand a great deal, rebuilding the entire spatial subdivision may become necessary. For large scenes this may involve a temporary drop in frame rate. For applications where this is unacceptable, it would be advisable to perform the rebuilding within a separate thread (rather than the display thread) and use double buffering of the scene to minimize the impact on the rendering threads.

---

<sup>3</sup>Result obtained on a 32-node SGI Origin 2000.

### III.4 Sample reuse techniques

As argued in previous chapters, interactive Whitted-style ray tracing has recently become feasible on high-end parallel machines [15, 18]. However, such systems only maintain interactivity for relatively simple scenes or small image sizes, due to the brute-force nature of these approaches. While keeping the algorithm as simple as possible is an important factor for their success, reasonably straightforward extensions have been devised to improve visual appearance for much larger image sizes and scene complexities. After a brief overview, one such system is explored further in this chapter.

By reusing samples instead of relying on brute force approaches, the limitations in scene complexity and image size can be overcome. There are several ways to reuse samples. All of them require interpolating between existing samples as the key part of the process. First, rays can be stored along with the color seen along them. The color of new rays can be interpolated from existing rays [3, 13]. Alternatively, the points in 3D where rays strike surfaces can be stored and then woven together as displayable surfaces [23]. This method was designed to display course results by a display processor while new samples are created by a rendering back-end which can consist of one or more renderers. As new results become available to the display processor, the image is refined and redisplayed. Finally, stored points can be directly projected to the screen, and holes can be filled in using image processing heuristics [28]. All techniques that re-use samples rely on the fact that the reprojection step is much cheaper than the generation of new samples and are therefore typically employed in cases where sample generation is too slow for creating interactive results. In the case of Simmons' work, this occurred because the lighting simulation is too complex for interactive display [23]. Walter's point reprojection algorithm is directed towards interactive display of scenes that are too complex to display interactively otherwise.

Another method to increase the interactivity of ray tracing is *frameless rendering* [4, 7, 18, 30]. Here, a master processor farms out single pixel tasks to be traced by the slave processors. The order in which pixels are selected is random or quasi-random. Whenever a renderer finishes tracing its pixel, it is displayed directly. As pixel updates are independent of their display, there is no concept of frames. During camera movements, the display will deteriorate somewhat, which is visually preferable to slow frame-rates in frame-based rendering approaches. It can therefore handle scenes of higher complexity than brute force ray tracing, although no samples are reused.

The main thrust of this chapter is the use of parallelism to increase data reuse and thereby increase allowable scene complexity and image size without affecting perceived update rates. The remainder of this chapter uses the *render cache* of Walter et al. [28] and applies to it the concept of frameless rendering. By distributing this algorithm over many processors we are able to overcome the key bottleneck in the original render cache work. We demonstrate our system on a variety of scenes and image sizes that have been out of reach for previous systems. The work described in this chapter is currently under submission for the IEEE 2001 Symposium on Parallel and Large-Data Visualization and Graphics [21].

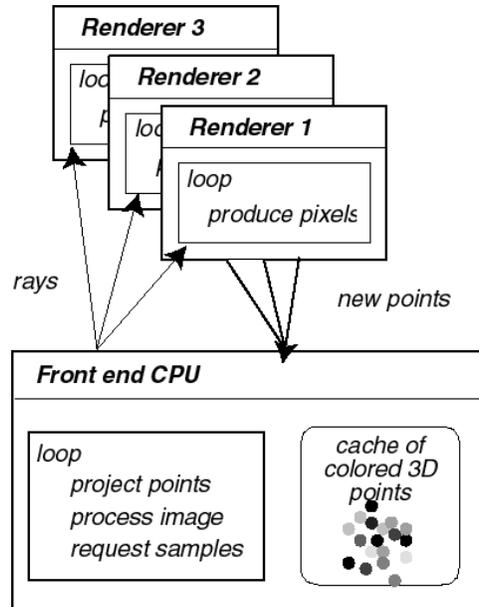


Figure III.8: *The serial render cache algorithm [28].*

### III.4.1 Render cache algorithm

The basic idea of the render cache is to save samples in a 3D point cloud, and reproject them when viewing parameters change [28]. New samples are requested all over the screen, with most samples concentrated near depth discontinuities. As new samples are added old samples are eliminated from the point cloud.

The basic process is illustrated in Figure III.8. The front-end CPU handles all tasks other than tracing rays. Its key data structure is the cache of colored 3D points. The front end continuously loops, first projecting all points in the cache into screen space. This will produce an image with many holes, and the image is processed to fill these holes in. This filling-in process uses sample depths and heuristics to make the processed image look reasonable. The processed image is then displayed on the screen. Finally, the image is examined to find “good” rays to request to improve future images. These new rays are traced by the many CPUs in the “rendering farm”. The current frame is completed after the front end receives the results and inserts them into the point cloud.

From a parallel processing point of view, the render cache has the disadvantage of a single expensive display process that needs to feed a number of renderers with sample requests and is also responsible for point reprojection. The display process needs to insert new results into the point cloud, which means that the more renderers are used, the heavier the workload of the display process. Hence, the display process quickly becomes a bottleneck. In addition, the number of points in the point cloud is linear in image size, which means that the reprojection cost is linear in image size.

The render cache was shown to work well on 256x256 images using an SGI Origin 2000 with 250MHz R10k processors. At higher resolutions than 256x256, the front end has too many pixels to reproject to maintain fluidity.

### III.4.2 Parallel render cache

Ray tracing is an irregular problem, which means that the time to compute a ray task can vary substantially depending on depth complexity. For this reason it is undesirable to run a parallel ray tracing algorithm synchronously, as this would slow down rendering of each frame to be as slow as the processor which has the most expensive set of tasks. On the other hand, synchronous operation would allow a parallel implementation of the render cache to produce exactly the same artifacts as the original render cache. We have chosen responsiveness and speed of operation over minimization of artifacts by allowing each processor to update the image asynchronously.

Our approach is to distribute the render cache functionality with the key goal of not introducing synchronization, which is analogous to frameless rendering. In our system there will be a number of renderers which will reproject point clouds and render new pixels, thereby removing the bottleneck from the original render cache implementation. Scalability is therefore assured.

We parallelize the render cache by subdividing the screen into a number of tiles. A random permutation of the list of tiles could be distributed over the processors, with each renderer managing its set of tiles independently from all other renderers. Alternatively, a global list of tiles could be maintained with each processor choosing the tile with the highest priority whenever it needs a new task to work on. While the latter option may provide better (dynamic) load balancing, we have opted for the first solution. Load balancing is achieved statically by ensuring that each processor has a sufficiently large list of tiles. The reason for choosing a static load balancing scheme has to do with memory management on the SGI Origin 3800, which is explained in more detail in Section III.4.3 and Appendix IIIA.

Each tile has associated with it a local point cloud and an image plane data structure. The work associated with a tile depends on whether or not camera movement is detected. If the camera is moving, the point cloud is projected onto the tile's local image plane and the results are sent to the display thread for immediate display. No new rays are traced, as this would slow down the system and the perceived smoothness would be affected. This is at the cost of a degradation in image quality, which is deemed more acceptable than a loss of interaction. It is also the only modification we have applied to the render cache concept.

If there is no camera movement, a depth test is performed to select those rays that would improve image quality most. Other heuristics such as an aging scheme applied to the points in the point cloud also aid in selecting appropriate new rays. Newly traced rays are both added to the point cloud and displayed on screen. The point cloud itself does not need to be reprojected.

The renderers each loop over their allotted tiles, executing for each tile in turn the following main components:

1. **Clear tile** Before points are reprojected, the tile image is cleared.

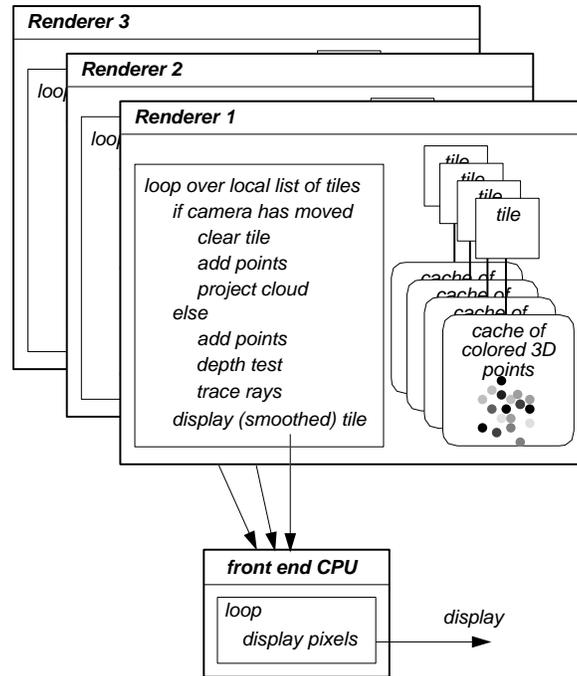


Figure III.9: *The parallel render cache algorithm.*

2. **Add points** Points that previously belonged to a neighboring tile but have been projected onto the current tile are added to the point cloud.
3. **Project point cloud** The point cloud is projected onto the tile image. Points that project outside the current tile are temporarily buffered in a data structure that is periodically communicated to the relevant neighboring tiles.
4. **Depth test** A depth test is performed on the tile image to determine depth discontinuities. This is then used to select new rays to trace.
5. **Trace rays** The rays selected by the depth test function, are traced and the results added to the local point cloud.
6. **Display tile** The resulting tile is communicated to the display thread. This function also performs hole-filling to improve the image's visual appearance.

If camera movement has occurred since a tile was last visited, items 1, 2, 3 and 6 in this list are executed for that tile. If the camera was stationary, items 1, 2, 3 and 6 are executed. The algorithm is graphically depicted in Figure III.9

While tiles can be processed largely independently, there are circumstances when interaction between tiles is necessary. This occurs for instance when a point in one tile's point cloud projects to a different tile (due to camera movement). In that case, the point is removed from the local point cloud and is inserted into the point cloud

associated with the tile to which it projects. The more tiles there are, the more often this would occur. This conflicts with the goal of having many tiles for load balancing purposes. In addition, having fewer tiles that are larger causes tile boundaries to be more visible.

As each renderer produces pixels that need to be collated into an image for display on screen, there is still a display process. This display thread only displays pixels and reads the keyboard for user input. Displaying an image is achieved by reading an array of pixels that represents the entire image, and sending this array to the display hardware using OpenGL. When renderers produce pixels, they are buffered in a local data structure, until a sufficient number of pixels has been accumulated for a write into the global array of pixels. This buffering process ensures that memory contention is limited for larger image sizes.

Finally, the algorithm shows similarities with the concept of frameless rendering, in the sense that tiles are updated independently from the display process. If the size of the tiles is small with respect to the image size, the visual effect is like that of frameless rendering. The larger the tile size is chosen, the more the image updating process starts to look like a distributed version of the render cache.

### III.4.3 Implementation details

The parallel render cache algorithm is implemented on a 32 processor SGI Origin 3800. While this machine has a 16 GB shared address space, the memory is physically distributed over a total of eight nodes. Each node features four 400 MHz R12k processors and one 2 GB block of memory. In addition each processor has an 8 MB secondary cache. Memory access times are determined by the distance between the processor and the memory that needs to be read or written. The local cache is fastest, followed by the memory associated with a processor's node. If a data item is located at a different node, fetching it may incur a substantial performance penalty.

A second issue to be addressed is that the SGI Origin 3800 may relocate a rendering process with a different processor each time a system call is performed. Whenever this happens, the data that used to be in the local cache is no longer locally available. Cache performance can thus be severely reduced by migrating processes.

These issues can be avoided on the SGI Origin 3800 by actively placing memory near the processes and disallowing process migration. This can, for example, be accomplished using the *dplace* library or the *mmap* system call (see also Section III.2.6). Associated with each tile in the parallel render cache is a local point cloud data structure and an image data structure which are mapped as close as possible to the process that uses it. Such memory mapping assures that if a cache miss occurs for any of these data structures, the performance penalty will be limited to fetching a data item that is in local memory. As argued above, this is much cheaper than fetching data from remote nodes. For this reason, using a global list of tiles as mentioned in the previous section is less efficient than distributing tiles statically over the available processors.

Carefully choreographing the mapping of processes to processors and their data structures to local memory enhances the algorithm's performance. Cache performance is improved and the number of data fetches from remote locations is minimized.

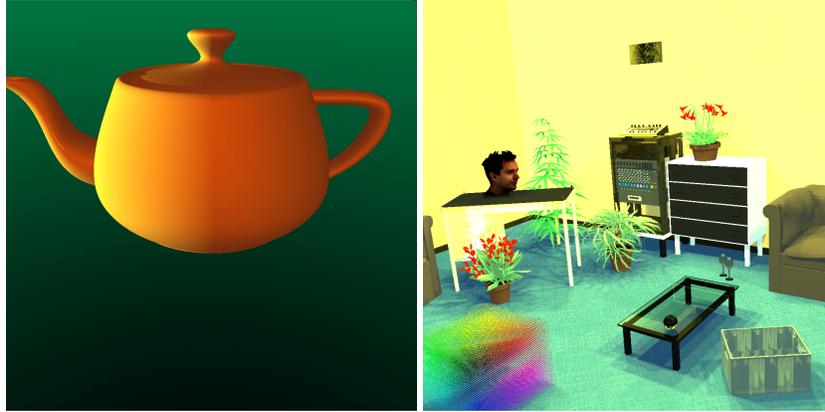


Figure III.10: *Test scenes. The teapot (top) consists of 32 bezier patches, while the room scene consists of 846,563 primitives and 80 point light sources.*

#### III.4.4 Measuring scalability

The main loops of the renderers consist of a number of distinct steps. During each iteration a subset of these steps is executed dependent on whether camera movement has occurred or not (see Section III.4.2). Standard speed-up measurements would under these circumstances produce unreliable results, since the measured speed-up would depend on how often the user moves the camera. The user cannot be expected to move the camera in exactly the same way for each measurement.

For this reason each of the steps making up the complete algorithm are measured separately. To assess scalability, the time to execute each step is measured, summed over all invocations and processors and subsequently divided by the number of invocations and processors. The result is expressed in events per second per processor, which for a scalable system should be independent of the number of processors employed. Hence, using more processors would then not alter the measurements. In case this measure varies with processor count, scalability is affected.

If the number of events per second per processor drops when adding processors, sublinear scalability is measured, whereas an increase indicates super-linear speed-up for the measured function. Also note that the smaller the number, the more costly the operation will be. Using this measure provides better insight into the behavior of the various parts of the algorithm than a standard scalability computation would give, especially since only a subset of the components of the render cache algorithm is executed during each iteration.

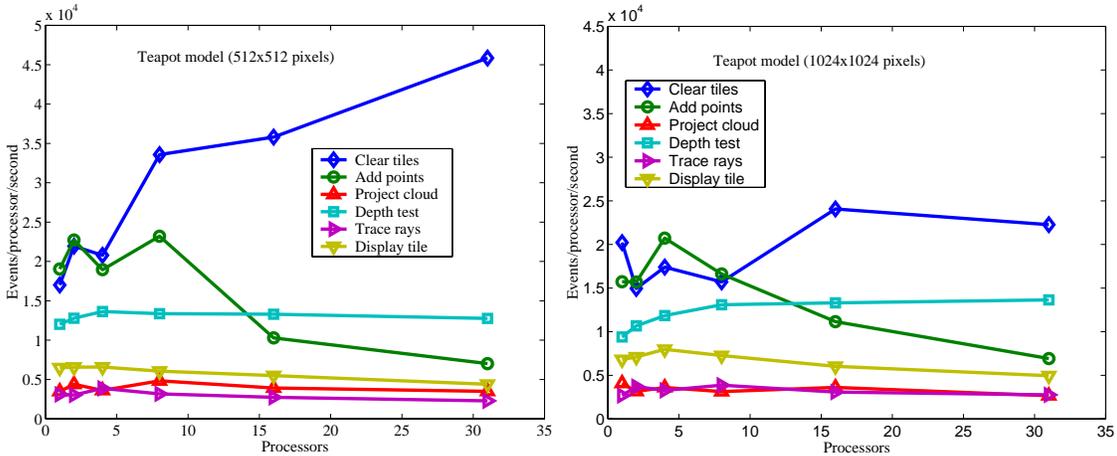


Figure III.11: Scalability of the render cache components for the teapot scene rendered at  $512^2$  pixels (left) and  $1024^2$  pixels (right). Negative slopes indicate sub-linear scalability, whereas horizontal lines show linear speed-ups.

### III.4.5 Results

Our implementation uses the original render cache code of Walter et al [28]<sup>4</sup>. Two test scenes were used: a teapot with 32 bezier patches<sup>5</sup> and one point light source, and a room scene with 846,563 geometric primitives and area light sources approximated by 80 point light sources (Figure III.10). For the teapot scene, the renderer is limited by the point reprojection algorithm, while for the room scene, tracing new rays is the slowest part of the algorithm. The latter scene is of typical complexity in architectural applications and usually cannot be interactively manipulated.

In the following subsection, the different components making up the parallel render cache are evaluated (Section III.4.5), the performance as function of task size is assessed (Section III.4.5) and the parallel render cache is compared with other methods to speed up interactive ray tracing (Section III.4.5).

#### Parallel render cache evaluation

The results of rendering the teapot and room models on different numbers of processors at a resolution of  $512^2$  and  $1024^2$  pixels are depicted in Figures III.11 and III.12.

While most of the components making up the algorithm show horizontal lines in these graphs, meaning that they scale well, the “Clear tiles” and “Add point” components show non-linear behavior. Clearing tiles is a very cheap operation which appears to become cheaper if more processors are used. Because more processors result in each

<sup>4</sup>The original code has since been improved (Walter, personal communication) but we have not ported that improved code. However, we expect that any improvements to the serial code would transfer to our parallel version since the serial code runs essentially as a black box.

<sup>5</sup>These bezier patches are rendered directly using the intersection algorithm from Parker et. al [18].

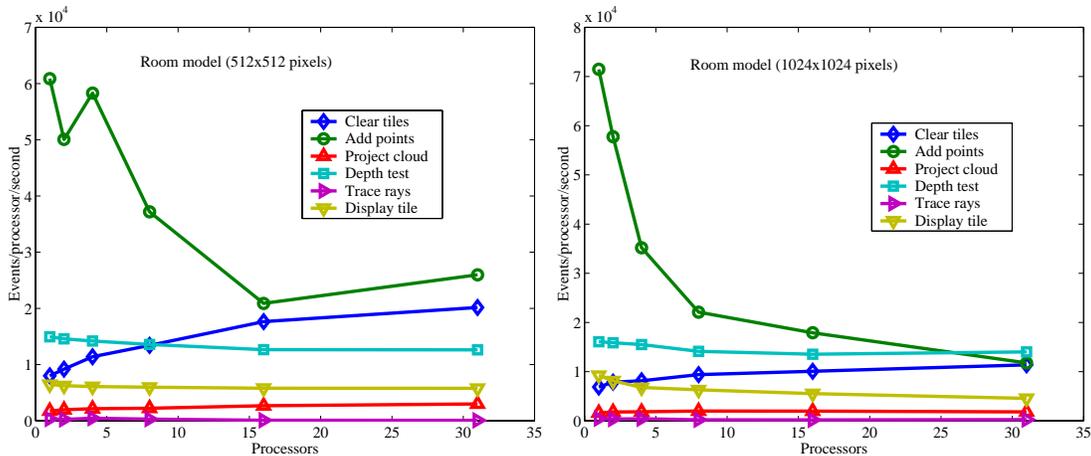


Figure III.12: Scalability for the room scene, rendered at  $512^2$  pixels (left) and  $1024^2$  pixels (right). Horizontal lines indicate linear scalability, whereas a fall-off means sub-linear scalability.

processor having to process fewer tiles, this super-linear behavior may be explained by better cache performance. This effect is less pronounced for the  $1024^2$  pixel renderings, which also points to a cache performance issue as here each processor handles more data.

The “Add point” function scales sub-linearly with the number of processors. Because the total number of tiles was kept constant between runs, this cannot be explained by assuming that different numbers of points project outside their own tile and thus have to be added to neighboring tiles. However, with more processors there is an increased probability that a neighboring tile belongs to a different processor and may therefore reside in memory which is located elsewhere in the machine. Thus projecting a point outside the tile that it used to belong to, may become more expensive for larger numbers of processors. This issue is addressed in the following section.

Note also that despite the poor scalability of “Add points”, in absolute terms its cost is rather low, especially for the room model. Hence, the algorithm is bounded by components that scale well (they produce more or less horizontal lines in plots) and therefore the whole distributed render cache algorithm scales well, at least up to 31 processors (see also Section III.4.5). In addition, the display of the results is completely decoupled from the renderers which produce new results and therefore the screen is updated at a rate that is significantly higher than rays can be traced and is also much higher than points can be reprojected. This three-tier system of producing new rays at a low frequency, projecting existing points at an intermediate frequency and displaying the results at a high frequency (on the Origin 3800 at a rate of around 290 frames per second for  $512^2$  images and 75 frames per second for  $1024^2$  images, regardless of number of renderers and scene complexity) ensures a smooth display which is perceived as interactive, even if new rays are produced at a rate that would not normally

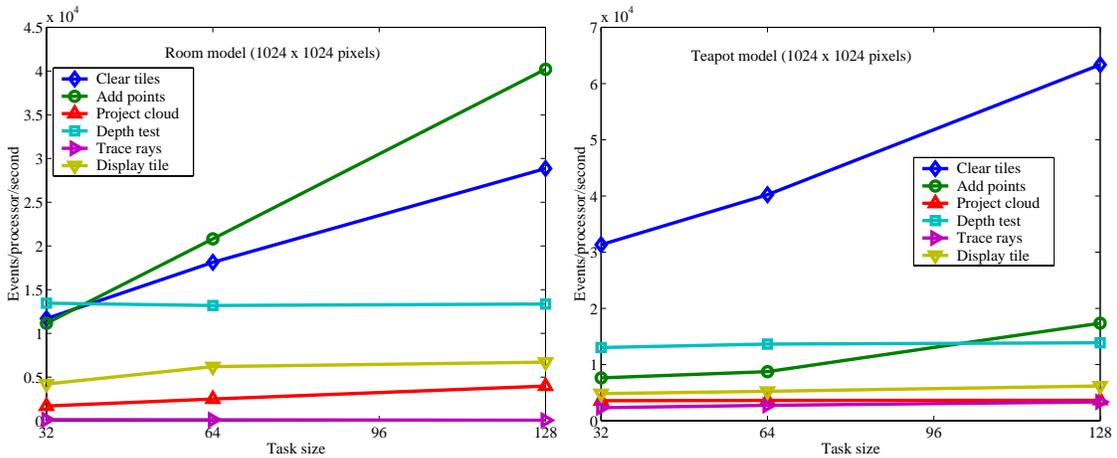


Figure III.13: Scalability for the room model (left) and teapot scene (right) as function of tile size ( $32^2$ ,  $64^2$  and  $128^2$  pixels per tile). The image size is  $1024^2$  pixels and for these measurements 31 processors were used. These graphs should be interpreted the same as those in Figures III.11 and III.12.

allow interactivity.

By abandoning ray tracing altogether during camera movement, the system shows desirable behavior even when fewer than 31 processors are used. For both the room scene and the teapot model, the camera can move smoothly if 4 or more processors are used. During camera movement, the scene deteriorates because no new rays are produced and holes in the point cloud may become visible. During rapid camera movement, tile boundaries may become temporarily visible. After the camera has stopped moving, these artifacts disappear at a rate that is linear in the number of processors employed. We believe that maintaining fluid motion is more important than the temporary introduction of some artifacts, which is why the distributed render cache is organized as described above.

For those who would prefer a more accurate display at the cost of a slower system response, it would be possible to continue tracing rays during camera movement. Although the render cache then behaves differently, the scalability of the separate components, as given in Figures III.11 and III.12, would not change. However, the fluidity of camera movement is destroyed by an amount dependent on scene complexity.

### Task size

In section III.4.2 it was argued that the task size, i.e. the size of the tiles, is an important parameter which defines both speed and the occurrence of visual artifacts. The larger the task size, the better artifacts become visible. However, at the same time, the reprojections that cross tile-boundaries are less likely to occur, resulting in higher performance. In Figure III.13 the scalability of the parallel render cache components as function of task size is depicted. Task sizes range from  $32^2$  pixels to  $128^2$  pixels

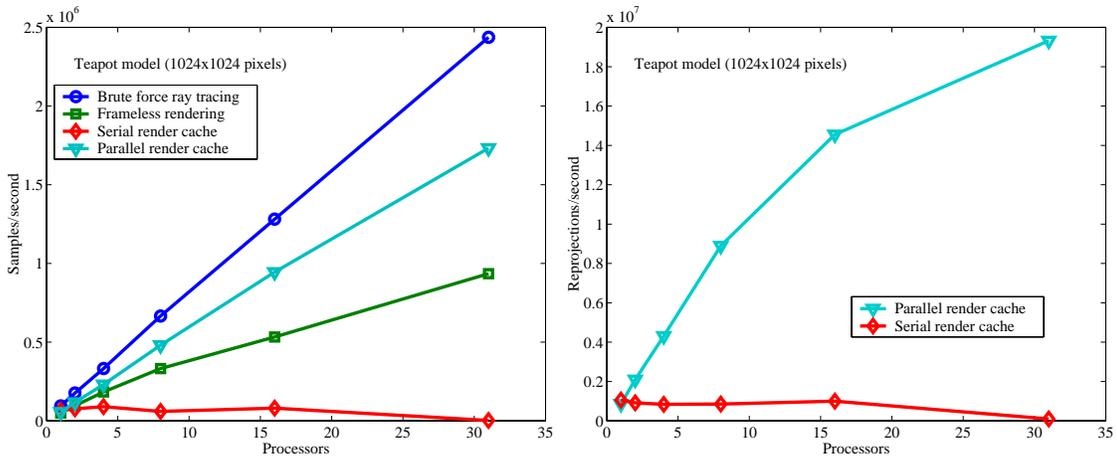


Figure III.14: *Samples per second (left) and point reprojections per second (right) for the teapot model.*

and the measurements were all obtained using 31 processors on  $1024^2$  images. Larger tile sizes are thus impossible, as the total number of tasks would become smaller than the number of processors. Task sizes smaller than  $32^2$  pixels resulted in unreasonably slow performance and were therefore left out of the assessment.

As in the previous section, the “Add points” and “Clear tile” components show interesting behavior. As expected, for larger tasks, the “Add points” function becomes cheaper. This is because the total length of the tile boundaries diminishes for larger task sizes, and so the probability of reprojections occurring across tile boundaries is smaller.

The “Clear tile” component also becomes less expensive for larger tiles. Here, we suspect that resetting one large block of memory is less expensive than resetting a number of smaller blocks of memory.

Although Figure III.13 suggests that choosing the largest task size as possible would be appropriate, the artifacts visible for large tiles are more unsettling than for smaller task sizes. Hence, for all other experiments presented in this paper, a task size of  $32^2$  pixels is used, which is based on an assessment of both artifacts and performance.

### Comparison with other speed-up mechanisms

In this section, the parallel render cache is compared with other state-of-the-art rendering techniques. All make use of the interactive ray tracer of Parker et. al. [18], either as a back-end or as the main algorithm. The comparison includes the original render cache algorithm [28], the parallel render cache algorithm as described in this paper, the interactive ray tracer (rtrt) without reprojection techniques and the interactive ray tracer using the frameless rendering concept [18]. In the following we will refer to the original render cache as “serial render cache” to distinguish it from our parallel render

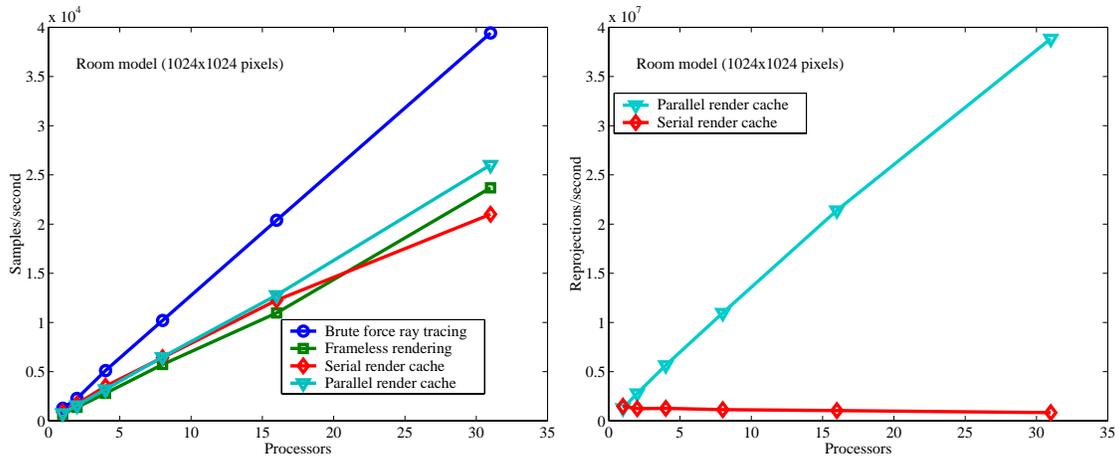


Figure III.15: *Samples per second (left) and point reprojections per second (right) for the room scene.*

cache implementation. All renderings were made using the teapot and room models (Figure III.10) at a resolution of  $1024^2$  pixels.

The measurements presented in this section consist of the number of new samples produced per second by each of the systems and the number of points reprojected per second (for the two render cache algorithms). These numbers are summed over all processors and should therefore scale with the number of processors employed. The results for the teapot model are given in Figure III.14 and the results for the room model are presented in Figure III.15.

The graphs on the left of these figures show the number of samples generated per second. All the lines are straight, indicating scalable behavior. In these plots, steeper lines are the result of higher efficiency and therefore, the real-time ray tracer would be most efficient, followed by the parallel render cache. The frameless rendering concept loses efficiency because randomizing the order in which pixels are generated destroys cache coherence. The parallel render cache does not suffer from this, since the screen is tiled and tasks are based on tiles. The serial render cache appears to perform well for complex scenes and poorly for simple scenes. For scenes that lack complexity, the point reprojection front-end becomes the bottleneck, especially since the image size chosen causes the point cloud to be quite large. Thus, the render cache front-end needs to reproject a large number of points for each frame and so constitutes a bottleneck.

Although the parallel render cache does not produce as many new pixels as the real-time ray tracer by itself does, this loss of efficiency is compensated by its ability to reproject large numbers of points, as is shown in the plots on the right of Figures III.14 and III.15. The point reprojection component of the parallel render cache shows good scalability, and therefore the goal of parallelizing the render cache algorithm is reached. The point reprojection part of the serial render cache does not scale because it is serial in nature.

### III.4.6 Discussion

While it is true that processors get ever faster and multi-processor machines are now capable of real-time ray tracing, scenes are getting more and more complex while at the same time frame sizes still need to increase. Hence, Moore's law is not likely to allow interactive full-screen brute-force ray tracing of highly complex scenes anytime soon.

Interactive manipulation of complex models is still not possible without the use of sophisticated algorithms that can efficiently exploit temporal coherence. The render cache is one such algorithm that can achieve this. However, for it not to become a bottleneck itself, the render cache functionality needs to be distributed over the processors that produce new samples. The resulting algorithm, presented in this paper, shows superior reprojection capabilities that enables smooth camera movement, even in the case where the available processing power is much lower than would be required in a brute force approach. It achieves this for scene complexities and image resolutions that are not feasible using any of the other algorithms mentioned in the previous section.

While smoothness of movement is an important visual cue, our algorithm necessarily produces other artifacts during camera motion. These artifacts are deemed less disturbing than jerky motion and slow response times. The render cache attempts to fill small holes after point reprojection. For larger holes, this may fail and unfilled pixels may either be painted in a fixed color, or can be left unchanged from previous reprojections. Either approach causes artifacts inherent to the algorithm and is present both in the original render cache and in our parallel implementation of it.

The parallel render cache produces additional artifacts due to the tiling scheme employed. During camera movement, tile boundaries may temporarily become visible, because there is some latency between points being reprojected from neighboring tiles and this reprojection becoming visible in the current tile. A further investigation to minimize these artifacts is in order, which we reserve for future work. Currently, the parallel render cache algorithm is well suited for navigation through highly complex scenes to find appropriate camera positions.

It has been shown that even with a relatively modest number of processors, the distributed render cache can produce smooth camera movement at resolutions typically sixteen times higher than the original render cache. The system as presented here scales well up to 31 processors. Its linear behavior suggests that improved performance is likely beyond 31 processors, although if this many processors are available, it would probably become sensible to devote the extra processing power to produce more samples, rather than increase the speed of reprojection.

### III.5 Summary and discussion

In this chapter we have explained some of the issues involved in implementing an interactive ray tracer. Using general purpose hardware, currently shared memory architectures appear the most attractive solution, since distributed memory architectures do not allow pixel data to be communicated to the display quickly enough. Distributed memory systems are better suited for coarse grain applications, such as ray tracing with diffuse interreflection [20].

Modern shared memory systems often have a single shared address space, but the memory is still physically distributed. Knowing the architecture for which code is written, may help reduce memory access bottleneck. Mapping memory close to the processes that most often access this memory, can be advantageous, especially if the volume of data that is read from memory is large compared with the local cache size. In that case, cache misses will cause local main memory to be read rather than remote main memory.

Other optimizations that are important include cache optimization. Important data structures should fit on a single cache line. Smaller data structures can be padded to occupy a single cache line. This has the advantage that when a cached item is swapped out because it has not been used for a while, it can not affect other data on the same cache line that is still in use.

While the basic interactive ray tracer described in this chapter is based on such low level optimizations, algorithmic extensions to this basic ray tracer can be employed to allow much larger images to be computed or have scenes of much higher complexity rendered interactively. Point reprojection techniques do not allow rays to be produced quicker, but rely on frame coherence by reusing samples computed for previous frames. For complex scenes, reprojection of existing results is much faster than tracing new rays and so point reprojection allows for smooth movement between camera points for scenes that are too complex for other algorithms to smoothly advance from one camera position to the next.

In this chapter extensions to spatial subdivisions are discussed as well. Normally, spatial acceleration structures are built as a preprocess and are therefore not flexible enough to accomodate interactively placed or moved objects. Extending grid and octree data structures to enable user interaction with the scene interactively have a modest impact on speed of rendering which is acceptable given their ability to allow objects to be moved within and even outside the extent of the scene.

Interactive ray tracing is now feasible and for certain types of application, such as interactive rendering of the visible female data set [19], it is a better choice than other forms of rendering, including z-buffer techniques. With increasing available computational power, the range of applications for interactive ray tracing is likely to grow and become possible on cheaper hardware.

## **Acknowledgements**

We would like to thank Brian Smits, Chuck Hansen and Pete Shirley for their help, support and involvement in the projects that are described in these course notes. In addition, we thank Bruce Walter, Steven Parker and George Drettakis for their render cache source code and John McCorquodale for valuable discussions regarding processor and memory placement issues. Thanks also to Silicon Graphics and Springer Verlag for their friendly cooperation and allowing us to republish appendices IIIA to IIIC. This work was supported by NSF grants NSF/ACR, NSF/MRI and by the DOE AVTC/VIEWS.

## Appendices

Appendices IIIA and IIIB contain extracts from the following SGI technical report, which are reproduced with permission from Silicon Graphics.

David Cortesi and Jeff Fier, *Origin2000 and Onyx2 Performance Tuning and Optimization Guide*, Silicon Graphics technical report, 1998, available from <http://techpubs.sgi.com:80/>

Appendix IIIC contains the following paper, which is reproduced with permission from Springer Wien - New York:

Reinhard, E., Smits, B., Hansen, C., *Dynamic Acceleration Structures for Interactive Ray Tracing*. in: *Rendering Techniques 2000* (Proceedings of the Eurographics Workshop in Brno, Czech Republic, 2000). pp299-306. Wien - New York: Springer. 2000

## **IIIA SGI Origin 2000**

In this section, the SGI Origin 2000 architecture is described, along with the R10k processor and its facilities for profiling. This will constitute basic knowledge for those who want to write optimized code on this hardware platform. The following subsections are extracted from the SGI document “Origin2000 and Onyx2 Performance Tuning and Optimization Guide” [6].

### **IIIA.1 The MIPS R10000 processor**

The R10000 has a two-level cache hierarchy. Located on the microprocessor chip are a 32 KB, two-way set associative level-1 instruction cache and a 32 KB, two-way set associative, two-way interleaved level-1 (L1) data cache. Off-chip is a two-way set associative, unified (instructions and data) level-2 (L2) cache. This secondary cache may range in size from 512 KB to 16 MB; the size of the secondary cache in the Origin 2000 is 4 MB for 195 MHz systems, and 1 MB for 180 MHz systems. The L1 instruction cache uses a line size of 64 bytes, while the L1 data cache has a line size of 32 bytes. The line size of the L2 cache may be either 64 or 128 bytes; in the Origin 2000 it is 128 bytes. Both the L1 data cache and the L2 unified cache employ a least recently used (LRU) replacement policy for selecting in which set of the cache to place a new cache line.

### **IIIA.2 Origin 2000 layout**

To understand how the Origin2000’s scalable shared memory multiprocessor (S2MP) architecture works, we first look at how the building blocks of an Origin system are connected. This is diagrammed in figure III.16. This figure represents Origin systems ranging from 2 to 16 processors. We start by considering the two-processor system in the upper left-hand corner. This is a single Origin 2000 node. It consists of one or two processors, memory, and a device called the hub. The hub is the piece of hardware that carries out the duties that a bus performs in a bus-based system; namely, it manages each processor’s access to memory and I/O. This applies to accesses that are local to the node containing the processor, as well as to those that must be satisfied remotely in multi-node systems.

The smallest Origin systems consist of a single node. Larger systems are built by connecting multiple nodes. A two-node system is shown in the upper middle of the figure. Since information flow in and out of a node is controlled by the hub, connecting two nodes means connecting their hubs. In a two-node system this simply means wiring the two hubs together. The bandwidth to local memory in a two-node system is double that in a one-node system: the hub on each of the two nodes can access its local memory independently of the other. Access to memory on the remote node is a bit more costly than access to local memory since the request must be handled by both hubs. A hub determines whether a memory request is local or remote based on the physical address of the data accessed.

When there are more than two nodes in a system, their hubs cannot simply be wired together. In this case, additional hardware is required to control information

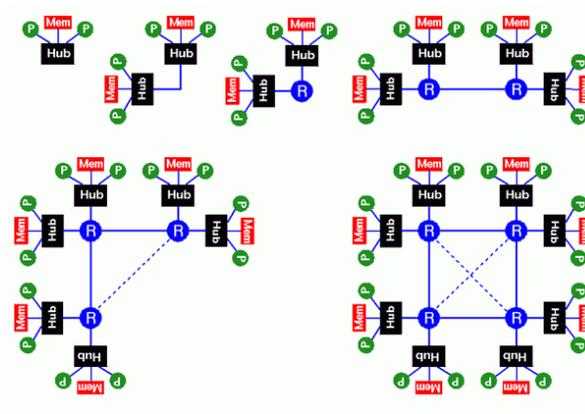


Figure III.16: Building blocks of the SGI Origin 2000.

flow between the multiple hubs. The hardware used for this in Origin systems is called a router. A router has six ports, so it may be connected to up to six hubs or other routers. In a two-node system, one may employ a router to connect the two hubs rather than wiring them directly together; this is shown adjacent to the other two-node configuration in the figure. These two different configurations behave identically, but because of the router in the second configuration, information flow between the two hubs takes a little more time. The advantage, though, of the configuration with the router is that it may be used as a basic building block from which to construct larger systems.

In the upper right corner of the figure, a four-node — or, equivalently, eight-processor — system is shown. It is constructed from two of the two-node-with-router building blocks. Here, the connection between the two routers allows information to flow and, hence, the sharing of memory between any pair of hubs in the system. Since a router has six ports, it is possible to connect all four nodes to just one router, and this one-router configuration can be used for small systems. That is a special case, and in general the two-router implementation is used since it conveniently scales to larger systems.

Two such larger systems are shown on the lower half of the figure; these are 12- and 16-processor systems, respectively. From these diagrams you can begin to see how the router configurations scale: each router is connected to two hubs, routers are then connected to each other forming a binary  $n$ -cube, or hypercube, where  $n$ , the dimensionality of the router configuration, is the base-2 logarithm of the number of routers. For the four-processor system,  $n$  is zero, and for the eight-processor system, the routers form a linear configuration, and  $n$  is one. In both the 12- and 16-processor systems,  $n$  is two and the routers form a square; for the 12-processor system, one corner of the square is missing. Larger systems are constructed by increasing the dimensionality of the router configuration and adding up to two hubs with each additional router. Systems with any number of nodes may be constructed by leaving off some corners of the  $n$ -dimensional hypercube. We will see these larger configurations later.

The key thing here is that the hardware allows the physically distributed memory of the system to be shared, just as in a bus-based system, but since each hub is connected to its own local memory, memory bandwidth is proportional to the number of nodes. As a result, there is no inherent limit to the number of processors that can be used effectively in the system. In addition, since the dimensionality of the router configuration grows as the systems get larger, the total router bandwidth also grows with system size (proportional to  $n^2n$ , where  $n$  is the dimensionality of the router configuration). Thus, systems may be scaled without fear that the router connections will become a bottleneck.

To allow this scalability, however, one nice characteristic of the bus-based shared memory systems has been sacrificed; namely, the access time to memory is no longer uniform: it varies depending on how far away the memory being accessed is in the system. The two processors in each node have quick access through their hub to their local memory. Accessing remote memory through an additional hub adds an extra increment of time, as does each router the data must travel through. But several factors combine to smooth out these nonuniform memory access (NUMA) times:

1. The hardware has been designed so that the incremental costs to access remote memory are not large. The choice of a hypercube router configuration means that the number of routers information must pass through is at most  $n + 1$ , where  $n$  is the dimension of the hypercube; this grows only as the logarithm of the number of processors. As a result, the average memory access time on even the largest Origin system is no more than the uniform memory access time on a Power Challenge 10000 system. We'll see a detailed table of these costs later.
2. The R10000 processors operate on data that are resident in their caches. If programs use the caches effectively, the access time to memory, whether it is local or remote, is unimportant since the vast majority of memory accesses are satisfied from the caches.
3. Through operating system support or programming effort, the memory accesses of most programs can be made to come primarily from local memory. Thus, in the same way that the caches can make local memory access times unimportant, remote memory access costs can be reduced to an insignificant amount.
4. The R10000 processors can prefetch data that are not cache resident. Other work can be carried out while these data move from local or remote memory into the cache; thus the access time can be hidden.

The architecture of the Origin 2000 system, then, provides shared memory hardware without the limitations of traditional bus-based designs.

## IIIB Profiling on the SGI Origin 2000

The hardware counters in the R10000 CPU make it possible to profile the behavior of a program in many ways without modifying the code. The software tools are `perfex`, which runs a program and reports exact counts of any two selected events from the R10000 counters. Alternatively, it time-multiplexes all 32 countable events and reports extrapolated totals of each. `Perfex` is useful for identifying what problem (for example, secondary data cache misses) is hurting the performance of your program the most. (see `timex` for simple timing functionality.) The `Perfex` functions are also available as callable library functions in `libperfex`. Similarly, for `speedshop`, the `ssapi` library is available. `Speedshop` (actually, the `ssrun` command), which runs a program while sampling the state of the program counter and stack, and writing the sample data to a file for later analysis. You select the timebase for the sampling and the particular type of information to be sampled. `SpeedShop` is useful for locating where in your program the performance problems occur. `Prof`, which analyzes a `Speedshop` data file and displays it in a variety of formats. `Dprof`, which, like `Speedshop`, samples a program while it is executing but records memory access information as a histogram file. It identifies which data structures in the program are involved in performance problems. Use these tools to find out what constrains the program and which parts of it consume the most time. Through the use of a combination of these tools, it is possible to identify most performance problems.

The profiling tools depend for most of their features on the R10000's performance counter registers. These on-chip registers can be programmed to count hardware events as they happen, for example, machine cycles, instructions, branch predictions, floating point instructions, or cache misses. There are only two performance counter registers. Each can be programmed to count machine cycles or 1 of 15 other events, for a total of 32 events that can be counted (30 of which are distinct). The specific events are summarized in Table III.1, which can be obtained by using the command `perfex -h`.

The counters are 64-bit integers. When a counter overflows, a hardware trap occurs. The kernel can preload a counter with  $2^{64} - n$  to cause a trap after  $n$  counts occur. The profiling tools use this capability. For example, the command `ssrun -gi_hwc` programs the graduated instruction counter (event 17) to overflow every 32 K counts. Each time the counter overflows, `ssrun` samples the program counter and stack state of the subject program. The reference page *r10k\_counters(5)* gives detailed information on how the counters can be accessed through the `/proc` interface. This is the interface used by the profiling tools. The interface hides the division of events between only two registers and allows the software to view the counters as a single set of thirty-two 64-bit counters. The operating system time-multiplexes the active counters between the events being counted. This requires sampling and scaling, which introduce some error when more than two events are counted. In general, it is better to access the counters through the profiling tools. A program that uses the counter interface directly cannot be profiled using `perfex` or using `ssrun` for any experiment that depends on counters. When a program must access counter values directly, the simplest interface is through `libperfex`, documented in the *libperfex* reference page.

0	Cycles	16	Cycles
1	Instructions issued to functional units	17	Instructions graduated
2	Memory data access (load, prefetch, sync, cacheop) issued	18	Memory data loads graduated
3	Memory stores issued	19	Memory data stores graduated
4	Store conditionals issued	20	Store conditionals graduated
5	Store conditionals failed	21	Floating point instructions graduated
6	Branches decoded	22	Quadwords written back from L1 cache
7	Quadwords written back from L2 cache	23	TLB refill exceptions
8	Correctable ECC errors on L2 cache	24	Branches mispredicted
9	L1 cache misses (instruction)	25	L1 cache misses (data)
10	L2 cache misses (instruction)	26	L2 cache misses (data)
11	L2 cache way mispredicted (instruction)	27	L2 cache way mispredicted (data)
12	External intervention requests	28	External intervention request hits in L2 cache
13	External invalidate requests	29	External invalidate request hits in L2 cache
14	Instructions done (formerly, virtual coherency condition)	30	Stores, or prefetches with store hint, to CleanExclusive L2 cache blocks
15	Instructions graduated	31	Stores, or prefetches with store hint, to Shared L2 cache blocks.

Table III.1: *Hardware counters of the R10000 processor.*

### IIIB.1 Performance analysis using **perfex**

The simplest profiling tool is **perfex**, documented in the *perfex* reference page. It runs a subject program and records data about the run, similar to **timex**:

```
% perfex [options] command [arguments]
```

The subject program and its arguments are given. **perfex** sets up the counter interface and forks the subject program. When the program ends, **perfex** writes counter data to standard output. **perfex** gathers its information with no modifications to your existing program. Although this is convenient, the data obtained come from the entire run of the program. If you only want to profile a particular section of the program, you need to use the library interface to **perfex**, *libperfex(3)*. To use this interface, insert a call to initiate counting into your program's source code and another to terminate it; a third call prints the counts gathered. The program must then be linked with the *libperfex* library:

```
% cc -o program -lperfex
```

Since you can use SpeedShop to see where in a program various event counts come from, *libperfex* is not described in detail. More information can be found in its reference page.

### IIIB.2 Absolute counts of one or two events

Use **perfex** options to specify what is to be counted. You can specify one or two countable events. In this case, the counts are absolute and accurate. For example, the command

### **% perfex -e 15 -e 21 adi2**

runs the subject program and reports the exact counts of graduated instructions and graduated floating point instructions. You use this mode to explore specific points of program behavior.

### **IIIB.3 Statistical counts of all events**

When you specify option -a (all events), **perfex** multiplexes all 32 events over the program run. Each count is active 1/16 of the time and then scaled by 16. The resulting counts have some statistical error. The error is small (and the counts sufficiently repeatable) provided that the subject program runs in a stable execution mode for a number of seconds. When the program runs for a short time, or shifts between radically different regimes of instruction or data use, the counts are less dependable and less repeatable. Nevertheless, **perfex -a** usually gives a good overview of program operation. Here is the **perfex** command line and output applied to a sample program called adi2:

#### **% perfex -a -x adi2**

WARNING: Multiplexing events to project totals--inaccuracy possible.

```
Time:          7.990 seconds
Checksum:      5.6160428338E+06
0 Cycles..... 1645481936
1 Issued instructions..... 677976352
2 Issued loads..... 111412576
3 Issued stores..... 45085648
4 Issued store conditionals..... 0
5 Failed store conditionals..... 0
6 Decoded branches..... 52196528
7 Quadwords written back from scache..... 61794304
8 Correctable scache data array ECC errors..... 0
9 Primary instruction cache misses..... 8560
10 Secondary instruction cache misses..... 304
11 Instruction misprediction from scache way prediction table.. 272
12 External interventions..... 6144
13 External invalidations..... 10032
14 Virtual coherency conditions..... 0
15 Graduated instructions..... 371427616
16 Cycles..... 1645481936
17 Graduated instructions..... 400535904
18 Graduated loads..... 90474112
19 Graduated stores..... 34776112
20 Graduated store conditionals..... 0
21 Graduated floating point instructions..... 28292480
22 Quadwords written back from primary data cache..... 32386400
23 TLB misses..... 5687456
24 Mispredicted branches..... 410064
25 Primary data cache misses..... 16330160
26 Secondary data cache misses..... 7708944
27 Data misprediction from scache way prediction table..... 663648
28 External intervention hits in scache..... 6144
29 External invalidation hits in scache..... 6864
30 Store/prefetch exclusive to clean block in scache..... 7582256
31 Store/prefetch exclusive to shared block in scache..... 8144
```

The -x option requests that **perfex** also gather counts for kernel code that handles exceptions, so the work done by the OS to handle TLB misses is included in these counts.

### IIIB.4 Analytic output with the -y option

The raw event counts are interesting, but it is more useful to convert them to elapsed time. Some time estimates are simple, for example, dividing the cycle count by the machine clock rate gives the elapsed run time (1645481936 / 195 MHz = 8.44 seconds). Other events are not as simple and can only be stated in terms of a range of times. For example, the time to handle a primary cache miss varies depending on whether the needed data are in the secondary cache, in memory, or in the cache of another CPU. Analysis of this kind can be requested using **perfex -a -x -y**. When you use -a , -x, and -y, **perfex** collects and displays all event counts, but it also displays a report of estimated times based on the counts. Here is an example, again, of the program adi2:

**% perfex -a -x -y adi2**

WARNING: Multiplexing events to project totals--inaccuracy possible.

Time: 7.996 seconds  
 Checksum: 5.6160428338E+06

Based on 196 MHz IP27				
Event Counter Name	Counter Value	Typical Time (sec)	Minimum Time (sec)	Maximum Time (sec)
0 Cycles .....	1639802080	8.366337	8.366337	8.366337
16 Cycles .....	1639802080	8.366337	8.366337	8.366337
26 Secondary data cache misses .....	7736432	2.920580	1.909429	3.248837
23 TLB misses .....	5693808	1.978017	1.978017	1.978017
7 Quadwords written back from scache .....	61712384	1.973562	1.305834	1.973562
25 Primary data cache misses .....	16368384	0.752445	0.235504	0.752445
22 Quadwords written back from primary data cache .....	32385280	0.636139	0.518825	0.735278
2 Issued loads .....	109918560	0.560809	0.560809	0.560809
18 Graduated loads .....	88890736	0.453524	0.453524	0.453524
6 Decoded branches .....	52497360	0.267844	0.267844	0.267844
3 Issued stores .....	43923616	0.224100	0.224100	0.224100
19 Graduated stores .....	33430240	0.170562	0.170562	0.170562
21 Graduated floating point instructions .....	28371152	0.144751	0.072375	7.527040
30 Store/prefetch exclusive to clean block in scache .....	7545984	0.038500	0.038500	0.038500
24 Mispredicted branches .....	417440	0.003024	0.001363	0.011118
9 Primary instruction cache misses .....	8272	0.000761	0.000238	0.000761
10 Secondary instruction cache misses .....	768	0.000290	0.000190	0.000323
31 Store/prefetch exclusive to shared block in scache .....	15168	0.000077	0.000077	0.000077
1 Issued instructions .....	673476960	0.000000	0.000000	3.436107
4 Issued store conditionals .....	0	0.000000	0.000000	0.000000
5 Failed store conditionals .....	0	0.000000	0.000000	0.000000
8 Correctable scache data array ECC errors .....	0	0.000000	0.000000	0.000000
11 Instruction misprediction from scache way prediction table ..	432	0.000000	0.000000	0.000002
12 External interventions .....	6288	0.000000	0.000000	0.000000
13 External invalidations .....	9360	0.000000	0.000000	0.000000
14 Virtual coherency conditions .....	0	0.000000	0.000000	0.000000
15 Graduated instructions .....	364303776	0.000000	0.000000	1.858693
17 Graduated instructions .....	392675440	0.000000	0.000000	2.003446
20 Graduated store conditionals .....	0	0.000000	0.000000	0.000000
27 Data misprediction from scache way prediction table .....	679120	0.000000	0.000000	0.003465
28 External intervention hits in scache .....	6288	0.000000	0.000000	0.000000
29 External invalidation hits in scache .....	5952	0.000000	0.000000	0.000000
Statistics				
Graduated instructions/cycle .....		0.222163		
Graduated floating point instructions/cycle .....		0.017302		
Graduated loads & stores/cycle .....		0.074595		
Graduated loads & stores/floating point instruction .....		5.422486		
Mispredicted branches/Decoded branches .....		0.007952		
Graduated loads/Issued loads .....		0.808696		
Graduated stores/Issued stores .....		0.761099		
Data mispredict/Data scache hits .....		0.078675		
Instruction mispredict/Instruction scache hits .....		0.057569		
L1 Cache Line Reuse .....		6.473003		
L2 Cache Line Reuse .....		1.115754		
L1 Data Cache Hit Rate .....		0.866185		
L2 Data Cache Hit Rate .....		0.527355		

Time accessing memory/Total time.....	0.750045
L1--L2 bandwidth used (MB/s, average per process).....	124.541093
Memory bandwidth used (MB/s, average per process).....	236.383187
MFLOPS (average per process).....	3.391108

”Maximum,” ”minimum,” and ”typical” time cost estimates are reported. Each is obtained by consulting an internal table which holds the maximum, minimum, and typical costs for each event, and multiplying this cost by the count for the event. Event costs are usually measured in terms of machine cycles, and so the cost of an event generally depends on the clock speed of the processor, which is also reported in the output. The maximum value in the table corresponds to the worst-case cost of a single occurrence of the event. Sometimes this can be a pessimistic estimate. For example, the maximum cost for graduated floating point instructions assumes that every floating point instruction is a double-precision reciprocal square root since it is the most costly R10000 floating point instruction. Because of the latency-hiding capabilities of the R10000, the minimum cost of virtually any event could be zero since most events can be overlapped with other operations. To avoid simply reporting minimum costs of zero, which would be of no practical use, the minimum time reported by **perfex -y** corresponds to the best-case cost of a single occurrence of the event. The best-case cost is obtained by running the maximum number of simultaneous occurrences of that event and averaging the cost. For example, two floating point instructions can complete per cycle, so the best case cost is 0.5 cycles per floating point instruction. The typical cost falls somewhere between minimum and maximum and is meant to correspond to the cost you see in average programs. **perfex -y** prints the event counts and associated cost estimates sorted from most costly to least costly. Although resembling a profiling output, this is not a true profile. The event costs reported are only estimates. Furthermore, since events do overlap with one another, the sum of the estimated times will usually exceed the program’s run time. This output should only be used to identify which events are responsible for significant portions of the program’s run time and to get a rough idea of what those costs might be. In the example above, the program spends a significant fraction of its time handling secondary cache and TLB misses. To make a significant improvement in the run time of this program, the tuning measures need to concentrate on reducing those cache misses. In addition to the event counts and cost estimates, **perfex -y** also reports a number of statistics derived from the typical costs. The meaning of many of the statistics is self-evident, for example, Graduated instructions/cycle. Below is a list of those statistics whose definitions require more explanation:

**Data mispredict/Data scache hits** The ratio of the counts for data misprediction from scache way prediction table and secondary data cache misses.

**Instruction mispredict/Instruction scache hits** The ratio of the counts for instruction misprediction from scache way prediction table and secondary instruction cache misses.

**L1 Cache Line Reuse** The number of times, on average, that a primary data cache line is used after it has been moved into the cache. It is calculated as graduated loads plus graduated stores minus primary data cache misses, divided by primary data cache misses.

**L2 Cache Line Reuse** The number of times, on average, that a secondary data cache line is used after it has been moved into the cache. It is calculated as primary data cache misses minus secondary data cache misses, divided by secondary data cache misses.

**L1 Data Cache Hit Rate** The fraction of data accesses that are satisfied from a cache line already resident in the primary data cache. It is calculated as  $1.0 - (\text{primary data cache misses} / \text{sum of graduated loads and graduated stores})$ .

**L2 Data Cache Hit Rate** The fraction of data accesses that are satisfied from a cache line already resident in the secondary data cache. It is calculated as  $1.0 - (\text{secondary data cache misses} / \text{primary data cache misses})$ .

**Time accessing memory/Total time** The sum of the typical costs of graduated loads, graduated stores, primary data cache misses, secondary data cache misses, and TLB misses, divided by the total program run time. The total program run time is calculated by multiplying cycles by the time per cycle (inverse of the processor's clock speed).

**L1–L2 bandwidth used (MB/s, average per process)** The amount of data moved between the primary and secondary data caches, divided by the total program run time. The amount of data moved is calculated as the sum of the number of primary data cache misses multiplied by the primary cache line size and the number of quadwords written back from primary data cache multiplied by the size of a quadword (16 bytes). For multiprocessor programs, the resulting figure is a per-process average since the counts measured by **perfex** are aggregates of the counts for all the threads. Multiply by the number of threads to get the total program bandwidth.

**Memory bandwidth used (MB/s, average per process)** The amount of data moved between the secondary data cache and main memory, divided by the total program run time. The amount of data moved is calculated as the sum of the number of secondary data cache misses multiplied by the secondary cache line size and the number of quadwords written back from secondary data cache multiplied by the size of a quadword (16 bytes). For multiprocessor programs, the resulting figure is a per-process average since the counts measured by **perfex** are aggregates of the counts for all the threads. Multiply by the number of threads to get the total program bandwidth.

**MFLOPS (MB/s, average per process)** The ratio of the graduated floating point instructions and the total program run time. Note that although a multiply-add carries out two floating point operations, it only counts as one instruction, so this statistic may underestimate the number of floating point operations per second. For multiprocessor programs, the resulting figure is a per-process average since the counts measured by **perfex** are aggregates of the counts for all the threads. Multiply by the number of threads to get the total program rate.

These statistics give you a quick way to identify performance problems in your program. For example, the cache hit-rate statistics tell you how cache friendly your

program is. Since a secondary cache miss is much more expensive than a cache hit, the L2 Data Cache Hit Rate needs to be close to 1.0 to indicate that the program is not paying a large penalty for the cache misses. Values of  $\pm 0.96$  and above indicate good cache performance. Note that, for the above example, the rate is 0.53, further confirmation of the cache problems in this program.

### IIIB.5 Using SpeedShop

The purpose of profiling is to find out exactly where a program is spending its time, that is, in precisely which procedures or lines of code. Then you can concentrate your efforts on the (usually small) areas of code where there is the most to be gained. Profiling using the SpeedShop package supports these methods:

- Sampling, in which the subject program is frequently interrupted; the program counter (PC) and stack are recorded on each interruption. The more frequently the PC is found in a particular procedure, the more execution time that procedure costs.
- SpeedShop can sample on a variety of time bases: the system timer or any of the R10000 performance counters. Ideal counting, in which a copy of the subject program binary is modified with trap instructions at the end of each basic block. During execution, the exact number of uses of each basic block is counted.
- Exception trace, not really a profiling method, records only floating point exceptions and their locations.

The SpeedShop package has three parts:

- `ssrun` performs experiments and collects data.
- `prof` processes data and prepares reports.
- The `ssapi` interface allows you to insert caliper points into a program to profile specific sections of code or phases of execution.

These programs are documented in the following reference pages: *speedshop* documents the types of experiments, as well as a number of environment variables you can set. *ssrun* documents specific options of `ssrun`. You need both **speedshop(1)** and **ssrun(1)** to run an experiment. *prof* documents the report types and the options you use to get them. *ssapi* documents the three library calls you can use.

### IIIB.6 PC sampling profiling

The accuracy of sampling depends on the time base that sets the sampling interval. In each case, the time base is the independent variable and the program state is the dependent variable. Select from the sampling methods summarized in Table III.2. Each time base finds the program PC more often in the code that consumes the most units of that time base:

ssrun Option	Time Base		Comments
-usertime	30 ms timer		Fairly coarse resolution; experiment runs quickly and output file is small; some bugs noted in speedshop(1).
-pcsamp[x] -fpcsamp[x]	10 ms timer	1 ms timer	Moderately coarse resolution; functions that cause cache misses or page faults are emphasized. Suffix x for 32-bit counts.
-gi_hwc -fgi_hwc	32771 insts	6553 insts	Fine-grain resolution based on graduated instructions. Emphasizes functions that burn a lot of instructions.
-cy_hwc -fcy_hwc	16411 clocks	3779 clocks	Fine-grain resolution based on elapsed cycles. Emphasizes functions with cache misses and mispredicted branches.
-ic_hwc -fic_hwc	2053 icache miss	419 icache miss	Emphasizes code that doesn't fit in L1 cache.
-isc_hwc -fisc_hwc	131 scache miss	29 scache miss	Emphasizes code that doesn't fit in L2 cache. Should be coarse-grained measure.
-dc_hwc -fdc_hwc	2053 dcache miss	419 dcache miss	Emphasizes code that causes L1 cache data misses.
-dsc_hwc -fdsc_hwc	131 scache miss	29 scache miss	Emphasizes code that causes L2 cache data misses.
-tlb_hwc -ftlb_hwc	257 TLB miss	53 TLB miss	Emphasizes code that causes page faults.
-gfp_hwc -fgfp_hwc	32771 fp insts	6553 fp insts	Emphasizes code that performs heavy FP calculation.
-prof_hwc	user-set		Hardware counter and overflow value from environment variables.

Table III.2: *Sampling methods.*

- The time bases that reflect actual elapsed time (-usertime, -pcsamp, -cy\_hwc) find the PC more often in the code where the program spends elapsed time. The time may be spent in that code because it is executed a lot, or it might be spent there because those instructions are processed slowly owing to cache misses, memory contention, or failed branch prediction. Use these to get an overview of the program and to find major trouble spots.
- The time bases that reflect instruction counts (-gi\_hwc, -gfp\_hwc) find the PC more often in the code that actually performs the most instructions. Use these to find the code that could benefit most from algorithmic changes.
- The time bases that reflect data access (-dc\_hwc, -sc\_hwc, -tlb\_hwc) find the PC more often in the code that has to wait for its data to be brought in from another memory level. Use these to find memory access problems.
- The time bases that reflect code access (-ic\_hwc, -isc\_hwc) find the PC more often in the code that has to be fetched from memory when it is called. Use these to pinpoint functions that could be reorganized for better locality, or to see when automatic inlining has gone too far.

It is easy to perform an experiment. Here is the application of an experiment to program adi2:

```
% srrun -fpcsamp adi2
```

```
Time:      7.619 seconds
Checksum:  5.6160428338E+06
```

The output file of samples is left in a file with the default name of ./command.experiment.pid. In this case the name was adi2.fpcsamp.4885. It is often more convenient, however, to dictate the name of the output file. You can do this by putting the desired filename and directory in environment variables. Using this csh script you can run an experiment, passing the output directory and filename on the command line, for example

```
% srruno -d /var/tmp -o adi2.cy -cy_hwc adi2
```

```
ssrun -cy\_hwc adi2 .....
Time:      9.644 seconds
Checksum:  5.6160428338E+06
..... srrun ends.
-rw-r--r--  1 guest  guest  18480 Dec 17 16:25 /var/tmp/adi2.cy
```

### IIIB.7 Using prof

Regardless of which time base you use for sampling, you display the result using **prof**. By default, **prof** displays a list of procedures ordered from the one with the most samples to the least:

```
% prof adi2.fpcsamp.4885
```

```

-----
Profile listing generated Sat Jan  4 10:28:11 1997
with:      prof adi2.fpcsamp.4885
-----

```

```

samples  time      CPU      FPU      Clock    N-cpu  S-interval  Countsize
 8574    8.6s R10000 R10010 196.0MHz  1      1.0ms      2 (bytes)

```

Each sample covers 4 bytes for every 1.0ms ( 0.01% of 8.5740s)

```

-----
-p[rocedures] using pc-sampling.
Sorted in descending order by the number of samples in each procedure.
Unexecuted procedures are excluded.
-----

```

samples	time(%)	cum time(%)	procedure (dso:file)
6688	6.7s( 78.0)	6.7s( 78.0)	zsweep (adi2:adi2.f)
671	0.67s( 7.8)	7.4s( 85.8)	xsweep (adi2:adi2.f)
662	0.66s( 7.7)	8s( 93.6)	ysweep (adi2:adi2.f)
208	0.21s( 2.4)	8.2s( 96.0)	fake_adi (adi2:adi2.f)
178	0.18s( 2.1)	8.4s( 98.1)	irand_ (/usr/lib32/libftn.so:../../libF77/rand_.c)
166	0.17s( 1.9)	8.6s(100.0)	rand_ (/usr/lib32/libftn.so:../../libF77/rand_.c)
1	0.001s( 0.0)	8.6s(100.0)	__oserror (/usr/lib32/libc.so.1:oserror.c)
8574	8.6s(100.0)	8.6s(100.0)	TOTAL

This profile indicates that you should focus on the routine zsweep, since it consumes almost 80% of the run time of this program. For finer detail, use the -heavy option. This supplements the list of procedures with a list of individual source line numbers, ordered by frequency:

```

-----
-h[eavy] using pc-sampling.
Sorted in descending order by the number of samples in each line.
Lines with no samples are excluded.
-----

```

samples	time(%)	cum time(%)	procedure (file:line)
3405	3.4s( 39.7)	3.4s( 39.7)	zsweep (adi2.f:122)
3226	3.2s( 37.6)	6.6s( 77.3)	zsweep (adi2.f:126)
425	0.42s( 5.0)	7.1s( 82.3)	xsweep (adi2.f:80)
387	0.39s( 4.5)	7.4s( 86.8)	ysweep (adi2.f:101)
273	0.27s( 3.2)	7.7s( 90.0)	ysweep (adi2.f:105)
246	0.25s( 2.9)	8s( 92.9)	xsweep (adi2.f:84)
167	0.17s( 1.9)	8.1s( 94.8)	irand_ (../../libF77/rand_.c:62)
163	0.16s( 1.9)	8.3s( 96.7)	fake_adi (adi2.f:18)
160	0.16s( 1.9)	8.5s( 98.6)	rand_ (../../libF77/rand_.c:69)

```

45 0.045s( 0.5) 8.5s( 99.1)      fake_adi (adi2.f:59)
32 0.032s( 0.4) 8.5s( 99.5)      zswEEP (adi2.f:113)
21 0.021s( 0.2) 8.5s( 99.7)      zswEEP (adi2.f:121)
11 0.011s( 0.1) 8.6s( 99.8)      irand_ (././libF77/rand_.c:63)
 6 0.006s( 0.1) 8.6s( 99.9)      rand_ (././libF77/rand_.c:67)
 4 0.004s( 0.0) 8.6s(100.0)     zswEEP (adi2.f:125)
 1 0.001s( 0.0) 8.6s(100.0)     yswEEP (adi2.f:104)
 1 0.001s( 0.0) 8.6s(100.0)     yswEEP (adi2.f:100)
 1 0.001s( 0.0) 8.6s(100.0)     __oserror (oserror.c:127)

8574 8.6s(100.0) 8.6s(100.0)      TOTAL

```

From this listing it is clear that lines 122 and 126 warrant further inspection. Even finer detail can be obtained with the `-source` option, which lists the source code and disassembled machine code, indicating sample hits on specific instructions.

### IIIB.8 Ideal time profiling

The other type of profiling is called ideal time, or basic block, profiling. Basic block is a compiler term for a section of code that has only one entrance and one exit. Any program can be decomposed into basic blocks. To obtain an ideal profile, `ssrun` copies the executable program and modifies the copy to contain code that records the entry to each basic block. Not only the executable itself but also all dynamic shared objects (DSOs; for more information, see *dso(5)*) that it links to are copied and instrumented. The instrumented executable and libraries are statically linked and run:

```
% ssrun -ideal adi2
```

```

Beginning libraries
  /usr/lib32/libssrt.so
  /usr/lib32/libss.so
  /usr/lib32/libfastm.so
  /usr/lib32/libftn.so
  /usr/lib32/libm.so
  /usr/lib32/libc.so.1
Ending libraries, beginning "adi2"
Time:      8.291 seconds
Checksum:  5.6160428338E+06

```

The number of times each basic block was encountered is recorded. The output data file is displayed using `prof`, just as for a sampled run. The report ranks source and library procedures from most to least used:

```
% prof adi2.ideal.4920
```

```

Prof run at: Sat Jan 4 10:34:06 1997
Command line: prof adi2.ideal.4920

285898739: Total number of cycles
1.45867s: Total execution time
285898739: Total number of instructions executed
1.000: Ratio of cycles / instruction
196: Clock rate in MHz
R10000: Target processor modeled

```

Procedures sorted in descending order of cycles executed.  
 Unexecuted procedures are not listed. Procedures  
 beginning with \*DF\* are dummy functions and represent  
 init, fini and stub sections.

```
-----
```

	cycles(%)	cum %	secs	instrns	calls	procedure (dso:file)
68026368	(23.79)	23.79	0.35	68026368	32768	xsweep(adi2.pixie:adi2.f)
68026368	(23.79)	47.59	0.35	68026368	32768	ysweep(adi2.pixie:adi2.f)
68026368	(23.79)	71.38	0.35	68026368	32768	zsweep(adi2.pixie:adi2.f)
35651584	(12.47)	83.85	0.18	35651584	2097152	rand(/libftn.so.pixn32:../libF77/rand_.c)
27262976	( 9.54)	93.39	0.14	27262976	2097152	irand(/libftn.so.pixn32:../libF77/rand_.c)
18874113	( 6.60)	99.99	0.10	18874113	1	fake_adi(adi2.pixie:adi2.f)
11508	( 0.00)	99.99	0.00	11508	5	memset(/libc.so.1.pixn32:/slayer_xlv0/ficussg-nov05/work/irix/lib/libc/libc_n32_M4/strings/bzero.s)
3101	( 0.00)	99.99	0.00	3101	55	__flsbuf(/libc.so.1.pixn32: flsbuf.c)
2446	( 0.00)	100.00	0.00	2446	42	x_putc(/libftn.so.pixn32:../libI77/wsfe.c)
1234	( 0.00)	100.00	0.00	1234	2	x_WEND(/libftn.so.pixn32:../libI77/wsfe.c)
1047	( 0.00)	100.00	0.00	1047	1	f_exit(/libftn.so.pixn32:../libI77/close.c)
1005	( 0.00)	100.00	0.00	1005	5	fflush(/libc.so.1.pixn32:flush.c)
639	( 0.00)	100.00	0.00	639	4	do_fio64_mp(/libftn.so.pixn32:../libI77/fmt.c)
566	( 0.00)	100.00	0.00	566	3	wrt_AP(/libftn.so.pixn32:../libI77/wrtfmt.c)
495	( 0.00)	100.00	0.00	495	6	map_luno(/libftn.so.pixn32:../libI77/util.c)
458	( 0.00)	100.00	0.00	458	14	op_gen(/libftn.so.pixn32:../libI77/fmt.c)
440	( 0.00)	100.00	0.00	440	9	gt_num(/libftn.so.pixn32:../libI77/fmt.c)
414	( 0.00)	100.00	0.00	414	1	getenv(/libc.so.1.pixn32:getenv.c)

The `-heavy` option adds a list of source lines, sorted by their consumption of ideal instruction cycles. An ideal profile shows exactly and repeatedly which statements are most executed and gives you an exact view of the algorithmic complexity of the program. An ideal profile does not necessarily reflect where a program spends its time since it cannot take cache and TLB misses into account. Consequently, the results of the ideal profile are startlingly different from that of the PC sampling profile. These ideal results indicate that `zsweep` should take exactly the same amount of run time as `ysweep` and `xsweep`. These differences can be used to infer where cache performance issues exist. On machines without the R10000's hardware profiling registers, such comparisons are the only profiling method available to locate cache problems.

### IIIB.9 Operation counts

Since ideal profiling counts the instructions executed by the program, it can provide all sorts of interesting information about the program. Already printed in the standard **prof** output are counts of how many times each subroutine is called. In addition, you may use the `-op` option to **prof** to get a listing detailing the counts of all instructions in the program. In particular, this will provide an exact count of the floating point operations executed:

```
% prof -op adi2.ideal.4920
```

```
Prof run at: Wed Jan 15 14:42:54 1997
Command line: prof -op adi2.ideal.4920
```

```
285898739: Total number of cycles
1.45867s: Total execution time
285898739: Total number of instructions executed
1.000: Ratio of cycles / instruction
196: Clock rate in MHz
R10000: Target processor modeled
```

```
-----  
pixstats summary  
-----
```

```
56590456: Floating point operations (38.796 Mflops @ 196 MHz)  
105500230: Integer operations (72.3265 M intops @ 196 MHz)
```

```
.  
. .  
. . .
```

Note that this is different from what you get using **perfex**. The R10000 counter #21 counts floating point instructions, not floating point operations. As a result, in a program that executes a lot of multiply-add instructions — each of which carries out two floating point operations — **perfex**'s MFLOPS statistic can be off by a factor of two. Since **prof -op** records all instructions executed, it counts each multiply-add instruction as two floating point operations, thus providing the correct tally. The MFLOPS figure it calculates, however, is based on the ideal time; to calculate floating point performance, divide the number of floating point operations counted by **prof -op** by wall clock time. Either method of profiling, PC sampling or ideal, can be applied to multiprocessor runs just as easily as it is applied to single-processor runs; each thread of an application maintains its own histogram, and the histograms may be printed individually or merged in any combination and printed as one profile.

### III.B.10 Gprof

One limitation of the **prof** output for either PC sampling or ideal time is that the information reported contains no information about the call hierarchy. That is, if the routine `zsweep` in the above example were called from two different locations in the program, you would not know how much time results from the call at each location; you would only know the total time spent in `zsweep`. If you knew that, say, the first location was responsible for the majority of the time, this could affect how you tune the program. For example, you might try inlining the call into the first location, but not bother with the second. Or, if you wanted to parallelize the program, knowing that the first location is where the majority of the time is spent, you might consider parallelizing the calls to `zsweep` there rather than trying to parallelize the `zsweep` routine itself. Speed-Shop provides two methods of obtaining hierarchical profiling information. The first method, which is called `gprof`, is used in conjunction with the ideal time profile. To obtain the `gprof` information for the above example, simply add the flag `-gprof` to the **prof** command:

```
% prof -gprof adi2.ideal.4920
```

```
Prof run at: Wed Jan 15 16:52:09 1997  
Command line: prof -gprof adi2.ideal.4920
```

```
285898739: Total number of cycles  
1.45867s: Total execution time  
285898739: Total number of instructions executed  
1.000: Ratio of cycles / instruction  
196: Clock rate in MHz  
R10000: Target processor modeled
```

-----  
 Procedures sorted in descending order of cycles executed.  
 Unexecuted procedures are not listed. Procedures  
 beginning with \*DF\* are dummy functions and represent  
 init, fini and stub sections.  
 -----

	cycles(%)	cum %	secs	instrns	calls	procedure (dso:file)
	68026368 (23.79)	23.79	0.35	68026368	32768	xsweep(adi2.pixie:adi2.f)
	68026368 (23.79)	47.59	0.35	68026368	32768	ysweep(adi2.pixie:adi2.f)
	68026368 (23.79)	71.38	0.35	68026368	32768	zsweep(adi2.pixie:adi2.f)
	35651584 (12.47)	83.85	0.18	35651584	2097152	rand_(./libftn.so.pixmap32:../libF77/rand_.c)
	27262976 (9.54)	93.39	0.14	27262976	2097152	irand_(./libftn.so.pixmap32:../libF77/rand_.c)

All times are in milliseconds.

NOTE: any functions which are not part of the call  
 graph are listed at the end of the gprof listing  
 -----

index	cycles(%)	self self(%)	Kids Kids(%)	called/total called-self called/total	parents name children
[1]	285895481 (100.00%)	57 ( 0.00%) 50 3 2	285895424 (100.00%) 285895369 0 0	0 1/1 1/1 1/1	__start [1] main [2] __istart [112] __readenv_sigfpe [113]
[2]	285895419 (100.00%)	50 50 ( 0.00%) 18874113 205	285895369 285895369 (100.00%) 267020445 606	1/1 1 1/1 5/5	__start [1] main [2] fake_adi [3] signal [44]
[3]	285894558 (100.00%)	18874113 18874113 ( 6.60%) 68026368 68026368 68026368 35651584 28 26 22 22 23 114	267020445 267020445 (93.40%) 0 0 0 27262976 13486 5368 2610 2610 2428 44	1/1 1 32768/32768 32768/32768 32768/32768 2097152/2097152 2/2 2/2 1/1 1/1 1/1 2/2	main [2] fake_adi [3] zsweep [4] ysweep [5] xsweep [6] rand_ [7] s_wsfe64 [9] e_wsfe [17] dc_ficxr4v [25] dc_ficxr8v [24] s_stop [28] dtime_ [68]
[4]	68026368 (23.79%)	68026368 (100.00%)	0 ( 0.00%)	32768/32768 32768	fake_adi [3] zsweep [4]
[5]	68026368 (23.79%)	68026368 (100.00%)	0 ( 0.00%)	32768/32768 32768	fake_adi [3] ysweep [5]
[6]	68026368 (23.79%)	68026368 (100.00%)	0 ( 0.00%)	32768/32768 32768	fake_adi [3] xsweep [6]

This produces the usual ideal time profiling output, but following that is the hierarchical information. There is a block of information for each subroutine in the program. A number, shown in brackets (e.g., [1]), is assigned to each routine so that the information pertaining to it can easily be located in the output. Let's look in detail at the block of information provided; we'll use fake\_adi [3] as an example. The line beginning with the number [3] shows, from left to right, the:

- Number of cycles consumed by this routine and the routines it calls (its descen-

dants)

- Number of cycles spent inside the routine, but not in any of its descendants
- Number of cycles spent in its descendants
- Total number of times the routine was called
- Name of the routine, fake\_adi [3]

Above this line are lines showing which routines fake\_adi [3] was called from. In this case, it is only called from one place, main [2], so there is just one line (in general, there would be one line for each routine which calls fake\_adi [3]). This line shows

- The proportion of the cycles spent inside fake\_adi [3] as a result of the call from main [2];
- The proportion of time spent in fake\_adi [3]'s descendants as a result of the call from main [2];
- How many calls there are to fake\_adi [3] in main [2] / the total number of calls to fake\_adi [3] from all places in the program

Since fake\_adi [3] is only called once, all the time in it and its descendants is the result of this one call. Below the line beginning with the number [3] are all the descendants of fake\_adi [3]. For each descendant you see:

- The proportion of the descendant's cycles spent inside it.
- The proportion of the descendant's cycles spent in its descendants (i.e., fake\_adi [3]'s grandchildren).
- How many calls to the descendant there are in fake\_adi [3] / the total number of calls to the descendant from all places in the program.

This block of information allows you to determine not just which subroutines but which paths in the program are responsible for the majority of time. The only limitation is that gprof reports ideal time, so cache misses are not represented.

### **III.B.11 Usertime profiling**

gprof only reports ideal time. To get hierarchical profiling information that accurately accounts for all the time in the program, the way PC sampling does, use usertime profiling. For this type of profiling, the program is sampled during the run. At each sample, the location of the program counter is noted and the entire call stack is traced to record which routines have been called to get to this point in the program. From this, a hierarchical profile is constructed. Since unwinding the call stack is an expensive operation, the sampling period for usertime profiling is relatively long: 30 milliseconds. usertime profiling is performed with the following command:

**% srun -usertime adi2**

The output is written to a file called `./adi2.usertime.pid`, where `pid` is the process ID for this run of `adi2`. The profile is displayed using **prof** just as for PC sampling and ideal time profiling:

**% prof adi2.usertime.19572**

The output is as follows:

```
-----
-----
Profile listing generated Wed Jan 15 16:57:10 1997
with:      prof adi2.usertime.19572
-----
-----

      Total Time (secs)      : 9.99
      Total Samples          : 333
      Stack backtrace failed: 0
      Sample interval (ms)   : 30
      CPU                    : R10000
      FPU                    : R10010
      Clock                  : 196.0MHz
      Number of CPUs         : 1

-----
-----

index  %Samples      self descendent  total      name
[1]    100.0%      0.00            9.99      333      __start
[2]    100.0%      0.00            9.99      333      main
[3]    100.0%      0.09            9.90      333      fake_adi
[4]     80.8%      8.07            0.00      269      zsweep
[5]     7.5%       0.75            0.00      25       xsweep
[6]     6.9%       0.69            0.00      23       ysweep
[7]     3.9%       0.12            0.27      13       rand_
[8]     2.7%       0.27            0.00      9        irand_
```

The information is less detailed than that provided by `gprof`, but when combined with `gprof`, you can get a complete hierarchical profile for all routines which have run long enough to be sampled.

## **IIIC Dynamic Acceleration Structures for Interactive Ray Tracing**

Acceleration structures used for ray tracing have been designed and optimized for efficient traversal of static scenes. As it becomes feasible to do interactive ray tracing of moving objects, new requirements are posed upon the acceleration structures. Dynamic environments require rapid updates to the acceleration structures. In this paper we propose spatial subdivisions which allow insertion and deletion of objects in constant time at an arbitrary position, allowing scenes to be interactively animated and modified.

### **IIIC.1 Introduction**

Recently, interactive ray tracing has become a reality [15, 18], allowing exploration of scenes rendered with higher quality shading than with traditional interactive rendering algorithms. A high frame-rate is obtained through parallelism, using a multiprocessor shared memory machine. This approach has advantages over hardware accelerated interactive systems in that a software-based ray tracer is more easily modified. One of the problems with interactive ray tracing is that previous implementations only dealt with static scenes or scenes with a small number of specially handled moving objects. The reason for this limitation is that the acceleration structures used to make ray tracing efficient rely on a significant amount of preprocessing to build. This effectively limits the usefulness of interactive ray tracing to applications which allow changes in camera position. The work presented in this paper is aimed at extending the functionality of interactive ray tracing to include applications where objects need to be animated or interactively manipulated.

When objects can freely move through the scene, either through user interaction, or due to system-determined motion, it becomes necessary to adapt the acceleration methods to cope with changing geometry. Current spatial subdivisions tend to be highly optimized for efficient traversal, but are difficult to update quickly for changing geometry. For static scenes this suffices, as the spatial subdivision is generally constructed during a pre-processing step. However, in animated scenes pre-processed spatial subdivisions may have to be recalculated for each change of the moving objects. One approach to circumvent this issue is to use 4D radiance interpolants to speed-up ray traversal [2]. However, within this method the frame update rates depend on the type of scene edits performed as well as the extent of camera movement. We will therefore focus on adapting current spatial subdivision techniques to avoid these problems.

To animate objects while using a spatial subdivision, insertion and deletion costs are not negligible, as these operations may have to be performed many times during rendering. In this paper, spatial subdivisions are proposed which allow efficient ray traversal as well as rapid insertion and deletion for scenes where the extent of the scene grows over time.

The following section presents a brief overview of current spatial subdivision techniques (Section IIIC.2), followed by an explanation of our (hierarchical) grid modifications (Sections IIIC.3 and IIIC.4). A performance evaluation is given in Section IIIC.5, while conclusions are drawn in the final section.

## III.C.2 Acceleration Structures for Ray Tracing

There has been a great deal of work done on acceleration structures for ray tracing [12]. However, little work has focused on ray tracing moving objects. Glassner presented an approach for building acceleration structures for animation [11]. However, this approach does not work for environments without *a priori* knowledge of the animation path for each object. In a survey of acceleration techniques, Gaede and Günther provide an overview of many spatial subdivisions, along with the requirements for various applications [9]. The most important requirements for ray tracing are fast ray traversal and adaptation to unevenly distributed data. Currently popular spatial subdivisions can be broadly categorized into bounding volume hierarchies and voxel based structures.

Bounding volume hierarchies create a tree, with each object stored in a single node. In theory, the tree structure allows  $O(\log n)$  insertion and deletion, which may be fast enough. However, to make the traversal efficient, the tree is augmented with extra data, and occasionally flattened into an array representation [25], which enables fast traversal but insertion or deletion incur a non-trivial cost. Another problem is that as objects are inserted and deleted, the tree structure could become arbitrarily inefficient unless some sort of rebalancing step is performed as well.

Voxel based structures are either grids [1, 8] or can be hierarchical in nature, such as bintrees and octrees [10, 26]. The cost of building a spatial subdivision tends to be  $O(n)$  in the number of objects. This is true for both grids and octrees. In addition, the cost of inserting a single object may depend on its relative size. A large object generally intersects many voxels, and therefore incurs a higher insertion cost than smaller objects. This can be alleviated through the use of modified hierarchical grids, as explained in Section III.C.4. The larger problem with spatial subdivision approaches is that the grid structure is built within volume bounds that are fixed before construction. Although insertion and deletion may be relatively fast for most objects, if an object is moved outside the extent of the spatial subdivision, current structures would require a complete rebuild. This problem is addressed in the next section.

## III.C.3 Grids

Grid spatial subdivisions for static scenes, without any modifications, are already useful for animated scenes, as traversal costs are low and insertion and deletion of objects is reasonably straightforward. Insertion is usually accomplished by mapping the axis-aligned bounding box of an object to the voxels of the grid. The object is inserted into all voxels that overlap with this bounding box. Deletion can be achieved in a similar way.

However, when an object moves outside the extent of the spatial subdivision, the acceleration structure would normally have to be rebuilt. As this is too expensive to perform repeatedly, we propose to logically replicate the grid over space. If an object exceeds the bounds of the grid, the object wraps around before re-insertion. Ray traversal then also wraps around the grid when a boundary is reached. In order to provide a stopping criterion for ray traversal, a logical bounding box is maintained which contains all objects, including the ones that have crossed the original perimeter. As this scheme does not require grid re-computation whenever an object moves far away, the

cost of maintaining the spatial subdivision will be substantially lower. On the other hand, because rays now may have to wrap around, more voxels may have to be traversed per ray, which will slightly increase ray traversal time.

During a pre-processing step, the grid is built as usual. We will call the bounding box of the entire scene at start-up the 'physical bounding box'. If during the animation an object moves outside the physical bounding box, either because it is placed by the user in a new location, or its programmed path takes it outside, the logical bounding box is extended to enclose all objects. Initially, the logical bounding box is equal to the physical bounding box. Insertion of an object which lies outside the physical bounding box is accomplished by wrapping the object around within the physical grid, as depicted in Figure III.17 (left).

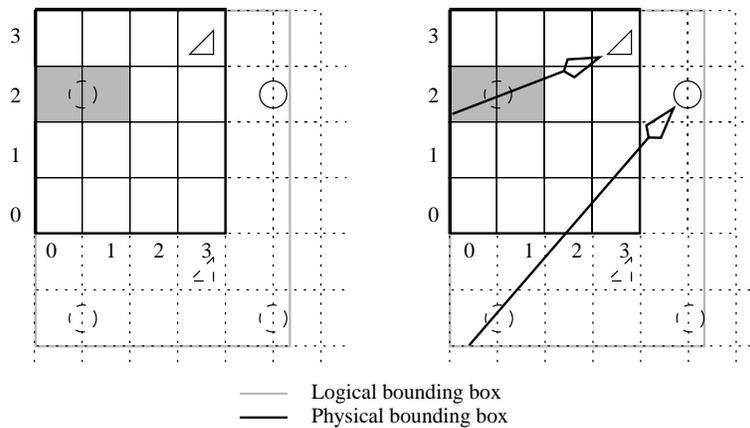


Figure III.17: Grid insertion (left). The sphere has moved outside the physical grid, now overlapping with voxels (4, 2) and (5, 2). Therefore, the object is inserted at the location of the shaded voxels. The logical bounding box is extended to include the newly moved object. Right: ray traversal through extended grid. The solid lines are the actual objects whereas the dashed lines indicate voxels which contain objects whose actual extents are not contained in that voxel.

As the logical bounding box may be larger than the physical bounding box, ray traversal now starts at the extended bounding box and ends if an intersection is found or if the ray leaves the logical bounding box. In the example in Figure III.17 (right), the ray pointing to the sphere starts within a logical voxel, voxel (-2, 0), which is mapped to physical voxel (0, 2). The logical coordinates of the sphere are checked and found to be outside of the currently traversed voxel and thus no intersection test is necessary. The ray then progresses to physical voxel (1, 2). For the same reason, no intersection with the sphere is computed again. Traversal then continues until the sphere is intersected in logical voxel (4, 2), which maps to physical voxel (0, 2).

Objects that are outside the physical grid are tagged, so that in the above example, when the ray aimed at the triangle enters voxels (0, 2) and (1, 2), the sphere does not have to be intersected. Similarly, when the ray is outside the physical grid, objects that are within the physical grid need not be intersected. As most objects will initially

lie within the physical bounds, and only a few objects typically move away from their original positions, this scheme speeds up traversal considerably for parts of the ray that are outside the physical bounding box.

When the logical bounding box becomes much larger than the physical bounding box, there is a tradeoff between traversal speed (which deteriorates for large logical bounding boxes) and the cost of rebuilding the grid. In our implementation, the grid is rebuilt when the length of the diagonals of the physical and logical bounding boxes differ by a factor of two.

Hence, there is a hierarchy of operations that can be performed on grids. For small to moderate expansions of the scene, wrapping both rays and objects is relatively quick without incurring too high a traversal cost. For larger expansions, rebuilding the grid will become a more viable option.

This grid implementation shares the advantages of simplicity and cheap traversal with commonly used grid implementations. However, it adds the possibility of increasing the size of the scene without having to completely rebuild the grid every time there is a small change in scene extent. The cost of deleting and inserting a single object is not constant and depends largely on the size of the object relative to the size of the scene. This issue is addressed in the following section.

#### **III.C.4 Hierarchical grids**

As was noted in the previous section, the size of an object relative to each voxel in a grid influences how many voxels will contain that object. This in turn negatively affects insertion and deletion times. Hence, it would make sense to find a spatial subdivision whereby the voxels can have different sizes. If this is accomplished, then insertion and deletion of objects can be made independent of their sizes and can therefore be executed in constant time. Such spatial subdivisions are not new and are known as hierarchical spatial subdivisions. Octrees, bintrees and hierarchical grids are all examples of hierarchical spatial subdivisions. However, normally such spatial subdivisions store all their objects in leaf nodes and would therefore still incur non-constant insertion and deletion costs. We extend the use of hierarchical grids in such a way that objects can also reside in intermediary nodes or even in the root node for objects that are nearly as big as the entire scene.

Because such a structure should also be able to deal with expanding scenes, our efforts were directed towards constructing a hierarchy of grids (similar to Sung [27]), thereby extending the functionality of the grid structure presented in the previous section. Effectively, the proposed method constitutes a balanced octree.

Object insertion now proceeds similarly to grid insertion, except that the grid level needs to be determined before insertion. This is accomplished by comparing the size of the object in relation to the size of the scene. A simple heuristic is to determine the grid level from the diagonals of the two bounding boxes. Specifically, the length of the grid's diagonal is divided by the length of the object's diagonal, the result determining the grid level. Insertion and deletion progresses as explained in the previous section.

The gain of constant time insertion is offset by a slightly more complicated traversal algorithm. Hierarchical grid traversal is effectively the same as grid traversal with the following modifications. Traversal always starts at a leaf node which may first

be mapped to a physical leaf node as described in the previous section. The ray is intersected with this voxel and all its parents until the root node is reached. This is necessary because objects at all levels in the hierarchy may occupy the same space as the currently traversed leaf node. If an intersection is found within the space of the leaf node, then traversal is finished. If not, the next leaf node is selected and the process is repeated.

This traversal scheme is wasteful because the same parent nodes may be repeatedly traversed for the same ray. To combat this problem, note that common ancestors of the current leaf node and the previously intersected leaf node, need not be traversed again. If the ray direction is positive, the current voxel's number can be used to derive the number of levels to go up in the tree to find the common ancestor between the current and the previously visited voxel. For negative ray directions, the previously visited voxel's number is used instead. Finding the common ancestor is achieved using simple bit manipulation, as detailed in Figure III.18.

```

bitmask = (raydir_x > 0) ? x : x + 1
forall levels in hierarchical grid
{
    cell = hgrid[level][x>>level][y>>level][z>>level]
    forall objects in cell
        intersect(ray, object)
    if (bitmask & 1)
        return
    bitmask >>= 1
}

```

Figure III.18: Hierarchical grid traversal algorithm in C-like pseudo-code. The bitmask is set assuming that the last step was along the x-axis.

As the highest levels of the grid may not contain any objects, ascending all the way to the highest level in the grid is not always necessary. Ascending the tree for a particular leaf node can stop when the largest voxel containing objects is visited.

This hierarchical grid structure has the following features. The traversal is only marginally more complex than standard grid traversal. In addition, wrapping of objects in the face of expanding scenes is still possible. If all objects are the same size, this algorithm effectively defaults to grid traversal. Insertion and deletion can be achieved in constant time, as the number of voxels that each object overlaps is roughly constant<sup>6</sup>.

### III.C.5 Evaluation

The grid and hierarchical grid spatial subdivisions were implemented using an interactive ray tracer [18], which runs on an SGI Origin 2000 with 32 processors. For evaluation purposes, two test scenes were used. In each scene, a number of objects were animated using pre-programmed motion paths. The scenes as they are at start-up are depicted in Figure III.21 (top). An example frame taken during the animation is given for each scene in Figure III.21 (bottom). All images were rendered on 30 processors at a resolution of  $512^2$  pixels.

<sup>6</sup>Note that this also obviates the need for mailbox systems to avoid redundant intersection tests.

To assess basic traversal speed, the new grid and hierarchical grid implementations are compared with a bounding volume hierarchy. We also compared our algorithms with a grid traversal algorithm which does not allow interactive updates. Its internal data structure consists of a single array of object pointers, which improves cache efficiency on the Origin 2000.

From here on we will refer to the new grid implementation as ‘interactive grid’ to distinguish between the two grid traversal algorithms. As all these spatial subdivision methods have a user defined parameter to set the resolution (voxels along one axis and maximum number of grid levels, respectively), various settings are evaluated. The overall performance is given in Figure III.19 and is measured in frames per second.

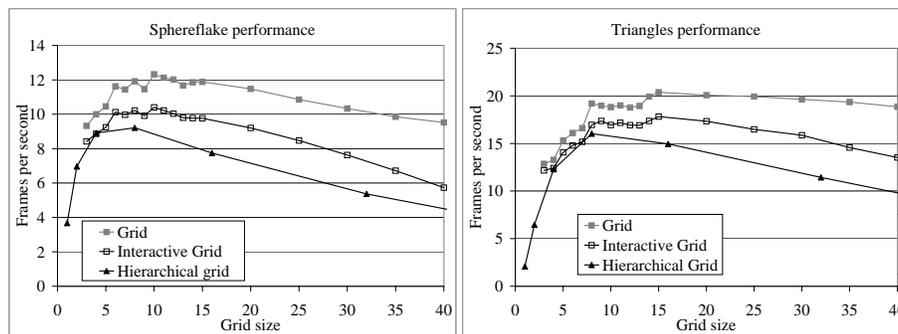


Figure III.19: Performance (in frames per second) for the grid, the interactive grid and the hierarchical grid for two static scenes. The bounding volume hierarchy achieves a frame rate of 8.5 fps for the static sphereflake model and 16.4 fps for the static triangles model.

The extra flexibility gained by both the interactive grid and hierarchical grid implementations results in a somewhat slower frame rate. This is according to expectation, as the traversal algorithm is a little more complex and the Origin’s cache structure cannot be exploited as well with either of the new grid structures. The graphs in Figure III.19 show that with respect to the grid implementation the efficiency reduction is between 12% and 16% for the interactive grid and 21% and 25% for the hierarchical grid. These performance losses are deemed acceptable since they result in far better overall execution than dynamically reconstructing the original grid. For the sphereflake, all implementations are faster, for a range of grid sizes, than a bounding volume hierarchy, which runs at 8.5 fps. For the triangles scene, the hierarchical grid performs at 16.0 fps similarly to the bounding volume (16.4 fps), while grid and interactive grid are faster.

The non-zero cost of updating the scene effectively limits the number of objects that can be animated within the time-span of a single frame. However, for both scenes, this limit was not reached. In the case where the frame rate was highest for the triangles scene, updating all 200 triangles took less than 1/680th of a frame for the hierarchical grid and 1/323th of a frame for the interactive grid. The sphereflake scene costs even less to update, as fewer objects are animated. For each of these tests, the hierarchical

grid is more efficiently updated than the interactive grid, which confirms its usefulness.

The size difference between different objects should cause the update efficiency to be variable for the interactive grid, while remaining relatively constant for the hierarchical grid. In order to demonstrate this effect, both the ground plane and one of the triangles in the triangle scene was interactively repositioned during rendering. The update rates for different size parameters for both the interactive grid and the hierarchical grid, are presented in Figure III.20 (left). As expected, the performance of the hierarchical grid is relatively constant, although the size difference between ground plane and triangle is considerable. The interactive grid does not cope with large objects very well if these objects overlap with many voxels. Dependent on the number of voxels in the grid, there is one to two orders of magnitude difference between inserting a large and a small object. For larger grid sizes, the update time for the ground plane is roughly half a frame. This leads to visible artifacts when using an interactive grid, as during the update the processors that are rendering the next frame temporarily cannot intersect this object (it is simply taken out of the spatial subdivision). In practice, the hierarchical grid implementation does not show this disadvantage.

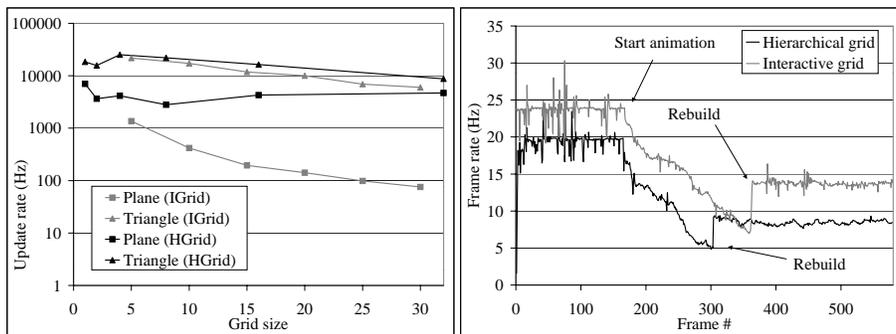


Figure III.20: Left: Update rate as function of (hierarchical) grid size. The plane is the ground plane in the triangles scene and the triangle is one of the triangles in the same scene. Right: Frame rate as function of time for the expanding triangle scene.

The time to rebuild a spatial subdivision from scratch is expected to be considerably higher than the cost of re-inserting a small number of objects. For the triangles scene, where 200 out of 201 objects were animated, the update rate was still a factor of two faster than the cost of completely rebuilding the spatial subdivision. This was true for both the interactive grid and the hierarchical grid. A factor of two was also found for the animation of 81 spheres in the sphereflake scene. When animating only 9 objects in this scene, the difference was a factor of 10 in favor of updating. We believe that the performance difference between rebuilding the acceleration structure and updating all objects is largely due to the cost of memory allocation, which occurs when rebuilding.

In addition to experiments involving grids and hierarchical grids with a branching factor of two, tests were performed using a hierarchical grid with a higher branching factor. Instead of subdividing a voxel into eight children, here nodes are split into 64 children (4 along each axis). The observed frame rates are very similar to the hierar-

chical grid. The object update rates were slightly better for the sphereflake and triangle scenes, because the size differences between the objects matches this acceleration structure better than both the interactive grid and the hierarchical grid.

In the case of expanding scenes, the logical bounding box will become larger than the physical bounding box. The number of voxels that are traversed per ray will therefore on average increase. This is the case in the triangles scene<sup>7</sup>. The variation over time of the frame rate is given in Figure III.20 (right). In this example, the objects are first stationary. At some point the animation starts and the frame rate drops because the scene immediately starts expanding. At some point the expansion is such that a rebuild is warranted. The re-computed spatial subdivision now has a logical bounding box which is identical to the (new) physical bounding box and therefore the number of traversed voxels is reduced when compared with the situation just before the rebuild. The total frame rate does not reach the frame rate at the start of the computation, because the objects are more spread out over space, resulting in larger voxels and more intersection tests which do not yield an intersection point.

Finally, Figure III.22 shows that interactively updating scenes using drag and drop interaction is feasible.

## III.6 Conclusions

When objects are interactively manipulated and animated within a ray tracing application, much of the work that is traditionally performed during a pre-processing step becomes a limiting factor. Especially spatial subdivisions which are normally built once before the computation starts, do not exhibit the flexibility that is required for animation. The insertion and deletion costs can be both unpredictable and variable. We have argued that for a small cost in traversal performance flexibility can be obtained and insertion and deletion of objects can be performed in constant time.

By logically extending the (hierarchical) grids into space, these spatial subdivisions deal with expanding scenes rather naturally. For modest expansions, this does not significantly alter the frame rate. When the scenes expand a great deal, rebuilding the entire spatial subdivision may become necessary. For large scenes this may involve a temporary drop in frame rate. For applications where this is unacceptable, it would be advisable to perform the rebuilding within a separate thread (rather than the display thread) and use double buffering to minimize the impact on the rendering threads.

## Acknowledgements

Thanks to Pete Shirley and Steve Parker for their help and comments and to the anonymous reviewers for their helpful comments. This work was supported by NSF grants CISE-CCR 97-20192, NSF-9977218 and NSF-9978099 and by the DOE Advanced Visualization Technology Center.

---

<sup>7</sup>For this experiment, the ground plane of the triangles scene was reduced in size, allowing the rebuild to occur after a smaller number of frames.

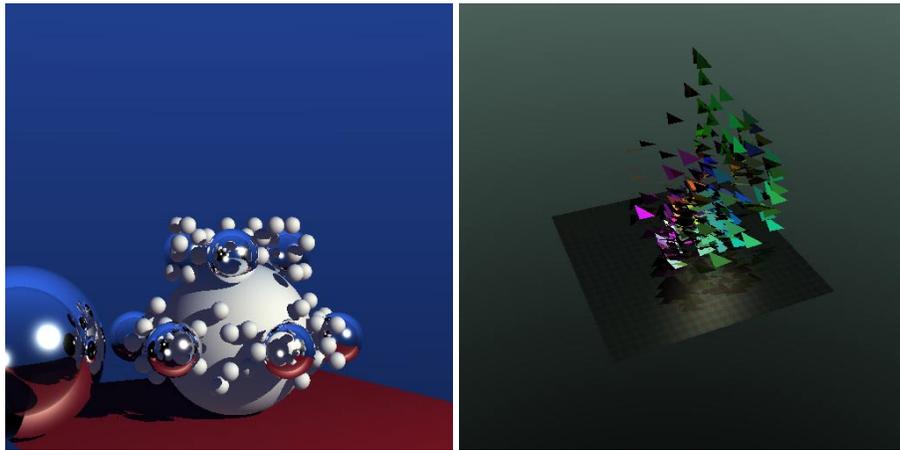
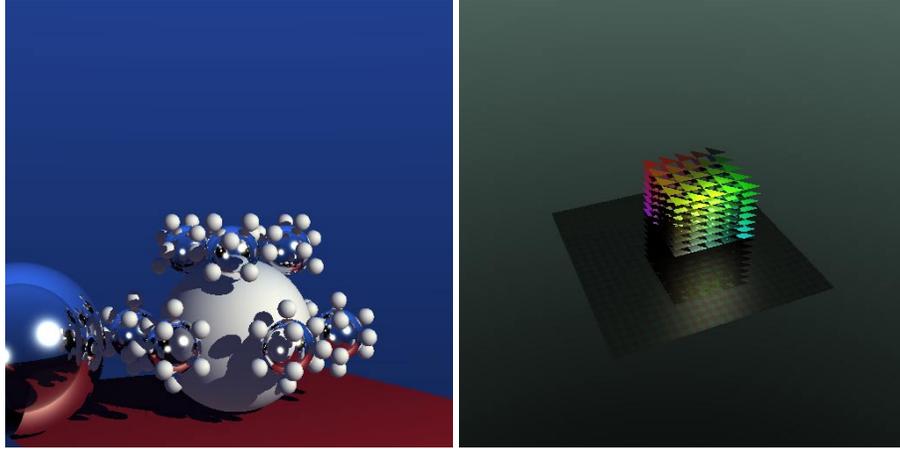


Figure III.21: Test scenes before any objects moved (top) and during animation (bottom).

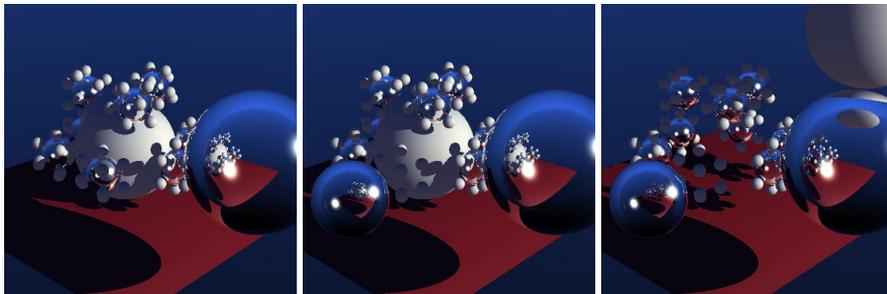


Figure III.22: Frames created during interactive manipulation.

# Bibliography

- [1] J. AMANATIDES AND A. WOO, *A fast voxel traversal algorithm for ray tracing*, in Eurographics '87, Elsevier Science Publishers, Amsterdam, North-Holland, Aug. 1987, pp. 3–10.
- [2] K. BALA, J. DORSEY, AND S. TELLER, *Interactive ray traced scene editing using ray segment tree*, in Rendering Techniques '99, D. Lischinski and G. W. Larson, eds., Eurographics, Springer-Verlag Wien New York, 1999, pp. 31–44.
- [3] ———, *Radiance interpolants for accelerated bounded-error ray tracing*, ACM Transactions on Graphics, (1999).
- [4] G. BISHOP, H. FUCHS, L. MCMILLAN, AND E. J. SCHER ZAGIER, *Frameless rendering: Double buffering considered harmful*, in Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994), A. Glassner, ed., Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH, ACM Press, July 1994, pp. 175–176. ISBN 0-89791-667-0.
- [5] J. G. CLEARY AND G. WYVILL, *Analysis of an algorithm for fast ray tracing using uniform space subdivision*, The Visual Computer, (1988), pp. 65–83.
- [6] D. CORTESI AND J. FIER, *Origin2000 (tm) and onyx2 (tm) performance tuning and optimization guide*, tech. rep., Silicon Graphics, 1998. Available from <http://techpubs.sgi.com:80/>.
- [7] R. A. CROSS, *Interactive realism for visualization using ray tracing*, in Proceedings Visualization '95, 1995, pp. 19–25.
- [8] A. FUJIMOTO, T. TANAKA, AND K. IWATA, *ARTS: Accelerated ray tracing system*, IEEE Computer Graphics and Applications, 6 (1986), pp. 16–26.
- [9] V. GAEDE AND O. GÜNTHER, *Multidimensional access methods*, ACM Computing Surveys, 30 (1998), pp. 170–231.
- [10] A. S. GLASSNER, *Space subdivision for fast ray tracing*, IEEE Computer Graphics and Applications, 4 (1984), pp. 15–22.
- [11] A. S. GLASSNER, *Spacetime ray tracing for animation*, IEEE Computer Graphics and Applications, 8 (1988), pp. 60–70.

- [12] A. S. GLASSNER, ed., *An Introduction to Ray Tracing*, Academic Press, San Diego, 1989.
- [13] G. W. LARSON AND M. SIMMONS, *The holodeck ray cache: An interactive rendering system for global illumination in non-diffuse environments*, ACM Transactions on Graphics, 18 (October 1999), pp. 361–368.
- [14] J. D. MACDONALD AND K. S. BOOTH, *Heuristics for ray tracing using space subdivision*, The Visual Computer, (1990), pp. 153–166.
- [15] M. J. MUUSS, *towards real-time ray-tracing of combinatorial solid geometric models*, in Proceedings of BRL-CAD Symposium, June 1995.
- [16] K. NAKAMARU AND Y. OHNO, *Breadth-first ray tracing utilizing uniform spatial subdivision*, IEEE Transactions on Visualization and Computer Graphics, 3 (1997), pp. 316–328.
- [17] K. NEMOTO AND T. OMACHI, *An adaptive subdivision by sliding boundary surfaces for fast ray tracing*, in Proceedings of Graphics Interface '86, M. Green, ed., May 1986, pp. 43–48.
- [18] S. PARKER, W. MARTIN, P.-P. SLOAN, P. SHIRLEY, B. SMITS, AND C. HANSEN, *Interactive ray tracing*, in Symposium on Interactive 3D Computer Graphics, April 1999.
- [19] S. PARKER, M. PARKER, Y. LIVNAT, P.-P. SLOAN, C. HANSEN, AND P. SHIRLEY, *Interactive ray tracing for volume visualization*, in IEEE Transactions on Visualization and Computer Graphics, July-September 1999.
- [20] E. REINHARD, A. CHALMERS, AND F. W. JANSEN, *Sampling diffuse inter-reflection within a hybrid scheduling ray tracer*, Journal of Parallel and Distributed Computing Practices, 3 (2001), pp. 11–19.
- [21] E. REINHARD, P. SHIRLEY, AND C. HANSEN, *Parallel point reprojection*. Submitted to IEEE 2001 Symposium on Parallel and Large-Data Visualization and Graphics.
- [22] E. REINHARD, B. SMITS, AND C. HANSEN, *Dynamic acceleration structures for interactive ray tracing*, in Proceedings of the 11th Eurographics Workshop on Rendering, Brno, Czech Republic, June 2000, pp. 299–306.
- [23] M. SIMMONS AND C. SÉQUIN, *Tapestry: A dynamic mesh-based display representation for interactive rendering*, in Proceedings of the 11th Eurographics Workshop on Rendering, Brno, Czech Republic, June 2000, pp. 329–340.
- [24] P. SLUSALLEK, P. HANRAHAN, S. PARKER, H. PFISTER, T. PURCELL, AND E. REINHARD, *Interactive ray tracing*, August 2001. Siggraph course #13.
- [25] B. SMITS, *Efficiency issues for ray tracing*, Journal of Graphics Tools, 3 (1998), pp. 1–14.

- [26] J. SPACKMAN AND P. WILLIS, *The SMART navigation of a ray through an octree*, Computers and Graphics, 15 (1991), pp. 185–194.
- [27] K. SUNG, *A DDA octree traversal algorithm for ray tracing*, in Eurographics '91, W. Purgathofer, ed., North-Holland, sept 1991, pp. 73–85. European Computer Graphics Conference and Exhibition; held in Vienna, Austria; 2-6 September 1991.
- [28] B. WALTER, G. DRETTAKIS, AND S. PARKER, *Interactive rendering using the render cache*, in Rendering Techniques '99, D. Lischinski and G. W. Larson, eds., Eurographics, Springer-Verlag Wien New York, 1999, pp. 19–30.
- [29] K.-Y. WHANG, J.-W. SONG, J.-W. CHANG, J.-Y. KIM, W.-S. CHO, C.-M. PARK, AND I.-Y. SONG, *Octree-r: An adaptive octree for efficient ray tracing*, IEEE Transactions on Visualization and Computer Graphics, 1 (1995), pp. 343–349.
- [30] E. S. ZAGIER, *Defining and refining frameless rendering*, Tech. Rep. TR97-008, UNC-CS, July 1997.

**Section IV**  
**The “Kilauea” Massively Parallel Ray  
Tracer**

**Toshi Kato**

# Parallel Rendering and the Quest for Realism: The “Kilauea” Massively Parallel Ray Tracer

Toshiaki Kato  
toshi@squareusa.com  
Square USA Honolulu Studio

Hitoshi Nishimura\*, Tadashi Endo†, Tamotsu Maruyama‡  
Jun Saito§, Per H. Christensen¶

<http://www.squareusa.com/kilauea/>

## Abstract

By taking advantage of the increased performance and decreased cost of personal computers today, and by making use of the massive computational resources gained from parallelizing such machines, we can create images of complexity and quality that has never before been possible. However, overcoming inherent difficulties of parallel processing requires various techniques and efforts, especially when developing a parallel renderer. In this article we describe our rendering research and development project, and explain the different approaches we attempted in order to solve these problems. This information should be useful when constructing a similar system in a similar environment. We hope this article will serve as an aid to correctly understand the possibilities of parallel processing on PC clusters, and make parallel processing a more realistic and feasible option for attacking new challenges.

## 1 What Is the Kilauea Project?

The Kilauea project, currently taking place within the R&D division of Square USA’s Honolulu Studio, is a project intended to develop a renderer that is completely different from the currently available renderers focusing on local illumination. The renderer is called Kilauea, named after an active volcano on the island of Hawaii. Kilauea is clearly distinguished from other renderers mainly due to two of its ultimate goals:

### 1. Global illumination

Demand for global illumination is increasing day by day. However, the amount of computation required for global illumination is much greater than that of a local illumination renderer.

---

\*Kilauea Developer. [hito@squareusa.com](mailto:hito@squareusa.com)

†Kilauea Developer. [tendo@squareusa.com](mailto:tendo@squareusa.com)

‡Kilauea Developer. [tamotu@squareusa.com](mailto:tamotu@squareusa.com)

§Kilauea Developer. [jun@squareusa.com](mailto:jun@squareusa.com)

¶Former Kilauea Developer. [per@pixar.com](mailto:per@pixar.com)

## 2. Support for extremely complex scenes

The scenes being rendered are also becoming more complex, and rendering performance naturally decreases as scene complexity increases.

In addition to increasing the required computation power, complex scenes also severely impact the memory requirements.

These two goals are fundamentally the ultimate goals of any photorealistic renderer, and achieving them clearly requires many innovative ideas different from the conventional methodology. With these goals in mind, the Kilauea project was started in order to create a test bed designed for conducting many technical experiments. Currently the development team consists of five members. Code is entirely in C++ and the total source code size is over seven hundred thousand lines, including test programs for all components.

## 2 Basic Idea

This section explains the basic approach of the Kilauea project.

The global illumination solution of Kilauea is based on the photon map method (including final gathering), instead of other approaches such as radiosity. As a result, ray tracing was absolutely necessary as the underlying algorithm in the rendering engine. These decisions are based on the fact that the ability to render effects such as motion blur and global illumination from both specular and diffuse components was mandatory. The photon map method allows a straightforward implementation of these requirements. As for the computational power, the Kilauea project's basic premise is to use parallel processing as a solution.

Parallel processing is also the Kilauea project's solution to the problem with the large memory space required for rendering complex scenes. In other words, if a single machine cannot store the entire scene data, it is distributed among two machines. If two is not enough, then three. In this manner, the entire scene is contained in memory using as many machines as necessary.

The Kilauea project took this approach because storing the entire scene data in memory is much more convenient from a ray tracing standpoint. The ray tracing optimization in terms of memory usage is implemented based on the assumption that the scene is distributed among multiple machines. Being able to render the same scene with less memory is important, since this means that more complex scenes can be rendered with the same amount of memory. Furthermore, the ability to distribute the scene across multiple machines is more important in eliminating the constraint of maximum memory size of one machine. Because Kilauea can distribute the scene data across multiple machines, there is no software upper limit to the supported scene size, despite using a simple ray tracing algorithm. The upper limit is determined by the sum of the memories of the multiple machines that Kilauea has access to. It provides an environment where users can just throw in more machines as they become necessary.

The initial goals of this project implied that we had to seriously devote ourselves to parallel processing.

We ended up using two different levels of parallel processing. One is a shared memory type of parallel processing within a multi-CPU environment, and another is message passing based parallel processing over a cluster of PCs. The current implementation of Kilauea is able to perform almost all of its computation in parallel.

### 3 System Design

This section describes an overview of Kilauea from the system design standpoint.

#### 3.1 Hardware environment

Kilauea is implemented with two levels of parallel processing. The first is parallel processing within a multi-CPU machine, and the other is parallel processing across multiple machines. These mechanisms are designed with the following hardware environment in mind:

1. multi-CPU machine (Linux)
2. PC cluster
3. 100Base-T connection (via a switching hub)

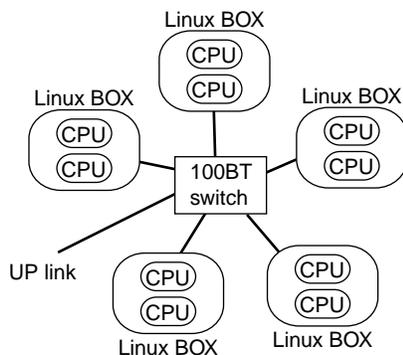


Figure 1: Kilauea hardware environment

PCs have become extremely inexpensive recently. A 1 GHz dual processor machine with 1 gigabyte of RAM can be bought for around \$1500 (as of Jan. 2001) if building from parts. Prices are dropping every month, so the same \$1500 will buy a much more powerful machine today. The machines run Linux, and individual machines are connected with 100Base-T connection via a switching hub (figure 1).

Figures 2 and 3 are the photographs of the render farm that the Kilauea team uses for experimentation. We are using machines ranging from Pentium III 450MHz single CPU with 512 megabyte of RAM, to a Pentium III 1 GHz dual CPU with 1 gigabyte of RAM in our experiments. In total, we have 37 machines with 60 CPUs available for development.

#### 3.2 Pthreads

Building a dual-CPU Linux box is very inexpensive these days. To take advantage of this, Kilauea is designed for parallel processing on a multi-CPU machine. The Pthread library is used in the actual implementation. Parallel processing within a single machine is a shared memory implementation.

#### 3.3 Message passing

Pthreads allows parallel processing implementation within a single machine, but some form of data transmission will be necessary for parallel processing between individual machines. We have considered several methodologies, with the emphasis on adopting the most widely accepted standard. We considered PVM and others at one time, but since the general trend seemed to be



Figure 2: Kilauea render farm A



Figure 3: Kilauea render farm B

converging towards MPI, we adopted MPI<sup>1</sup> without hesitation and started developing with it. We actually ended up doing a variety of experimentation on the MPI implementation itself. This will be described in detail in section 6.2.

### 3.4 Front-end process

Kilauea is a rendering engine. It has no GUI of any sort. It also has no features for animation or model data creation. Normally Maya<sup>2</sup> serves as a front-end to Kilauea. Animation, modeling, and other scene data are created/edited within Maya. This information is then sent to Kilauea for rendering. Kilauea is a collection of processes running on multiple machines, but to a Maya user it appears to be a single external rendering process (figure 4).

A file format called ShotData is used for sending data from Maya to Kilauea. This is the equivalent of rib files for RenderMan. All the scene data is converted to this ShotData format before being sent to Kilauea. The important features of the ShotData format are:

1. multiple frame storage
2. incremental data storage
3. binary format

Details on the ShotData format are discussed in section 4.

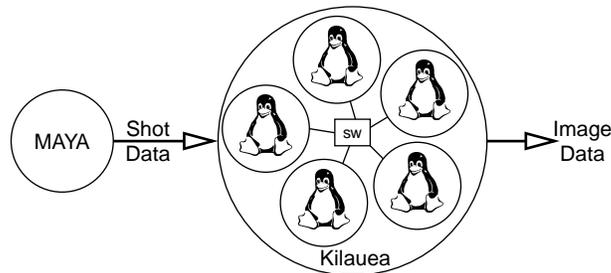


Figure 4: Maya and Kilauea interaction

### 3.5 Launching Kilauea

Kilauea is a system in which processes running on multiple machines communicate with each other in order to generate the final image. The boot process of Kilauea needs to be controlled between multiple machines in some standardized way. Kilauea daemons manage this by residing in Kilauea server machines and launching processes as requested, while communicating with each other (figure 5).

This Kilauea daemon also handles the termination of the Kilauea processes and plays a role in the fail-safe mechanism. The fail-safe mechanism is discussed further in section 6.9.3.

### 3.6 Single executable binary

Kilauea runs across multiple machines, but the executable binary which runs on all machines is exactly the same. One of the benefits in doing so is the simplicity of maintenance in freezing a specific version of the binary by just copying a single file.

<sup>1</sup><http://www-unix.mcs.anl.gov/mpi/>

<sup>2</sup><http://www.aliaswavefront.com/>

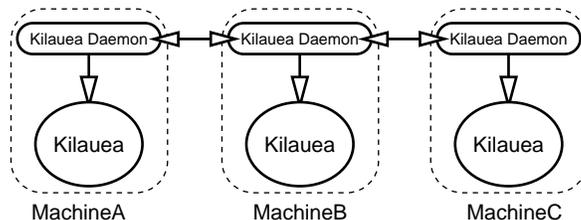


Figure 5: Booting Kilauea through Kilauea daemon

### 3.7 Multi-frame rendering

The Kilauea system is designed to render multiple frames at a time. This design has the advantage of reducing the time required to construct the data for the next frame by loading only the changes between the current frame and the next. For example, the scene construction time is greatly reduced if a very complex background scene is read in once and shared among multiple frames, rather than reading it in every frame. This becomes more advantageous as the amount and the complexity of static data in the animation increase.

Not having to reread data is even more useful when recomputing the same scene. Consider the situation where the user is experimenting with the shading parameter of some object's surface. Kilauea only needs to receive the modified shading parameter from outside (from Maya, for example) and immediately restart a rendering. In this case all the geometry data in the scene can be reused, reducing the time required for communication and rendering preprocessing to almost zero.

In order to take advantage of these traits, Kilauea is designed to stay resident as a process until it is explicitly shut down from outside.

### 3.8 Global illumination renderer

Kilauea has the following remarkable characteristics derived from its design as a global illumination renderer.

#### 3.8.1 Parallel ray tracer

Kilauea generates images by performing ray tracing in parallel. Kilauea executes two types of parallel processing depending on the scene size.

If the scene data fits completely within a single Kilauea process, each process will contain a complete copy of the scene. In this case, parallelizing ray tracing is very simple.

If the scene is too large to fit within a single Kilauea process, then multiple Kilauea processes are used to store the entire scene. Ray tracing is not so straightforward in this case. This will be discussed in detail in section 5.

Currently, Kilauea does not perform optimizations such as the ones specific to primary rays from the screen space. This is because the time spent on primary ray computation is only a very small percentage of the global illumination rendering. This type of optimization is one of the planned enhancements in the future.

#### 3.8.2 Antialiasing

Aliasing problems in screen space and along the time axis are all handled by supersampling. Because Kilauea uses ray tracing as its underlying algorithm, supersampling is the simplest method that can handle all cases.

### 3.8.3 Global illumination

Kilauea renders global illumination with the combination of final gathering and photon maps<sup>34</sup>. This allows global illumination effects such as caustics and color bleeding. The photon map method also has other benefits such as the relative ease of implementing motion blur with acceptable quality.

### 3.8.4 Volume shading

Kilauea allows the user to define a volume container within the scene. When a ray enters the volume container, ray marching is used to perform volume rendering. Global illumination may be performed on volume elements by using volume photon maps<sup>5</sup>.

### 3.8.5 Sub-surface scattering

Kilauea supports sub-surface scattering, which is necessary for special types of shading like a human skin. Sub-surface scattering also uses volume photon maps<sup>6</sup>.

## 4 The ShotData File Format

Kilauea uses the ShotData format to receive data from a front-end such as Maya. This section lists and explains the properties of ShotData.

### 4.1 Binary data format

ShotData is a binary data format. This makes the file size smaller and increases the read/write performance of the file. ShotData absorbs the byte order issues between different CPU architectures, allowing Intel Linux executables to read data from an SGI.

### 4.2 Version-free file format

File format specifications frequently undergo changes or additions as a result of debugging or the addition of new features. This is an annoying problem in an ordinary binary format. Because ShotData has a version-free data structure, it can adapt to these changes in a very flexible way.

ShotData is composed of various commands. For example, a command specifying a vertex is then followed by the values of the  $(x, y, z)$  coordinates. Every element of the scene can be described by these combinations of commands. There are two main problems with this approach alone. The first problem occurs when the number of parameters for a command needs to be modified for some reason. In this case, new executables must still be able to read any previous data that was previously dumped. The other problem occurs when adding new commands. Especially because this is a binary format, special mechanisms are necessary for skipping commands that the parser does not know about. To solve these problems, ShotData embeds information about commands and their parameter within itself. The first problem is circumvented by not allowing a previously defined command to be redefined. Whenever a command's parameter need to be changed, a new command is added instead and new data is dumped using the new command. Because command definition information is written into the ShotData, the second problem can be dealt with by skipping. The parser looks at the command definition block and parse only the commands which

---

<sup>3</sup>SIGGRAPH 2000 Course Note. "A Practical Guide to Global Illumination Using Photon Maps". [6]

<sup>4</sup>SIGGRAPH 2001 Course Note. "A Practical Guide to Global Illumination Using Photon Mapping". [7]

<sup>5</sup>Henrik Wann Jensen, Per H. Christensen. "Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps". Computer Graphics (SIGGRAPH '98 Proceedings), pages 311-320, July 1998, Orlando, Florida. [3]

<sup>6</sup>Julie Dorsey, Alan Edelman, Henrik Wann Jensen, Justin Legakis, Hans Kohling Pedersen. "Modeling and Rendering of Weathered Stone", Computer Graphics (SIGGRAPH '99 Proceedings), pages 225-234, August 1999, Los Angeles, California. [2]

it can understand. By using this mechanism, ShotData that was dumped a long time ago can be read by the latest Kilauea executable and vice versa.

### 4.3 Multi-frame format

ShotData contains data spanning multiple frames. ShotData can just contain a single frame, but generally an animation sequence of some length is recorded. The advantage of using this will become clearer after introducing the incremental data format described below.

### 4.4 Incremental data format

When recording multiple frames, ShotData defines the scene as incremental data between the frames. The current frame contains only the difference from the previous frame. Since the first frame obviously does not have a previous frame, it contains full information for the frame. When there is an object in the scene that neither moves nor deforms, the description for that object is only contained in the first frame. This tremendously compresses the data size for a sequence of animation.

A frame that contains all the data for the frame is referred to as a full dump frame. Thus the first frame of ShotData is always a full dump frame.

Because ShotData takes the form of incremental data, it can very efficiently handle the reading of frames in a sequence. When rendering only a specific frame in the middle of a sequence, having to read the incremental data from the top to the target frame is inefficient. Such requests probably will happen frequently. In order to handle these situations, full dump frames can be embedded at user-specified intervals (figure 6).

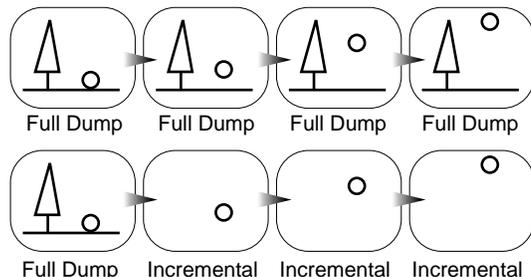


Figure 6: ShotData construction by incremental algorithm

### 4.5 Frame boundaries

Inside Kilauea, frame numbers take on a real number value, and frame  $N$  actually specifies a range between frame  $N - 0.5$  and frame  $N + 0.5$ . All frame numbers in ShotData are specified in this way. Motion blur shutter timing is specified by a moment value which is normalized to be between 0.0 and 1.0. 0.0 is equivalent to  $N - 0.5$ , and 1.0 is equivalent to  $N + 0.5$ . When motion blur is turned off, the image generated uses the values of frame  $N$ , in other words, shutter timing 0.5 (figure 7).

### 4.6 Motion blur information

To allow Kilauea to more correctly compute motion blur, incremental data for every subframe is written out to the ShotData.

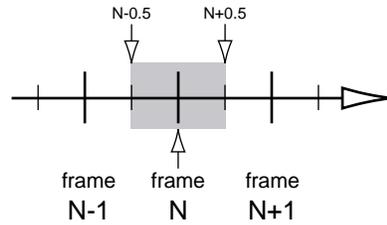


Figure 7: Frame number within Kilauea

Kilauea interpolates the subframe data internally to generate a motion blurred image. Motion blur precision can be increased by increasing the number of these subframes, which can be user-specified. Around three subframes is sufficient for most cases.

For example, when the number of subframes is one (motion blur off), a single frame is composed of one subframe. In this case the subframes will fall exactly at integer frame numbers. For example, when outputting the frame data between frame one and three, data at frames 1.0, 2.0, and 3.0 is the output.

If three subframes are specified, each frame will generate three subframes of data. When processing frames one to three, frame one generates data at 0.5, 1.0, and 1.5. Frame two generates data at 1.5, 2.0, and 2.5. Finally, frame three generates data at 2.5, 3.0, and 3.5. This data is finally represented as seven subframes at 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, and 3.5.

Users may only specify an odd number of subframes — i.e., 1,3,5,7,9... This way, the subframes always come to the center of the frame. Turning motion blur off is equivalent to simply rendering the subframe in the middle. This ensures the generation of consistent images between when motion blur is turned on/off.

Interpolation between subframes is done by using Catmull-Rom splines. This ensures that interpolated data go through the control points and the values at the subframes are always used. Linear interpolation of subframes is also available, if desired.

## 4.7 String table

Scene data usually contain a lot of string data. Strings are frequently used for object and light source names, shader names attached to the object surfaces, etc. The representation of these strings has a large impact on the file size because they are almost always reused between multiple frames. ShotData creates a lookup table of strings used in the scene. These strings are internally represented as indices to the string table. ShotData contains a string table block, which is also represented as incremental data between frames. By using this method, every string is represented as an ID in the string table, contributing significantly to file size reduction. Any processing that occurs within or after the parsing of the ShotData can check for identical strings by simply comparing this ID. This speeds up string search operations, which frequently occur.

## 4.8 Multiple ShotData files

Essentially, the upper bound of the scene complexity Kilauea can handle only depends on the total number of machines that Kilauea has access to. Therefore, a scene which is difficult to handle on a single machine can be rendered as long as Kilauea can read the data into multiple machines. However, Kilauea depends completely on the front-end environment (Maya) for the scene creation and therefore the scene size is practically limited to the size of memory available to the front-end. In order to work around this issue, Kilauea can read multiple ShotData simultaneously. The scene is divided according to some convention and converted into multiple ShotData. Each ShotData can be created independently, using the maximum memory allowed by the front-end environment,

and multiple ShotData are finally combined in three dimensions. In the actual rendering stage, Kilauea does not care whether a scene is composed of single or multiple ShotData, and simply renders it as one large scene.

Figure 8 shows the three typical ways in which shot data is created and used between Maya and Kilauea.

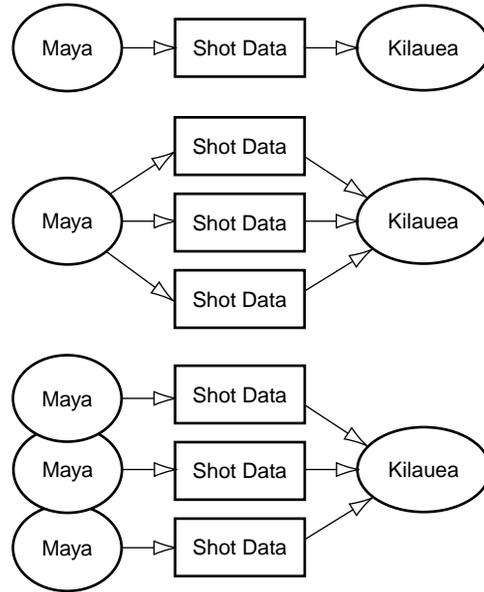


Figure 8: Multiple ShotData files

The multiple ShotData idea was originally designed as a way to load such complex scenes into Kilauea, but there are three properties which turned out to be useful when we actually started testing the system.

The first property is that since the scene can be converted to ShotData in pieces, more than one person can work on a single scene on different machines.

The second property is that ShotData can be created incrementally as work progresses, which means that there is no need to regenerate a ShotData for completed parts of the scene when doing test rendering. Once a part of a scene with director’s approval is converted into ShotData, the ShotData can be reused. This makes it easier to think about caching ShotData.

The third property is that by managing special-case data as separate ShotData, implementation can be made easier.

For example, consider the case of creating a special geometry data for “an object which undergoes deformation using cyclic animation.” This can be done by adding a new framework to describe this object within the ShotData, but by taking advantage of how ShotData works, implementation can be made simpler than that. “An object which undergoes deformation using cyclic animation” itself is represented as a sequential data of subframes which can be described as incremental data. The only missing piece is the length of the animation within the scene; all the other elements can share almost all the existing mechanisms of the ShotData format. By taking advantage of the ShotData format and interpreting the data differently, implementation can be made easier.

Rendering multiple ShotData combined within Kilauea does not suffer from any of the problems associated with rendering a large scene in separate layers. Objects casting shadows onto each other and color bleeding between objects naturally work perfectly fine. The human labor of dividing the scene into layers can be completely offloaded to machines.

ShotData can describe every element of the scene. This offers a great freedom in how the scene is divided and converted into ShotData. Users have complete freedom over dividing the scene on the basis of the elements composing the scene (such as light sources, objects, and shading information), or the progress of work (every ShotData has its own lights, objects, and shading information).

#### 4.9 Overwrite rule

When rendering with multiple ShotData and there happens to be the same data with different contents, a decision has to be made on which one to use. For example, an inconsistency occurs if there are two ShotData files and an object called A with different attributes exists in both files. Kilauea handles this problem by following an overwrite rule. When combining multiple ShotData within Kilauea, whether the existing object A will be overwritten by the new object A, or vice versa, is user-specified. By using this functionality, it is possible to create a ShotData that only contains the update information of a preexisting ShotData, applying a patch to the scene at loading time. This further broadened the beneficial aspects of the ShotData caching which was mentioned in section 4.8. This can be a very powerful feature in a production pipeline where every second counts.

### 5 Parallel Ray Tracing

Kilauea performs all hidden surface removal by pure ray tracing. Ray tracing fundamentally needs to reference all the geometry data in the scene. The algorithm simply consumes as much memory as the scene size, thus it poorly handles extremely large scenes. Kilauea overcomes this weakness in the ray tracing algorithm by sharing data among multiple machines. This is a brute-force approach, but it is a reasonable decision considering the price drops of PCs today.

In a simple scene where the entire scene data can be contained within a single process, each node computes ray tracing completely independent of each other, resulting in a direct linear speedup (figure 9).

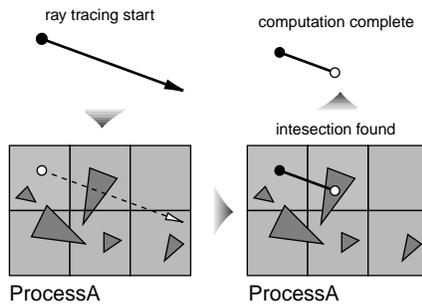


Figure 9: Ray tracing where one process contains the entire scene

The parallel ray tracing methodology when the entire scene is too complex to be stored in one machine is a highly debatable subject with various suggested algorithms. The outline on how Kilauea tackles this problem is provided here.

Assume that a certain scene is too large that two machines are required to store it. Several possible methods to execute ray tracing in this situation basically boil down to these two ideas:

1. Transfer object data among machines
2. Transfer ray data among machines

In the first approach, as machine A computes the ray's path, it will eventually need to retrieve the scene information from other machines (figure 10). This transmitted scene information must be stored in memory efficiently. There should be no extra memory to waste in the system since that is the reason why the scene was distributed to multiple machines in the first place. For the efficiency, geometry data should probably be separated into the one that stays resident and the one that is managed by some clever memory cache considering certain geometry size.

There has been several proposals on clever techniques to strategically distribute the scene data among multiple machines<sup>7</sup>. Nevertheless, performance of this approach essentially depends on how fast the geometry data can be transmitted between machines.

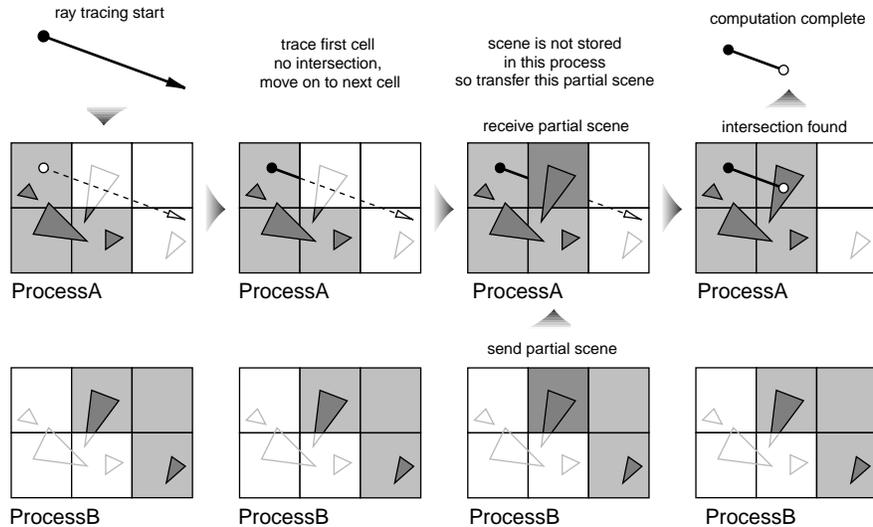


Figure 10: Ray tracing by transferring partial scene data

The second approach is the completely opposite approach. Instead of geometry information, ray information is exchanged between machines. If ray tracing cannot be resolved within the partial scene managed by one machine, the ray information is sent to another machine. The receiver of the ray then continues the computation (figure 11).

Generally speaking, the ray data is much simpler than the geometry data, making ray transmission a much more attractive solution. The disadvantage of exchanging ray data is that the amount of transmitted data depends on the quality of the image being rendered, the properties of the scene each machine is handling, and the paths of the rays. In a two-machine configuration, if ray tracing cannot be resolved in one, the other must have the solution. If the scene is distributed to more than two machines, the ray tracing may have to be performed over multiple levels (figure 12).

Kilauea ended up adopting the method which transfers the ray data instead of the geometry data. We created a prototype using the first method earlier in our development, but in the end we were unable to suppress the geometry transmission overhead to an acceptable level.

This geometry transmission overhead is due to the way Kilauea stores the geometry information. Kilauea restricts the type of primitive that can be rendered only to triangles. Any other primitive is internally converted to triangles. This geometry data is stored in a data structure called accel grid (see section 6.5.3) to speed up ray tracing. As scenes become more complex, required memory size increases. The maximum scene data size that one machine can hold is closely related to the

<sup>7</sup>Erik Reinhard. "Scheduling and Data Management for Parallel Ray Tracing". Ph.d. dissertation, University of Bristol Department of Computer Science, October, 1999. [5]

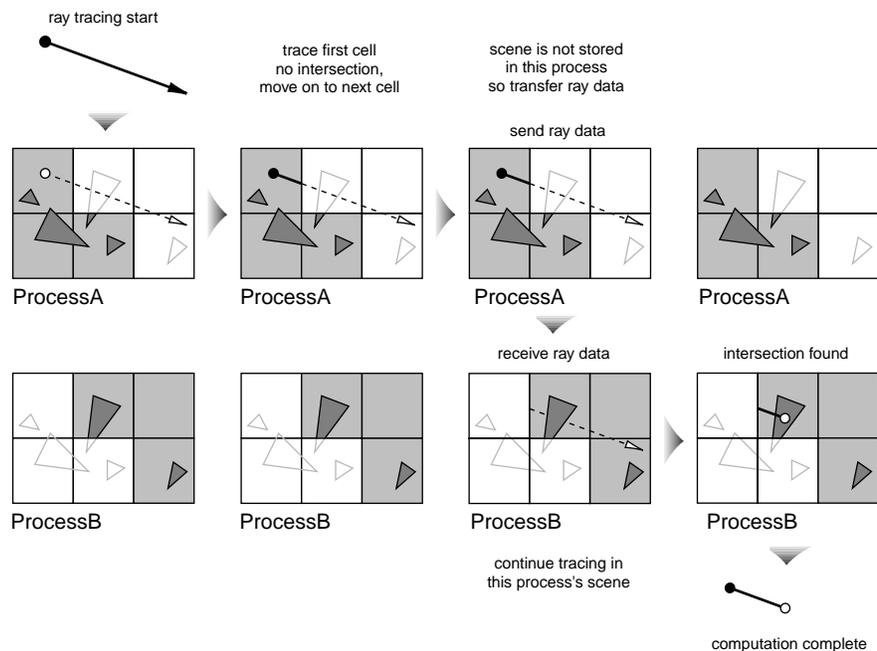


Figure 11: Ray tracing by transferring ray data

accel grid data size. Kilauea manages all the vertices of triangles separately from the triangle itself, allowing sharing of vertices and thus conserving memory as much as possible.

This is an excellent idea to save memory, but it becomes a problem when trying to send only a portion of the scene. Vertices are shared on a per-object basis and do not pay much attention to dividing the scene into parts. Shared vertex data in this form perform inefficiently in exchanging the scene partially.

On the other hand, exchanging rays between machines is a lot simpler than exchanging geometry data. This is clear from the fact that rays are basically independent of each other. The problem is then reduced to the speed of sending a large amount of ray data.

Kilauea's approach distributes geometry data statically across the different machines, and only the ray data is exchanged.

However, this does not solve all the problems. If the scene data is distributed among a very large number of machines, the ray may have to be transmitted from machine to machine. This is especially apparent when there are larger number of machines. The situation also depends on how the geometry data is distributed across the machines.

These situations depend on the state of the ray itself. Certain rays may be able to get a result on the first machine, while others may have to go though all the machines before finding a solution. By dynamically assigning what route the ray should take at runtime, this overhead might be minimized. This can be done by making a table of scenes that each machine has, and sending the rays using the table as a guide. Even if this portion can be optimized enough, however, the load may still be concentrated on one machine in situations where space traversal requests concentrate on a certain area in space.

Discussions above lead to the conclusion that an ideal method is the one which statically distribute geometry data across the machines, require no complex routing of rays, and ensure that the load does not concentrate on a single machine.

The stability was the top agenda in coming up with such a method. That is, an algorithm which only works efficiently for certain scene types but raises fatal problems in others is avoided.

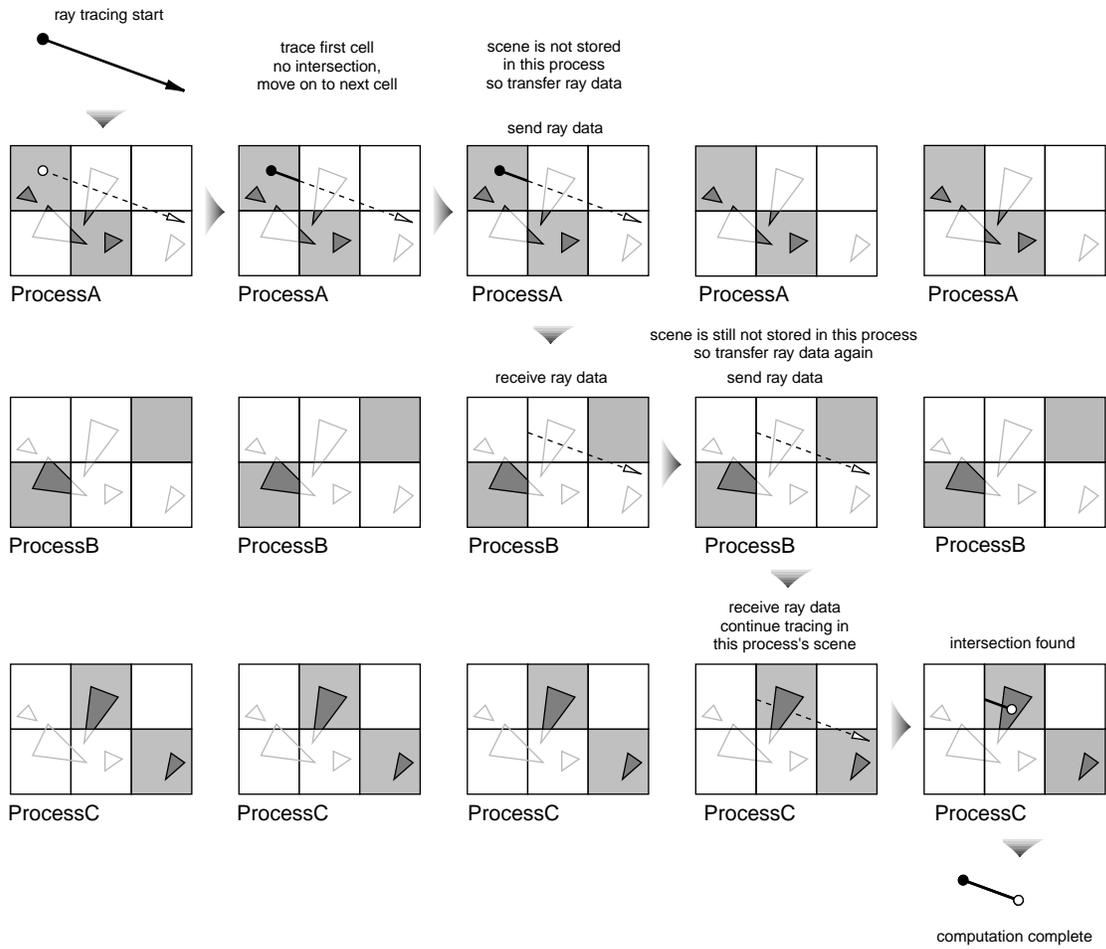


Figure 12: Ray tracing over multiple processes

Instead, we prefer a robust algorithm that works decently on all scenes.

Once again consider the case where the scene is distributed across the two machines in a completely random manner on a per-primitive basis. Ideally, each machine has close to half the scene's geometry data, and the spatial density in the machines is almost identical. The ray tracing performance characteristics on the two machines are statistically close to being equal, given enough rays (figure 13).

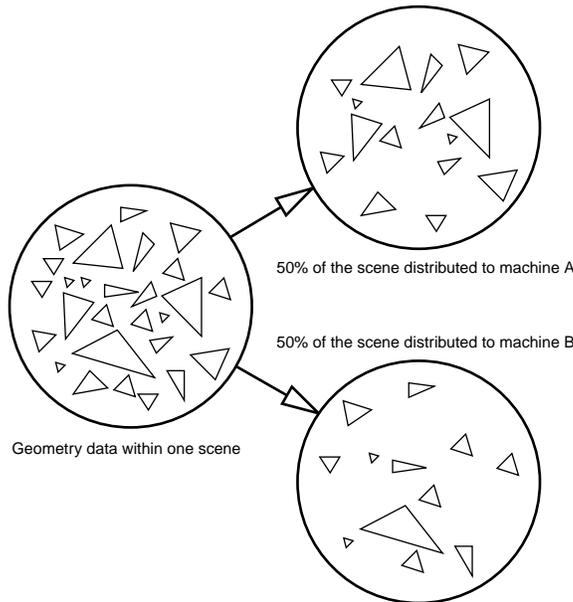


Figure 13: Scene data distributed to two machines

Ray tracing computation on scenes divided in this fashion is executed as follows (figure 14). First, the two machines perform the ray tracing on the exact same ray. Each machine independently performs the space traversal computation. Each of the ray tracing operations can be completely parallelized. Because the two machines share no scene data, the ray tracing results are obviously different. The results are compared on one of the machines and the closer intersection point is determined as the final result, discarding the other result.

The comparison of the space traversal requests can be performed in parallel as well. Two machines first get two completely different ray tracing requests. These rays are duplicated to two equivalent requests in both machines. Each machine sends one request to itself, the other to the other machine, and each machine performs the ray tracing independently. The computed result is sent back to the sender if the request originated on the other machine, and the returned result is compared with the result of the originating machine. The comparison finds out the closest intersection point as the final result, and one ray tracing is complete. As described in figure 15, the distribution of rays, the ray tracing, and the comparison of the results can all be executed in parallel.

This method may appear to be performing unnecessary computation, but it does have the advantage of being extremely stable under any pattern and/or size of the scene data. Even if the objects are distributed across a very large number of machines, it can be executed in the exact same manner. At this moment, experiments show that this method works reliably with up to seventy-five million triangle scene data, and there is no reason why it should not work for larger scenes.

A structure called accel grid is used to speed up the space traversal in each machine. This

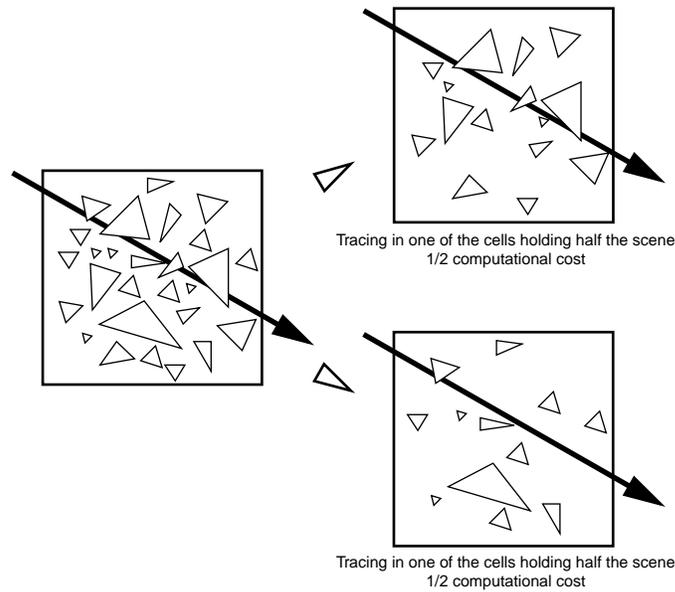
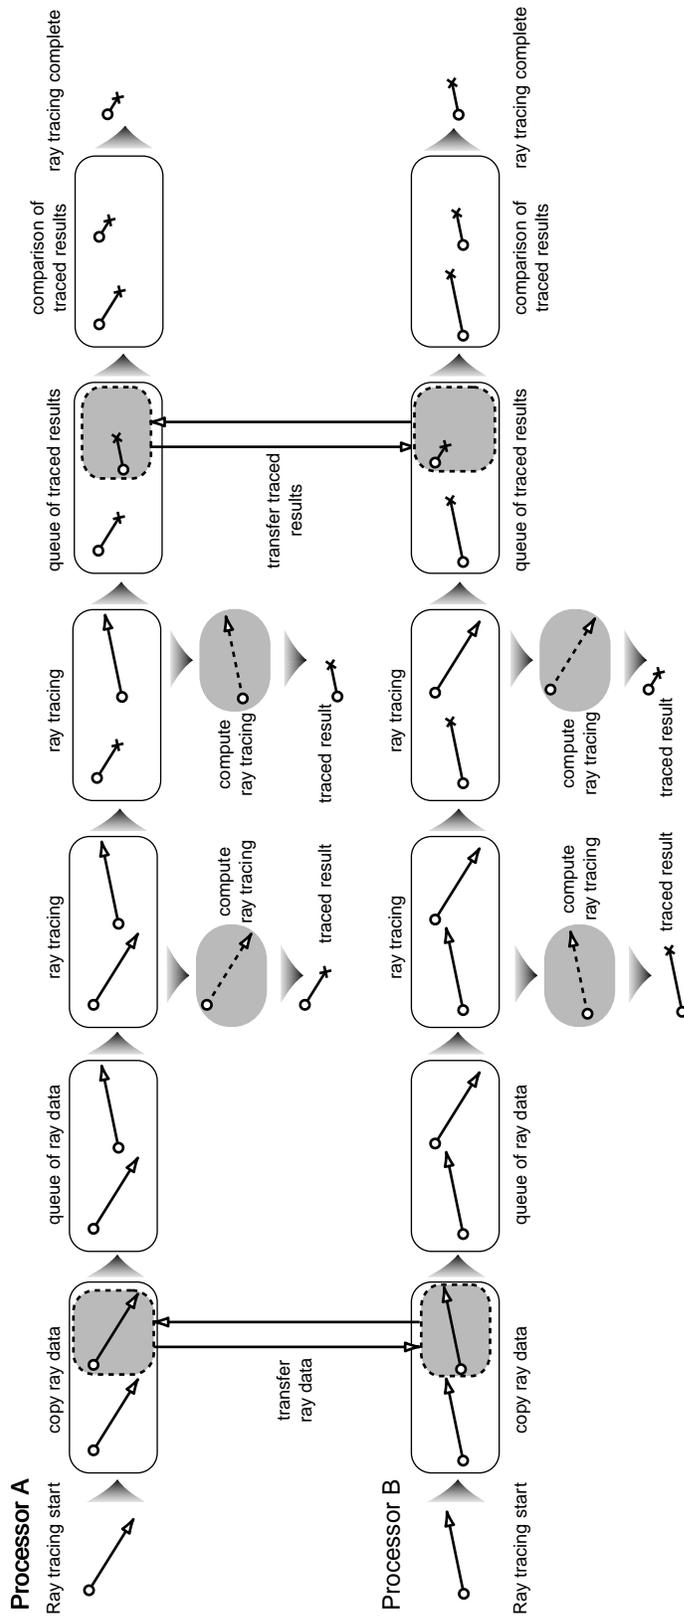


Figure 14: Performance of parallel intersection calculation to multiple objects within one cell

method accelerates the space traversal by dividing the entire scene into uniform cells. Ray tracing is done by traversing these cells, narrowing down the objects to perform the final intersection calculation. This method performs intersection calculation for every primitive in a cell and returns the closest one as the result. The parallel space traversal in Kilauea is equivalent to performing this in-cell intersection detection in parallel.

Each ray goes through a space traversal on multiple machines, including the originating machine. This process is performed in almost the same computation time across all machines. Every ray trace request resolves in one level of ray transmission and computation. This characteristic indicates that the method is far more reliable and tractable than the one where a ray must be transmitted over unpredictable number of levels to resolve. Experiments show that the load average using this method across multiple machines turns out to be almost equal, meaning that complex processing for load balancing is not required.

There are problems, on the other hand. The method may perform lots of wasteful computation that gets discarded eventually. Consider a scene in which the geometry is divided into spaces A and B. Inside A's accel grid, one triangle is placed directly in front of the camera. The scene is uniformly distributed (no data is shared), so this triangle does not exist in B's accel grid. In a such case, space traversal in A hits this big triangle and finishes immediately. B however has no way of knowing this. B must continue its space traversal further, until it resolves somehow. B performs far more computation than A in order to find an intersection point that gets thrown away. Clearly, reducing this wasteful calculation will be the key improvement to this parallel ray tracing. We are currently experimenting several ways to improve the algorithm, but unfortunately we have not implemented them in Kilauea yet.



Parallel execution of ray distribution, ray tracing, and comparison of results

Figure 15: Parallel execution of ray distribution, ray tracing, and comparison of results

## 6 Implementation

This section describes the various implementation issues in Kilauea.

### 6.1 Low-level data structure

Kilauea is a multi-level pipeline engine with large amounts of computation. Computation at each pipeline stage is implemented as an independent thread, and the engine is designed to let data flow through these threads. Low-level data structures which form the core part of this engine, and their programming policy, are discussed below.

#### 6.1.1 Queue data structure

One of the most frequently used and extremely important data structures in Kilauea is the queue. Kilauea has several special purpose low-level queue classes to simplify coding for specific needs. These queues are closely tied to the memory management scheme which is explained in section 6.1.2. Some techniques to solve the problems often encountered when developing specialized queues are provided in this section as well.

##### 1. Simple doubly-linked list

A linked list data structure is the basis of all queue structures. The Kilauea library defines a generic doubly-linked list base class. In most cases the lists need to be doubly-linked, though for some cases a separate single-linked list class is prepared for full optimization (figure 16).

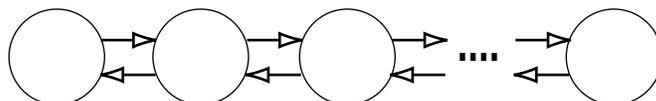


Figure 16: Simple doubly-linked list

##### 2. Multi-thread safe linked list

Whenever data needs to be exchanged between different threads using a queue, a mutual exclusion lock mechanism is necessary to insert or remove items to/from a queue. Multi-thread safety is basically achieved by just locking the queue every time insertion or removal operations to the queue occurs. This however can hurt the performance by causing frequent lock collisions between threads. In order to deal with this situation, the Kilauea library provides several multi-thread safe linked lists optimized with attention to these special cases (figure 17):

- (a) Insertion and removal can be restricted to a single thread, respectively
- (b) Insertion is done from multiple threads, but removal is restricted to a single thread
- (c) Insertion is restricted to a single thread, but removal is done from multiple threads
- (d) Insertion and removal are both done from multiple threads

The development stance of Kilauea is to first derive a queue class from the most appropriate one of the existing base classes. If further performance increase is desired in certain situations, they are individually optimized. However, the generic multi-thread safe list classes are optimized well enough that in most cases they can be used without any modification.

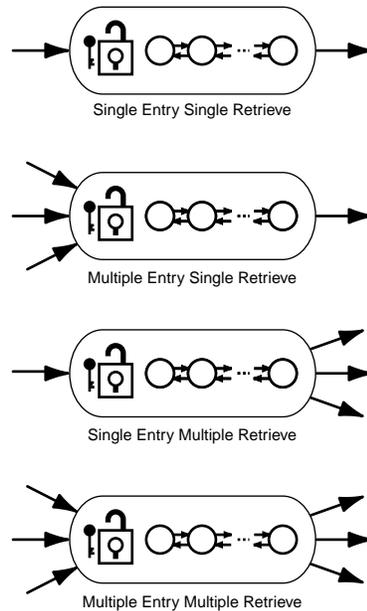


Figure 17: Multi-thread safe linked list

### 3. Lock problem

If accesses to the queue from multiple threads occur relatively infrequently, mutual exclusion locks do not cause too many problems. However, frequent collision of blocking locks have undesirable effects such as imposing a burden on thread scheduling and causing race conditions when acquiring the lock. One possible solution for getting around this problem is the use of a non-blocking lock (try-lock).

Figure 18 illustrates how this try-lock mechanism avoids lock collisions. Before inserting data to a queue, a thread first try-locks to check if it can acquire the lock. If try-lock is successful, the thread continues the insertion as in the case of an ordinary blocking lock. If try-lock fails, i.e. some other thread already owns the lock, then the thread gives up and inserts the data into a local list, hoping that it will be inserted in the following attempts. If try-lock succeeds in the next attempt, the thread first inserts the data from this thread-local list, and then insert the new data. The thread-local list continues to grow if the attempt fails again.

This mechanism allows increased independency of the thread execution by avoiding blocking when the lock cannot be obtained. Taking this idea further, the frequency of try-lock failures can be greatly reduced by not attempting try-lock every time.

However, queues using such mechanisms are essentially more complex than a simple thread-safe list structure, leading to a performance disadvantage. Storing data into a thread-local list in order to reduce the number of try-locks causes a latency in the data exchange between threads. Excessive use of this mechanism results in a serious degradation of performance.

The use of this mechanism should be restricted only to the case where a normal multi-thread safe linked list structure fails to perform well. It offers a very powerful optimization solution for some special purposes, however.

### 4. Condition wait

If a thread attempts to take data out of the queue but the queue is empty, the thread sleeps in Pthread's condition wait mechanism. The sleeping thread wakes up when another thread

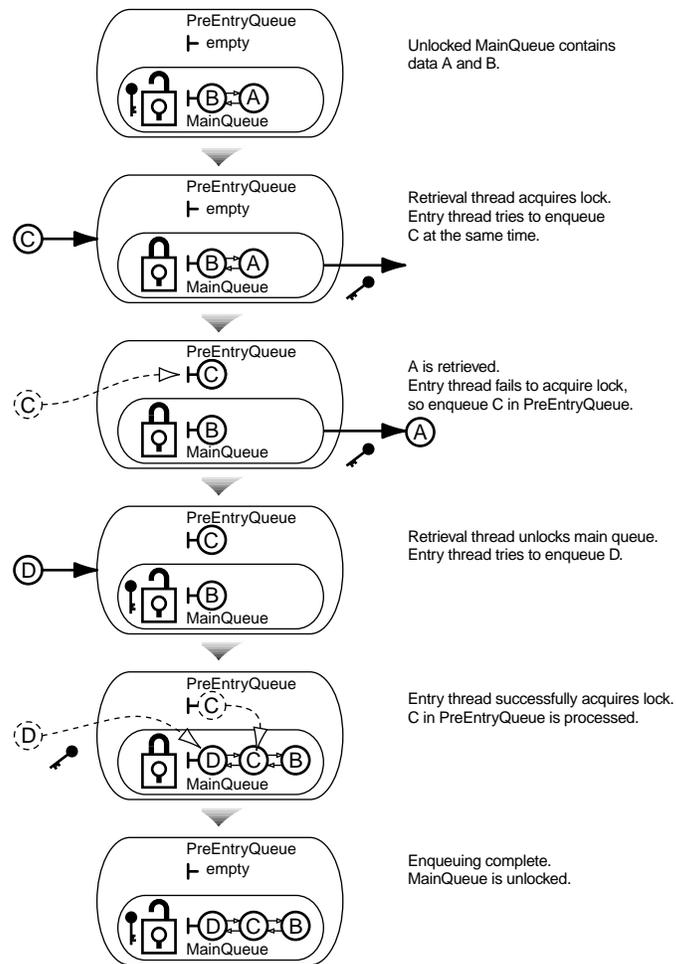


Figure 18: Use of try-lock and local queue to avoid lock collisions

inserts new data into the queue. Figure 19 shows this mechanism in action. Condition wait allows the CPU resource to be used more efficiently. However, frequent sleeps and wake-ups inflict a major performance overhead. The key to maximum performance is to keep the threads running at all times as much as possible and to sleep and wake up threads at appropriate times.

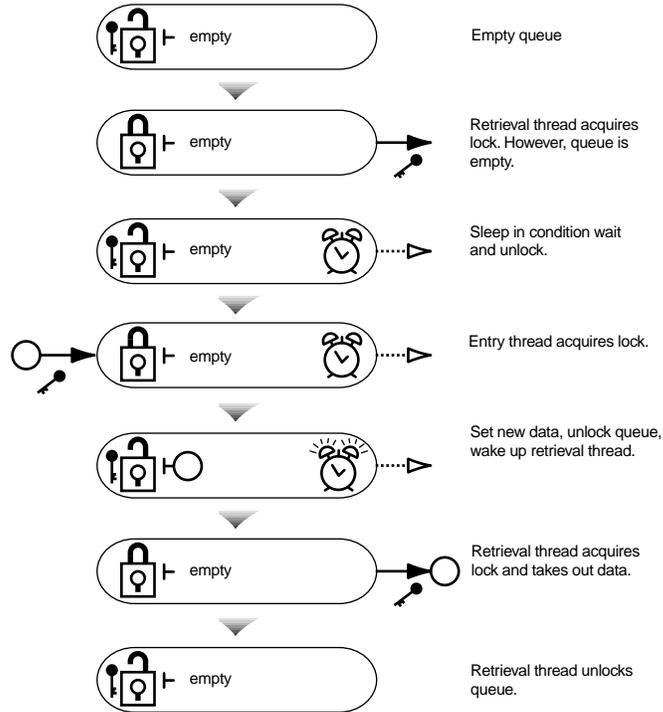


Figure 19: Condition wait

More information on how to manage threads and queues can be found in literatures such as [4].

### 6.1.2 Memory management

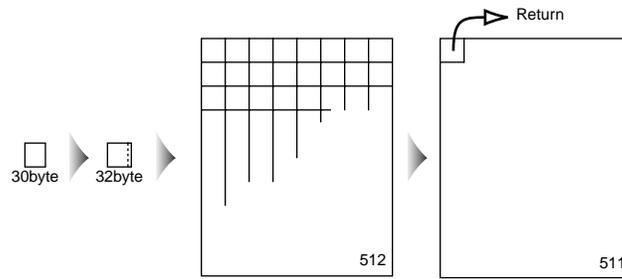
Special attention to memory management is vital when programming in a multi-threaded environment. Kilauea has a built-in memory manager to cope with issues such as multi-thread safety and fragmentation.

#### 1. Multi-thread safety

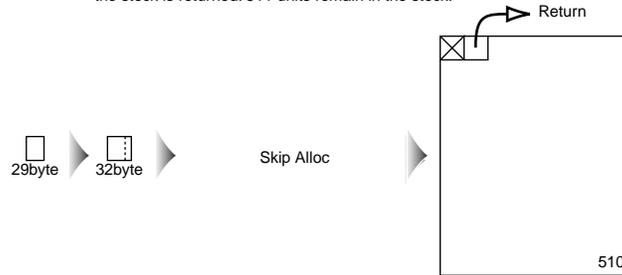
`malloc()` is usually used when dynamically allocating memory, but this causes serious problems in a multi-threaded environment. Kilauea has its own multi-thread safe `malloc()` with automatic locking. The C++ `new` operator is also overloaded to use this `malloc()`, so all normal construction using `new` is executed in a multi-thread safe manner.

#### 2. Dynamic memory allocation

Implementing original memory allocation routines instead of using the existing ones is very effective when the memory allocation pattern is predictable. In Kilauea, the memory allocation request size is usually extremely small, mostly under 16 kilobytes. Kilauea's own



A request for 30 bytes is treated as a 32 byte request. First 512 units of 32 byte memory block are allocated. The top one in the stock is returned. 511 units remain in the stock.



Next request for 29 bytes is also treated as a 32 byte request. Stock of 511 memory units already exists so a memory allocation does not take place. The topmost memory unit of the stock is returned. 510 units remain in the stock.

Figure 20: Memory management

memory manager is optimized with the focus on such small memory allocations. Every memory request is first categorized as either over 16 kilobytes or under 16 kilobytes. Requests under 16 kilobytes are then grouped into requests of predetermined size. Requests over 16 kilobytes are considered to be a request of continuous 16 kilobyte blocks. Figure 20 explains how the Kilauea memory manager handles requests under 16 kilobytes.

For every memory request, the very first one reserves the entire 16 kilobytes, which is then broken down into series of equally sized memory blocks. The next request is allocated from within the stock of memory that has just been created. The memory manager does not allocate any new memory area until all memory in the stock is used up.

Consider the case where an allocation of 30 bytes takes place. According to a predefined table for memory allocation size, this request is internally processed as a request for 32 bytes. When the first 32 byte request comes in, the memory allocation library allocates 16 kilobytes from the system. The 16 kilobyte block is then used as 512 units of 32 byte blocks. The first block from this 512 blocks is returned as the result of the memory request.

Assume that the next allocation request is 29 bytes. Internally this request is again treated as a request for 32 bytes. Since the memory manager already has 511 units of 32 bytes blocks internally, no new memory is allocated from the system, and the next available block is simply returned.

Whenever memory is freed, the freed memory is returned to the internal stock of memory.

If the stock is exhausted, a new memory space of 16 kilobytes is allocated and used in the same way. When every element of this 16 kilobyte block is freed, the entire 16 kilobyte block is freed in such a way that it can be used for memory requests of a different size.

In conjunction with the previously mentioned lock mechanism for multi-thread safety, threads

can independently manipulate memory in 16 kilobyte units, allowing a finer management of heap memory.

### 3. Thread local memory control

Frequent allocation of numerous equal-sized blocks of data occurs inside Kilauea. These memory requests repeatedly create a data structure of a certain size and then free it, so in most cases objects do not really have to be **new**ed and **delete**d. Once an object that has been allocated by **new** is no longer necessary, it is linked to an internal stock list and completely recycled. If correctly implemented, this recycling mechanism not only reduces the cost of construction and destruction, but also eliminates or significantly reduces the number of locks necessary in a multi-thread environment, thus increasing the independency and efficiency of each thread.

In some cases, an object that is **new**ed by one thread is **delete**d by another thread. Returning the object to the original thread and having the thread which **new**ed the object **delete** it solves the complications most of the time, as illustrated in figure 21.

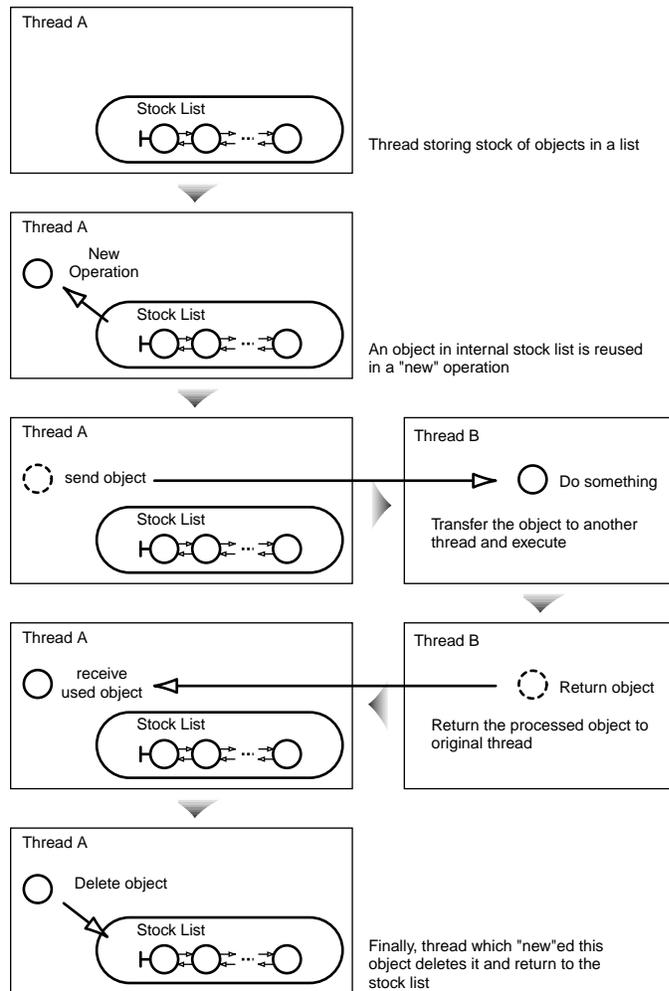


Figure 21: Thread-local memory control

The Kilauea library has a thread local memory management for random-sized memory requests in the same way as our implementation of `malloc()`. Using this memory management, developers do not need to pay special attention to make sure that memory is local to the thread. They can instead concentrate on the essence of the algorithm itself.

#### 4. Memory verification mechanism

Even when writing a sequential program, memory-related problems cause a lot of headaches at times. When writing a multi-thread program, the problem is even more serious. Such bugs are often very difficult to locate. Kilauea's memory manager has an integrated memory verification mechanism in order to easily spot memory problems. This is possible only because Kilauea has its own implementation of dynamic memory management routines such as `malloc()`.

When the memory verification function is called, tables constructed and managed internally are checked from various aspects for any inconsistency. This is done in a mutually exclusive way to all memory requests, thus changing the various timings such as thread scheduling. Therefore it is not effective in all situations. However, the verification failure most surely indicates the existence of a critical problem in the program. Problems such as mismatched `new(malloc)` and `delete(free)`, accidental reuse of memory that is currently being used, and freeing of memory that has already been freed, are relatively easy to spot using this memory verification.

In some cases, developers want to find out the details about the types of memory usage patterns. To satisfy such demand, the memory manager can log the memory management history to a file for later analysis. Doing so obviously puts a heavy load on the system and the thread scheduling and performance are significantly affected. Nevertheless, by carefully using this logging facility, developers were able to obtain valuable memory usage information and figure out the optimal memory management scheme.

#### 6.1.3 Thread tree structure

Kilauea uses the Pthread library for multi-threaded execution. The Kilauea library has a C++ wrapper class for starting and stopping threads. This class maintains every thread created by a process in a tree structure, and this tree can handle such operations as stopping all the threads that are subtrees of a certain node, or checking the execution status of a certain thread. Programmers are able to very easily create and control threads without knowing about the actual Pthread functions. Also, the program is able to monitor the status of a thread at any time in the program execution, based on this tree structure information. Monitoring of threads such as whether a thread has started executing or already completed executing can be done all by using this thread tree structure (figure 22).

## 6.2 MPI (Message Passing Interface) layer

The MPI layer handles the exchange of data between multiple machines. This layer defines the message passing information description and actually performs the message passing.

In the initial stage of development, we adopted MPI which was starting to become widely used at the time. At that point (around 1999), we were developing on the assumption that an external library will be used for this portion, but there were two problems.

#### 1. Reliability

At the time, we were using an MPI implementation which was available from SGI for SGI machines. With this MPI implementation, messages frequently did not arrive at the recipient or were lost in transit when a large number of machines were involved. In most cases it was a problem in our application, but we still had tough time dealing with the program dying

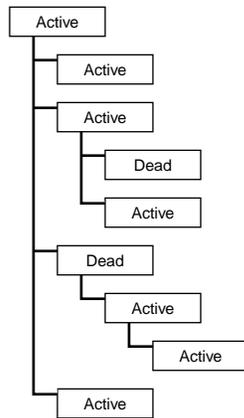


Figure 22: Thread tree structure

suddenly with a broken pipe. Also, the maximum number of nodes were restricted to around one hundred nodes at the time, which did not meet our requirements.

## 2. Multi-thread safety

The MPI implementation at the time was not multi-thread safe. This was a serious restriction for our application, since the renderer was initially designed to be multi-threaded with messages flying all over the place from each thread independently of each other.

After considering these issues, we have decided to implement a subset of MPI on our own. We based this decision on the following discussions inside the team:

1. MPI's spec covers a wide area, but our application only uses a small subset of it.
2. There is a specific pattern to the data size of messages being used inside Kilauea, and we want to specifically optimize for that.
3. We need a multi-thread safe MPI implementation.
4. We want to do a dynamic node configuration.

Kilauea's MPI implementation currently partially supports item 4, but it has complete support for 1, 2, and 3.

In implementing our own MPI subset, we took care not to change the MPI API whenever possible. This is to leave open the possibility of returning to a standard MPI distribution at some point in the future. Message creation uses the same methodology as in standard MPI. Message tags, ranks, etc. all follow MPI's philosophy.

Our implementation has the following characteristics:

1. Point-to-point communication only. Broadcast is not used.
2. Designed with multi-threaded environment in mind, and is implemented in a multi-thread safe manner.
3. Socket `send` and `recv` within the MPI layer is a multi-threaded implementation. Also, socket communication uses TCP instead of UDP.
4. Based on the analysis of message sizes required for Kilauea, the memory management is optimized for a specific data size, resulting in a speed up.

5. Simplified implementation based on the assumption that the byte order will be consistent within a single running system.
6. If the recipient of a message is the same process as the sender, then socket communication is bypassed for speed.

Kilauea currently performs all message passing using our implementation of MPI. The portability of this layer is extremely high, and we have been able to easily port this code to all our target hardware environments.

## 6.3 Tcl command interface

Kilauea is just a renderer, but it is far more complex than other renderers designed as a sequential program, in that it runs on a cluster of PCs communicating with each other. It is designed as a test bed for experimenting new features designed for very ambitious goals. Developers needed to be able to implement different features and easily test them individually. Also, since Kilauea stays resident like a daemon process once launched, having a way to control it from something like a console was useful.

Users also wanted to be able to control the various features of Kilauea from a command, and be able to combine these commands together.

For these reasons, Kilauea has a Tcl command interface. Tcl was preferred over other interpreters for its ease of integration.

### 6.3.1 Changes to Tcl itself

All Kilauea controls are processed through the Tcl interpreter. Kilauea opens a console displaying a command prompt, constantly waiting for Tcl command input.

The Tcl command interpreter was implemented as one of the threads inside Kilauea. Tcl is not intended to be used from threads this way and Kilauea needed to work around this problem.

1. Thread-safe Tcl

Thread safety here refers to using Tcl from inside Kilauea. From Kilauea's point of view, the Tcl interpreter is just one thread out of many running inside the entire Kilauea system. All threads inside Kilauea are managed in a tree structure, so any thread that the Tcl interpreter spawns needs to be put into this thread management tree. In the actual implementation, Tcl's thread initialization mechanism is modified to support this.

2. Sharing data such as variables between multiple Tcl interpreters

Spawning another instance of Tcl interpreter from the main Tcl interpreter can be very useful at times. The original Tcl interpreter had problems with data sharing in this case. For example, imagine the case of opening up a socket port to connect Kilauea and other programs such as Maya or telnet (we refer to this kind of connection as a back door). For such cases, the Tcl interpreter had to be modified to share the Tcl internal data in a thread-safe manner.

### 6.3.2 Kilauea command

Every Kilauea command is accessible as a Tcl command. Most commands are for fine-tuning specific functions inside Kilauea.

Kilauea developers add new Tcl commands to test new features. By doing this, specific functions can be tested quickly and individually, rather than performing complete rendering tests in the entire Kilauea. This increases the development productivity by conducting series of tests and perfecting

each component faster. After an individual feature has been tested thoroughly, this Tcl command is then used in the final rendering sequence.

General rendering sequences are represented by a combination of Tcl commands. A group of functions called the script library combines the individual low-level Kilauea commands, allowing users to manage Kilauea at a more abstract level. This system may seem very strange compared to other renderers. However, this is a very reasonable and flexible solution considering the architecture of Kilauea.

There are multiple units of computation inside Kilauea which are referred to as tasks (explained in the next section). For example, STask is responsible for shading and ATask is responsible for ray tracing. Kilauea can control which task to be launched on which machine with what kind of parameters, all from a Tcl script. Trying out all the combinations and various adjustments as many times as desired from a Tcl script is far more productive than modifying C++ source files and recompiling every time to conduct a test. Once a desirable configuration is found from these tests and decided to be adopted for general rendering, this configuration is embedded in a function inside a Tcl script. This allows flexibility in modifying settings on the fly and enhancing features.

Kilauea's front-end process composes and sends these Tcl commands to the Kilauea Tcl interpreter through a socket connection. Therefore users do not need to worry about Tcl scripting at all under normal circumstances. Power users may prefer to directly create Tcl scripts for full control. Tcl scripts can also provide a temporary workaround for a bug.

Kilauea usually reads in a scene from ShotData, but it is also possible to describe a scene completely from within Tcl scripts. This is clearly not a normal way to create a scene, and is very difficult to do. However this is effective as an emergency procedure at times.

### 6.3.3 Incremental rendering

Since the Tcl interpreter is implemented as one independent thread within Kilauea, Tcl commands may be issued while Kilauea is running some rendering task. This mechanism is used to stop Kilauea's rendering task in the middle. Being able to enter a command through the Tcl interpreter and dump the internal state of Kilauea is very useful for debugging at times.

Users are not only limited to viewing internal values. Changing values in the middle of rendering can be used for purposes outside of debugging. For example, users may change the material parameters or light sources dynamically. We are currently implementing a GUI for fine-tuning of material parameters and incremental rendering without rereading the scene.

### 6.3.4 Command interface to C++ objects

Kilauea handles almost everything in parallel by using threads, and each thread is implemented as a group of objects of a specific class.

Being able to send commands to these objects at run time will be effective for adding new features and debugging. To support this, we have implemented a unified interface for sending text command strings to a specified object. For commands such as passing parameters at initialization which are not executed often and small in data transmission size, raw text messages are directly sent to individual objects, where the messages are parsed and interpreted. This allowed us to implement new commands quickly without writing complex code for MPI style message passing.

## 6.4 Rank and task

The concept of "rank" and "task" needs to be explained first when describing the structure of Kilauea. Kilauea's message passing basically follows the MPI standard. The meaning of rank is equivalent to that of MPI's. When Kilauea is using three machines and one Kilauea process is running on each machine, there are total of three Kilauea processes running within the entire Kilauea system. Kilauea assigns rank IDs to these processes using a convention. In this case the ID takes a value of zero to two. Rank ID can be viewed as a unique ID for specifying a process

which reside within the Kilauea system. Rank also has a one-to-one correspondence with a Unix process (figure 23).

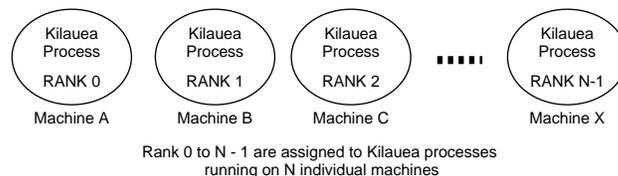


Figure 23: Rank ID

Inside Kilauea, there is a concept of multiple “tasks” which are units of some computation classified according to their function. Basically a task is a set of computation which has been classified and grouped together. Section 6.4.1 explains the roles of different tasks in Kilauea.

For example, WTask is a task that reads ShotData, expanding the incremental data and building the scene data necessary for rendering. ATask performs ray tracing on a scene. STask processes the shading.

The task assignment to ranks can be defined in a very flexible manner in Kilauea.

A task handler which is running inside each rank controls the execution of multiple tasks in parallel, using multiple threads. This means that multiple tasks can be assigned to a single rank.

The task configuration inside Kilauea can be separated from the correct execution of individual tasks. This property is very effective for moving development along flexibly. An optimal task configuration for a certain scene complexity can be adjusted just by modifying a Tcl script, which means that Kilauea can be easily adapted to the needs of a specific scene in a movie production.

A task handler running on each rank provides a general method for controlling a specific task by message passing of text command data, as shown in figure 24. Control commands can be sent to and results received from specific tasks without writing new code (also see section 6.3.4). The ability to stop and diagnose a specific task is also available within this general framework (see section 6.1.3).

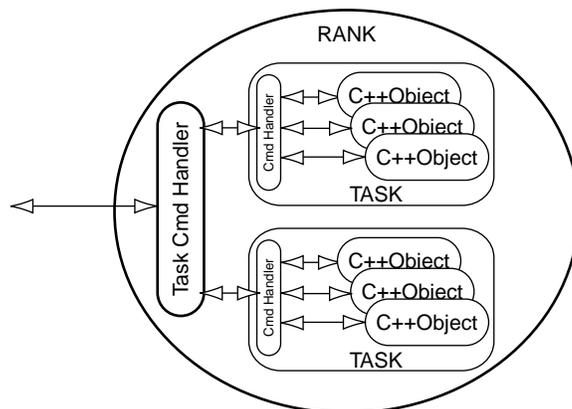


Figure 24: Task handler and text command interface

### 6.4.1 Task types

The following tasks currently exist inside Kilauea:

- MTask (Master Task)
- BTask (Binary Task)
- WTask (World Task)
- TTask (Tile Task)
- STask (Shading Task)
- RTask (Ray Task)
- ATask (Accel grid Task)
- ETask (Emit Task)
- LTask (Lookup Task)
- PTask (Photon Task)
- OTask (Output Task)

Multiple instances of each task may exist within a Kilauea system, and one task master is required for each kind of task.

Figure 25 shows the case where there are N STasks managed by one SMTask, which is the task master for STasks, inside the Kilauea system. All management of STasks are done through this SMTask, and individual STasks may not be controlled directly from other tasks. Other tasks will not even be able to find out how many STasks currently exist within the Kilauea system unless they query the SMTask. In the actual rendering, STask configuration details are queried ahead of time; data transmission destination is determined based on this information and directly processed for each STask.

There is one exception to the task master rule. This is MTask (Master Task). This task bundles all the task masters and controls them. There is no equivalent of an MMTask because only one instance of MTask can exist within the Kilauea system.

All task controls begin with the MTask. MTask sends control data to each task master, which then sends the data to the corresponding Tasks. This data path is for the task control only. In the actual rendering, data will be exchanged directly.

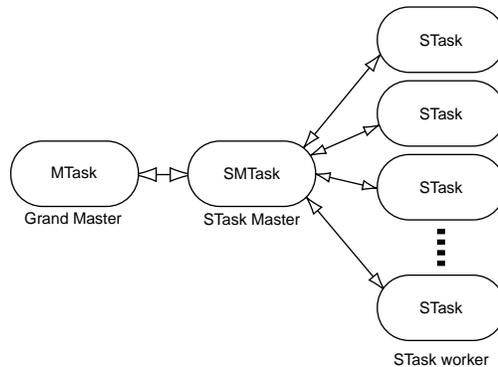


Figure 25: Task control inside Kilauea

The role of each task in Kilauea is explained below.

- MTask (Master Task)

MTask manages the information of all the task masters within the Kilauea system. MTask has an integrated Tcl interpreter, and is able to open a command console. Direct control of everything inside Kilauea is possible by entering Tcl commands through this console.

All control of tasks are executed with this MTask as a starting point.

Only one instance of MTask exists within a Kilauea system.

- BTask (Binary Task)

This task reads in external data files. Inside Kilauea, other tasks do not open and read from a file directly. All file I/O goes through this BTask, and data is then sent to the appropriate tasks via message passing. In most normal production work, the scene data is managed by a central server. When the Kilauea system is running and multiple machines are running in parallel, having every machine access the central server is not desirable. Restricting access to one machine and thus narrowing problems down to that machine and the central server is safer. When a very large number of file I/O occurs, creating problems between the central server and the machine running the BTask, augmenting the file I/O and network hardware performance of that machine solves the problem. BTask may be preferentially assigned to a machine customized for I/O performance to minimize the problem in communication with the outside network.

Data that is read in will eventually be sent to the appropriate tasks. This transmission is necessary for Kilauea system to render an image, and nothing can be done about the amount of data being sent here. However, transmission efficiency is easily improved by intelligently compressing data within the BTask. Also, because the network packets generated for this transmission do not need to go outside the Kilauea system, we can control the network congestion appropriately by working out the physical machine connections via a switching hub.

There are instances where data reading efficiency can be increased without generating actual file I/O by controlling the data caching inside BTask. Take the case of reading in the ShotData. Many file pointer seeks will be necessary to read in a ShotData. By cleverly using cached data to handle these seeks, it is possible to read data without generating an actual file I/O.

It is also possible to launch multiple BTasks within the Kilauea system. This allows file I/O to be handled by multiple BTasks instead of a single BTask, allowing them to be processed in parallel.

For these reasons Kilauea have adopted the design of using this BTask for all file I/O.

- WTask (World Task)

ShotData is represented as incremental data. WTask handles the computation necessary to interpret this data correctly and construct the scene data for the specified frame.

WTask handles the reading of multiple ShotData, as well as controlling the data overwrite management. The built scene is then sent to ATask. WTask is also responsible for distributing the scene across multiple ATasks.

Because WTask is in charge of building the scene data, it is idle during the actual rendering stage until the scene data for the next frame needs to be built.

It is also possible to launch multiple WTasks. In this case, each instance is in charge of building a part of the scene data. When multiple WTasks are launched, all the processing beginning from the incremental data analysis through the building of scene data and finally the transmission of data to ATasks, will be executed in parallel.

- TTask (Tile Task)

TTask creates the schedule for pixel sampling in screen space, and shoots rays accordingly. Even though the task is named “tile,” the scheduling is not restricted to just rectangular tiles. The scheduling unit is represented as an abstract data and implementation of a new scheduling algorithm is relatively simple.

TTask also handles the adaptive sampling in screen space. This is very effective in normal ray tracing. TTask also performs adaptive sampling along the time axis, which is necessary for motion blur.

Multiple TTasks are usually executed to shoot primary rays and collect their result in parallel.

- STask (Shading Task)

STask performs the shading computation on an object surface. Kilauea uses a method called SPOTEngine (described in section 6.6.3) for shading computation, and STasks control this SPOTEngine. This Task only handles the computation on the object surface, and does not perform further ray tracing from the surface. In addition, STask handles the texture data sampling and filtering for shading.

STask is not just limited to surface shaders. Every shader implemented in Kilauea is executed in this STask. Details about the various shaders are explained in section 6.6.4.

Multiple STasks are usually launched to handle shading computation in parallel.

- RTask (Ray Task)

RTask compares the ray tracing results returned from ATasks. A single ray trace is complete after RTask returns the result of the comparison.

- ATask (Accel grid Task)

ATask is the ray tracing engine. This task performs ray tracing, using a data structure called accel grid for optimization (see section 6.5).

ATask contains the geometry data of the scene, needed for ray tracing.

Ray tracing is computed in parallel by executing multiple ATasks.

- ETask (Emit Task)

Kilauea uses photon maps for global illumination. This task handles the photon shooting from light sources used in the photon map method.

By launching multiple ETasks, photon shooting is executed in parallel.

- LTask (Lookup Task)

LTask handles the photon lookup operation during the shading computation of object surface.

Photon maps can be distributed over multiple machines, and this task also handles the lookup of distributed photon maps (see section 6.7.2).

By executing multiple LTasks, photon lookup can be processed in parallel.

- PTask (Photon Task)

PTask contains the photon map data. This task stores photons, preprocesses them before rendering, and performs photon searching at rendering time. (see section 6.7).

By executing multiple PTasks, these operations can be processed in parallel.

This Task performs the equivalent of ATask for photon maps.

- **O**Task (Output Task)

O

Because Kilauea executes all sampling operations in parallel, there is no guarantee that the sampling results will return to O

Holding the entire image in memory may not be possible when rendering a large image. O

This task itself does not have the ability to display an image. It can however send the data via a socket to an external image viewer.

Currently only one O

### 6.4.2 Task configuration

The only task existing in the system after Kilauea boots up is M

Kilauea starts rendering as soon as the task configuration is determined. Currently, task configuration cannot be modified during the rendering phase. Runtime modification of the task configuration has not yet been discussed since there has not been any situation where this feature was necessary. This is listed as one of the enhancements for the future, however.

Many serious problems occur if dynamic configuration of tasks takes place at runtime. There are many dependencies between tasks while rendering, so an appropriate mechanism will have to be implemented in order to handle the dynamic modification of these dependencies. A flexible modification mechanism of inter-task dependencies is required in order to handle dynamic task configuration.

### 6.4.3 Task grouping

The scene data which is too large to fit in the memory of one machine is shared across multiple machines. This is implemented by grouping several A

During ray tracing, one ray sent to an A

Creating multiple A

Such task grouping can be applied to P

#### 6.4.4 Data flow between tasks

Rendering within Kilauea is handled by passing data between tasks. Within each task, data is repeatedly fetched from an input queue, processed, and then the result is written out to the output queue. Figure 26 shows how the data flows between the tasks.

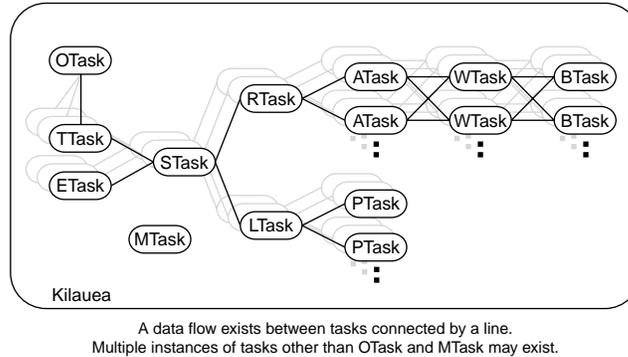


Figure 26: Kilauea internal data flow

Instead of sending individual data separately between tasks, they are often packeted and sent out in one shot to reduce the communication overhead if the tasks reside in different processes. The decision on what data to packet is solely dependent on the type of inter-task message passing.

### 6.5 Details of ray tracing

This section describes the details of Kilauea’s ray tracing engine.

#### 6.5.1 Engine design policy

Currently all hidden surface removal inside Kilauea is done by ray tracing. No optimization such as screen space scan line algorithm for primary rays is implemented. In a global illumination renderer, the percentage cost of primary ray computation is very small and optimizing this does not contribute much to the final performance. It is on the list of things to develop, but its priority is very low.

All rays of all generations are processed with exactly the same algorithm. In addition to the direct illumination rays, Kilauea also has photon tracing and final gather rays. These special rays are treated in the same way. Kilauea’s ray tracing engine requires a uniform performance regardless of the position or the direction of the ray. Optimizing just for a particular position or direction causes other computations to become more complex, especially shading.

#### 6.5.2 Primitives

There is a lot of discussion regarding optimal primitive types for ray tracing. In Kilauea, every primitive is converted to triangles in order to simplify the internal computation and use a unified algorithm. The ray tracing computation only considers intersection with triangles. Restricting primitives in this way puts less stress on developers for feature enhancements and maintenance.

The data structure of triangles changes depending on whether rendering with motion blur or not.

If there is no motion blur, i.e. triangles exist statically in the scene, then triangles are simply stored in the accel grid, which is explained in section 6.5.3. Intersection with the ray is simply an intersection with the stored triangle.

If there is motion blur, the path of the motion is computed and the data is stored in all accel grids where the motion passes. The path of motion blur is computed by sampling the coordinates of the triangle vertices at some intervals while the shutter is open, and connecting the points with Catmull-Rom splines. The volume constructed in this way is divided into parts along the time axis while the shutter is open, and the partial volumes of each part are stored in the accel grid. When computing the actual intersection with the triangle, it is moved to the position at the moment (time) of the ray. The triangle is moved by simply moving the vertices along the spline. After the triangle is moved to that moment, the intersection point is computed. In the case of motion blur, primitives are inside the accel grid as small blocks of space over the time axis of when the shutter is open, so it is possible to remove objects which will obviously not intersect at that moment.

### 6.5.3 Accel grid

Kilauea uses a hierarchical uniform grid to optimize the space traversal of the ray. As figure 27 shows, it is a data structure of space divided into equal parts, where each element (cell) is further subdivided recursively. If the division size is two, then the data structure is equivalent to an octree.

The reason why Kilauea uses a hierarchical grid instead of an octree is that the cost for space traversal is smaller than an octree, with the sacrifice of using extra memory.

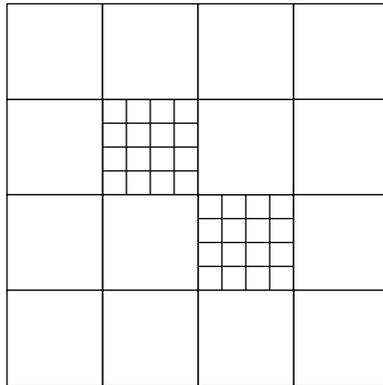


Figure 27: Hierarchical uniform grid

Accel grid does not need to be modified dynamically while rendering a specific frame, but is completely reconstructed when rendering the next frame. Accel grid does its own memory management in order to avoid problems like memory fragmentation when reconstructing a large accel grid.

When the scene data is small enough to fit completely within a single machine, the accel grids in each machine are identical when performing parallel ray tracing in multiple ranks. In this case, each rank independently performs the exact same computation to generate the same accel grid. Ideally, performing the same computation should be avoided to speed up the accel grid construction using parallel processing. But in that case, partial accel grids must be transmitted to other ranks as they are being computed. Since accel grids internally reference data using pointers in a complex way, this is not simple. In any case, the time required to construct the accel grid is only a small part of the entire rendering computation, so we do not think this is a very critical issue for now.

If the scene does not fit in a single rank, then the scene is distributed over multiple ranks (see section 5). Each rank independently generates accel grids in parallel. There is a high level of parallelization here, and we can expect a speedup directly proportional to the number of machines. This is because no data is exchanged between ranks in the creation of accel grid. However there are two problems.

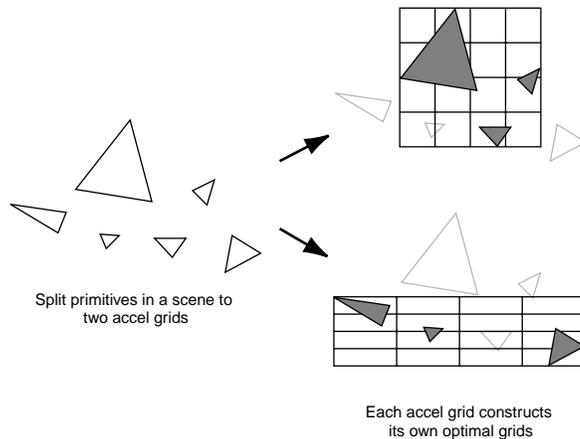


Figure 28: Problem with parallel accel grid construction

The first problem is space correspondence. Independently creating accel grids for scenes divided in multiple ranks implies that there can be almost no correspondence between the created grids. Each accel grid should generate an optimized accel grid for its particular data. This causes problems when accel grids generated this way need some sort of correspondence (figure 28).

As mentioned in section 5, Kilauea currently exchanges no geometry data whatsoever between the ranks. Therefore specifying some part of space, removing that part, or inserting that part will never take place. But if we ever want to do this, having differently structured accel grids in two different ranks is a problem.

The second problem is in the precision of space traversal. This is a problem in the precision of the accel grid. For example, the precision of values used to describe the size and position of the grids is an issue. We construct accel grids as a hierarchical uniform grid, but traversing inside the grids is done using incremental computation whenever possible. The precision of the traversal is related closely to the precision of the incremental computation, and ultimately the precision of the size and position of the grids.

This implies that ray tracing results could differ for the exact same data, depending on how the accel grids were constructed.

Because our ray tracing algorithm executes the traversal of the exact same ray on multiple accel grids, discrepancies in the space traversal results could be fatal. We must somehow ensure that the results are computed at the same precision.

We handle this problem as follows. The size and position of accel grids are specified using an IEEE single precision floating point number (32 bits). We precompute the space that can be represented by this number into a table, and make sure that grid boundaries only fall on values in this table. By doing this, we have made it possible to keep the size and position information of accel grids completely consistent no matter what kind of data is created in what order. By doing this, we have ensured that the final ray tracing result is always done at the same precision (figure 29).

## 6.6 Shading computation

Shading computation inside Kilauea is separated from the rest of the system such as parallel ray tracing. This allows the shading computation implementation to be improved independently.

Kilauea’s shading engine handles the various issues of parallel processing all “under the hood”; shader writers do not need to think about parallel processing and can write their shaders as if Kilauea was an ordinary sequential ray tracer.

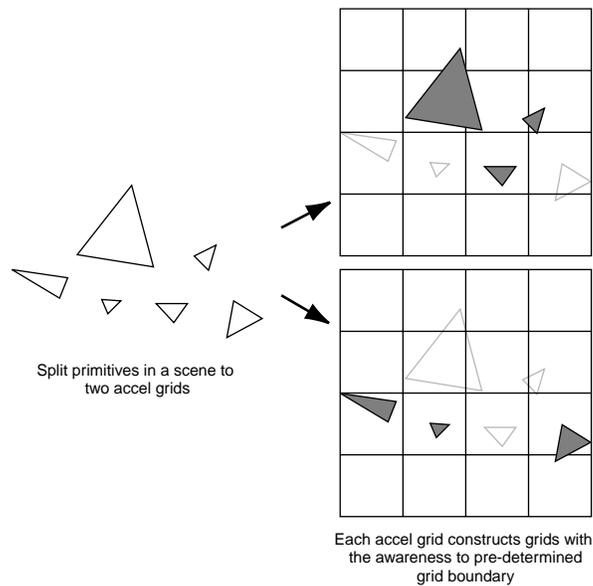


Figure 29: Avoiding problem with parallel accel grid construction

### 6.6.1 Issues with parallel shading computation

Inside Kilauea, shading computation on the object surface and space traversal using ray tracing are processed separately in different threads. They may even be processed by different machines (depending on the configuration of tasks). This raises situations where something that can easily be done on a sequential shader is extremely difficult.

Say that some shading computation is being done on some object surface. This object has a specular reflection, so the color of reflected ray must be taken into account for computing the final color of this surface.

In a normal sequential ray tracer, shading computation is done on the object surface. The reflection ray is then computed, and a new ray is shot. The new ray is processed recursively, before any other computation takes place. This result is then used in the surface computation, and the final color of the object is determined. However, this cannot be done easily in Kilauea.

Ray tracing must be done in order to compute reflection. Since the computation on the surface and the ray tracing are done in different tasks within Kilauea, ray data must be exchanged between the tasks in some way. In the worst case, the desired task may be on a different rank, so data must be transmitted using message passing.

Ray tracing is then executed on the transmitted ray, performing a new shading computation which intersects another object. The computed color is then returned to the original object surface and added to the surface color, determining the final color of the surface. The problem here is that since the ray tracing is being done in a different task, computation on the object surface cannot proceed until the ray tracing resolves. If the surface computation stops until the color of the ray comes back, CPU cycles are wasted, seriously degrading parallel processing performance (figures 30 and 31).

### 6.6.2 Overview of solving the parallel shading problem

There are mainly two solutions to this problem:

1. Do not allow shaders that depend on ray colors

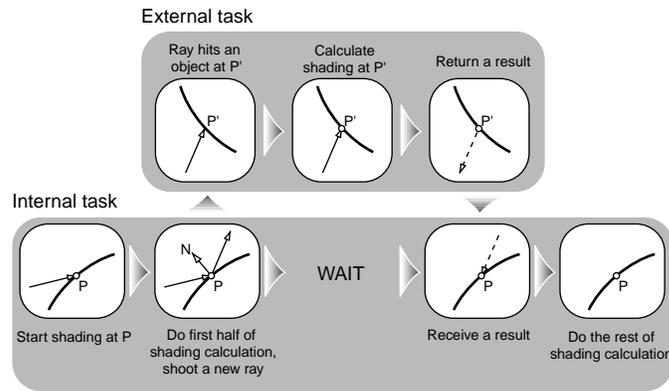


Figure 30: Problem in parallel shading calculation

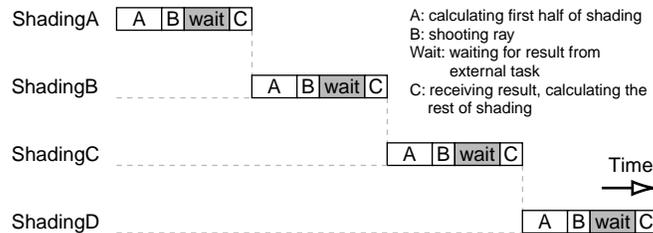


Figure 31: Apparently wasteful CPU processing

## 2. Somehow keep the CPU busy

The first solution gets around the problem by changing the way in which shaders are written. For example, by restricting shaders to not perform a computation based on the color of the returned ray. However, we have given up on this approach because being able to easily write shaders is extremely important for a production renderer.

The second method allows users to write a shader in the same way as a sequential ray tracer without having to use a special way of thinking. There are several ideas for accomplishing this.

One idea is to send the intermediate results attached to the ray whenever a new ray needs to be cast. Computation proceeds as the required elements become available, and the solution can be found when all the elements are available. This appears to avoid making the CPU idle.

This idea may seem to work, but there are problems with implementing it. The information added to the ray will most likely grow very large, especially for complex shaders. This is not desirable for Kilauea, since Kilauea is designed to frequently exchange rays between ranks.

Also, in some cases the distributed computation must be collected in one place. One example is when computing the shadow contributions by collecting multiple shadow rays shot from the object surface. To do this computation using this method, multiple shadow rays must be sequentially evaluated. Multiple shadow rays are independent, so the computation should be parallelized.

Because of these reasons, we have adopted the following method (figure 32).

1. If a new ray needs to be shot while computing a surface, shoot the ray and wait until the result comes back.
2. But this will cause the CPU to be idle, so start a new shading computation.

3. Usually, there are a large number of shading computations going on inside Kilauea, so execute whatever shading computations that can be processed in turn, in order to keep the CPU from being idle.
4. Some shading computations will shoot a ray and wait.
5. The first ray comes back while doing all this.
6. The returned result is used in the original computation, and the computation continues.
7. Finally all the computation completes, and the colors are finalized.

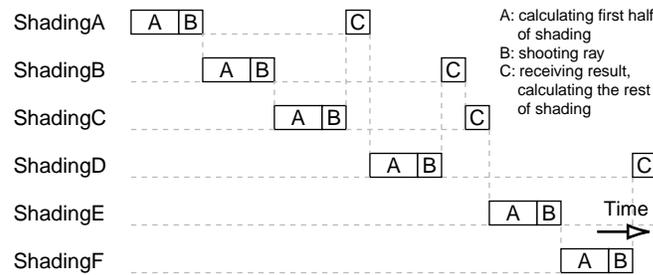


Figure 32: Keep CPU busy by optimally scheduling the shading

This method takes advantage of the fact that there are many independent shading computations going on, and controls them so that the CPU does not become idle. While the previous method tries to solve the problem by adding more information to the rays — i.e.; send more messages, this idea does not change the amount and the size of the messages. Instead, the execution order of the computations are changed. On the other hand, this method requires much more memory in order to manage the execution sequence.

We have designed a shading engine based on this idea that is general enough. Memory management has been optimized as much as possible by writing special-purpose mechanisms.

Currently all shaders run on top of this engine. We call this engine the SPOTEngine (Shading Parallel Object Task Engine).

### 6.6.3 SPOTEngine

The SPOTEngine is explained here in terms of shader creation and computation.

1. Decomposition of shading computation

When implementing a shader, the shader writer breaks down the shader computation into functions performed by an external task and functions performed by the shader itself. There are only two kinds of computation performed by an external task:

- (a) Ray tracing
- (b) Photon lookup

The shader computation is in a wait state whenever these two tasks are being processed.

Once the shader computation is decomposed into these two elements, any part of the computation which does not depend on outside tasks can be performed sequentially. We refer to these parts as SPOT (Shading Parallel Object Task, pronounced es-pot).

Computation to be performed after receiving data from outside tasks can also be considered SPOTs for the same reason.

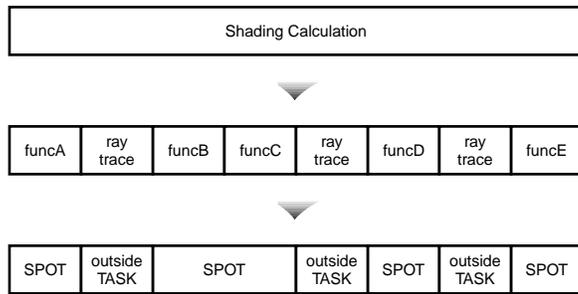


Figure 33: Decomposing shading computation

Therefore, computation before and after external tasks can always be combined into a SPOT. An SPOT that receives data from an outside task and acts on it is referred to as receiving SPOT.

Dividing shader computation into SPOT execution units easily allows the SPOTEngine to determine under what condition some computation can proceed with its execution (figure 33).

Ultimately, one shader is composed of multiple SPOTs where the data flows through to perform a shading computation.

## 2. SPOT network

SPOTEngine internally treats an SPOT as some unit of execution that cannot be decomposed further. SPOTs can define data along with computation (method). This represents exactly what an object is in an object-oriented philosophy (figure 34).

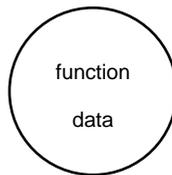


Figure 34: SPOT

Actual shaders are represented as a network structure of SPOTs that pass data to each other and perform the computation defined for each SPOT. The network can be structured in a very flexible way, and all shaders are described with this network. In other words, Kilauea executes shader computation by passing data in sequence along this SPOT network (figure 35).

At actual rendering time, there are as many SPOTs waiting for data from outside tasks, or sending data to outside tasks (i.e. ray tracing requests or photon lookup requests) as there are a number of shading computation being performed inside the SPOTEngine. SPOTEngine's job is to appropriately pick out a SPOT that can be executed and process it. Shading computation inside Kilauea proceeds by processing massive amounts of SPOTS quickly. SPOTEngine's performance determines the speed of shader computation and required memory.

## 3. SPOT (Shading Parallel Object Task)

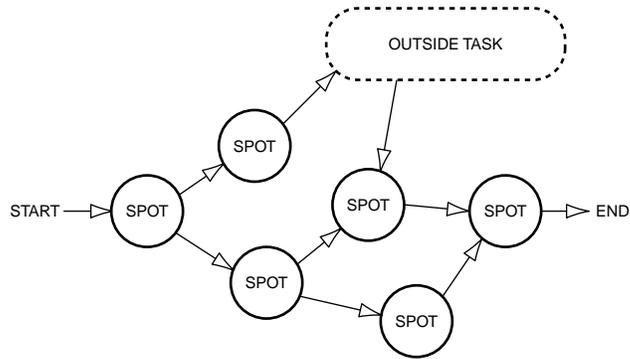


Figure 35: SPOT network

As previously mentioned, SPOT is the smallest execution unit inside the SPOTEngine. The characteristics of SPOTs which are required to build the SPOT network will now be described. SPOT can define data entry points called “slots.” (figure 36)

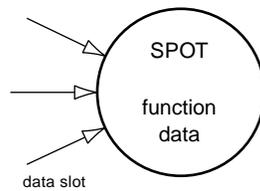


Figure 36: SPOT slot

These slots are used to get data from other SPOTs, or receive data from outside tasks. These slots are basically used to describe the network structure of SPOTs. Most SPOTs exist somewhere on the network, so they usually have at least one or more slots, but it is possible in special cases to define a SPOT with zero slots.

An SPOT also has an execution status. Currently there are four kinds of execution status (figure 37):

- (a) Waiting
- (b) Active
- (c) Dead
- (d) Static

A waiting condition means that the SPOT is waiting for values to be inserted to its slots. The method defined in the SPOT cannot be executed yet.

An active condition means that all the slots are filled, and that the method defined in the SPOT can be executed.

Dead means that the SPOT has finished its execution.

Static is a special case, used for defining a SPOT that exist regardless of the other conditions mentioned above.

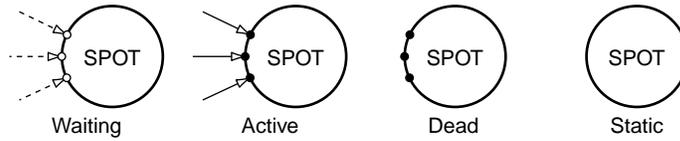


Figure 37: SPOT Condition

In the network that uses slots, it is easy to determine which SPOTs can be executed by monitoring the SPOT's status. SPOTEngine tracks the condition of every SPOT that it manages, and executes the methods defined in the SPOTs as they become Active.

In an actual shader implementation, SPOTs derive from the basic SPOT class, implementing the required computation as a method defined in the subclass. Computation such as sending data to other SPOTs, sending data to external tasks, and receiving data from other tasks, are all performed by calling the various methods implemented in the basic SPOT class. Programmers need not worry about these details when implementing a new SPOT.

#### 4. SPOTSpace

Inside the SPOTEngine, multiple SPOTs exist connected in a network structure. SPOTEngine's main job is to execute the SPOTs that are ready to run. To handle this more efficiently, SPOTs related to each other are grouped into an SPOTSpace.

Specifically, one shading computation unit (all computation on a particular surface) composes one SPOTSpace, and all SPOTs required for that surface are constructed within that SPOTSpace. When this shading computation is complete, SPOTSpace ends its purpose and is destroyed (figure 38).

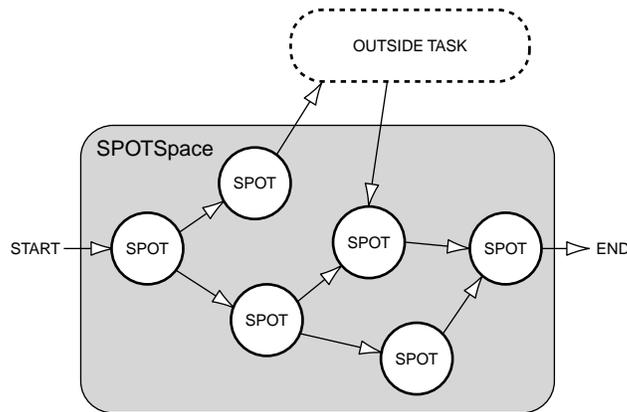


Figure 38: One shading calculation

Shading computations inside Kilauea are basically independent of each other. Some shading computations cannot refer to an intermediate result of another. For example, shading computations on two different pixels on screen are processed as completely independent computations, and one cannot refer to the results of other while computing. The sequence in which shading computation is performed depends on the particular state of parallel processing at the time, and for this reason, it is not possible to guarantee the availability of data (figure 39).

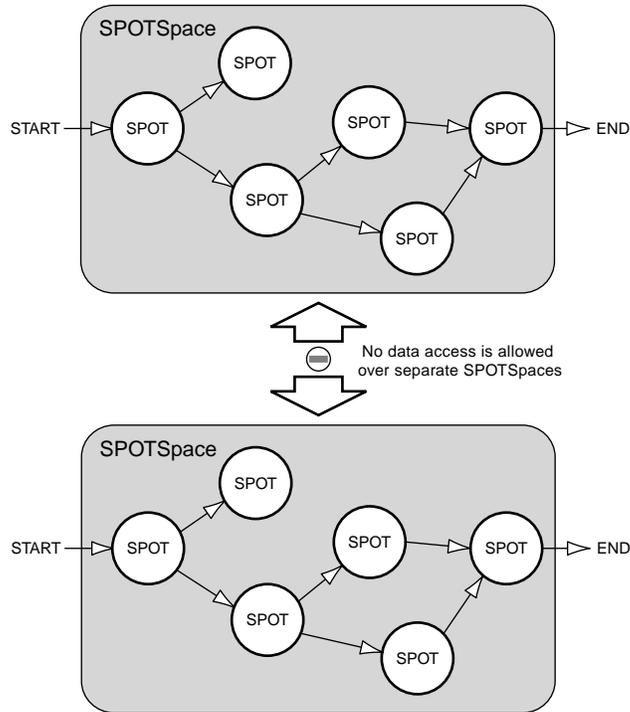


Figure 39: Prohibition of data exchange between SPOTS in SPOTS spaces

SPOTEngine uses SPOTS spaces to make one shading computation completely independent, and to help handle the memory management of one shading request appropriately.

At rendering time, multiple SPOTS spaces exist inside the SPOTEngine, and a network of SPOTS exist within the SPOTS space. SPOTS in the network that are ready to be executed are processed as they become ready.

#### 5. PSA (Pot Space Address)

Consider the case in which one SPOT shoots a ray and another SPOT is waiting for its result. The shot ray goes through an appropriate path and is sent to the accel grid for space traversal. If some intersection point is found as a result, then a new shading computation begins at that point. The color at that point is then determined, returned to the original SPOT, and computation proceeds according to the SPOT network inside that SPOTS space.

To make this work, one important mechanism needs to be implemented. That is the mechanism for correctly returning the result to the original SPOT that requested the color information. In order to meet this requirement, we need an address information that uniquely points to a single SPOT within the Kilauea system.

We call this address the PSA (Pot Space Address). All data going to an outside task use PSA to specify the destination SPOT where the result of computation should be sent.

Usually inside Kilauea system, there are multiple STasks in multiple ranks. Inside an STask is an SPOTEngine. Inside an SPOTEngine are a large number of SPOTS spaces. Inside an SPOTS space are a large number of SPOTS. PSA is an address that can uniquely specify a single SPOT out of the very large number of SPOTS, from when the Kilauea system is launched to when it is shut down.

By using the PSA one can guarantee that data will be sent without errors. PSA information forms the root of Kilauea’s shading computation scheme

## 6. Multi-threaded implementation

SPOTEngine itself, like other parts of Kilauea, is a multi-threaded implementation.

SPOTSpaces do not need to communicate with each other, and it is simple to find parallelism in the evaluation of the SPOT network in SPOTSpace. So there is no reason why we shouldn’t try a multi-threaded implementation. We have divided the computation inside SPOTEngine across several threads.

## 7. SPOT garbage collection

The massive number of shading computations are represented inside SPOTEngine by a large number of SPOTSpaces and the network of SPOTs inside the SPOTSpace. Memory used by SPOTs must be freed as soon as they perform their computation and become dead. Otherwise the available memory will be exhausted.

However, freeing of memory must be done carefully in a multi-threaded SPOTEngine.

If the thread which allocated the memory is different from the thread freeing it, memory management must be coded to expect accesses from two different threads, and be coded in a thread-safe manner. This makes coding more complicated, and is bug-prone. This also makes optimization and modification more difficult. One solution is to always return memory to the thread that allocated the memory, and let the allocating thread handle the release. SPOTEngine manages memory in this way, and increases efficiency by performing thread-local memory management.

Detection of dead SPOTs is one of the features implemented in the SPOT class itself, and subclasses of SPOT used for writing shaders do not need to handle this. Similarly, SPOTEngine directly controls the freeing of dead SPOTs, optimally handling memory release automatically.

Thus, shader writers do not need to take into consideration the small details of multi-threaded memory management — the SPOTEngine handles everything automatically.

### 6.6.4 Shader types

The functionality required for shading computation is classified into several types of shaders in Kilauea. Currently Kilauea has the following shaders:

- Emit shader
- Surface shader
- Light shader
- Background shader
- Volume shader

All of these shaders are implemented as SPOT networks, and SPOTEngine executes them.

All shader types are described below.

#### 1. Emit shader

Any kind of computation involving the shooting of rays is classified as an emit shader. Inside Kilauea, there are two main cases where this shader is necessary.

The first case is when controlling the rays being shot from the camera in normal ray tracing. Each camera has one emit shader. By writing an emit shader that takes into account the

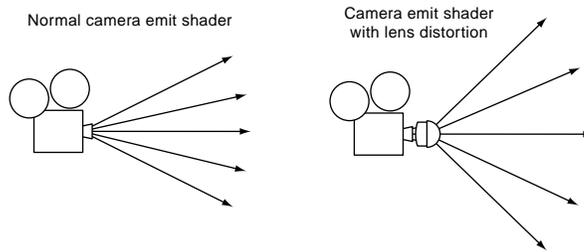


Figure 40: Camera emit shader

characteristics of the camera, it is simple to implement features such as a lens distortion. Separate emit shaders will be created for different lens characteristics (figure 40).

The other case is the controlling of photon tracing. In order to create a photon map, photon tracing must be done. This shoots photons of appropriate number and intensity from light sources in the scene, tracing how these photons spread in the scene. Emit shaders control this basic framework. However, the photon emit function implemented in each of the light sources is in charge of taking into consideration the emission characteristics of each light source and shooting the photons appropriately. This function exists inside the light shader. Therefore, emit shaders in the photon tracing only control the higher levels of how multiple light sources in the scene shoot the photons, and are not directly involved in the shooting of the photons. Currently, one kind of emit shader for photon shooting suits all our needs, and there is no need to develop new emit shaders for photon tracing.

## 2. Surface shader

Surface shaders describe how shading computation should be done on an object surface. When a ray hits some object surface as a result of space traversal, a new surface shader is called at that point. A surface shader only describes the computation that will be done on the surface. If necessary, it will shoot shadow rays, reflection rays, refraction/transparent rays, etc. SPOT already implements the necessary functions for shooting these rays, so new shaders simply call these functions.

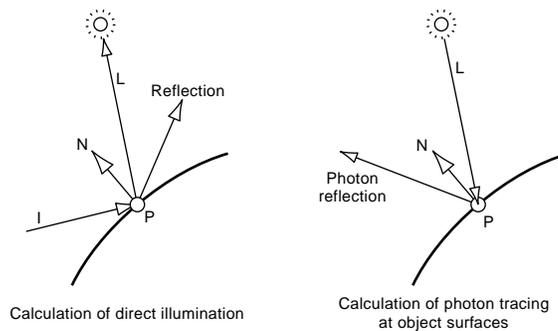


Figure 41: Two roles of surface shader

When creating a surface shader that takes global illumination into account, final gathering and photon map searching must be done on the object surface. It is also simply a matter of calling the functions that SPOT already implements.

Whether global illumination is considered in the shading computation of object surface depends on the kind of surface shader assigned to that object.

This means that in order to correctly perform global illumination on the entire scene, every object in the scene must be assigned a surface shader that takes global illumination into account.

Surface shaders also define a function that decides how the shader will behave at the time of photon tracing. This function describes how a photon will behave on the object surface when photon tracing is done. New photons may have to be shot from the object surface inside this function. This is also performed by calling functions already implemented in SPOT.

Basically, surface shaders independently define the following two computations:

- (a) Direct illumination on the object surface
- (b) Photon tracing computation on the object surface

These are closely related, and correct computation cannot be done if there are any inconsistencies here. However, the two functions can be defined independently in order to allow more freedom in shading description even at the cost of unnatural computation (figure 41).

### 3. Light shader

Light shaders define the computation for light sources. Basically they define the following:

- (a) Light source computation for direct illumination
- (b) Behavior of photon emission for photon tracing

These two elements may be described independently, but the simulation of light will not be physically correct unless its attributes are consistent. For the same reason as in the surface shader, we allow them to be defined independently for more freedom in lighting description (figure 42).

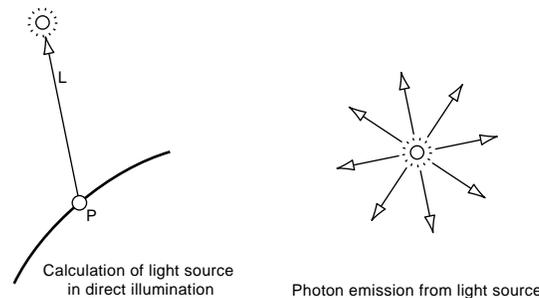


Figure 42: Two roles of light shader

When tracing photons, the emit shader calls the second mechanism of the light shader. Photon emission from this light source is controlled according to the definition of the emission characteristics, light intensity, color, etc.

In the direct illumination computation, the surface shader calls the first mechanism. The color on the surface is computed based on the relationship between the surface shader and the light.

Light shaders also take care of the shadow computation. The light shaders are called from the object surface, and first shoot a shadow ray onto itself to see if it is a shadow. If not,

the light color computation is executed. Shader writers do not have to think about these details when writing new light shaders because this mechanism is implemented in the shader template described below.

Light shaders can describe many different variations of light. Currently, area lights are described in addition to the common point, directional, and spot lights.

Sky lights are also implemented as a light shader denoting a special kind of light source. This light source computes the light coming from the sky, and it is used when shooting photons that take into account light from the sky. This is necessary when the image takes such light into account.

#### 4. Background shader

This shader is called whenever the ray performs space traversal but fails to hit anything.

Currently background shaders are used only when computing direct illumination. Photons that do not hit anything at photon tracing time disappear into infinity.

#### 5. Volume shader

This is called when a ray performs space traversal and enters a volume container. Kilauea uses ray marching for volume computation, and volume shaders define what kind of computation will be performed for each step of ray marching.

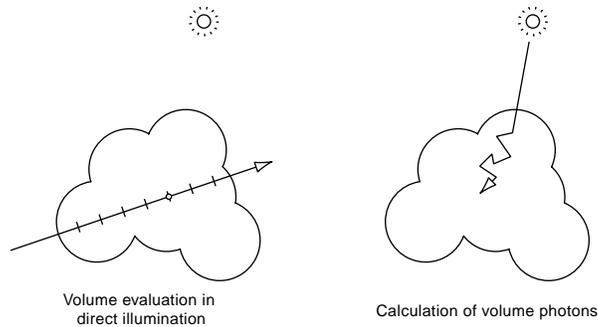


Figure 43: Two roles of volume shader

Volume shaders, just like the surface shaders, have two roles in computation:

- (a) Direct illumination in the volume
- (b) Photon tracing in the volume

The second computation defines the behavior of photons for volume photon map creation.

### 6.6.5 Relationship of shaders

Figure 44 illustrates how the data flows between each shader in a general shading computation.

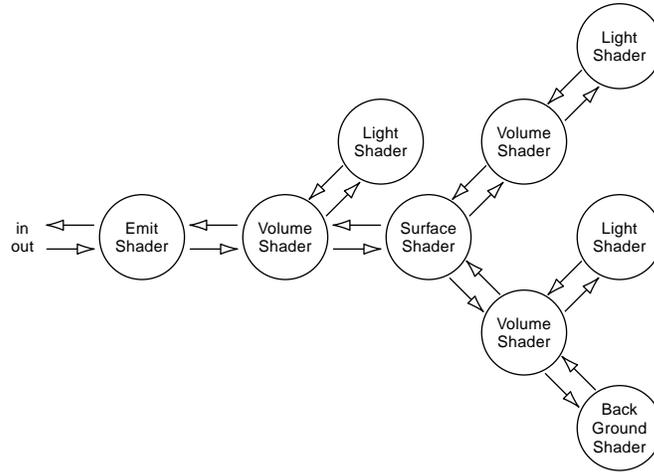


Figure 44: Relationship of shaders

### 6.6.6 Shader implementation techniques

Implementation-level techniques used for Kilauea shader development are described here.

#### 1. Shader template

If the shading computation on an object surface can be processed by a local illumination model, then each shader simply performs independent computation. Each shader simply determines the color somehow. When writing a shader that takes global illumination into consideration, however, shaders must perform the various global illumination computations.

In the case of Kilauea, final gathering and photon lookup is necessary for global illumination. Final gathering must be executed in the surface shader of the object surface taking into account global illumination. Basically, the same algorithm can be applied to all surface shaders that take global illumination into account.

All surface shaders are implemented using SPOT, so final gathering and photon lookup are treated as computation on the SPOTs. Therefore, every surface shader which takes global illumination into account must have final gathering and photon lookup implemented.

The shader template is prepared to simplify the task of writing shaders supporting global illumination. The shader template basically bundles everything that does not depend on the surface characteristics, and includes controls for final gather and photon maps. By using this shader template to write new surface shaders, it is possible to easily write a global illumination shader simply by writing the parts which depend on the object surface characteristics.

It is obviously possible to write a global illumination surface shader without using this shader template. In this case, the new shader will simply have to control everything itself.

Implementation-wise, this shader template is also composed of the SPOT network. One shader template exists for each shader type.

#### 2. Texture sampling

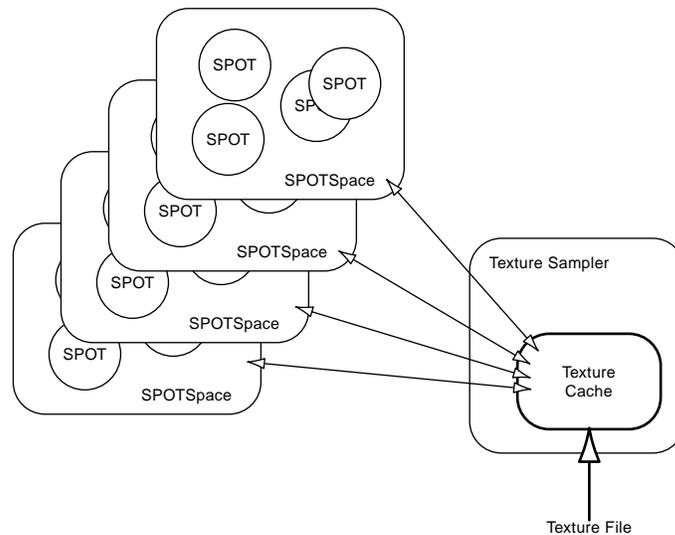


Figure 45: Texture sampling function

Texture data sampling can be done from every SPOT. The actual sampling operation is managed separately within the SPOTEngine, and things like the partial caching of texture data when multiple sampling operations are done on the same texture, is automated. It is also possible to appropriately control the memory required for sampling multiple separate textures. Basically, it takes the approach of caching part of an image in memory, so users are able to have complete control over how much memory will be used for the texture cache.

Because this functionality is implemented in the base SPOT class, all SPOT classes which inherit the base SPOT class can perform texture sampling.

## 6.7 Photon map method

Kilauea uses the photon map method (including final gathering) for global illumination computation. Kilauea's implementation of the photon map method is explained here.

The photon map method generates global illumination images in two passes. The first pass is called photon tracing, where photons are emitted from light sources, traced through the scene, and stored in photon maps.

The second pass performs the actual rendering, where the photon maps are looked up to compute the global illumination component. Kilauea's implementation simply adopts the general photon map method. The only difference is that our implementation takes advantage of parallel processing, performing parallel photon tracing and parallel photon lookups.

### 6.7.1 Parallel photon tracing

Parallelizing photon tracing is extremely simple in an environment where ray tracing is already parallelized. Photon tracing is purely ray tracing from an algorithmic standpoint. In our case ray tracing is already parallelized, so there is no problem in the parallelizing of the photon tracing.

The strategy for saving photon data in multiple photon maps needs careful consideration. When there are multiple PTasks in different ranks, each PTask stores photon map data. There are two cases to be considered here depending on how photons are stored in each photon map:

1. The photon map fits in the memory of one machine

## 2. The photon map does not fit in the memory of one machine

The first case is simple. The same photon data are stored in every photon map and every photon map ends up with the exact same data.

The second case is similar to the idea of parallel ray tracing. The distribution strategy of photons across multiple machines follows the method used for distribution of the accel grids.

If two machines share the memory to store the photon map, each machine will contain approximately 50% of the total photons. The distribution and density of the photons in two machines should be almost equivalent. If there are more than two machines, the photons are distributed uniformly among photon maps, with a similar distribution and density (figure 46).

The processing of each photon in photon tracing is completely independent from other photons, so the parallel processing performance of this portion has the same characteristics as Kilauea's parallel ray tracing performance. However, there are some different characteristics from the final ray tracing. Ray tracing is called from within a shading computation, and it resolves when color values are returned at the end. In the case of photon tracing, there is no need to return a value once the data is stored in the photon map. Because this is essentially a one-way operation, it has an even higher degree of parallelism, very suitable for multi-CPU or PC cluster environment.

To speed up the rendering stage, the irradiance estimate at the location of each photon is precalculated after the photon tracing<sup>8</sup>. This pre-computation phase can be processed in parallel for each distributed photon map.

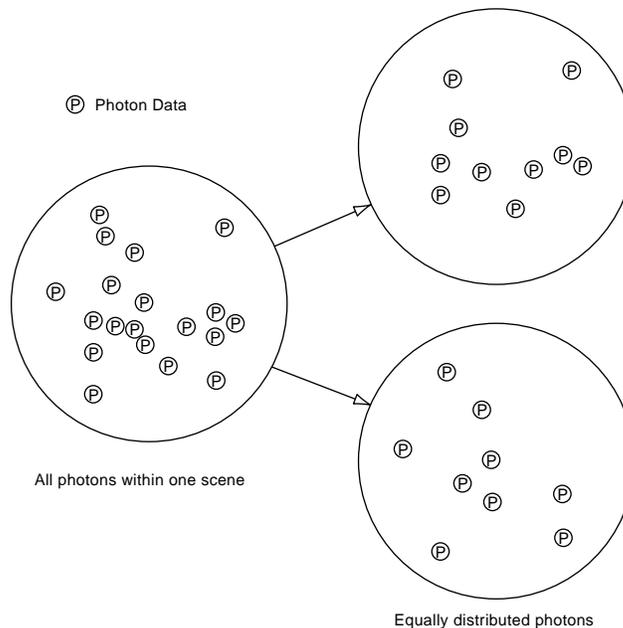


Figure 46: Distribution of photon data

### 6.7.2 Parallel photon lookup

Kilauea looks up photons during shading computation to take global illumination components into account. This lookup operation searches the already constructed photon maps and retrieves the necessary values. Photon lookup is handled in a task outside of SPOTEngine, just like ray tracing

<sup>8</sup>Per H. Christensen. "Faster Photon Map Global Illumination". Journal of Graphics Tools, volume 4, number 3, pages 1-10. ACM, 1999. [1]

is handled outside of SPOTEngine. Photon lookup requests are sent to a photon map inside the PTask, and the lookups are performed there. The results are returned to the SPOTSpace that requested the lookups, and the shading computation continues from there.

If all the photons in the scene can be handled in a single photon map, then it is just a matter of doing the lookup in that photon map. If the photon map for the entire scene is distributed across several photon maps, the same kind of strategy as parallel ray tracing is used (figure 47).

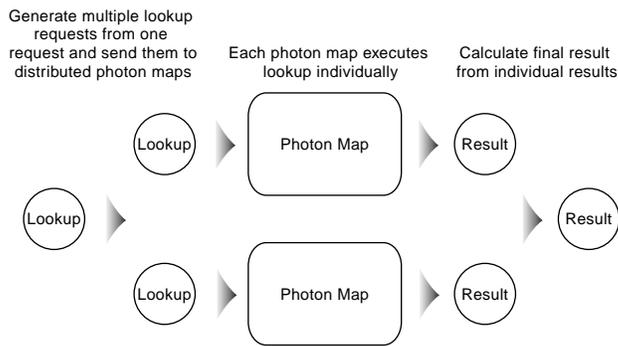


Figure 47: Parallel photon lookup

First, an identical photon lookup request is created and sent to each of the distributed photon maps. Each request executes the lookup operation for each photon map, and returns the precomputed irradiance. Based on the results, the final photon lookup result is computed.

Each photon lookup is handled in the PTask, and computing the final result from the returned results is done in LTask.

Photon lookups in multiple photon maps basically have the exact same properties as ray tracing in multiple accel grids. Implementation and multi-threading can be done using the exact same ideas. In the case of Kilauea, parallel ray tracing implementation and parallel photon lookup implementation have a one-to-one correspondence.

There are fundamental differences as well. In the case of ray tracing, only the closest intersection point is used and others are thrown out. In the case of photon lookups, results from every photon map are used in the computation of the final photon lookup result.

Please refer to the SIGGRAPH 2001 course notes on photon maps (course 38), particularly the note “A few photon map tricks,” for further details on parallel photon maps.

## 6.8 Things to note in shading computation

Some issues which came up while developing the Kilauea system is discussed here.

### 6.8.1 Auto-cruise control of queues

All shading computation is done using the SPOTEngine. SPOTEngine can be abstracted as an extremely deep pipeline. The queue data structure which handles this pipeline could exhaust all the available memory at times. The problem manifests itself when the execution speed of two stages in the pipeline become excessively unbalanced for some reason. If a later pipeline stage becomes too slow, the data stored in the queue will accumulate.

Most of these problems can be predicted at the initial design phase of the algorithm, and can be circumvented by evenly distributing the computational load among multiple threads. There is still a possibility that memory will be wasted due to some external cause, however. In extreme cases, CPU power is used to allocate more memory for the queue data, which further aggravates the situation. If the socket connection is hindered for some reason even for a short amount of time,

data that was supposed to leave for outside ranks quickly can no longer do so, eventually causing the kinds of problems mentioned above.

Another cause of dramatic changing in the data flow through the pipeline is that the computation of each SPOT is completely dependent on the shader. SPOTEngine constantly faces the above danger because the computational load of some SPOT is unknown until it is actually executed.

The current implementation of SPOTEngine is able to resist the worst situation of explosive use of memory by the pipeline queue for almost every SPOT. SPOTEngine also has a built-in “auto-cruise” control of queues to avoid such situation.

When some method is inserting its computed result into the queue to be handled by the next stage in the pipeline, the insertion operation is blocked if the queue is longer than a certain length. The next stage in the pipeline will process items from the queue while the previous stage is blocked, eventually allowing the value to be added to the queue (figure 48).

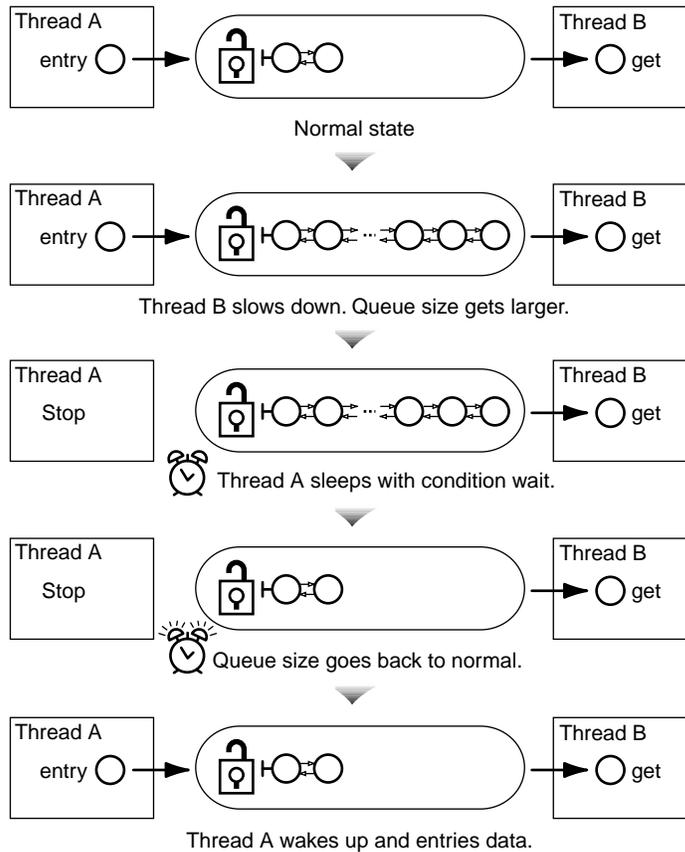


Figure 48: Auto-cruising queue data flow

Ideally, the program logic should be constructed in such a way that this kind of mechanism is unnecessary, but this is an effective technique for avoiding the worst possible situation.

### 6.8.2 Thread priority boost

One solution to the problem mentioned above is to change the execution priority of the threads. This is generally referred to as the priority boost of the thread. This sounds like an appealing idea, though experiments show that controlling thread scheduling with priority does not work as expected. First of all, the Pthread library that Kilauea currently uses does not allow changing

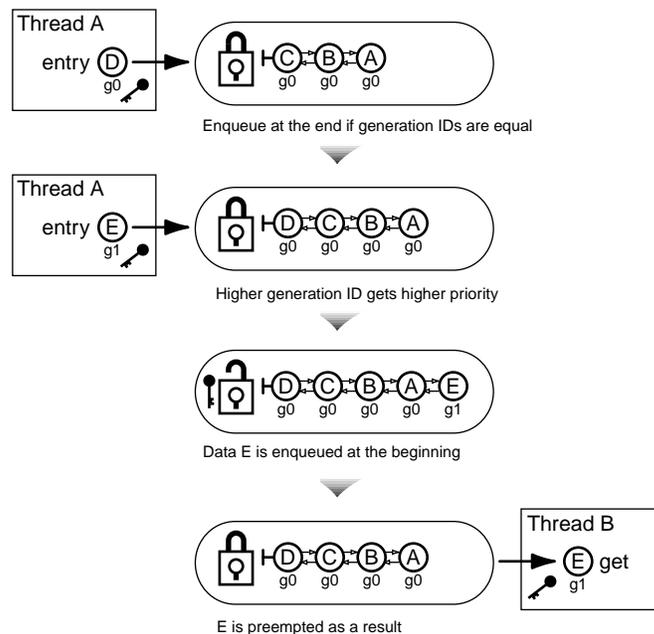


Figure 49: Passing rule

individual thread’s priority. In our experiments on an SGI, controlling the thread scheduling was very difficult even with full control over the individual thread’s priority. For these reasons, Kilauea does not use any thread priority boosting.

### 6.8.3 Passing rule inside the queue

In a similar problem as the one mentioned above, there are times when the memory in SPOTEngine is excessively used, and the SPOTEngine cannot continue its job because it is too busy trying to allocate memory. The engine easily encounters this situation if all shading requests are processed with equal priority.

Consider shooting a new ray from the object surface when shading. This ray hits another object, shoots another ray, and so on at multiple levels. Normally SPOTEngine handles shading requests as they enter the queue. All requests actually start from primary rays shot from the camera. SPOTEngine executes these requests as they go in the queue, and these requests generate new rays, which request new shading computation at their intersection points. If the new shading request is inserted to the end of the queue, and it cannot execute until the requests earlier in the queue have been completed, the SPOTEngine will have to store all in-progress SPOTs by using an extremely large amount of memory. In other words, once all primary ray computations are complete, the secondary ray computation begins, then the tertiary ray computation begins. This is often referred to as a breadth-first strategy. To get around this problem, we take the approach of adding generation IDs to all data that goes through the SPOTs. A computation with a larger generation ID unconditionally preempts — has a higher priority than — a computation with a lower generation ID, as illustrated in figure 49. This method, often referred to as a depth-first strategy, avoids the problem of an explosive increase of the queue length inside the SPOTEngine.

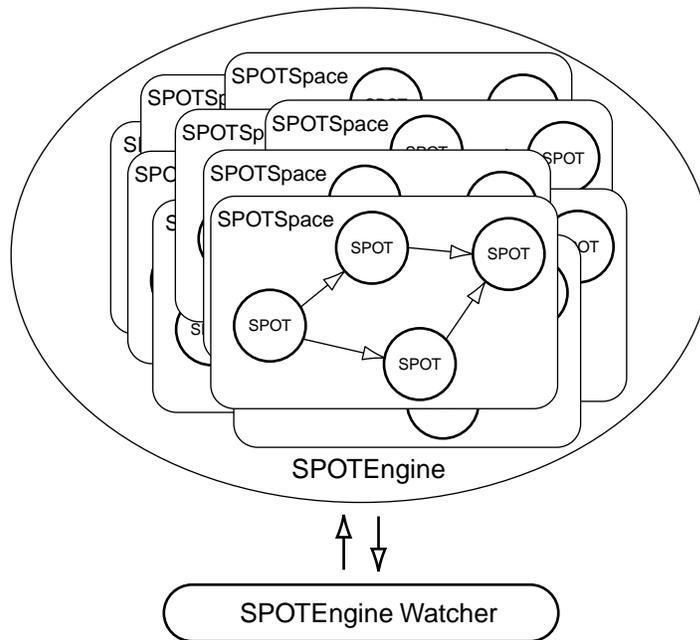


Figure 50: SPOTEngine watcher thread

#### 6.8.4 Latency and dynamic adaptive sampling

As mentioned previously, the shading computation inside Kilauea is a pipeline with a very large number of stages. Compared with a sequential implementation, this parallel shading engine suffers from a much longer latency in computing the shading results.

Dynamic adaptive sampling has trouble dealing with this latency in the shading engine. Dynamic adaptive sampling determines the next action based on the previous sampling results. In the initial stage of rendering, there are enough sampling requests to keep the pipeline busy all the time. However, the CPU gradually becomes more and more idle as the computation gets near the end. Currently the only solution is to shorten the time required to process one request. This problem is alleviated to some extent by improving the latency of SPOTEngine, but a fundamental solution has not been discovered yet.

#### 6.8.5 Profiling

Being able to learn how the shading computation is processed within the SPOTEngine and pinpointing the bottleneck is very vital in improving the system. However, in the case of our SPOTEngine which is implemented in a multi-threaded manner, we were unable to find effective tools for analyzing the performance. SPOTEngine performs self-analysis and self-diagnosis as it executes its shading computations. The self-diagnosis/analysis mechanism is a separate thread, collecting various statistical information from the internals of the SPOTEngine. This information is very helpful in determining whether the SPOTEngine is running soundly or some potential problem is disturbing the performance (figure 50). This mechanism can be started and stopped as necessary from the Tcl script.

## 6.9 Development in general

Kilauea is the result of cooperative effort of multiple developers. This section discusses our development environment and policy.

### 6.9.1 Debugging, debuggers, and development style

Debugging is a very difficult process. We have not found any debuggers or debugging methods that perfectly match our requirements. Our main debugger is gdb, and it is only a tool for debugging within one process, down to threads. When debugging something related to message passing between multiple ranks, we cannot use any debugger whatsoever. Launching two Kilauea processes on two machines with gdb in most cases is not effective either.

In order to avoid such situations as much as possible, our development style first focuses on perfecting each module independently. Kilauea has several mechanisms for helping this. Kilauea is an extremely large system as a renderer, but all computation is divided into tasks, making individual testing of tasks simpler. In some cases, parts of rendering computation can be simulated inside Kilauea, or simple scene data can be created inside Kilauea. Each task has a general way of sending text commands through the master task in a unified way, thereby making implementation of debugging and testing methods simple. Multiple developers are currently involved in the project, and in most cases development effort is divided on a per-task basis. Individual programmers can perform tests individually and check code into the master source tree when it is relatively bug-free. This development style is working very well.

### 6.9.2 Message passing protocol

Kilauea adopted MPI's philosophy for message passing between multiple ranks, but someone must determine how and what data is transmitted. One idea is to standardize the details of binary data transmission protocol for specific tasks, documenting and implementing them after the programmers have agreed on it. This is not the approach we take because we want to avoid the overhead of having to communicate all the time in order to decide on transmission protocols that may change every day. We only define a class for transmission and its API methods, completely leaving actual transmission method and protocol to the programmer in charge of implementing that class. The programmer may change and/or improve the transmission protocol as often as he desires, but under the rule that the class API does not change. In this way, we have tried to allow each other's development to proceed smoothly by removing as many dependencies on each other as possible.

### 6.9.3 Inherent difficulty of debugging parallel computation

Debugging parallel programs is unquestionably much more difficult than debugging sequential programs. Bugs in parallel processing are often not reproducible. Most reproducible bugs are likely to be not bugs related to the parallel processing, but rather very simple errors in logic.

When development gets to a certain stage, most of the simple bugs are fixed, leaving only the complex and difficult bugs that are deeply rooted in the parallel processing.

Sometimes bugs may be difficult to reproduce if they only occur after ten hours of continuous rendering. Cores dumped by Kilauea often do not have any useful information, or are very difficult to analyze. Even worse, Kilauea may just hang up in a thread deadlock. Tracking down this sort of bug is a very tedious process.

Instability is a fatal flaw in applications such as a renderer. In an ordinary sequential renderer we should improve stability by fixing all the bugs as they appear, but in the case of our parallel renderer, this kind of policy may just put a hopeless load on the developers without improving the situation at all. Of course, bugs should be located and fixed, so we can't just leave them alone.

For Kilauea, we have decided to come up with a fault-tolerant mechanism for getting around fatal bugs. We ended up with this conclusion after facing the reality that we can not keep the code

100% bug-free, even though we should definitely fix bugs.

Consider the case where Kilauea is rendering multiple frames over many hours. Some potential bug is lurking in the system, and it is extremely difficult to fix because it is deeply rooted in the parallel processing logic. The worst symptom of this bug is that it stops the rendering. Imagine that one starts a rendering job at night, only to find out next morning that it has frozen in the middle of the first frame. Kilauea observes this phenomenon as some sample requests never returning the results for some reason. Kilauea has the mechanism to monitor this problem by setting timeout for every sampling computation. For samples that time out, Kilauea determines that fatal errors which lock up the computation occurred and attempts to reschedule those samples.

If after several attempts they keep timing out, then Kilauea finally gives up and shuts down all Kilauea processes. At this time, Kilauea returns a status code saying that it shut down due to some internal trouble.

The Kilauea daemons which monitor the Kilauea system analyze the status code, and understand that Kilauea shut itself down due to internal troubles. In this case, the Kilauea daemons reboot all Kilauea processes, and instruct them to restart the computation starting from the frame with the problem. The restarted Kilauea may be able to compute the frame correctly this time. The more the bug is based on some complex timing issues of parallel processing, the more likely rendering would succeed on the second try.

This fault-tolerant mechanism is rarely triggered now because Kilauea is becoming more stable day by day. However, this kind of mechanism should exist considering the difficulty of debugging a parallel program.

Apparently, the mechanism is no use against simple bugs which are easily reproduced. It is only effective for getting around bugs whose cause is difficult to find and fix.

## 7 Rendering Examples

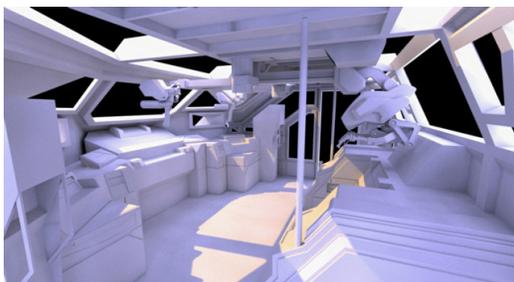
Details of the rendered images will be explained in the slides.



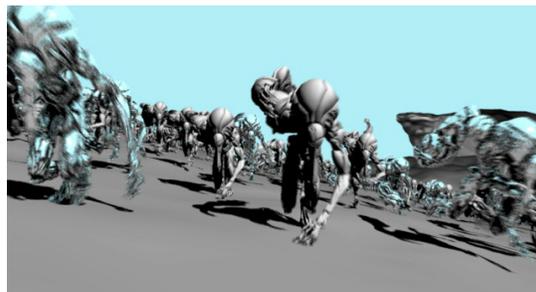
Cornel Box with Stanford Bunny 208,387 triangles, 1 area light



Running Phantom 55,336 triangles, skylight and 1 point light



Gondola 72,637 triangles, skylight and 1 point light



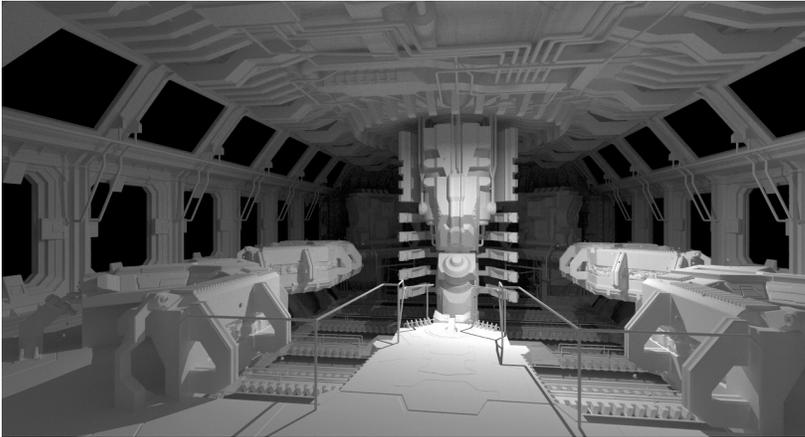
Phantom Crowd 25,340,322 triangles 1 point light



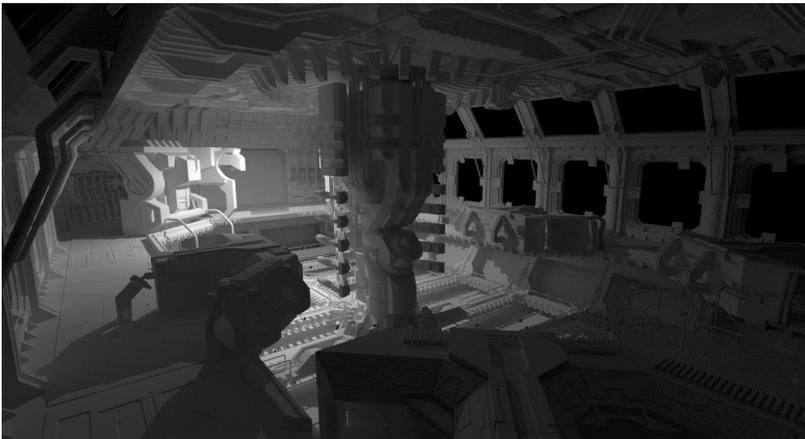
Japanese Tea Room 37,664 triangles, 2 area lights



MCB 2,156,176 triangles, 12 area lights



SPC  
899,720 triangles,  
1 point light



Jeep 1,999,302 triangles, skylight and 1 area spot light

## 8 Future Plans and Tasks

Here is a brief list of planned additions and enhancements to Kilauea in the near future.

### 8.1 Optimization

Kilauea has distinguishing features such as message passing and SPOTEngine that do not exist in ordinary sequential ray tracers. Thus we must admit that running on a single machine, Kilauea has a speed disadvantage against a simple ray tracer. One technical goal is to make this penalty as small as possible.

Optimization is being done from two perspectives. One is the optimization of SPOTEngine internals, and the other is the optimization of parallel ray tracing. These two both form the root of all computation in Kilauea, so they are constantly being improved to be more efficient.

As far as SPOTEngine goes, the first implementation was very experimental, leaving lots of room for improvement. We are currently implementing a new version based on the problems and characteristics discovered through experimentation.

We are also conducting some experiments on improving the parallel ray tracing algorithm, looking for an efficient implementation with more parallelism.

Another future topic is improving the efficiency of final gathering and photon mapping.

### 8.2 Stability

We are currently stabilizing Kilauea. Specifically, these include improvements to the Kilauea daemon and modifying the internal computation to safer logics.

### 8.3 User interface

Kilauea is completely controlled by Maya which serves as the front-end, and all GUI is created with Mel or Maya plug-ins. We are working to provide a more comfortable environment. Another topic for the future is to create an interactive environment for determining the material properties of objects, taking advantage of Kilauea's ability to change elements of the scene while rendering.

### 8.4 Other plans for the future

- Implement a shading compiler
- Dynamic reconfiguration of tasks at runtime
- Support for more primitives
- etc.

## 9 Acknowledgments

Our sincere thanks go to Kazuyuki Hashimoto, Kaveh Kardan, Shiro Kawai, and to the entire Square USA R&D team for their advice and ideas concerning the Kilauea Project.

We also thank Jack Liao for his great assistance for the creation of the test images, and Junichi Kimura for his assistance with the writing of this course note.

## References

- [1] Per H. Christensen. “Faster Photon Map Global Illumination”. *Journal of Graphics Tools*, volume 4, number 3, pages 1-10. ACM, 1999.
- [2] Julie Dorsey, Alan Edelman, Henrik Wann Jensen, Justin Legakis, Hans Kohling Pedersen. “Modeling and Rendering of Weathered Stone”. *Computer Graphics (SIGGRAPH '99 Proceedings)*, pages 225-234, August 1999, Los Angeles, California.
- [3] Henrik Wann Jensen, Per H. Christensen. “Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps”. *Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 311-320, July 1998, Orlando, Florida.
- [4] Bradford Nichols, Dick Buttlar, Jacqueline Proulx Farrell, Jackie Farrell. “Pthreads Programming: A POSIX Standard for Better Multiprocessing”. O'Reilly & Associates, 1996.
- [5] Erik Reinhard. “Scheduling and Data Management for Parallel Ray Tracing”. Ph.d. dissertation, University of Bristol Department of Computer Science, October, 1999.
- [6] SIGGRAPH 2000 Course Note. “A Practical Guide to Global Illumination Using Photon Maps”.
- [7] SIGGRAPH 2001 Course Note. “A Practical Guide to Global Illumination Using Photon Mapping”.



Toshiaki Kato  
toshi@squareusa.com



Hitoshi Nishimura  
hito@squareusa.com



Tadashi Endo  
tendo@squareusa.com



Tamotsu Maruyama  
tamotu@squareusa.com



Jun Saito  
jun@squareusa.com



Per H. Christensen  
per@pixar.com