

Visual Supercomputing

— Technologies, Applications and Challenges

K Brodrie^b, J Brooke^c, M Chen^{d†}, D Chisnall^d, A Fewings^a, C Hughes^a, N W John^a, M W Jones^d, M Riding^c and N Roard^d

^a School of Informatics, University of Wales **Bangor**, UK

^b School of Computing, University of **Leeds**, UK

^c Manchester Computing, University of **Manchester**, UK

^d Department of Computer Science, University of Wales **Swansea**, UK

Abstract

If we were to have a Grid infrastructure for visualization, what technologies would be needed to build such an infrastructure, what kind of applications would benefit from it, and what challenges are we facing in order to accomplish this goal? In this report, we make use of the term ‘visual supercomputing’ to encapsulate a subject domain concerning the infrastructural technology for visualization. We consider a broad range of scientific and technological advances in computer graphics and visualization, which are relevant to visual supercomputing. We identify the state of the art technologies that have prepared us for building such an infrastructure. We examine a collection of applications that would benefit enormously from such an infrastructure, and discuss their technical requirements. We propose a set of challenges that may guide our strategic efforts in the coming years.

Keywords: visualization, visual supercomputing, infrastructure, high performance computing, Internet-based visualization, virtual environments, large scale data visualization, computational steering, mission-critical visualization, mobile visualization, Grid computing, autonomic computing.

1. Introduction

Today there is a variety of computational resources available to visualization. While a huge number of users are contented with the visualization capabilities provided through modern desktop computers and powerful 3D graphics accelerators, many are still relying on high performance computing facilities to visualize very large data sets or to achieve real-time performance in rendering a complex visualization. In some areas, users have already demanded visualization capabilities to be provided through mobile computing systems, such as PDAs (Personal Digital Assistants), most of which are yet to benefit from powerful 3D graphics accelerators. As the size of visualization data (e.g., in visual data mining),

the complexity of visualization algorithms (e.g., with volumetric scene graphs), and demand for instant availability of visualization (e.g., for virtual environments) continues to grow, it is unlikely that visualization users can be served adequately, at least in the coming years, by an infrastructure largely based on desktop computers.

Inevitably, this leads to a series of questions that we must ask ourselves:

- What would be an adequate infrastructure?
- In what way do the computational requirements of visualization differ from other software technologies?
- Is it desirable or feasible to bring a range of technologies under one management (not necessarily under one roof)?
- If it were feasible to build such an infrastructure, what would be an appropriate virtual machine interface for the infrastructure?
- How should users’ experience be managed when they access visualization resources in the infrastructure?

[†] Contact Author: Min Chen, Department of Computer Science, University of Wales Swansea, Singleton Park, Swansea SA2 8PP, United Kingdom. Tel: +44 1792 295663. Fax: +44 1792 295708. Email: m.chen@swansea.ac.uk

In fact, the computer graphics and visualization community has been seeking answers for these questions for the past few decades. The community has invested a huge amount of effort in developing specialized graphics hardware, has always been among the first to deploy the latest technologies for high performance computing, and has accumulated large volumes of research outputs in parallel, distributed, and web-based techniques for visualization. Recently, the community has shown equally great enthusiasm to embrace the cluster, Grid and mobile technologies. However, in general, the community tended to address these questions mainly from the perspective of visualization technology. With the rapid expansion of the visualization user community, there is an urgent need to examine these questions from the perspective of end-users, for instance, surgeons, field workers, and fraud detectives.

The authors of this report are engaged in a collaborative project, e-Viz [e-V], to develop an advanced infrastructure for managing a variety of visualization tasks. In this report, we trace the historic route of deploying advanced computing technologies for visualization, and survey a broad collection of scientific and technological developments, including theories, algorithms, hardware, software and services, for visualization. We utilize the term *Visual Supercomputing* to encapsulate a subject domain concerning such an infrastructure for visualization, and outline the user requirements by considering a range of applications. We present an overview of the state of the art of technologies in hardware and software for visualization, and the impacts of the Internet, Grid and mobile technologies to visualization. We highlight those latest developments that are relevant, or potentially relevant, to visualization. We propose a set of technical challenges in realizing a visual supercomputing infrastructure that manages visualization tasks in complex networked computing environments, as well as manages users experience in accessing and interacting with visualization resources. We believe that *autonomic computing* can play an integral role in the evolutionary development of such an infrastructure.

Our survey comes at a timely moment in considering the relationship between visual supercomputing and *Grid computing*. There is now a growing body of experience in adapting applications to a Grid environment. What is emerging is a consensus that the original idea of a computational Grid that behaved like a utility Grid for computation is perhaps oversimplified. There may be several different structures for Grids depending on whether the resources aggregated in the Grid are to serve large-scale computation, large-scale data handling, complex data sources (e.g., bioinformatics databases) or perhaps to integrate business processes. In this, the visual supercomputing paradigm presents novel challenges to the Grid concept. A number of pioneering projects, described in this report, have been testing the implications of a Grid for various visualization applications and have raised many technical issues including real-time processing, syn-

chronicity of resource allocation and interactivity between clients and Grid services.

This report is organized as follows. In Section 2, we will give a more precise definition of the term *Visual Supercomputing* and outline its technical scope. In Section 3, we will review major scientific and technological developments by following the arrivals of different computing technologies, and identify the state of the art technologies that have prepared us for building an infrastructure for visual supercomputing. In Section 4, we will examine a collection of applications that would benefit enormously from such an infrastructure, and discuss their technical requirements. In Section 5, we will propose a set of challenges that may be used to guide our strategic efforts in the coming years. These will be followed by a summary of our conclusions in Section 6.

2. Visual Supercomputing

In this section, we first define the term ‘Visual Supercomputing’. We examine its relevance to the three semantic contexts of visualization. We then outline the technical scope of visual supercomputing from the perspectives of *applications*, *users* and *systems* respectively.

2.1. Definition

Definition. *Visual supercomputing is concerned with the infrastructural technology for supporting visual and interactive computing in general, and visualization in particular, in complex networked computing environments.*

In this report, we are focusing only on the subject domain of visualization, though most of the discussions can be extrapolated to other subject domains involving visual and interactive computing, such as computer-aided design, computer animation, and computer vision.

As an *infrastructural technology*, visual supercomputing encompasses a large collection of hardware technologies and software systems for supporting the computation and management of visualization tasks. It focuses on generic technologies for managing the specification, execution and delivery of visualization tasks. It addresses issues such the scheduling of visualization tasks, hardware and software configurations, parallel and distributed computation, data distribution, communications between different visualization tasks, and communications between visualization tasks and their couplings such as computation tasks or data collection tasks. In addition, it provides infrastructural support to users’ interaction with visualization systems, and manages users’ experience in accessing and interacting with visualization resources. Nevertheless, visual supercomputing does not concern a specific algorithm and technique for processing a specific type of data in order to generate visualization results.

We give an explicit emphasis on *complex networked computing environments*, as this report is not intended as only a survey of the technologies that have been developed so far in the context of visual supercomputing, but as a report on technologies that are in place as well as that are desirable for a future infrastructure. No doubt such an infrastructure must take *web computing*, *Grid computing* and *mobile computing* into account. Hence it has to provide comprehensive support to visualization tasks in complex networked computing environments.

The best way to capture our imagination of a visual supercomputing environment is to consider a global Grid infrastructure for visualization. The above-mentioned technical features of visual supercomputing have clearly set it apart from the traditional subject domains such as *hardware architectures for visualization*, *parallel and distributed computation for visualization*, *web-based visualization*, and *collaborative visualization*. While the advances in these traditional subject domains will have significant influence in shaping the infrastructure of visual supercomputing, we need not only integrate these technical advances together in an environment, but also bring in, and develop new, technologies for significantly improving the quality of services (QoS) of such an infrastructure and users' experience. This will become apparent in Sections 4 and 5.

2.2. Semantic Contexts

The gerund 'visualizing' refers to a process that extracts meaningful information from data, and constructs a visual representation of the information. In the field of visualization, this *process* is commonly considered in three different but interrelated semantic contexts as illustrated in Figure 1.

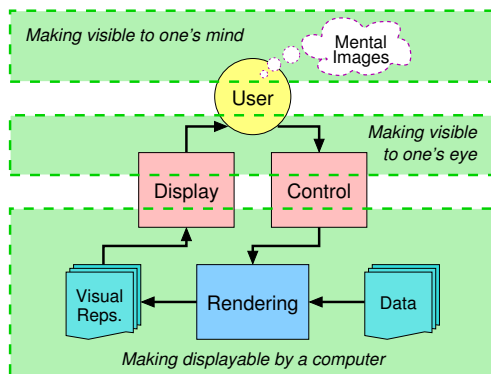


Figure 1: Three semantic contexts of visualization.

- *Making displayable by a computer.* This is concerned with the algorithmic and computational process of extracting information and rendering a visual representation of the information. In this semantic context, a visual supercomputing infrastructure should address issues such as allo-

cating and scheduling computational resources for visualization tasks, managing data distribution, and providing mechanisms for inter-process, and inter-task communications within an infrastructure.

- *Making visible to one's eyes.* This is concerned with the process of specifying meaningful information, designing appropriate visual representations, and conveying visual representations to viewers. In this semantic context, a visual supercomputing infrastructure should address issues related to the interaction between users and their visualization tasks, which can be conducted in a variety of forms, including interactive virtual environments, Internet-based collaborative environments, mobile visualization environments, and so on.
- *Making visible to one's mind.* This is concerned with users' thought process and cognitive experience of interpreting received information (not necessarily in a visual form) in one's mind and converting the information to knowledge in pictorial representations. In this semantic context, it is neither desirable nor perhaps feasible for a visual supercomputing infrastructure to manage the thought process of a user.

This section presents further detailed discussions on how visual computing relates to the above three semantic contexts, and provides rationalization for focusing infrastructural support on the processes of *making displayable by a computer* and *making visible to one's eye*, with the state of the art technologies in visual supercomputing. It also argues for the necessity for introducing gradually new capabilities in a visual supercomputing environment, to support the process of *making visible to one's mind*.

2.3. Application Perspective

The demands for visualization multiply in every direction, and there is an increasing number of new applications, resulting in new, and sometimes conflicting, requirements. For example:

- In some applications (e.g., bioinformatics), the size of datasets to be processed continues to grow, while in others (e.g., mobile visualization), a careful control of data size is absolutely necessary.
- In many applications (e.g., those involving 3D virtual environments), users still have plenty of appetite for photo-realistic visualization at an interactive speed, while in others (e.g., visual data mining), schematic visual representations and non-photo-realistically rendered images are often able to convey more information.
- In many applications (e.g., virtual endoscopy), interactive visualization can now be achieved with modern personal computers, hence small integrated systems provide a high degree of independence to users who operate in various practical situations. Meanwhile, other applications (e.g., those centralized around one or more data warehouses)

require a substantial amount of computation for visualization to be closely coupled with the source of data. Some applications, which have distributed data sources or dynamic data sources, demand a more complex computational model.

From the perspective of applications, an important requirement for a visual supercomputing infrastructure is *choice*, that is, it has to provide a large collection of platforms, methods, mechanisms and tools to serve different applications, as well as offer each individual application a diverse selection of means to accomplish a visualization task.

In Section 4, we will consider several major applications, which collectively characterize the main requirements for a visual supercomputing infrastructure.

2.4. User Perspective

Visualization users are no longer limited to scientists and engineers. At the same time, a visualization process often requires a high degree of domain knowledge about the application concerned. While the diversity of applications demands a visual supercomputing environment to provide a large collection of platforms, methods, mechanisms and tools, users require the service to be tailored to individual needs, and to be delivered in a seamless manner. Many users, especially those less technically oriented, would very much hope for a secretary-like visualization service, where they simply submit the data, give instructions and receive results. Although to get appropriate results may require a few feedback loops, many users certainly do not wish to get involved in choosing hardware, programming parallelism, organizing storage for input and output data, and so on. Further more, like a secretary, perhaps a visual supercomputing infrastructure should accumulate knowledge about various entities in the environments, profiling hardware capabilities, software usage, users' preference, etc. and gradually improving its quality of services to individual users.

Recent developments in business computing, such as electronic customer relationship management (e-CRM) [PL03, LaR00], has shown that it is possible to provide users with better quality of services with appropriate technologies that are capable of collecting and processing users' experience. The emergence of *autonomic computing* [KC03] is gathering further momentum in developing self-managed services in a complex infrastructure (see also Sections 3.6.2 and 5). Therefore a visual supercomputing infrastructure should have the responsibility for managing:

- visualization resources,
- visualization processes,
- source data and resultant data,
- users' interaction and communication,
- users' experience in accomplishing a visualization tasks.

2.5. System Perspective

From the system perspective, a visualization task is a kind of computation task, which exhibits a specific class of characteristics. The infrastructure of visual supercomputing is built upon a range of underlying technologies, including computer hardware, operating systems, programming languages, data warehouses, communications, world-wide-web, Grid computing, knowledge-based systems, and standardization. It is neither sensible nor feasible for the visualization community to attempt to provide solutions in all these aspects. However, it is necessary for the construction of such an infrastructure to bring in the latest advances in other fields of computing and communications, and moreover, to influence the developments in these fields.

In the following section, we thereby examine in detail the major advances and the state of the art in the relevant fields.

3. Technologies of Visual Supercomputing

The technological infrastructure for visualization has heavily depended on high performance computing environments until recently. In this section, we examine how the advances in computing and communication technologies have shaped, and reshaped, the foundation of visual supercomputing. Obviously, it is not possible to provide a comprehensive coverage for the large number of visualization works that have impacted upon the development of visual supercomputing. We hence focus on the contributions, in connection with each major technological advance, which are particularly relevant to the state of the art of visual supercomputing. For further historic details, readers are encouraged to refer to several excellent surveys [Whi96, Han96, BS01, Vit01, EE02, KMS02, BWD*03] and some major publications [DEH89, The89, Gre91] [Whi92, SSC02, Kou03].

3.1. The Era of Supercomputers

Elwald and Mass's vector graphics library for Cray-1 [EM78] represents the earliest efforts for providing visualization capability to support scientific computation on supercomputers. Since then, there has been a huge volume of publications devoted to parallel architectures and parallel algorithms for computer graphics and visualization. While most of these architectures are no longer in action, and many of these algorithms have difficulties to benefit from modern hardware, the research in the era of supercomputers has provided us with a collection of abiding concepts which can still be entrusted to serve modern visual supercomputing environments.[†]

[†] Note that some of the works described in this section were developed much later than the actual 'era of supercomputers'. As they represent some fundamental concepts and methodologies, we have conveniently placed in this section.

3.1.1. Models of Parallel Computation

Since the creation of the very first computer, there have been ever-increasing demands for processing power. Although Moore's Law [Moo65], which suggests that processor power doubles every 18 months, has been satisfied for the last 39 years, today's seemingly powerful desktop computers still cannot meet the requirements of many scientists and engineers who seek to model, compute and visualize even larger and more complex problems. Hence, there has always been, and will continue to be, a need for parallel computation. There are three paradigms for parallel computation:

- *Functional parallelism* splits up the process of computation by dividing an algorithm into separately functional sections and distributing these sections onto different processors organized along a logical pipeline. Data is passed from one processor to another to be computed. The parallelism is achieved when different parts of data are processed concurrently by different functional sections on different processors. In many special purpose graphics hardware systems, including commodity graphics cards, a graphics rendering pipeline is partially realized using functional parallelism.
- *Data parallelism* splits up the process of computation by dividing the data amongst the processors, all of which perform more or less the same algorithmic function. The parallelism is achieved when multiple streams of the data are computed in parallel. Some graphics hardware, such as the SGI InfiniteReality, makes use of data parallelism at individual stages of a graphics pipeline. A large collection of parallel visualization algorithms have been designed based on data parallelism.
- *Farm parallelism* splits up the process of computation into 'tasks', each of which is essentially a portion of data coupled with a functional operation to be performed. The tasks are kept in a queue, and are distributed to a 'worker' processor whenever one becomes available (i.e., idle). Many modern parallel visualization algorithms (e.g., [PMS*99]) have employed farm parallelism to optimize processor utilization.

In 1972, Flynn's taxonomy [Fly72] redefined parallel architectures and, whilst it may be a little outdated now, it is still generally appropriate and widely used. Flynn suggested four categories of parallel machines, namely SISD (*Single Instruction stream, Single Data stream*), SIMD (*Single Instruction stream, Multiple Data stream*), MISD (*Multiple Instruction stream, Single Data stream*), and MIMD (*Multiple Instruction stream, Multiple Data stream*).

In 1978, Fortune proposed the PRAM (*Parallel Random Access Machine*) model [FW78], which is an idealized parallel machine of p processors sharing an unbounded global memory and a common clock. PRAM architectures are essentially synchronous shared-memory MIMD systems, which are further categorized into four subclasses according to whether a memory location can be read or written

concurrently. By not considering synchronization problems and communication issues, the model focuses on the actual parallelization of a problem. As interconnection of the processors and memory is a fundamental factor in classical parallel architecture, other models of parallel computations were proposed, including UMA (*Uniform Memory Access*) and NUMA (*Non-Uniform Memory Access*) [Ski88, Buy99]. UMA systems are better known as SMP (*Symmetric Multi-Processor*) systems, where all processors are able to access all shared memory in the same, consistent time period. NUMA systems have differing access times for processors depending on the locality of the memory being accessed. Hence, NUMA systems can be larger and more distributed [ZQ91].

There are two principal memory structures that can operate in both UMA and NUMA systems. Firstly, in a *distributed memory* system, each processor has private access to its own fast, local memory, but must use some form of message passing over the interconnection to access the memory of another processor. A typical example is the Cray T3D. Distributed memory systems are generally regarded as difficult to program and debug, but they can scale to many thousands of processors [Tea02]. This is in contrast to *shared memory* systems, where all processors can access all memory directly via a shared bus (normally in a UMA system) or a complex switched interconnection network (normally in a NUMA system). Both of these require synchronization functions in order to safely handle contention for shared data. In hardware specifically designed for shared memory purposes, extra cache memories are often present along with a Cache Coherency Protocol to ensure consistency between local cache and global shared memory [TOIA95]. Volume visualization often relies on memory systems supplying conflict-free simultaneous access to multiple voxel values in a volume dataset [RPSC99].

Another major consideration in parallel computation is *granularity*, which is often used to indicate, intuitively, the size of parallel tasks in relation to the whole computation requirement. Granularity of a parallel architecture is defined as the ratio of the number of processors to the computation capacity of each processor. Granularity of a parallel algorithm is measured as the ratio of the time required for a basic communication operation to that for a basic computation. Different applications suit fine- or coarse-grain parallelization. Finer granularity brings greater potential for parallelism but increases the overhead of synchronization and communication. Over the last two decades, researchers in graphics and visualization have developed a large collection of parallel visualization hardware and algorithms of a wide range of granularity.

3.1.2. Models of Inter-process Communications

Since the first multitasking systems, it has been necessary to provide means for concurrent processes to communicate. In

parallel and distributed systems, inter-process communication introduces delays, which may affect the efficiency of a parallel algorithm significantly.

Shared memory architectures rely on low-latency ($< 1\text{ms}$) communication between processing units and memory via a *dynamic interconnection network* [Tan01]. Several parallel architectures, such as Cray Y-MP, utilized a *crossbar switching* network, which uses a grid of switching elements to connect p processors to q memory banks. A simple but less scalable alternative is a *bus-based* network, with which p processors connect to q memory banks by sharing a common data path. Multi-stage interconnection networks are a class of networks that offer more scalable performance than bus networks and more scalable costs than crossbar networks. A typical configuration is the *omega network* where p processors connect to p memory banks via p stages, and each stage is an interconnection pattern connecting p inputs to p outputs.

Distributed memory architectures typically involve *static interconnection networks*, which may be of a variety connection topologies [Fen81]. In such architectures, some kind of mechanism for *message passing* [Sil99] or remote procedure calls (RPC) [BN84] is required. The former enables data communication between remote processes, and the latter facilitates server-client communication by allowing a client to activate pre-defined remote procedures at a server and exchange data in a manner similar to conventional subroutine calls. An object-oriented approach to inter-process communication enables a process to send data as well as operations to remote processes, hence significantly improve the flexibility and scalability in dynamic management of parallel computation tasks. *Common Object Request Broker Architecture* (CORBA) [Bak97] provides such an inter-process communication in UNIX-like systems, while Microsoft Windows incorporates such features into DCOM as an operating system service. Some systems, such as Globe [vSTKS99] allow a single object to be distributed across a wide area network.

A number of modern parallel environments provide programmers with high level programming interfaces for managing inter-process communications. These include coordination-based middleware such as Jini [OW00], and document-based middleware such as Globus [FK97]. This enables application developers to focus on the contents of the communications, and many have adopted the XML standard for defining the syntax of the contents [BWD*02]. Although there is significant overhead in parsing transferred data when compared with binary encodings, the XML standard facilitates integration of different protocols and extension of existing protocols.

3.1.3. Performance Metrics for Parallel Systems

Many different metrics have used to measure the performance of parallel systems and algorithms. The primary objective of using p processors in parallel to solve a problem

of size n is the multiplication of the amount of processing power, commonly measured in terms of MIPS (*millions of instructions per second*) or FLOPS (*floating-point operations per second*). However, as previously outlined, it is not possible to parallelize all problems perfectly without introducing any additional costs.

One widely used performance metric is the *speedup* [JW00] which measures the ratio of *sequential run time* on a single processor to *parallel run time* taken on p processors. Speedup can never exceed the number of processors p , though in practice *superlinear speedup* (speedup $> p$) may sometimes be observed due to non-optimal *sequential run time*. Amdahl's Law [Amd67] describes another upper bound of speedup in relation to the *problem size*. Given a problem of size w that has a sequential fraction of size w_s and a parallel fraction of size $w_p = w - w_s$, the upper bound of speedup is w/w_s , regardless the number of processors.

It is also important to measure a parallel system with the *efficiency* metric, which is defined as the ratio of speedup to the number of processors; and the *cost* metric, which is the product of parallel run time and the number of processors used. One design goal for a parallel algorithm is to achieve a *cost-optimal* system, the *cost* of which is proportional to the execution time of the fastest-known sequential algorithm. The main obstacle to achieving a *cost optimal* parallel system is the *overhead* resulting from parallelization, which is usually caused by inter-process communication, extra computation (e.g., initialization, distributed data management), and idle waiting (e.g., due to load imbalance, task synchronization).

Increasing the number of processors reduces efficiency, while increasing the size of the computation increases total speedup hence efficiency. One of the most important metrics is *scalability* [Kri89], which measures the capability of a parallel system to maintain efficiency by increasing problem size and speedup in proportion to the number of processors.

The *isoefficiency function* of a parallel system describes the problem size w required, as a function of the *overhead* and *efficiency*, p , in order to maintain a constant level of efficiency. Isoefficiency has been found useful in characterizing the scalability of a variety of parallel algorithms. A small *isoefficiency function* indicates a highly scalable parallel system [GGK93].

Gustafson et al. [GMB88] studied the *time-constrained scalability* which is the core issue in some applications, such as weather forecasting, where it is necessary to fix the parallel run time, and to scale the problem size according to the number of available processors. They also examined the *memory-constrained scalability*, focusing on the largest problem that can fit the available memory in a parallel system.

3.1.4. Parallel Programming Paradigms

It is generally accepted that there are three primary programming paradigms for developing parallel applications, namely *message passing*, *shared-address-space* and *data parallel* paradigms. The first two are sometimes collectively referred to as the *control-parallel paradigm*[‡]

Message passing is a widely adopted programming paradigm. Although it is commonly associated with MIMD computers, it is universal enough to run on SIMD systems and uniprocessor systems as well as cluster systems and symmetric multi-processor systems. It requires the programmer to ‘manually’ specify subtasks to be executed in parallel, start and stop their execution, and coordinate their interaction and synchronization.

Message Passing Interface (MPI) [Sil99] is perhaps one of the most popular programming environments for developing parallel applications in C/C++ and Fortran. It is hardware independent, and provides a set of library interface standards for managing dynamic processes and message communications. It is supported by most of the hardware platforms available today.

The *Parallel Virtual Machine* (PVM), which was first developed at Oak Ridge National Laboratory (ORNL), is another implementation of the message passing paradigm. Using the notion of a virtual machine, PVM enables programmers to treat a set of heterogeneous computers as a single parallel computer. Although MPI is believed to be faster within a large multiple processor system, PVM still scores highly due to its fault tolerance and recovery [GKP96]. Other vendor independent libraries for the messaging passing paradigm include EXPRESS, P4 and PICTL.

The shared-address-space paradigm aims to provide programmers with a virtual shared memory machine, which can be built upon distributed as well as shared memory architectures. An programming environment for this paradigm normally include primitives for creating processes, allocating shared variables, managing mutual exclusion and facilitating synchronization. Managing mutual exclusion during concurrent memory accesses is critical to the correctness of parallel programs in this paradigm [GBD*94].

Linda [CG86] is a coordination language, in the form of extensions to C and Fortran respectively, for supporting shared-address-space programming. SR [AO93] is a language that supports both shared-address-space paradigm and messaging passing paradigm. X3H5 is an ANSI standard for shared-address-space programming in the context of single program and multiple data stream (SPMD).

[‡] The terms *control parallelism*, *functional parallelism* and *task parallelism* are often used in an interchangeable manner, while each places different emphasis on aspects of parallel computation.

The *data parallel* paradigm provides programmers a collection of virtual processors. Hence it facilitates a high level abstraction in developing parallel applications, hiding the architectural features of the underlying hardware. Data are distributed among virtual processors. It enables programmers to focus on data parallelism within a parallel algorithm. The parallelization of a computation task is usually realized by an appropriate compiler which must map virtual processors onto physical processors [KGGK93].

Many languages were developed for supporting the data parallel paradigm in the late 1980s and early 1990s, including the CM-2 family (i.e., C*, CM-Fortran and *Lisp by Thinking Machine Co.), MP-2, Dataparallel C, DINO and PC++ [KGGK93].

One important strand of the data parallel paradigm is *dataflow computation* [Sha85, Sha92], in which operations are executed in an order determined by the data interdependencies and the availability of resources. The execution can be activated by the availability of input data (i.e., data-driven) or by requirements for specific output data (i.e., demand driven). The concept of dataflow computing facilitates a functional specification of a computation task and the permitted freedom as well as constraints in its parallelization.

This concept has also played a more significant role in visualization (see also 3.2). Systems, such as OpenDX [AT95], AVS [UFK*89], IRIS Explorer [Fou95], SCIRun [SGPJ98] and DDV [MH99], are dataflow-based *modular visualization environments*. They provide a network of modules as the specification of a visualization task, which in principle can support dataflow parallelism [SG93]. As most networks normally define a coarse-grain dataflow, and most available modules cannot handle partial datasets, these environments offer only limited data parallelism under a centralised executive [ABM*01]. Both AVS and OpenDX can achieve control parallelism with remote modules. SCIRun provides threaded-task and data parallelism on shared-memory multiprocessors. DDV enables a pipelined-based, demand-driven execution that requires the minimum amount of input data to produce the results.

Stream-based computation, inspired by some parallel hardware architectures, represents a combination of simple control parallelism and simple data parallelism. Chromium [HHN*02] provides a collection of pluggable *stream processing units*, and allows streams of OpenGL commands (which contain mostly data) to be processed in parallel. Moreland and Thompson [MT03] recently described a new set of VTK parallel rendering components built on the top of Chromium for supporting ‘cluster to wall’ visualization.

3.1.5. Design Methods for Parallel Visualization

Parallel and distributed computation in visualization is broadly divided into two fundamental categories — *object space* and *image space* [Gre91]. ‘Object space parallel’

refers to the decomposition of a visualization task by dividing up input data into a collection of smaller components, each being processed by a computation node. Algorithms in this category are also known as *sort-last* [MMEF94], reflecting the need for sorting graphics primitives generated by different computation nodes at the image composition stage of a graphics pipeline. 'Image space parallel' refers to the decomposition of a visualization task into a collection of sub-tasks, each responsible for a small portion of pixels in the visualization image to be synthesized. Algorithms in this category are also known as *sort-first*, reflecting the need for organizing (or 'sorting') data according to the target sub-images prior to their entering into the graphics pipeline.

There is always a need in any parallel implementations to keep a balance between two, often conflicting, requirements, namely *data locality* and *load balance*. The former helps reduce the communication overhead, whilst the latter attempts to minimise the idle time of the processors involved.

Data partitioning is important for any visualization tasks to be computed on parallel and distributed architectures. It is particularly critical for distributed memory architectures, such as Beowulf clusters, where partitioned data components are distributed to different processing nodes. Data or spatial coherence is often harnessed by partitioning algorithms to ensure data locality while minimizing the amount of data residing on each node [MC94]. Further consideration includes image and frame coherence [GR95], and overlapping and exchange of boundary data [MLM*03]. In general, sophisticated partitioning methods are largely datatype-dependent, though they can sometimes also be architecture-dependent.

Data partitioning and distribution schemes may be classified according to division criteria (e.g., image-space [MC94], object-space [WH94], or hybrid methods [Lac96]), or organization of data replication, which may be in one the following three forms:

- *Complete Data Replication*, in which each node holds all data locally. This allows simple parallelization, classically image-space parallelization, through the same sequential algorithm on all nodes and minimizes communication overhead during processing. This technique is only effective for processing read-only data. It does not scale well as the cost of initial data distribution is a function of both the size of data and the number of nodes. The demand for large memory in each computation node is often difficult to meet.
- *Block Replication*, in which a dataset is typically partitioned into blocks or slices based on the 'physical' organization of the data. This meets the basic needs of object-space parallelization, and replicates a small proportion of an input dataset on each processor. For example, a *regular block decomposition method* may divide a volume dataset into equally-sized regular blocks. As equally-sized blocks do not ensure an equal amount of workload in each block, this sometimes leads to difficulties in load balance. An *ir-*

regular block decomposition method is often employed to produce blocks that contain similar workloads.

- *Structured or Hierarchical Partitioning*, in which one or more higher level structures are superimposed upon the raw dataset, facilitating data decomposition based on the 'logical' organization of the data. An *occupancy map* [MDKH01, KWH03] is one of the simplest forms of such structures, which employs a binary flag to indicate whether or not a block of data is of any interest to the rendering algorithm. A relatively more complex approach is the *Kd-tree Partitioning* [Ben75], which is used for partitioning k-dimensional space into sub-volumes along planes through the dataset. Another commonly used approach is *Octree subdivision* [DT81, CA86], which recursively divides the object-space (or an octant) into eight octants. Such a structure can be used to organize the data according to various attributes, including spatial occupancy and workload [VA88]. While most structured partitioning takes places in the object-space, many of these methods can also serve image-space parallelization as they can facilitate efficient view-dependent data fetch [LP02], and combined image and data coherence. Recently, scene graphs were used as a hierarchical structure for managing sort-first, distributed memory parallel visualization [BHPB03], and facilitating real-time virtual reality applications [NLSG03].

Load balancing is normally addressed by appropriate task assignment methods, which are typically classified by its run-time behaviour. *Static task assignment* (e.g., [WH94, MPK94]) pre-determines the workload of each processor according to the predicated workload of each sub-task and processing power of each computation node. Though it requires the pre-processing of task assignment, it demands less communication overhead and little cost in run-time monitoring and scheduling. It usually facilitates efficient data partition and distribution by taking data coherence into account in task assignment.

Dynamic task assignment (e.g., [Lac96]) maintains a pool of tasks, which are often of small and varying workloads. Whenever a processor is free, it is assigned a new task from the pool. This procedure repeats until the pool is empty. This method is particularly effective in heterogeneous environments, where the available computation capacity of each processor is somehow difficult to predict. It is often associated with image-space parallelization where the workload of each sub-task is difficult to predict.

Image composition, which transforms parallel streams into a useful output (usually a single image), is often a bottleneck in algorithms, especially sort last algorithms, in high performance visualization. Many classical implementations use the *direct send* method, in which each processor sends its rendered pixels directly to the processor responsible for image composition. However this simple method suffers from the problem of link contention with a large com-

munication overhead. Lee et al. [LRN96] suggested a parallel compositing algorithm to avoid link contention by routing messages along pre-defined grid paths in a mesh network. Ma et al. proposed to organize message paths in the form of a binary-tree (also by [WH94]), together with a binary swap algorithm for improving processor utilization. Recently, Stompel et al. [SML*03] presented a scheduled linear image compositing algorithm, as a highly optimized direct send method, offering better scaling on larger numbers of processors.

3.2. The Arrival of Graphics Workstations and Modular Visualization Environments

The arrival of graphics workstations in the late 1970s changed the face of visual computing. Up to that point, graphics was a speciality, provided in the form of a graphics terminal connected over a relatively slow communication line to a time-sharing processor. Suddenly the processor was co-located with the display, and so interaction became much more dynamic. Moreover this development coincided with the emergence of network-based windowing systems. This was significant to visualization users, who benefited from not only the WIMP-based user interface, but also from the interactive graphics capability that allowed visualization tasks to be carried out on the desktop.

However it took some time before visualization software emerged to support these new opportunities. In the late 1980s, the performance of workstations reached a point where interactive 3D visualization was feasible, and this performance leap was accompanied by new algorithmic developments such as Marching Cubes for iso-surfacing [LC87], and ray casting for volume rendering [Lev88]. A number of products started to appear, first AVS [UFK*89] and aPE [Dye90], and soon followed by IRIS Explorer [Fou95], Khoros [RASW91] and IBM Visualization Data Explorer [AT95], and more recently TGS Amira [TGS]. These are known as *modular visualization environments*, since applications are composed by wiring modules together in a dataflow network, using a visual programming paradigm. They are designed to suit end-users with limited programming knowledge and enable them to interrogate interactively a dataset via its visual representation.

Often these modular visualization environments were developed in the first instance as software tools to accompany and promote particular graphics workstations. Thus AVS was developed as a tool for use with Ardent workstations, and later Stardent; IRIS Explorer was developed to enhance the promotion of SGI workstations. The cost was typically very low, if not free. It is interesting that the software in most cases has lived rather longer than the hardware it was designed to support. For instance, responsibility for the development of AVS and IRIS Explorer was passed to software vendors in the 1990s, NAG Ltd. in the case of IRIS Explorer. Khoros was recently renamed as VisiQuest and

marketed by AccuSoft. IBM Visualization Data Explorer became OpenDX as IBM decided to make it an open source product. AVS, IRIS Explorer, Khoros and OpenDX remain vibrant products today.

In the late 1990s, relatively expensive graphics workstations were gradually replaced with modern personal computers equipped with commodity graphics cards. This has certainly created new demands for visualization tools from users of all types of occupations, for instance, security officers, and stock-brokers. It has also introduced a new dilemma as to the best way to provide users with visualization capabilities, and the role of modern personal computers equipped with powerful graphics hardware in the infrastructure of visual supercomputing. Undeniably, it is a formidable argument that a future visual supercomputing infrastructure should be based on all these personal computers, either loosely or tightly connected.

3.3. From Special Purpose Hardware to General Purpose Hardware

Many graphics and visualization tasks are computationally intensive, and continuing efforts have been made to offload the tasks performed by different parts of a graphics pipeline onto special purpose hardware. These efforts are exemplified by several most quoted developments, which include:

- *The video random-access memory (VRAM)* [PNG83], which provides an effective solution to improve the size and access of the frame-buffer required by almost every graphics pipeline.
- *Graphics processors*, such as Intel's i860 [GKB89], which led to an era when graphics processing units (GPUs) facilitated firstly window-based user interfaces to the desktop computers, followed by computer games, interactive 3D graphics, and interactive visualization toolkits.
- *Multi-processor graphics architectures*, such as Silicon Graphics' POWER IRIS [Ake89], which distributed the computational costs to a number of subsystems, each serving a set of special purpose operations, such as geometric manipulation, scan-conversion, and visibility determination.
- *Texture mapping hardware*, which has provided computer graphics and visualization with low cost pseudo-photorealism. In addition, such hardware has played a significant role in the development of visualization algorithms, and has been effectively deployed to accelerate a range of visualization tasks, including texture-based volume rendering [CCF95, WE98], flow visualization [RSHT99, TvW03, LBS03], splatting [SMM*97] and point-based rendering [PZvBG00]. Both 2D and 3D texture mapping techniques benefit from hardware support, but only high performance workstations currently offer 3D texture mapping hardware. For example, datasets of up to 512 Mbytes in size can be loaded into the dedi-

cated 3D texture memory available on a SGI InfiniteReality graphics pipe and then be visualized using hardware accelerated direct volume rendering. This is adequate to cope with the size of datasets from most medical scanners currently operating in hospitals around the world.

The latest generations of commodity graphics cards, such as the NVidia GeForce and ATI Radeon families, are allowing more and more applications to take advantage of graphics hardware. Demanding visualization techniques such as volume rendering and ray casting have already been successfully implemented, for example, [MHS99, RSEB*00, RGW*03, MLM*03, DPH*03]. With their generous memory capability and sophisticated numerical processing power, these cards have also been utilized for many circumstances other than graphics and visualization. Their affordability and extensive availability on almost all desktop computers, allows them to become more general purpose than ever before. There are limitations on what can be achieved today, however. For example, the size of the volume that can be manipulated is limited by the amount of dedicated graphics memory available on the card, and this can easily become a bottleneck when dealing with large datasets. Texture data must be fetched via the *accelerated graphics port* (AGP) from the main memory of the PC, and this prevents interactive performance from being achieved. Sophisticated partitioning of the data can be applied as a pre-processing stage to help overcome this limitation [CKS03]. However, it will be the replacement of the AGP with technology based on the new PCI-express standard that will eventually overcome this bandwidth bottleneck [WST03].

Among all of the increasingly 'general purpose' cards, one stands out as a piece of truly special purpose hardware; that is, the TeraRecon VolumePro, which delivers high-quality and real-time volume rendering capability [PHK*99]. Built upon the results of earlier research [PK96], the commercial VolumePro card currently available for PCs can deliver up to 30 frames per second for a 512^3 voxel dataset.

While there has been a surge of interest in transferring more computational costs from a visualization algorithm to a commodity graphics card, there has also been effort put into building high performance architectures that benefit from the collective power of an array of graphics cards. Several recent developments have demonstrated how graphics hardware of a PC cluster can accelerate a graphics and visualization task [WPLM01, MLM*03], implementing either image-space (sort-first) or object-space (sort-last) parallelism (see Section 3.1.5).

WireGL [HEB*01] was the first of a new breed of graphics software specifically designed to make use of such cluster systems, and it delivered general-purpose rendering capabilities through its support of sort-first rendering to tiled displays. The design of WireGL evolved into Chromium [HHN*02], which is a stream-oriented frame-



Figure 2: A large-scale, front projected, semi-immersive virtual environment.

work for processing streams of OpenGL commands on parallel architectures such as clusters. It can support sort-first, sort-last and hybrid parallelization strategies through the use of stream processing units. Integration between Chromium and visualization software such as VTK and OpenRM was recently reported [BHPB03, MT03]. The popularity of cluster computing has already led to a number of commercial products, including software products such as Modviz [Mod04] and hardware products such as the Sun Fire Visual Grid system [Sun04] and IBM DeepView [KKV*02].

The latest developments in graphics hardware have certainly suggested a modern approach to the architectural design for visual supercomputing, aiming at gaining the collective power from a large number of CPUs and GPUs simultaneously. No doubt, cluster computing is set to become a formidable technology in a visual supercomputing infrastructure.

3.4. The Drive for Virtual Reality

Immersive and semi-immersive virtual environments (Figure 2) represent a major technical drive in computer graphics and visualization, and have helped push a range of hardware and software technologies forward. Such a virtual environment enables users to be immersed inside a computer generated environment with a sense of spatial presence and often coupled with physical presence. For many visualization applications, virtual environments can provide users with realistic experiences in 'interrogating', 'navigating within', 'feeling' and 'manipulating' data via its visual representation.

3.4.1. Hardware Technologies

Although conventional computer displays and input devices can offer the most basic means for graphical interaction, they do not provide a sense of immersion, which is highly desirable in complex visualization tasks. Such tasks may require the user to have a better spatial awareness, better physical control in direct manipulation, better interaction with other users in the same virtual world, or better association with the real world.

Several techniques have been developed to enable users visually immersed in a virtual world with 3D stereoscopic views and volumetric views [BC03, SSC02, SC02]. These include:

- *Head mounted display* — It mounts a visual display in front of each eye. It is limited to one user at a time, and requires some form of cabled connection to the computer which could be cumbersome.
- *Projection-based display* — It provides stereoscopic views by projecting two different series of images, one for each eye, and allows several users to share the same visualization at a time [SMDW98]. Users typically gain stereoscopic experience using *shutter glasses* (i.e., in *active stereo*), or polarizing glasses (i.e., in *passive stereo*).
- *Autostereoscopic display* — It does not require the user to wear special glasses. One of such techniques is *volumetric display* [BS02], which allows users to view a 3D dataset directly. *Parallax techniques*, including *hologram*, *parallax barrier display*, *lenticular display* and *holographic stereogram* [Hal97], facilitate stereoscopic vision with motion parallax. A special layer (e.g., for example, a horizontal array of cylindrical lenses in *lenticular display*) is placed in front of interleaved images of a 3D object from different viewing angles. When the viewer moves, the a different image is picked up by the display, and the object is perceived to have rotated.

In addition to stereoscopic displays, one growing trend is building very large high resolution displays, involving, for instance, 63 million pixels [MT03]. Such a display can create an unusual sensation of *presence*, and involvement, enabling a team of users to interrogate a high fidelity model in its totality.

Techniques are available for users to interact with a virtual world with 3D input devices, some of which facilitate users' experience of physical immersion [Sto00]. These techniques include:

- *3D mouse* — It is a low cost hand-held device that provides a tracker sensor and a set of buttons. By changing the orientation of the mouse, the user can exert navigation control or apply direct manipulation in a virtual environment.
- *Interactive glove* — It is worn by the user and transducers sewn into the finger joints, which can be used to tell the computer the physical characteristics of the fingers when

they are bent. This allows the computer to identify when an object is being picked up, although the user would have no real sense of holding the object [SSC02].

- *Force feedback devices* — They are able to give the user a feel of physically interacting with virtual objects, and are often referred to as *active haptic devices*. One of the available techniques is the *Phantom-like haptic device* [Sto00], which involves a stylus fixed to a base, and can provide force resistance according to users' input actions and physical attributes defined with the object being manipulated. It can produce realistic feeling of the shape and textures of a solid object and the physical property of a deformable object.

3.4.2. Resourcing a Virtual Environment

The computational resources required to generate and interact with a virtual environment can be very different depending upon what is being simulated. A single desktop computer, or a cluster, with a £1,000 graphics card can be sufficient. However, many high performance applications are looking to the Grid and parallel computing to provide high quality graphics and resource-intensive data processing.

One of the most successful implementations of a virtual environment is the CAVE (*Cave Automatic Virtual Environment*) [CNSDe92]. It provides the illusion of immersion by projecting stereo images on the walls and floor of a room-sized cube. Simply by wearing lightweight stereo glasses, multiple users can enter and walk freely inside the CAVE. A head tracking system continuously adjusts the stereo projection to the current position of the main viewer. The technology of the CAVE and other large scale visualization environments has developed greatly over the last decade. Reconfigurable environments are providing even greater flexibility today. CAVE has been deployed in numerous visualization applications around the world. Such an immersive virtual environment requires the use of a high performance computer, for example, a SGI Onyx 3400 with 12 CPUs and 3 graphics pipes for CAVE. Special purpose software is also needed to manage the virtual environment, such as the open source DIVERSE [KASK02].

A related development to DIVERSE is the *Resource Aware Visualization Environment* (RAVE) [RAV04], which supports a collaborative visualization environment that scales from immersive platforms, to non-immersive but network-enabled platforms, including PCs, and PDAs. RAVE is 'resource-aware' so that the rendering platform and the visual representation will be determined dynamically by factors such as the client capabilities and the network bandwidth.

The Grid is becoming more and more important in the field of visualization, particularly when computational resources required for real time interaction in a virtual environment are not locally available. Also the popular component-based programming paradigm, which has been adopted

by many visualization systems such as VTK, AVS and OpenDX, can make use Grid resources. This allows different computation steps of a visualization pipeline to be distributed around the globe [SB03].

3.4.3. Collaborative Virtual Environments

Collaborative Virtual Environments allow multiple users to interact with each other and objects in a shared virtual environment. The users are usually also represented in the virtual environment by embodying themselves in virtual actors. Many of such environments are distributed systems, providing remote users with a sense of common presence.

Examples of collaborative virtual environments include DVIE [CH93], MASSIVE [GB95], VRML-extension [Bro95], COVEN [NT96], DEVRL [SUB*96], and VirtuOsi [BBG*94]. Many of them have focused on 3D virtual worlds, while others have attempted to address a wide range of issues related to collaborative virtual environments, such as avatar design (body-like figures), users' awareness, dynamic behaviours, system scalability, human factors and interest management. However, most of these environments were built independently on a project-based infrastructure by assembling different technologies together in an *ad hoc* manner.

Environments, such as those listed above, demand a noticeable amount of computational resources, complex distributed data management, dynamic resource allocation, as well as a variety of graphics support, it is only appropriate for the future development of such environments to be built upon a visual supercomputing infrastructure, which can facilitate computation, communication, graphics, data management, interaction management and interest management in a consistent and coherent manner.

3.4.4. Augmented Reality

Augmented Reality (AR) is an extension of the traditional virtual environment technology, which aims to immerse a user inside a virtual world completely. By contrast, AR allows the user to see the real world, whilst supplementing it with virtual objects which are superimposed within the real world [Azu97].

Most AR technologies have been based upon the use of some form of transparent display which is positioned between the real world and the eyes of the user [Pin01]. The most basic method is by overlaying computer graphics onto a 2D tabletop surface. In order to align the computer graphics with the physical reality, cameras are used to track the movements of the users vision and allow the graphics to be realigned [LaR98]. Rekimoto et al. [RS99] developed an InfoTable which combined a set of cameras for identifying real objects that were placed on the tabletop and an LCD projector for adding useful information to the known objects. The development of a collaborative AR environment has also been reported, in which several users can be

tracked and see the same virtual objects from different perspectives [SFH*00].

Potential AR applications include medical visualization, maintenance and repair, annotation, robot path planning, entertainment and aircraft navigation [Azu97]. Several AR techniques have now been shown to add value to the information available to doctors in the medical world. 3D medical datasets of a patient can be rendered in real time and overlaid onto the patient, allowing the doctor virtually to see inside the patient [GLPW*96]. This technique could also be used for medical training. Some examples of deploying this technology can be found in a recent survey by Vidal et al. [VBB*04].

One approach to facilitating interaction in an AR environment is to use *Tiles* as a reference between the virtual object and the real world [PTB*02]. Through a head mounted camera, the computer can identify the uniquely labelled *Tiles* and superimpose other graphics onto each *Tile*. ARToolKit is a software library for building AR applications [BKP01], which has been successfully in several example applications [The04]. The *Tile* approach was extended to become a *Personal Interaction Panel* (PIP), which provided a two handed 'pen and pad' interface for AR applications and allowed users to interact with virtual controls overlaid onto the panel [SFH*00].

The Eurographics Association sponsored an annual medical prize, acknowledging research utilizing computer graphics within the medical field [Uni04]. In 2003 the prize went to an AR application for Liver Surgery Planning, which utilized a PIP and a tracked pen to allow doctors to examine a patients liver. The PIP was used to allow doctors to specify cross sections of the liver that they wanted to examine.

Desktop PCs are continuing to increase in power and the latest range of GPUs are capable of meeting the requirements for many virtual environments. However as many immersive virtual environments consume a substantial amount of computational resources, particularly when handling large datasets that associated with some grand challenge problems, there remains a need for a visual supercomputing infrastructure.

3.5. The Ever-growing World Wide Web

The world wide web has made navigating 3D virtual worlds a readily-accessible technology, through programming environments such as VRML, X3D and Java3D. It provides a generic framework, under which it is possible to deliver visualization services to every corner of the globe. Interestingly, the web itself is becoming one of the focus points in information visualization as its complex infrastructure, its highly dynamic traffic, and its enormous amount of contents present serious challenges to the state of the art visualization technology.

The initial seminal work by Ang et al. [AMD94] demonstrated that the Web had a role in visualization. They associated visualization data with a MIME-type and this launched a helper application on the client side when the browser (i.e., Mosaic in 1994) downloaded the data. This data-driven approach has subsequently been rarely used, but it did show that the web can be an infrastructure for carrying out visualization, not just publishing previously created visualizations.

Two distinct approaches have emerged. In the *server-side approach*, the user submits a request from a web page, specifying the data to be visualized and the technique to be used. The request is processed on the server, and result returned as an image or a 3D VRML world. An early example, using CGI-scripting and IRIS Explorer, was developed by Wood et al. [WBW96]. Engel et al. [EWE99] exploit this for isosurface extraction.

In the alternative, *client-side approach*, Java applets can be used to provide simple visualizations on demand. An early example is by Michaels and Bailey [MB97]. This client-side approach has not gained wide popularity, perhaps because the security restrictions on Java applets prohibit the processing of local data. Thus many applets tend to be educational demonstrations rather than real services.

We can expect the server-side approach to be the forerunner of serious attempts at visualization web services. Prototypes are being developed using SOAP/XML (for Java-based services) and gSOAP (for services based on C/C++ and even Fortran). In terms of visual supercomputing, this may offer an attractive approach with its simplicity via a browser interface and power via remote server processing.

3.5.1. Collaborative Visualization

The Internet has also encouraged and facilitated collaborative visualization where geographically distributed users can work together as a team. Three distinct approaches have emerged: *display sharing*, where a single application runs, but the interface is shared; *data sharing*, where data is distributed to a group of users to visualize as they wish; and *full collaboration* in which the participants are able to program the way they collaborate.

Display sharing is supported by conferencing technology such as Microsoft NetMeeting, and the non-proprietary VNC [VNC]. This technology uses efficient compression on the frame buffer so that screen updates can be feasibly transmitted to a group of users. *Data sharing* has been exploited in collaborative environments such as CUMULVS [NER04] and in pV3 [pV304], where data from parallel computations is made available to multiple viewers. Another example of data sharing is provided by COVISE [COV] where geometry is made accessible to a group, each person in the group being able to render as they please. The most flexible approach is *full collaboration*, epitomized by the COVISA extension of IRIS Explorer developed by Wood et al. [WWB97]. In this

approach, each collaborator runs their own dataflow pipeline to create a visualization, but can export data and parameter settings to other users, and likewise import data and parameters. Although developed for IRIS Explorer, the idea can be exploited in any dataflow environment.

Figure 3 shows COVISA in action. It demonstrates the sort of application where collaboration can be useful: two doctors (Bone and Blood), each with their own speciality, can collaborate over the network. Bone looks at CT, and Blood at SPECT, but the two modalities can also be combined and this combined visualization viewed by both. Shared pointers allow discussion of significant features. The whole process is supported by video conferencing facilities: either desktop based using for example VRVS, or room-based using the AccessGrid [Acc03].




Figure 3: A collaborative visualization environment, where Dr. Bone is collaborating with Dr. Blood at a remote site to look at CT and SPECT data together.

There are many significant issues in the design of collaborative visualization systems: technical issues such as heterogeneity of visualization systems and of operating systems (collaboration between different visualization systems is hard because of lack of standardized data formats); and social issues such as privacy and floor control.

Both web-based and collaborative visualization have presented a visual supercomputing environment with the requirement for two essential services. As the web is likely to be the dominant information highway in the near future, it is inevitable that a visual supercomputing infrastructure will deliver a substantial amount of its services through the web. Web-based visualization and collaborative visualization will continue to challenge the underlying technologies of a visual supercomputing infrastructure.

3.6. The Beginning of Grid Computing and Autonomic Computing

3.6.1. Grid Computing

The Grid, as described by Ian Foster, one of the central characters in its development, is a distributed computing infrastructure for 'co-ordinated resource sharing' [FK98b]. The Grid is a virtual organization, which is composed of

autonomous organizations that maintain various local policies and software for controlling their resources. This distinguishes Grid computing from cluster computing, and introduces a great deal of complexity to the software engineering needed to provide the services and resources. Though it is generally agreed upon what actual services are required in a Grid infrastructure [SGPJ98], there are as of yet no agreed standards. A great deal of experimentation is being carried out to determine the best way to provide middleware services that 'glue' differing underlying systems together. The Grid middleware sits between users' applications and remote computing resources. It is generally accepted that the following key issues and services must be addressed within the Grid middleware:

- networking quality of service (QoS),
- resource co-scheduling,
- load balancing,
- message passing,
- file transfer mechanisms,
- data security, integrity and coherence.
- authentication.

Initially Grid middleware was built as a layer on top of services and protocols common in the Unix world, e.g., `ssh`, `ftp` and `LDAP`. The Globus Project [FK98a] developed a reference implementation of Grid protocols by providing all of the services and capabilities to construct a computational grid. This reference implementation has now become a *de facto* standard in the form of the *Globus Meta-computing Toolkit* [FK97], which contains a range of tools for resource allocation and process management, unicast and multicast communication services, authentication and related security services, distributed access to structure and state information, monitoring of health and status of system components, remote access to data via sequential and parallel interfaces.

The Grid infrastructure must be able to support a vast range of applications, allowing its services to be incorporated into the applications using a mix-and-match approach. An important aspect of the Globus Toolkit (GT) is that it separates local and global services. Local services are kept simple to allow deployment, and global services are built on top of local services. The *Metacomputing Directory Service* is provided as part of the toolkit to discover available resources and services. This allows resources to be added and removed dynamically and enables the Grid to recover if a failure was to occur.

Several other middleware developments took place in the same time frame as Globus but based on different principles. UNICORE (*The UNiform Interface to COmputing RESources*) [ES01] facilitated seamless access to computing resources and to integrate legacy applications for which source code was not available. This development was further extended into the European Grid infrastructure via the EUROGRID [EUR] and GRIP [GRI] projects. ICENI (*The Imperial College e-Science Networked Infrastructure*) provided

an environment for deploying software components over a federated pool of resources, with rich metadata structures for describing the characteristics of the components. Legion [CKKG99] federated computing resources as a virtual supercomputer. Condor [Con], which predates the idea of a Grid, is Grid-like in that it locates and federates resources to perform application tasks mainly in numerical computation. Codine, which was designed for scheduling tasks across a distributed infrastructure, was developed into Sun Grid Engine [Sun].

Several existing Grid infrastructures have been implemented and are already being used for research. These include the UK e-Science Grid, NASA's Information Power Grid (IPG) and the European Data Grid. In addition, there is the AccessGrid [Acc03], which is not an infrastructure for computation, but an IP-based conferencing infrastructure for supporting large scale collaborative activities.

In 2002, an alliance was formed between the Globus Project and IBM to promote an Open Grid Services Architecture (OGSA) [FCKT02]. This changed the delivery of Grid middleware to a *web services framework* and all Grid resources were virtualized as Grid services accessed via web service standard interfaces written in WSDL (*Web Services Description Language*). The emphasis has been shifted from interfacing with resource control mechanisms to hosting environments for the Grid Services defined by OGSA. Thus Apache Axis and Microsoft .NET became the tooling environments. Globus Toolkit 2 (GT2) was re-factored via tools such as Axis to create Globus Toolkit 3 (GT3), which was intended to preserve the functionality of GT2 in the new Grid Service environment.

In March 2003, a working group of the Global Grid Forum produced the first draft of the specification of Open Services Grid Infrastructure (OGSI), which involved moving key web services standards beyond what has been defined by the relevant standards groups. Because of the growing strain between the business and commerce users of web services and the scientific users of Grid computing, Globus and IBM recently announced that they were moving from OGSI to a newly proposed standard WS-RF (*Web Services Resource Framework*), which exposes resources (as in GT2) but now in the context of web services. The concept of inherent state in the former OGSI Grid Service has been abandoned, as web services are stateless entities and have difficulties in handling stateful management such as (e.g., a job queue). A new version of the Globus Toolkit (i.e., GT4) is to be released to replace GT3 later in 2004. In the meantime, WS-RF implementations are being developed based on .NET (by University of Virginia) and Perl (by University of Manchester) respectively.

The main problem with applying the Grid methodology, and any of the above implementations or proposed standards, to visual supercomputing is the need for interactivity with components running on the Grid. While users' interactive

intervention is an integral part of many visualization tasks, it does not always fit naturally with the idea of virtual 'visualization' resources. Some sophisticated middleware components are therefore required. One interesting attempt is the development of an *Interactive Access* plug-in to the UNICORE client [Sne01], which allows end-users to interact, via the UNICORE middleware, with simulation processes running at multiple locations.

The development of the Grid has laid critical foundations for a visual supercomputing infrastructure. It is highly possible that the development of visual supercomputing can be piggybacked on that of the Grid, and can learn a great deal from the evolution of the Grid technology. However it is also important to recognize that visualization is not just another computation process and hence a visual supercomputing infrastructure is not just a subset of the Grid infrastructure for launching computational tasks.

The great emphasis on a *web services framework* in Grid computing indicates that techniques developed for web-based visualization may have a more generic use in Grid computing. The server-based visualization services (see Section 3.5) allow visualization to be delivered in a coarse-grained visualization service, taking input data and sending back images. However, it poses a much more difficult challenge to deliver a sophisticated service, similar to modular visualization environments (see Section 3.2), that could access Grid resources using the web services paradigm.

3.6.2. Autonomic Computing

A Grid infrastructure, or more generally, a pervasive infrastructure, will be considerably complex, and the difficulties in managing such an infrastructure raise a serious question as to whether it is adequate for it to be managed by human administrators, and whether it requires a much more system-level automation than what is currently implemented. Researchers and developers in many fields, such as distributed systems, data communications, Internet technology, Grid computing, agent technology, database systems, expert systems and business management systems, are embracing the concept of *autonomic computing* in managing large and complex infrastructures and services.

Autonomic computing [KC03] refers to computing systems which possess the capability of self-knowing and self-management. Such a system may feature one or more of the following attributes:

- *Self-configuring* — The system can integrate new and existing components without low-level intervention from an administrator.
- *Self-optimizing* — The system can continually try alternate configurations to determine if the current one is optimal.
- *Self-healing* — The system can detect, and recover from, failure of components, hardware or software.
- *Self-protecting* — The system can detect attempts to compromise it, perhaps from hackers or viruses, and react accordingly.

A noticeable amount of research effort in autonomic computing has been placed on the self-management of system infrastructure and business services. Examples of this include self-configuration in patching management [DRD*04] and in Grid service composition [AP03], self-optimization in power management [KAH04], business objectives management [AGL*04], and resource management in dealing with network traffic spikes [NCFC04], and self-healing in online service management [CZL*04] and distributed software systems [Min03].

Efforts have also been made to broaden the scope of autonomic computing, addressing a wide range of related research issues, such as economic models [ERA*03], physiological models [LI03], interaction law [MU00], preference specification [WTKD04], ontology [Lin03, TT03], human-computer interaction [AHP*03], and so forth.

Though the development of generic software environments for autonomic applications is still in its infancy, several attempts were made, which include projects such as QADPZ [Con03], AUTONOMIA [DHX*03] and Almaden OptimalGrid [DLJ03].

QADPZ (*Quite Advanced Distributed Parallel System*) [Con03] provides an open source framework for managing heterogeneous distributed computation in a network of desktop computers using autonomic principles. In QADPZ, the system complexity is hidden in the middleware layer, facilitating self-knowledge, self-configuration, self-optimization and self-healing.

AUTONOMIA [DHX*03] is a prototype software development environment that provides application developers with tools for specifying and implementing autonomic requirements in network applications and services. It features an *application management editor* for requirements specification, a *mobile agent system* as a uniform execution interface to underlying hardware and operating systems, an *autonomic middleware service* for managing autonomic services, an *application delegated manager* as a broker between components and resources in the context of Jini lookup service [OW00], and a *fault handler* for self-healing.

OptimalGrid is a self-configuring, self-healing and self-optimizing grid middleware, using a set of distributed whiteboards for communication between the different nodes. A computational problem is expressed using Original Problem Cells (OPCs), which describe the connectivity of the cells with their neighbours and the calculations to be performed using the neighbours information. OPCs are aggregated in collections which are themselves part of Variable Problem Partitions (VPP), assigned to grid nodes. The OptimalGrid system is then able to self-configure, using a list of available compute nodes with their characteristics, and can optimize

the repartition of OPCs after each computation cycle. As the communication history between nodes is saved in the whiteboards, if a node is lost the system is able to recover and catch up with the computation, rather than restarting the entire problem. The use of these different autonomic features permits to deliver a grid system more robust and easier to use. Future plans include integrating support for the Open Grid Services Architecture (OGSA) [FCKT02].

By mimicking the behaviour of the human autonomic system especially in dealing with homeostasis, autonomic computing is believed to be a solution to the increasing administrative complexity of computing infrastructures. Hence, no visual supercomputing infrastructure can afford to ignore this emerging technology.

4. Applications of Visual Supercomputing

If we were to have a Grid for visualization, what kind of applications would benefit from it, and perhaps more importantly, how would these applications necessitate specific requirements for such an infrastructure? Shalf and Bethel recently outlined a futuristic scenario depicting how a geophysics researcher and her international collaborators may benefit from grid-based computation and visualization. They concluded that the current state of visualization is not grid ready [SB03]. In this section, we examine several traditional and newly emerged application areas, and discuss their requirements, especially those difficult to be met by the state of the art visualization environments.

4.1. Visual Data Mining and Large-scale Data Visualization

Never before in history have we had such capability for generating, collecting and storing digital data. Data repositories at terabyte level are becoming a common place in many applications, including bioinformatics, medicine, remote sensing and nano-technology. In some applications, such as network traffic visualization [Kou99] and video visualization [DC03] (Figure 4), we are encountering the scenario that dynamic data streams are almost temporally unbounded. Many visualization tasks are evolving into visual data mining processes [KMS02].

These applications are placing a huge strain on the existing visualization environments, and challenging the state of the art technologies in many ways. They demand a variety of infrastructural supports, such as,

- for providing sufficient run-time storage space to active visualization tasks;
- managing complex data distribution mechanisms for parallel and distributed processing;
- choosing the most efficient algorithm according to the size of the problem;
- facilitating the search through a huge parameter space for the most effective visual representation.



Figure 4: Video visualization needs to deal with data streams of an arbitrarily large size. Stream-based rendering can be effectively deployed to visualize video streams.

Data management is the very first issue in handling large datasets. Many visualization processes involve datasets that are much too large for the internal memory of a computer, and have to rely on external disk storage, usually under the virtual memory management of an operating system. The external disk access can become a serious bottleneck in terms of rendering speed. *Out-of-core algorithms* (also known as *external memory algorithms*) [Vit01] are designed to solve a variety of batch and interactive computational problems by minimizing disk I/O overhead. Various out-of-core visualization algorithms have been proposed to handle large structured and unstructured 3D datasets, for instance, in the context of (i) isosurface extraction [CMPS96, CS97, SH00], (ii) terrain rendering [LP02], (iii) streamline visualization [USM97], (iv) mesh simplification [Lin00], (v) rendering time-varying volume data [SCM99], (vi) rendering unstructured volumetric grids [FS01], and (vii) ray tracing [PKGH97]. While some algorithms rely little on internal memory (e.g., [CS97, FS01]), others utilize preprocessed data structures, such as octree [USM97] and indexing [SCM99] to optimize disk I/O operations. Kurc et al. [KÇC*01] recently reported their experience in visualizing large volume datasets using *Active Data Repository*, which is composed of a set of modular services and a unified interface for supporting the management of, and mapping between, in-core and out-core data.

There has been a similar amount of research effort, if not more, for developing techniques that synthesize a visualization image using less than the full dataset. Two commonly used approaches for determining a subset of data to be visualized are *multi-resolution* and *view-dependent* data organization.

Multi-resolution data organization makes use of various hierarchical spatial structures to manage *levels-of-details*

(LODs) of a graphical model or scene. Such structures facilitate real-time rendering by allowing an appropriate LOD to be selected according to the requirements of interactivity and the constraints of computational resources. In computer graphics and visualization, there exists a large collection of works based on this approach. For example, octrees and min-max indexing were used for isosurface extraction [WvG92, Jon95, CMPS96]. Laur and Hanrahan [LH91] utilized an octree for progressive refinement in splatting. Wilhelms et al. [WvGTG96] employed a k-D tree for direct rendering irregular and multiple volumetric grids.

View-dependent data organization makes use of the concepts and algorithms of hidden surface removal, and prioritizes geometrical primitives according to their visibility to the viewer. For example, Livnat and Hansen [LH98] proposed a view-dependent isosurfacing algorithm. LaMar et al. [LHJ99] prioritized volume data based on its proximity to the viewer. Other view-dependent works include visible set estimation [KS00], visibility-based prefetching [CKS03], and view-dependent progressive rendering [NR03].

While it is necessary to deal with problems arising from very large datasets, it is equally important to improve our capability for managing inter-related datasets in order to generate more meaningful visualization. In computer graphics and computer aided design, *scene graphs*, built-upon the concept of constructive solid geometry, have played an indispensable role in combining simple objects into a complex object and bringing many objects together into a scene. It is common for graphics systems to support scene graphs, for instance, in RenderMan, OpenGL, OpenRM, VRML, Java3D, POV-ray and Open Scene Graph. However, support for combinational modelling in visualization systems [BHPB03, NLSG03] is largely based on surface-based scene graphs, relying on image-space composition. Early research efforts for modelling complex visualizations involving multiple datasets were focused on voxelization [WK93]. In order to address the problems associated with voxelization [Jon96], such as excessive data size and data degeneration, Chen and Tucker [CT00] outlined the concept of *constructive volume geometry* for combining volumetric datasets and procedurally defined scalar fields. vlib [WC01], an open source volume graphics API, offers volumetric scene graphs as its fundamental data structure, and provides a discrete ray tracer for direct rendering *volumetric scene graphs*.

In large-scale data visualization, high performance rendering techniques, such as massively parallel rendering (e.g., [MP01]), progressive rendering (e.g., [LH91]) and stream-based rendering (e.g., [HHN*02]), are essential to the process of *making displayable by a computer* (Figure 1). However, facing very large datasets, *making meaningful information visible to one's eyes* is often more critical in visualization. With very large datasets, 'meaningful information' is often featured in a visualization at a sub-pixel level, in a large amount or in four or higher dimensions. This

challenges us to develop visualization techniques into tools for visual data mining [KMS02].

A popular approach to the handling of a huge amount of visual information is the use of *focus and context* techniques, which highlights a 'focus' in detail and depict its 'context' with less details to provide an overview. Focus and context techniques such as *fish-eye views* [SB92], *perspective wall* [MrC91], *hyperbolic space* [Mun98] and *rubber sheets* [SSTR93], have been deployed extensively in information visualization. This approach has also been employed in scientific visualization, deformation-based volume visualization by Kurzion and Yagel [KY95], distortion viewing by Carpendale et al. [CCF97], non-photorealistic rendering by Treavett and Chen [TC00], magnification lens by LaMar et al. [LHJ01], two-level rendering by Hadwiger et al. [HBH03], digital dissection in cardiac visualization by Chen et al. [CTCH03].

Data mining should be closely coupled with visualization [Won99]. Interactive visualization is an indispensable tool in many data mining activities [HAC*99, HKW99], interactive visualization of large datasets not only demands sufficient computational resources, but also requires effective interactive techniques for data exploration, view navigation, data segmentation, data filtering, data fusion and direct manipulation [KMS02].

Perhaps one of the main challenges in the coming year is *computer-assisted design of visual representations*. Many techniques in information visualization enable automated placement of information in a visualization, for instance, *treemap* [Shn92] and *Sunburst* [SZ00] in hierarchy visualization, *recursive pattern* [KKA95] and *circle segments* [AKK96] in time-series visualization, and *spring models* [TK98] and *Kohonen networks* [Koh97] for self-organization and self-optimization in the entire information space. In volume visualization, initial attempts have been made to automate the specification of transfer functions. Marks et al. [MAB*97] proposed a *design galleries* approach to the problem, while Kindlmann and Durkin [KD98] developed a semi-automatic method for generating transfer functions.

The problems surrounding large scale data visualization are collectively becoming an infrastructure issue, as it is unlikely an individual technique can provide a satisfactory solution alone. To process such large amounts of data at the speed required, it is necessary for a visual supercomputing infrastructure to provide dedicated computational resources and application software systems. It is no doubt useful for the infrastructure to select appropriate modelling, processing and rendering techniques according to the available resources and interaction requirements. It is also highly desirable for the infrastructure to offer a wide range of tools for visual data mining as such activities are often unplanned and the effectiveness of a particular tool cannot always be pre-determined.

4.2. Scientific Computation and Computational Steering

Problem Solving Environments (PSEs) are ‘computer systems that provide all of the computational facilities necessary to solve a target class of problems’ [GHR94]. For example, Cactus, is an open source PSE, which was originally designed to provide a framework for solving Einstein’s Equations, and gradually evolved into a ‘unified modular and parallel computational framework for physicists and engineers’ [ABD*01]. While PSEs have been successfully deployed to model many problems in science, engineering and finance, new problems, including a number of grand challenge problems, continue to be formulated.

In scientific modelling and simulation, it is rare to get a correct model without a complex feedback loop involving specification, modelling, computation, visualization and optimization. Upson et al suggested such a computation cycle [UFK*89]. Marshall et al. [MKDY90] identified three modes of combining simulation and visualization, namely *post-processing*, *tracking* and *steering*.

- In *post-processing*, visualization is merely a post-processing stage of simulation and cannot directly influence (or even abort) the simulation. This asynchronous working requires simulation to complete before visualization begins, and so there is no opportunity to effect any control on the simulation through the visualization. A benefit however is that the scientists can take as long as they want in visualizing the results, as the time scale for visualization is independent of that for the simulation.
- In *tracking*, the simulation and visualization are coupled, but there is no concept of the user altering the simulation on the basis of the visualization, other than the user hitting the abort key!
- In *steering*, the control parameters of the simulation are exposed, and can be manipulated as it runs. The model was expressed as dataflow. This concept was extended by Brodlić et al. [BBB*93] to allow an audit trail of checkpoint information to be stored in a tree structure, called *History Tree*. This generalized steering to facilitate a ‘reverse gear’. Simulation and visualization were seen as separate processes, linked through a manager. Recently Zhou et al. [ZCW02] proposed an approach towards automatic steering based on comparative visualization involving both experimental and computational results.

Building a generic computational steering environment is a non-trivial task. A significant development in this area was carried out by van Liere, van Wijk and Mulder in a series of works [vWvL97, vLMvW97, MvWvL99]. The key innovation was to build steering widgets, which sat within the visualization, to enable direct manipulation of the simulation.

A major software advance was made with SCIRun from Utah, which was a dataflow environment specially designed for steering. It facilitated the interactive construc-

tion, debugging and steering of large-scale scientific computations [SGPJ98].

CUMULVS (*Collaborative User Migration, User Library for Visualization and Steering*) [NER04] was developed as part of the large US DOE ACTS project which aims to provide tools for scientific programmers. It is a software framework for linking steering and visualization services with parallel simulation. It provides two libraries: one for the application, the other for the steering and visualization front-end. It is collaborative in the sense that multiple remote viewers can connect to a simulation.

Recently the RealityGrid project (Figure 5) have built some impressive demonstrations of steering Lattice-Boltzmann simulations, which are massive Grid applications, involving collections of machines across the world, and are state of art in what can be achieved on a global scale [BCH*03].

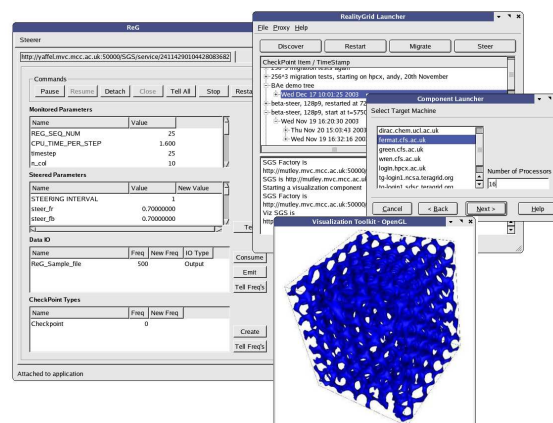


Figure 5: A computational steering environment developed in the RealityGrid project.

On a smaller scale, the gViz e-science project [WBW03] has studied two approaches to computational steering. One extends IRIS Explorer to run in secure distributed fashion across Grid machines, so an IRIS Explorer session spans the internet. The simulation runs inside IRIS Explorer. The other is very similar to RealityGrid in building an API for steering, and decoupling the simulation code from the visualization.

The close-coupling between computation and visualization in computational steering has highlighted the need for advanced inter-process and inter-task management in an visual supercomputing infrastructure. This challenges the underlying technologies of visual supercomputing, requires further advances in fields such as operating systems (e.g., for process management and migration), and programming environments (e.g., for component-based programming, dynamic integration management).

As scientific modelling and simulation usually involves

many repetitive steps in a feedback loop, there is a great scope for a visual supercomputing infrastructure to collect performance data in such a feedback loop, and transform the data into the knowledge, which can be used to offer users appropriate guidance, identify the best configuration, automate part of the process, and hence provide a higher quality of services. Such an approach has been extensively deployed in business, a conceptually similar situation, where customers are involved in a repetitive process loop, and data measuring various attributes of the process can easily be collected and analyzed. There are some successful examples where the quality of services has been improved.

4.3. Mission Critical Visualization

This category of visualization requires the real time processing of large datasets, possibly from diverse sources, that can then be fed into an interactive visualization environment. Typically, such a system provides decision support tools to the end user. Application areas exist in defence and intelligence, law enforcement, healthcare and social services, scientific research and education, transportation and communication, and energy and the environment. A mature example of mission critical systems are training simulators such as flight simulators, which have used custom built hardware to train pilots for many years both in routine flying and critical incident handling [Sog02].

Medical simulators are expected to be the next major application to benefit from simulator technology, but based on commodity graphics hardware (see Section 3.3). Clinicians are also using intra-operative surgical planning tools and neurosurgeons, for example, have been utilizing image guidance for the last decade [PDe96]. The military is another large market for mission critical visualization. For example, the US Fleet Numerical Meteorology and Oceanography Centre was recently tasked with supplying military forces deployed in the Persian Gulf with highly accurate meteorological information critical to conducting land, sea and air operations.

A characteristic of mission critical visualization has been the requirement for specialized and often expensive equipment. Until recently, growth has been restricted to niche areas and little work has been published on the optimization and scheduling problems of the visualization task. Grid and cluster based computing, however, are now providing an infrastructure for further exploitation and visualization will be a key component of future work.

For example, as shown in Figure 6, in a system for delivering interactive volume interrogation of patient data in the operating theatre [MJ03], visualization tasks were carried out on a server over a mile away from the hospital and then delivered across the data network. Applications such as this raise many issues including: how to guarantee a minimum bandwidth required for both data communication and



Figure 6: Visualization-guided surgery is a typical application of mission-critical visualization.

data processing; the use of redundancy for both communication and computation to ensure a reliable delivery of visualization; and the handling of secure information. Synchronization algorithms as well as data distribution techniques must also be considered when making use of multiple compute resources [Fuj01]. Those are exactly the issues to be addressed by a visual supercomputing infrastructure.

4.4. Mobile Visualization

Ubiquitous computing is capturing our imagination of a global infrastructure that supports not only networks of desktop computers and high performance computers, but also a huge number of wearable and mobile computing devices [Wei93]. The prospect for integrating mobile devices into the visualization pipeline and its applications offers new opportunities for accessing, interrogating and manipulating data remotely.

Izadi et al. [ICRS02] proposed the FUSE system as a development tool for collaborative systems across multiple platforms. Lamberti et al. [LZS*03] demonstrated a mobile graphical interactive rendering task running on a PDA which is provided by a remote graphics workstation. Wolf et al. [WCHS02] proposed the Smart Pointer as a role for PDA devices, where it either presents a subset of the visualization when part of a larger visualization environment (such as a CAVE) or it aims to provide the same overall image as other (desktop) clients, both approaches using a remote visualization server. Hartling et al. [HBCN02] presented a middleware system, Tweek, which displays a 2D GUI to a virtual environment using a PDA. The user may interact with the virtual environment via the PDA. D'Amora and Bernardini [DB03] developed a PDA 3D viewer that can access a remote database of CAD models. Apart from the technical aspect, human factor issues in using PDAs for visualization need to be addressed [PRM00].

We have categorized the demands upon both the mobile device and the visualization service into the following classes ordered according to their communication requirements:

- *Remote scheduling* — A device, such as a PDA, can be used to monitor the account status of the user on a visualization server. The user should be able to consult their account, see the current state of any jobs, and perform basic management tasks, such as *start*, *stop*, *hold* and *remove*. This requires a low bandwidth duplex channel for textual communications.
- *Remote monitoring* — Higher level monitoring functions can take advantage of the colour displays on the device. Users may query their account to retrieve still images which are visualizations of their data. They may (pre)select parameters for rendering (such as rendering method and transfer function), and be presented with the image. Such parameters may be used to assist with scheduling decisions. This class requires a duplex channel with a higher bandwidth downstreaming traffic.
- *Remote steering* — A remote user can be notified on job (or intermediate result) completion, and may view a visualization of the result. Some limited interaction with the visual representation is possible as the user's feedback can be used to generate modifications to the current job. This is most useful for checking intermediate results during batch mode without having to be tied down to one location. Some steering of the simulation is possible as jobs can be stopped and restarted from a recent state with new parameters. The bandwidth requirement is higher as the wait time for several images may be undesirable. The computational demands on the PDA are higher due to the need to zoom, pan, and interact with the data. At this stage, transmission and interaction with small 3D models may be desirable and possible.
- *Remote visualization* — The user interacts freely with the simulation, using the visualization to explore all aspects of their data. This places a high demand on the PDA as well as the server. The visualization could be in the form of a sequence of images generated by the server and transmitted compressed to the PDA, or the server could send a stream which could be processed by the limited graphics hardware available on the PDA. User interface widgets could be overlaid over the data, and the user will send interaction data back to the server in order to steer the simulation. Some frame loss, and some pauses in results are inconvenient but not critical.

Mobile visualization (Figure 7) introduces an interesting design problem for a visual supercomputing environment. It reminds us of the desktop technology two decades ago when low resolution displays and limited computation resources were supported by main frame computers. However, it also exhibits a completely new scenario where the requests for visualization, or managing visualization tasks, can come from anywhere with often unreliable communication chan-



Figure 7: Mobile technology has offered an exciting scope for developing new visualization applications.

nels in terms of both bandwidth and security. The infrastructural support to mobile visualization may significantly broaden the application scope of visualization, and transform this largely laboratory-based technology to a pervasive technology.

5. Challenges in Visual Supercomputing

The above discussions have clearly indicated the need for encompassing a large collection of infrastructural issues related to the management of visualization tasks in a common framework, for which we have introduced the subject domain of *visual supercomputing*. The requirements from applications, such as visual data mining, computational steering, mission-critical visualization and mobile visualization, have indicated a high research priority to the infrastructure of visual supercomputing. While such an infrastructure can benefit from the state of the art technologies in visualization, we are still facing many new challenges in order to realize a well-designed, serviceable and cost-effective infrastructure for visual supercomputing.

Hoare outlined a set of criteria for a grand challenge in computer science [HMTB02]. According to these criteria, building a visual supercomputing infrastructure can be considered as a grand challenge in the field of visualization. It raises a series of scientific questions such as:

- *Architectural Design* — Would it be desirable or feasible to build an infrastructure for visual supercomputing based on that of the Grid? How would it accommodate the different needs for centralized, distributed or independent services from various applications? How would such an infrastructure provide generic support to the management of visualization data, distributed visual data mining, very large scale data visualization, mission-critical visualization and mobile visualization?
- *Technology Deployment* — Should special purpose graphics hardware form the central core of a visual supercomputing environment? If so, what would be the relation-

ship between such central hardware and graphics hardware available on personal computers? What would different hardware attributes impact upon visualization algorithms, and how would visualization tasks be managed to take such attributes into account?

- *Quality of Service* — How would a visual supercomputing infrastructure provide seamless services to many users and for many applications, instead of just another ‘remote login’ service? What would be the role of the infrastructure in managing interaction, data and knowledge about users’ experience? In what way could users benefit from a knowledge-based infrastructure?

One emerging strategy for developing complex computing infrastructure is *autonomic computing* [KC03] (see also Section 3.6.2), which seeks inspiration in self-adaptive biological systems and self-governing social and economic systems.

Adapting the deployment model, proposed by IBM [IBM04], for the gradual evolution of complex system-wide self-managing environments, one can envisage a similar five-level deployment model for visual supercomputing, which can be developed evolutionarily.

- *Level 1: Basic* — At this level, a visual supercomputing infrastructure is an integrated *system* platform that provides visualization applications with necessary computation and communication resources. Typically, users are fully involved in identifying appropriate tools, locating computation resources, and managing data distributions. It is often necessary for users to navigate themselves through complicated technical obstacles, such as networking, security, parallel computing, data replication, and so forth.
- *Level 2: Managed* — At this level, a visual supercomputing infrastructure will have a managed *service* layer between the user interface and the system platform. The service layer is aware of the availability and ontology of data and resources, and can provide services to various visualization applications according to dynamic requirements of users and applications as well as dynamic states of the system platform. To a large extent, the development of the Grid technology is aiming at the delivery of a general-purpose infrastructure. To managing visualization applications effectively, it is necessary to incorporate more advanced service features into the Grid technology for supporting a variety of visualization needs such as interactive, distributed, mobile, and mission-critical applications in a more transparent manner.
- *Level 3: Predictive* — At this level, a visual supercomputing infrastructure will have an *information* layer between the user interface and the service layer, which collects, monitors and correlates various user interaction data and system performance data. It provides users with analytical data, which may indicate the quality of visualization results, effectiveness of visualization tools, and so on. In

addition, this layer can enable faster and better task specification by reporting potential problems and recommending suitable tools and visual representations. It is at this level, the infrastructure starts to manage users’ experience in carrying out visualization tasks.

- *Level 4: Adaptive* — At this level, a visual supercomputing infrastructure will have an *adaptation* layer between the information layer and the service layer. Based on the information collected, the adaptation layer has the functionality for self-configuring and self-optimizing the computational requirements of a visualization task, as well as the functionality for self-managing the system platform and various visualization services dynamically. It is at this level, visualization users can be largely freed from software management, and are able to focus on their core business, that is, visualization.
- *Level 5: Autonomic* — At this level, the traditional user interface in a visual supercomputing infrastructure will be replaced by an intelligent user interface, for instance ‘a virtual secretary’, which is capable of transform information to knowledge and provides users with a wide range assistance. Such assistance may include specifying visualization tasks, scheduling inter-dependent jobs, organizing raw data and visualization results, managing security, checking the quality of the service and results, and arranging the sharing of the data with other users.

Figure 8 illustrates evolutionary advance of the infrastructure through the five levels. In this deployment model, each layer is merely a conceptual placeholder for a collection of functional components (e.g., services, tools, agents, databases, knowledge-bases, and so on). It is not necessary for the development and deployment of each level to follow a temporal order. Nor is it desirable to make each layer a centralized bottleneck in the process of visualization. It is most likely that the infrastructure will be realised with a large number of autonomous, interacting, self-governing functional components.

Building a visual supercomputing infrastructure is no doubt an ambitious grand challenge. However, we have already had a solid foothold at Level 1, and are rapidly approaching Level 2. A noticeable amount of research effort is being made to develop system-level autonomic computing techniques in many fields, including distributed systems, data communications, Internet technology, Grid computing, agent technology, database systems and business management systems. Some of such effort can be viewed as ‘horizontal’ deployment of autonomic computing at the *system layer* and *service layer* of a visual supercomputing infrastructure (Figure 8), while others can provide new concepts, methods and tools for the development of the *intelligent user interface*, *information and knowledge layer* and *adaptation layer*. Hence, we believe that having such a visual supercomputing infrastructure is a realistic challenge.




Figure 8: developing new visualization applications.

6. Conclusions

In this STAR report, we have outlined an agenda for *visual supercomputing*, which defines a subject domain concerning the infrastructural technology for visualization. We have considered a broad range of scientific and technological advances in computer graphics and visualization, which are relevant to visual supercomputing. We have identified the state of the art technologies that have prepared us for building such an infrastructure. We have examined a collection of applications that would benefit enormously from such an infrastructure, and discussed their technical requirements. We have proposed a set of challenges that may guide our strategic efforts in the coming years. In particular, we have highlighted the integral role of *autonomic computing* in the gradual evolution of an infrastructure for visual supercomputing.

Acknowledgements

This work is conducted under the framework of the e-Viz project, which is jointly funded by EPSRC ICT programme and e-Science Programme with grants GR/S46567, GR/S46574 and GR/S46581. Special acknowledgements to Chris Hughes whose initial review motivated this report, Nigel John who suggested making this as a group exercise for the e-Viz project, and Min Chen who coordinated the preparation of this report. We are thankful to John Sharp, University of Wales Swansea, for his valuable advice on data flow computing.

References

- [ABD*01] ALLEN G., BENDER W., DRAMLITSCH T., GOODALE T., HEGE H.-C., LANFERMANN G., MERZKY A., RADKE T., SEIDEL E., SHALF J.: Cactus tools for grid applications. *Cluster Computing* 4, 3 (2001), 179–188.
- [ABM*01] AHRENS J., BRISLAWN K., MARTIN K., GEVECI B., LAW C. C., PAPKA M.: Large-scale data visualization using parallel data streaming. *IEEE Computer Graphics and Applications* 21, 4 (July/August 2001), 34–41.
- [Acc03] Access Grid documentation. <http://www.accessgrid.org>, December 2003.
- [AGL*04] AIBER S., GILAT D., LANDAU A., RAZINKOV N., SELA A., WASSERKRUG S.: Autonomic self-optimization according to business objectives. In *Proc. 1st International Conference on Autonomic Computing* (2004), pp. 206–213.
- [AHP*03] ANDERSON S., HARTSWOOD M., PROCTER R., ROUNCEFIELD M., SLACK R., SOUTTER J., VOSS A.: Making autonomic computing systems accountable: the problem of human computer interaction. In *Proc. 14th International Workshop on Database and Expert Systems Applications* (2003), pp. 718–724.
- [Ake89] AKELEY K.: The silicon graphics 4D/240GTX superworkstation. *IEEE Computer Graphics and Applications* 9, 4 (1989), 71–83.
- [AKK96] ANKERST M., KEIM D. A., KRIEGEL H. P.: Circle segments: a technique for visual exploring large multidimensional data sets. In *Proc. IEEE Visualization '96* (1996), p. Hot Topic Session.
- [Amd67] AMDAHL G.: Validity of the single processor approach to achieving large-scale computing capabilities. In *Proc. AFIPS Conference* (1967).
- [AMD94] ANG C. S., MARTIN D. C., DOYLE M. D.: Integrated control of distributed volume visualization through the world wide web. In *Proc. IEEE Visualization 94* (1994), Bergeron R. D., Kaufman A. E., (Eds.), pp. 13–20.
- [AO93] ANDREWS G. R., OLSSON R. A.: *The SR Programming Language: Concurrency in Practice*. Benjamin/Cummings, 1993.
- [AP03] AGARWAL M., PARASHAR M.: Enabling autonomic compositions in grid environments.

- In *Proc. 4th International Workshop on Grid Computing* (2003), pp. 34–41.
- [AT95] ABRAM G., TREINISH L.: An extended dataflow architecture for data analysis and visualization. *ACM SIGGRAPH Computer Graphics* 29, 2 (1995), 17–21.
- [Azu97] AZUMA R. T.: A survey of augmented reality. *Presence: Teleoperators & Virtual Environments* 6, 4 (1997), 355–85.
- [Bak97] BAKER S.: *CORBA Distributed Objects*. Addison Wesley, 1997.
- [BBB*93] BRODLIE K., BRANKIN L., BANECKI G., GAY A., POON A., WRIGHT H.: GRASPARC – a problem solving environment integrating computation and visualization. In *Proc. IEEE Visualization '93* (1993), pp. 102–109.
- [BBG*94] BENFORD S. D., BOWERS J., GRAY S., RODDEN T. R., RYGOL M., STANGER V.: The VirtuOsi project. In *Proc. VR'94 (London Virtual Reality Expo 1994)* (February 1994).
- [BC03] BURDEA G., COIFFET P.: *Virtual Reality Technology*, 2nd ed. Wiley, 2003.
- [BCH*03] BROOKE J. M., COVENEY P. V., HARTING J., JHA S., PICKLES S. M., PINNING R. L., PORTER A. R.: Computational steering in RealityGrid. In *Proc. the UK e-Science All Hands Meeting 2003* (2003).
- [Ben75] BENTLEY J. L.: Multidimensional binary search trees used for associative searching. *Communications of the ACM* 18, 9 (1975), 509–517.
- [BHPB03] BETHEL E. W., HUMPHREYS G., PAUL B., BREDERSON J. D.: Sort-first, distributed memory parallel visualization and rendering. In *Proc. IEEE Symposium on Parallel and Large Data Visualization and Graphics* (2003), pp. 41–50.
- [BKP01] BILLINGHURST M., KATO H., POUPYREV I.: The MagicBook: A transitional ar interface. *Computers and Graphics* (November 2001), 745–753.
- [BN84] BIRRELL A. D., NELSON B. J.: Implementing remote procedure calls. *ACM Transactions on Computer Systems* (1984), 39–59.
- [Bro95] BROLL W.: Interacting in distributed collaborative VE. In *Proc. VRAIS '95* (1995).
- [BS01] BARTZ D., SILVA C.: Rendering and visualization in parallel environments. In *Eurographics 2001 Tutorials*. 2001.
- [BS02] BLUNDELL B. G., SCHWARZ A. J.: The classification of volumetric display systems: characteristics and predictability of the image space. *IEEE Transactions on Visualization and Computer Graphics* 8, 1 (2002), 66–75.
- [Buy99] BUYYA R. (Ed.): *High Performance Cluster Computing – Volume 1: Architectures and Systems*. Prentice Hall, 1999.
- [BWD*02] BRODLIE K. W., WOOD J., DUCE D., GALLOP J. R., GAVAGHAN D., GILES M., HAGUE S., WALTON J., RUDGYARD M., COLLINS B., IBBOTSON J., KNOX A.: XML for visualization. In *Euroweb* (2002).
- [BWD*03] BRODLIE K., WOOD J., DUCE D., GALLOP J., WALTON J.: Distributed and collaborative visualization. In *Eurographics 2003 State of the Art Report*. 2003.
- [CA86] CHIEN C. H., AGGARWAL J. K.: Volume/surface octrees for the representation of three-dimensional objects. *Computer Vision, Graphics and Image Processing* 36, 1 (1986), 100–113.
- [CCF95] CABRAL B., CAM N., FORAN J.: Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proc. ACM/IEEE Symposium on Volume Visualization* (1995), pp. 91–98.
- [CCF97] CARPENDALE M. S. T., COWPERTHWAIT D. J., FRACCHIA F. D.: Extending distortion viewing from 2D to 3D. *IEEE Computer Graphics and Applications* 17, 4 (1997), 42–51.
- [CG86] CARRIERO N., GELERNTER D.: Linda in context. *Communications of the ACM* 32 (1986), 444–458.
- [CH93] CARLSON C., HAGSAND O.: DIVE: a platform for multi-user virtual environments. *Computers and Graphics* 17, 6 (1993).
- [CKKG99] CHAPIN S. J., KATRAMATOS D., KARPOVICH J., GRIMSHAW A. S.: The Legion resource management system. In *Job Scheduling Strategies for Parallel Processing*, Feitelson D. G., Rudolph L., (Eds.). Springer Verlag, 1999, pp. 162–178.
- [CKS03] CORRÊA W. T., KLOSOWSKI J. T., SILVA C. T.: Visibility-based prefetching for interactive out-of-core rendering. In *Proc. IEEE Symposium on Parallel and Large-Data Visualization and Graphics* (2003), pp. 2–8.

- [CMPS96] CIGNONI P., MONTANI C., PUPPO E., SCOPIGNO R.: Optimal isosurface extraction from irregular volume data. In *Proc. IEEE Symposium on Volume Visualization* (1996), pp. 31–38.
- [CNSDe92] CRUZ-NEIRA C., SANDIN D. J., DEFANTI T. A., ET AL: The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM* 35, 6 (1992), 65–72.
- [Con] Condor. <http://www.cs.wisc.edu/condor>.
- [Con03] CONSTANTINESCU Z.: Towards an autonomic distributed computing environment. In *Proc. 14th International Workshop on Database and Expert Systems Applications* (2003), pp. 699–703.
- [COV] COVISE. <http://www.vircinity.com/>.
- [CS97] CHIANG Y.-J., SILVA C. T.: I/O optimal isosurface extraction. In *Proc. IEEE Visualization '97* (1997), pp. 293–300.
- [CT00] CHEN M., TUCKER J. V.: Constructive volume geometry. *Computer Graphics Forum* 19, 4 (2000), 281–293.
- [CTCH03] CHEN M., TUCKER J. V., CLAYTON R. H., HOLDEN A. V.: Constructive volume geometry applied to visualization of cardiac anatomy and electrophysiology. *International Journal of Bifurcation and Chaos* 13, 12 (2003), 3591–3604.
- [CZL*04] CHEN M., ZHENG A. X., LLOYD J., JORDAN M. I., BREWER E.: Failure diagnosis using decision trees. In *Proc. 1st International Conference on Autonomic Computing* (2004), pp. 36–43.
- [DB03] D'AMORA B., BERNARDINI F.: Pervasive 3D viewing for product data management. *IEEE Computer Graphics and Applications* 23, 2 (2003), 14–19.
- [DC03] DANIEL G. W., CHEN M.: Video visualization. In *Proc. IEEE Visualization 2003* (Seattle, WA, October 2003), pp. 409–416.
- [DEH89] DEW P. M., EARNSHAW R. A., HEYWOOD T. R. (Eds.): *Parallel Processing for Computer Vision and Display*. Addison-Wesley, 1989.
- [DHX*03] DONG X., HARIRI S., XUE L., CHEN H., ZHANG M., PAVULURI S., RAO S.: Autonomia: an autonomic computing environment. In *Proc. IEEE International Conference on Performance, Computing, and Communications* (2003), pp. 61–68.
- [DLJ03] DEEN G., LEHMAN T., JAUFMAN J.: The almaden optimalgrid project. In *Autonomic Computing Workshop – Proc. 5th International Workshop on Active Middleware Services* (2003).
- [DPH*03] DEMARLE D. E., PARKER S., HARTNER M., GRIBBLE C., HANSEN C.: Distributed interactive ray tracing for large volume visualization. In *Proc. IEEE Symposium on Parallel and Large-Data Visualization and Graphics* (2003), pp. 88–94.
- [DRD*04] DUNAGAN J., ROUSSEV R., DANIELS B., VERBOWSKI A. J. C., WANG Y.: Towards self-managing software patching process using black-box persistent-state manifestes. In *Proc. 1st International Conference on Autonomic Computing* (2004), pp. 106–113.
- [DT81] DOCTOR L. J., TORBORG J. G.: Display techniques for octree-encoded objects. *IEEE Computer Graphics and Applications* 3, 1 (1981), 29–38.
- [Dye90] DYER D. S.: A dataflow toolkit for visualization. *IEEE Computer Graphics and Applications* 10, 4 (1990), 60–69.
- [e-V] e-Viz. www.eviz.org.
- [EE02] ENGEL K., ERTL T.: Interactive high-quality volume rendering with flexible consumer graphics hardware. In *Eurographics 2002 State of the Art Reports*. 2002.
- [EM78] ELWALD R., MASS L.: A high performance graphics system for the Cray-1. *ACM SIGGRAPH Computer Graphics* (1978), 82–86.
- [ERA*03] EYMANN T., REINICKE M., ARDAIZ O., ARTIGAS P., FREITAG F., NAVARRO L.: Self-organizing resource allocation for autonomic network. In *Proc. 14th International Workshop on Database and Expert Systems Applications* (2003), pp. 656–660.
- [ES01] ERWIN D. W., SNELLING D. F.: UNICORE: a Grid computing environment. *Lecture Notes in Computer Science* 2150 (2001), 825–834.
- [EUR] EUROGRID. <http://www.eurogrid.org>.
- [EWE99] ENGEL K., WESTERMANN R., ERTL T.: Isosurface extraction techniques for web-based volume visualization. In *Proc. IEEE Visualization '99* (1999), pp. 139–146.
- [FCKT02] FOSTER I., C. KESSELMAN J. N., TUECKE S.: The physiology of the Grid: an open grid services architecture for distributed system integration. In *Open Grid Services In-*

- frastructure WG (June 2002), Global Grid Forum.
- [Fen81] FENG T. Y.: A survey of interconnection networks. *IEEE Computers* (1981), 12–27.
- [FK97] FOSTER I., KESSELMAN C.: Globus: A metacomputing infrastructure toolkit. *International Journal Supercomputer Applications* 11, 2 (1997), 115–128.
- [FK98a] FOSTER I., KESSELMAN C.: The Globus project: a status report. In *Proc. Heterogeneous Computing Workshop* (1998), IEEE Press, pp. 4–18.
- [FK98b] FOSTER I., KESSELMAN C. (Eds.): *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
- [Fly72] FLYNN M. J.: Some computer organizations and their effectiveness. *IEEE Transactions on Computers C-21*, 9 (1972), 948–960.
- [Fou95] FOULSER D.: IRIS Explorer: A framework for investigation. *ACM SIGGRAPH Computer Graphics* 29, 2 (1995), 13–16.
- [FS01] FARIAS R., SILVA C. T.: Out-of-core rendering of large, unstructured grids. *IEEE Computer Graphics and Applications* 21, 4 (2001), 42–50.
- [Fuj01] FUJIMOTO R. M.: Advanced tutorials: Parallel simulation: parallel and distributed simulation systems. In *Proc. the 33rd Conference on Winter Simulation* (2001), IEEE Computer Society, pp. 147–157.
- [FW78] FORTUNE S., WYLLIE J.: Parallelism in random access machines. In *Proc. ACM Symposium on Theory of Computing* (1978), pp. 114–118.
- [GB95] GREENHALGH C. M., BENFORD S. D.: MASSIVE: a virtual reality system for teleconferencing. *ACM Transactions on Computer Human Interfaces* 2, 3 (1995), 239–261.
- [GBD*94] GEIST A., BEGUELIN A., DONGARRA J., JIANG W., MANCHEK R., SUNDERAM V.: *PVM Parallel Virtual Machine, A User's Guide and Tutorial for Networked Parallel Computing*. MIT Press, Cambridge, MA, 1994.
- [GGK93] GRAMA A. Y., GUPTA A., KUMAR V.: Isoefficiency: Measuring the scalability of parallel algorithms and architectures. *IEEE Parallel and Distributed Technology: Systems and Applications* 1, 3 (1993), 12–21.
- [GHR94] GALLOPOULOS E., HOUSTIS E. N., RICE J. R.: Computer as thinker/doer: Problem-solving environments for computational science. *IEEE Computational Science and Engineering* 1, 2 (1994), 11–23.
- [GKB89] GRIMES J., KOHN L., BHARADHWAJ R.: The intel i860 64-bit processor: a general-purpose cpu with 3d graphics capabilities. *IEEE Computer Graphics and Applications* 9, 4 (1989), 85–94.
- [GKP96] GEIST G. A., KOHLA J. A., PAPADOPOULOS P. M.: PVM and MPI: A comparison of features. *Calculateurs Paralleles* 8, 2 (1996), 137–150.
- [GLPW*96] GRIMSON W. E. L., LOZANO-PÉREZ T., WELLS III W. M., ETTINGER G. J., WHITE S. J., KIKINIS R.: An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization. *IEEE Transactions on Medical Imaging* 15, 2 (1996), 129–140.
- [GMB88] GUSTAFSON J. L., MONTRY G. R., BENNER R. E.: Development of parallel methods for a 1024 hypercubes. *SIAM Journal on Scientific and Statistical Computing* 9, 4 (1988), 609–638.
- [GR95] GOIL S., RANKA S.: Dynamic load balancing for ray-traced volume rendering on distributed memory machines. In *Proc. International Conference on High Performance Computing* (1995).
- [Gre91] GREEN S.: *Parallel Processing for Computer Graphics*. The MIT Press, Cambridge, MA, 1991.
- [GRI] GRIP. <http://www.grid-interoperability.org/>.
- [HAC*99] HELLERSTEIN J. M., AVNUR R., CHOU A., HIDBER C., OLSTON C., RAMAN V., ROTH T., HAAS P. J.: Interactive data analysis: The control project. *Computer* 32, 8 (1999), 51–58.
- [Hal97] HALLE M.: Autostereoscopic displays and computer graphics. *ACM SIGGRAPH Computer Graphics* 31, 2 (1997), 58–62.
- [Han96] HANSEN C.: Known and potential high performance computing applications in computer graphics and visualization. In *Proc. International Workshop on High Performance Computing for Computer Graphics and Visualisation* (1996), Chen M., Townsend P., Vince J. A., (Eds.), Springer, pp. 23–29.

- [HBCN02] HARTLING P., BIERBAUM A., CRUZ-NEIRA C.: Virtual reality interfaces using Tweek. In *SIGGRAPH 2002 Sketches*. ACM, 2002.
- [HBH03] HADWIGER M., BERGER C., HAUSER H.: High-quality two level volume rendering of segmented data sets on consumer graphics hardware. In *Proc. IEEE Visualization 2003* (2003), Seattle, WA, pp. 301–308.
- [HEB*01] HUMPHREYS G., ELDRIDGE M., BUCK I., STOLL G., EVERETT M., HANRAHAN P.: WireGL: a scalable graphics system for clusters. In *Proc. SIGGRAPH 2001* (2001).
- [HHN*02] HUMPHREYS G., HOUSTON M., NG R., FRANK R., AHERN S., KIRCHNER P. D., KLOSOWSKI J. T.: Chormium: a stream-processing framework for interactive rendering on clusters. In *Proc. SIGGRAPH 2002* (2002).
- [HKW99] HINNEBURG A., KEIM D. A., WAWRYNIUK M.: Hd-eye: Visual mining of high-dimensional data. *IEEE Computer Graphics and Applications* 19, 5 (1999), 22–31.
- [HMTB02] HOARE T., MILNER R., THOMAS M., BUNDY A.: Criteria for a grand challenge. <http://www.cra.org/Activities/grand.challenges/hoare.pdf> or http://www.nesc.ac.uk/esi/events/Grand_Challenges/criteria.html, 2002.
- [IBM04] IBM: Autonomic deployment model. <http://www-306.ibm.com/autonomic/levels.shtml>, 2004.
- [ICRS02] IZADI S., COUTINHO P., RODDEN T., SMITH G.: The FUSE platform: supporting ubiquitous collaboration within diverse mobile environments. *Automated Software Engineering* 9, 2 (2002), 167–186.
- [Jon95] JONES M. W.: *The Visualisation of Regular Three Dimensional Data*. 1995. PhD Thesis, University of Wales.
- [Jon96] JONES M. W.: The production of volume data from triangular meshes using voxelisation. *Computer Graphics Forum* 15, 5 (1996), 311–318.
- [JW00] JOGALEKAR P., WOODSIDE M.: Evaluating the scalability of distributed systems. *IEEE Transactions on Parallel and Distributed Systems* 11, 6 (2000), 589–603.
- [KAH04] KANDASAMY N., ABDELWAHED S., HAYES J.: Self-optimization in computer systems via on-line control: applications to power management. In *Proc. 1st International Conference on Autonomic Computing* (2004), pp. 54–61.
- [KASK02] KELSO J., ARSENAULT L. E., SATTERFIELD S. G., KRIZ R. D.: DIVERSE: a framework for building extensible and reconfigurable device independent virtual environments. In *Proc. IEEE Virtual Reality* (2002), pp. 183–190.
- [KC03] KEPHART J. O., CHESS D. M.: The vision of autonomic computing. *IEEE Computer* (2003), 41–50.
- [KÇC*01] KURE T., ÇATALYÜREK Ü., CHANG C., SUSSMAN A., SALTZ J.: Visualization of large data sets with the active data repository. *IEEE Computer Graphics and Applications* 21, 4 (2001), 24–33.
- [KD98] KINDLMANN G., DURKIN J.: Semi-automatic generation of transfer functions for direct volume rendering. In *Proc. IEEE Symposium on Volume Visualization* (1998), pp. 79–86.
- [KGGK93] KUMAR V., GRAMA A., GUPTA A., KARYPIS G.: *Introduction to Parallel Computing: Design and Analysis of Algorithms*. Benjamin, 1993.
- [KKA95] KEIM D. A., KRIEDEL H. P., ANKERST M.: Recursive pattern: a technique for visualizing very large amounts of data. In *Proc. IEEE Visualization '95* (1995), pp. 279–286.
- [KKV*02] KLOSOWSKI J. T., KIRCHNER P. D., VALUYEVA J., ABRAM G., MORRIS C. J., WOLFE R. H., JACKMAN T.: Deep view: high-resolution reality. *IEEE Computer Graphics and Applications* 22, 3 (2002), 12–15.
- [KMS02] KEIM D. A., MULLER W., SCHUMANN H.: Visual data mining. In *Eurographics 2002 State of the Art Reports*. 2002.
- [Koh97] KOHONEN T.: *Self-organizing Maps*, 2 ed. Springer, 1997.
- [Kou99] KOUTSOPIO E. E.: Visualizing large-scale telecommunication networks and services. In *Proc. IEEE Visualization* (1999).
- [Kou03] KOUTEK M.: *Scientific Visualisation in Virtual Reality: Interaction Techniques and Application Development*. Delft University of Technology, 2003.
- [Kri89] KRISHNAMURTHY E. V.: *Parallel Pro-*

- cessing: *Principles and Practice*. Addison-Wesley, 1989.
- [KS00] KLOSOWSKI J. T., SILVA C. T.: The prioritized-layered projection algorithm for visible set estimation. *IEEE Transaction on Visualization and Computer Graphics* 6, 2 (2000), 108–123.
- [KWH03] KANUS U., WETEKAM G., HIRCHE J.: Voxelcache: a cache-based memory architecture for volume graphics. In *Proc. ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware* (2003), pp. 76–83.
- [KY95] KURZION Y., YAGEL R.: Space deformation using ray deflectors. In *Proc. the 6th Eurographics Workshop on Rendering* (Dublin, Ireland, June 1995), pp. 21–32.
- [Lac96] LACROUTE P.: Analysis of a parallel volume rendering system based on the shear-warp factorisation. *IEEE Transactions on Visualization and Computer Graphics* 2, 3 (1996), 218–231.
- [LaR98] LAROSE D.: *A Fast, Affordable System for Augmented Reality*. Master's thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, April 1998.
- [LaR00] LAROW M.: The five engines of e-crm. *Computer Technology Review* (August 2000).
- [LBS03] LI G.-S., BORDOLOI U. D., SHEN H.-W.: Chameleon: an interactive texture-based rendering framework for visualizing three-dimensional vector fields. In *Proc. IEEE Visualization '03* (2003), pp. 241–248.
- [LC87] LORENSEN W., CLINE H.: Marching cubes: a high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics* 21, 4 (1987), 163–169.
- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 5 (1988), 29–37.
- [LH91] LAUR D., HANRAHAN P.: Hierarchical splatting: a progressive refinement algorithm for volume rendering. *ACM SIGGRAPH Computer Graphics* 25, 4 (1991), 285–288.
- [LH98] LIVNAT Y., HANSEN C.: View dependent isosurface extraction. In *Proc. IEEE Visualization '98* (1998), pp. 175–180.
- [LHJ99] LAMAR E., HAMANN B., JOY K.: Multiresolution techniques for interactive texture-based volume visualization. In *Proc. IEEE Visualization '99* (1999), pp. 355–362.
- [LHJ01] LAMAR E., HAMANN B., JOY K. I.: A magnification lens for interactive volume visualization. In *Proc. IEEE Pacific Conference on Computer Graphics and Applications* (2001), pp. 223–231.
- [LI03] LEE A., IBRAHIM M.: Emotional attributes in autonomic computing systems. In *Proc. 14th International Workshop on Database and Expert Systems Applications* (2003), pp. 681–685.
- [Lin00] LINDSTROM P.: Out-of-core simplification of large polygonal models. In *Proc. ACM SIGGRAPH 2000* (2000), pp. 259–262.
- [Lin03] LINN C.: Semantic reliability in distributed systems: ontology issues and system engineering. In *Proc. the IEEE/WIC International Conference on Web Intelligence* (2003), pp. 292–300.
- [LP02] LINDSTROM P., PASCUCCI V.: Terrain simplification simplified: A general framework for view-dependent out-of-core visualization. *IEEE Transactions on Visualization and Computer Graphics* 8, 3 (2002), 239–254.
- [LRN96] LEE T.-Y., RAGHAVENDRA C. S., NICHOLAS J. B.: Image composition schemes for sort-last polygon rendering on 2d mesh multicomputers. *IEEE Transactions on Visualization and Graphics* 2, 3 (1996), 202–217.
- [LZS*03] LAMBERTI F., ZUNINO C., SANNA A., FIUME A., MANIEZZO M.: An accelerated remote graphics architecture for PDAS. In *Proc. the 8th International Conference on 3D Web Technology* (2003), pp. 55–ff.
- [MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUMMLER W., RYALL K., SEIM J., SHIEBER S.: Design galleries: a general approach to setting parameters for computer graphics and animation. In *Proc. ACM SIGGRAPH '97* (1997), pp. 389–400.
- [MB97] MICHAELS C. K., BAILEY M. J.: VizWiz: A Java applet for interactive 3D scientific visualization over the web. In *Proc. IEEE Visualization '97* (1997).
- [MC94] MACKERRAS P., CORRIE B.: Exploiting data coherence to improve parallel volume rendering. *IEEE Parallel and Distributed Technology: Systems and Applications* 2, 2 (1994), 8–16.

- [MDKH01] MEISSNER M., DOGGETT M., KAMUS U., HIRCHE J.: Accelerating volume rendering using an on-chip SRAM occupancy map. In *Proc. IEEE International Symposium on Circuits and Systems* (2001), vol. 2, pp. 757–760.
- [MH99] MORAN P. J., HENZE C.: Large field visualization with demand-driven calculation. In *Proc. IEEE Visualization '99* (1999), pp. 27–33.
- [MHS99] MEISSNER M., HOFFMANN U., STRASSER W.: Enabling classification and shading for 3d texture mapping based volume rendering using OpenGL and extensions. In *Proc. IEEE Visualization '99* (1999).
- [Min03] MINSKY N. H.: On conditions for self-healing in distributed software systems. In *Autonomic Computing Workshop – Proc. 5th International Workshop on Active Middleware Services* (2003).
- [MJ03] MCCLOY R. F., JOHN N. W.: Remote visualization of patient data in the operating theatre during hepato-pancreatic surgery. In *Computer Assisted Radiology and Surgery*. Elsevier, 2003, pp. 53–58.
- [MKDY90] MARSHALL R., KEMPF J., DYER S., YEN C.: Visualization methods and simulation steering for a 3D turbulence model for Lake Erie. *ACM SIGGRAPH Computer Graphics* 24, 2 (1990), 89–97.
- [MLM*03] MURAKI S., LUM E. B., MA K.-L., OGATA M., LIU X.: A PC cluster system for simultaneous interactive volumetric modeling and visualization. In *Proc. IEEE Symposium on Parallel and Large-Data Visualization and Graphics* (2003), pp. 95–102.
- [MMEF94] MOLNAR S., M C., ELLSWORTH D., FUCHS H.: A sorting classification of parallel rendering. *IEEE Computer Graphics and Applications* 14, 4 (July 1994), 23–32.
- [Mod04] MODVIZ INC.: <http://www.modviz.com>, January 2004.
- [Moo65] MOORE G. E.: Cramming more components onto integrated circuits. *Electronics* 38, 8 (1965).
- [MP01] MA K.-L., PARKER S.: Massively parallel software rendering for visualizing large-scale data sets. *IEEE Computer Graphics and Applications* 21, 4 (2001), 72–83.
- [MPK94] MA K., PAINTER J. S., KROGH M. F.: Parallel volume rendering using binary swap composition. *IEEE Computer Graphics and Applications* 14, 4 (1994), 59–67.
- [MrC91] MACKINLAY J. D., ROBERTSON G. G., CARD S. K.: The perspective wall: detail and context smoothly integrated. In *Proc. ACM CHI '91* (1991), pp. 172–180.
- [MT03] MORELAND K., THOMPSON D.: From cluster to wall with VTK. In *Proc. IEEE Symposium on Parallel and Large-Data Visualization and Graphics* (2003), pp. 25–32.
- [MU00] MINSKY N., UNGUREANU V.: Law-governed interaction: a coordination and control mechanism for heterogeneous distributed systems. In *ACM Transactions on Software Engineering and Methodology* (2000).
- [Mun98] MUNZNER T.: Exploring large graphs in 3d hyperbolic space. *ACM SIGGRAPH Computer Graphics* 18 (1998).
- [MvWvL99] MULDER J., VAN WIJK J., VAN LIERE R.: A survey for computational steering environments. *Future Generation Computer Systems* 15, 2 (1999).
- [NCFC04] NORRIS J., COLEMAN K., FOX A., CANDEA G.: OnCall: defeating spikes with a free-market application cluster. In *Proc. 1st International Conference on Autonomic Computing* (2004), pp. 1198–205.
- [NER04] NERSC.: <http://acts.nersc.gov/cumulvs/main.html>, 2004.
- [NLSG03] NAEF M., LAMBORAY E., STAADT O., GROSS M.: The blue-c distributed scene graph. In *Proc. IEEE Virtual Reality 2003* (March 2003), pp. 275–276.
- [NR03] NORTON A., ROCKWOOD A.: Enabling view-dependant progressive volume visualization on the grid. *IEEE Computer Graphics and Applications* 23, 2 (2003), 22–31.
- [NT96] NORMAND V., TROMP J.: Collaborative virtual environments: the COVEN project. In *Proc. the Framework for Immersive Virtual Environments Conference, FIVE'96* (December 1996).
- [OW00] OAKS S., WONG H.: *Jini in a Nutshell*. O'Reilly & Associates, 2000.
- [PDe96] PETERS T., DAVEY B., ET AL: Three dimensional multimodal image guidance for neurosurgery. *IEEE Trans. on Med. Imaging* 15 (1996), 121–128.
- [PHK*99] PFISTER H., HARDENBERGH J., KNITTEL

- J., LAUER H., SEILER L.: The VolumePro real-time ray-casting system. In *Proc. 26th Conference on Computer Graphics and Interactive Techniques* (1999), pp. 251–260.
- [Pin01] PINHANEZ C.: Augmenting reality with projected interactive displays. In *Proc. International Symposium on Virtual and Augmented Architecture* (2001).
- [PK96] PFISTER H., KAUFMAN A.: Cube-4 – a scalable architecture for real-time volume rendering. In *Proc. ACM/IEEE Symposium on Volume Rendering* (1996), pp. 47–54.
- [PKG97] PHARR M., KOLB C., GERSHBEIN R., HANRAHAN P.: Rendering complex scenes with memory-coherent ray tracing. In *Proc. ACM SIGGRAPH '97* (1997), pp. 101–108.
- [PL03] PAN S. L., LEE J.-N.: Using e-CRM for a unified view of the customer. *Communications of the ACM* 46, 4 (2003), 95–99.
- [PMS*99] PARKER S., MARTIN W., SLOAN P.-P. J., SHIRLEY P., SMITS B., HANSEN C.: Interactive ray tracing. In *Proc. Symposium on Interactive 3D Graphics* (1999), pp. 119–126.
- [PNG83] PINKHAM R., NOVAK M., GUTTAG K.: Video RAM excels at fast graphics. *Electronic Design* 31, 17 (1983), 161–182.
- [PRM00] PASCOE J., RYAN N., MORSE D.: Using while moving: HCI issues in fieldwork environments. *ACM Transactions on Computer-Human Interaction* 7, 3 (2000), 417–437.
- [PTB*02] POUPYREV I., TAN D. S., BILLINGHURST M., KATO H., REGENBRECHT H., TETSUTANI N.: A survey of augmented reality. *IEEE Computer* (March 2002), 2–9.
- [pV304] pV3. <http://raphael.mit.edu/pv3/pv3.html>, 2004.
- [PZvBG00] PFISTER H., ZWICKER M., VAN BAAR J., GROSS M.: Surfels: surface elements as rendering primitives. In *Proc. 27th Conference on Computer Graphics and Interactive Techniques* (2000), pp. 335–342.
- [RASW91] RASURE J., ARGIRO D., SAUER T., WILLIAMS C.: A visual language and software development environment for image processing. *International Journal of Imaging Systems and Technology* (1991).
- [RAV04] RAVE. <http://www.wesc.ac.uk/projects/rave>, 2004.
- [RGW*03] ROETTGER S., GUTHE S., WEISKOPF D., ERTL T., STRASSER W.: Smart hardware accelerated volume rendering. In *Proc. Joint EUROGRAPHICS/IEEE-TCVG Symposium on Visualization* (2003).
- [RPSC99] RAY H., PFISTER H., SILVER D., COOK T. A.: Ray casting architectures for volume visualization. *IEEE Transactions on Visualization and Computer Graphics* 5, 3 (1999), 210–223.
- [RS99] REKIMOTO J., SAITOH M.: Augmented surfaces: A spatially continuous work space for hybrid computing environments. In *Proc. ACM CHI '99* (1999), pp. 378–385.
- [RSEB*00] REZK-SALAMA C., ENGEL K., BAUER M., GREINER G., ERTL T.: Interactive volume rendering on standard PC graphics hardware using multi-textures and multi-stage-rasterization. In *Proc. 2000 Eurographics/SIGGRAPH Workshop on Graphics Hardware* (2000), Addison-Wesley, pp. 109–118, 147.
- [RSHT99] REZK-SALAMA C., HASTREITER P., TEITZEL C., ERTL T.: Interactive exploration of volume line integral convolution based on 3d-texture mapping. In *Proc. Visualization '99* (1999), pp. 233–240.
- [SB92] SARKAR M., BROWN M.: Graphical Fish-Eye views of graphs. In *Proc. ACM CHI '92* (1992), pp. 83–91.
- [SB03] SHALF J., BETHEL E. W.: The Grid and future visualization systems architectures. *IEEE Computer Graphics and Applications* 23, 2 (2003), 6–9.
- [SC02] SHERMAN W. R., CRAIG A. B.: *Understanding Virtual Reality: Interface, Application and Design*. Morgan Kaufmann, 2002.
- [SCM99] SHEN H.-W., CHIANG L.-J., MA K.-L.: A fast volume rendering algorithm for time-varying fields using a time-space partitioning (tsp) tree. In *Proc. IEEE Visualization '99* (1999), pp. 371–378.
- [SFH*00] SCHMALSTIEG D., FUHRMANN A., HESINA G., ARI Z., ENCARNAC L., GERVAUTZ A., PURGATHOFER W.: The studierstube augmented reality project, 2000.
- [SG93] SONG D., GOLIN E.: Fine-grain visualization algorithms in dataflow environments. In *Proc. IEEE Visualization '93* (1993), pp. 126–133.
- [SGPJ98] S. G. PARKER M. MILLER C. D. H., JOHN-

- SON, C. R.: An integrated problem solving environment: the SCIRun computational steering system. In *Proc. the 31st Hawaii International Conference on System Sciences (HICSS-31)* (January 1998), IEEE Computer Society Press.
- [SH00] SUTTON P. M., HANSEN C. D.: Accelerated isosurface extraction in time-varying fields. *IEEE Transactions on Visualization and Computer Graphics* 6, 2 (2000), 98–107.
- [Sha85] SHARP J. A.: *Data Flow Computing*. Ellis Horwood, 1985.
- [Sha92] SHARP J. A. (Ed.): *Data Flow Computing: Theory and Practice*. Ablex Publishing, 1992.
- [Shn92] SHNEIDERMAN B.: Tree visualization with treemaps: a 2D space filling approach. *ACM Transaction on Graphics* 11, 1 (1992), 92–99.
- [Sil99] SILVA L. M. E.: High performance cluster computing: Programming and applications. vol. 2. Prentice Hall, 1999.
- [Ski88] SKILLICORN D. B.: A taxonomy for computer architectures. *IEEE Computer* 21, 11 (1988), 46–57.
- [SMDW98] SMITH S., MARSH T., DUKE D., WRIGHT P.: Drowning in immersion. In *Proc. UK-VRSIG '98* (1998).
- [SML*03] STOMPEL A., MA K.-L., LUM E. B., AHRENS J., PATCHETT J.: SLIC: Scheduled linear image compositing for parallel volume rendering. In *Proc. IEEE Symposium on Parallel and Large-Data Visualization and Graphics* (October 2003).
- [SMM*97] SWAN, II J. E., MUELLER K., MÜLLER T., SHAREEF N., CRAWFIS R., YAGEL R.: An anti-aliasing technique for splatting. In *Proc. IEEE Visualization '97* (1997), pp. 197–206.
- [Sne01] SNELLING D.: UNICORE: A grid computing environment. In *Proc. Europar 2001* (August 2001).
- [Sog02] SOGANDARES F.: Stone axes and warhammers: a decade of distributed simulation in aviation research. In *Proc. the 16th Workshop on Parallel and Distributed Simulation* (2002), pp. 125–132.
- [SSC02] SLATER M., STEED A., CHRYSANTHOU Y.: *Computer Graphics and Virtual Environments*. Addison Wesley, 2002.
- [SSTR93] SARKAR M., SNIBBE S. S., TVERSKY O., REISS S. P.: Stretching the rubber sheet. In *Proc. ACM Symposium on User Interface Software and Technology* (1993).
- [Sto00] STONE R. J.: Haptic feedback: A potted history from telepresence to virtual reality. In *Proc. 1st International Workshop on Haptic Human-Computer Interaction* (2000), Brewster S., Murray-Smith R., (Eds.), vol. LNCS 2058, Springer-Verlag, pp. 1–7.
- [SUB*96] SLATER M., USOH M., BENFORD S., SNOWDON D., BROWN C., RODDEN T., SMITH G., WILBUR S.: Distributed extensible virtual reality laboratory (DEVRL). In *Proc. 3rd Eurographics Workshop on Virtual Environments* (February 1996), Goebel M., (Ed.).
- [Sun] SUN MICROSYSTEMS INC.: <http://www.sun.com/solutions/infrastructure/grid>.
- [Sun04] SUN MICROSYSTEMS INC.: <http://www.sun.com/solutions/hpc/compute/visualgrid/>, January 2004.
- [SZ00] STASKO J., ZHANG E.: Focus+context display and navigation techniques for enhancing radial space-filling hierarchy visualization. In *Proc. IEEE Information Visualization 2000* (2000), pp. 57–65.
- [Tan01] TANENBAUM A. S.: *Modern Operating Systems, 2nd Ed.* Prentice Hall, 2001.
- [TC00] TREAVETT S. M. F., CHEN M.: Pen-and-ink rendering in volume visualization. In *Proc. IEEE Visualization 2000* (2000), pp. 203–210.
- [Tea02] TEAM T. B.: *An Overview of the Blue-Gene/L Supercomputer*. Tech. Rep. 0-7695-1523-X, IBM Research and IBM Rochester and the Lawrence Livermore National Laboratory, 2002.
- [TGS] TGS: Amira. <http://www.tgs.com>.
- [The89] THEOGARIS T.: *Algorithms for Parallel Polygon Rendering*. Springer-Verlag, New York, NY, 1989.
- [The04] THE HIGH PERFORMANCE COMPUTING CENTER IN STUTTGART: Visualization - augmented reality. <http://www.hlrs.de/organization/vis/ar/>, May 2004.
- [TK98] THEISEL H., KREUSELER M.: An enhanced spring model for information visualization. *Computer Graphics Forum* 17, 3 (1998), 335–343.

- [TOIA95] TERASAWA T., OGURA S., INOUE K., AMANO H.: A cache coherency protocol for multiprocessor chip. In *Proc. 7th IEEE International Conference on Wafer Scale Integration* (1995), pp. 238–247.
- [TT03] TZIALLAS G., THEODOULIDIS B.: Building autonomic computing systems based on ontological component models. In *Proc. 14th International Workshop on Database and Expert Systems Applications* (2003), pp. 674–680.
- [TvW03] TELEA A., VAN WIJK J. J.: 3D IBFV: hardware-accelerated 3D flow visualization. In *Proc. IEEE Visualization '03* (2003), pp. 233–240.
- [UFG*89] UPSON C., FAULHABER JR. T., KAMINS D., LAIDLAW D. H., SCHLEGEL D., VROOM J., GURWITZ R., VAN DAM A.: The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications* 9, 4 (1989), 30–42.
- [Uni04] UNIVERSITY OF WALES BANGOR: The EG2003 medical prize. <http://www.informatics.bangor.ac.uk/nigel/eg2003-mp/>, January 2004.
- [USM97] UENG S.-K., SIKORSKI C., MA K.-L.: Out-of-core streamline visualization on large unstructured meshes. *IEEE Transactions on Visualization and Computer Graphics* 3, 4 (1997), 370–380.
- [VA88] VEENSTRA J., AHUJA N.: Line drawings of octree-represented objects. *ACM Transactions on Graphics* 7, 1 (1988), 61–75.
- [VBB*04] VIDAL F. P., BELLO F., BRODLIE K., JOHN N. W., GOULD D., PHILLIPS R., AVIS N.: Medical visualization and virtual environments. In *Eurographics 2004 State of the Art Reports*. 2004.
- [Vit01] VITTER J. S.: External memory algorithms and data structures: dealing with massive data. *ACM Computer Survey* 33, 2 (2001), 209–271.
- [vLMvW97] VAN LIERE R., MULDER J., VAN WIJK J.: Computational steering. *Future Generation Computer Systems* 12, 5 (1997). Also available in *High-Performance Computing and Networking*, H. Liddell, A. Colbrook, B. Hertzberger and P. Sloot (eds.), Springer-Verlag, 1996.
- [VNC] VNC. <http://www.realvnc.com/>.
- [vSTKS99] VAN STEEN M., TANENBAUM A. S., KUZ A. S., SIPS H. J.: A scalable middleware solution for advanced wire-area web services. *Distributed Systems Engineering* 7 (1999), 34–42.
- [vWvL97] VAN WIJK J. J., VAN LIERE R.: An environment for computational steering. In *Scientific Visualization: Overviews, Methodologies, and Techniques*, Nielson G. M., Müller H., Hagen H., (Eds.). 1997.
- [WBW96] WOOD J. D., BRODLIE K. W., WRIGHT H.: Visualization over the world wide web and its application to environmental data. In *Proc. IEEE Visualization '96* (San Francisco, 1996), Yagel R., Nielson G. M., (Eds.), ACM Press, pp. 81–86.
- [WBW03] WOOD J., BRODLIE K., WALTON J.: gViz – visualization and steering for the Grid. In *Proc. e-Science All Hands Meeting* (September 2003).
- [WC01] WINTER A. S., CHEN M.: vlib: a volume graphics API. In *Proc. Volume Graphics 2001* (New York, 2001), Springer Wien New York, pp. 133–147. see also <http://vg.swan.ac.uk/vlib>.
- [WCHS02] WOLF M., CAI Z., HUANG W., SCHWAN K.: Smartpointers: personalized scientific data portals in your hand. In *Proc. the 2002 ACM/IEEE Conference on Supercomputing* (2002), IEEE Computer Society Press, pp. 1–16.
- [WE98] WESTERMANN R., ERTL T.: Efficiently using graphics hardware in volume rendering applications. In *Proc. 25th Conference on Computer Graphics and Interactive Techniques* (1998), pp. 169–177.
- [Wei93] WEISER M.: Some computer science issues in ubiquitous computing. *Communications of the ACM* 36, 7 (1993), 75–84.
- [WH94] WITTENBRINK C. M., HARRINGTON M.: A scalable MIMD volume rendering algorithm. In *Proc. 8th International Parallel Processing Symposium* (April 1994), pp. 916–922.
- [Whi92] WHITMAN S.: *Multiprocessor Methods for Computer Graphics Rendering*. AK Peters, Wellesley, MA, 1992.
- [Whi96] WHITMAN S. R.: A survey of parallel algorithms for graphics and visualization. In *Proc. International Workshop on High Performance Computing for Computer Graphics and Visualisation* (1996), Chen M.,

- Townsend P., Vince J. A., (Eds.), Springer, pp. 3–22.
- [WK93] WANG S., KAUFMAN A.: Volume sampled voxelization of geometric primitives. In *Proc. IEEE Visualization '93* (1993), pp. 78–84.
- [Won99] WONG P.: Visual data mining. *IEEE Computer Graphics and Applications* 19, 5 (1999), 20–21.
- [WPLM01] WYLIE B., PAVLAKOS C., LEWIS V., MORELAND K.: Scalable rendering on pc clusters. *IEEE Computer Graphics and Applications* 21, 4 (2001), 62–69.
- [WST03] WILEN A., SCHADE J. P., THORNBURG R.: *Introduction to PCI Express: A Hardware and Software Developer's Guide*. Intel Press, 2003.
- [WTKD04] WALSH W. E., TESAURO G., KEPHART J. O., DAS R.: Utilities functions in autonomic systems. In *Proc. 1st International Conference on Autonomic Computing* (2004), pp. 70–77.
- [WvG92] WILHELMS J., VAN GELDER A.: Octrees for faster isosurface generation. *ACM Transactions on Graphics* 11, 3 (1992), 201–227.
- [WvGTG96] WILHELMS J., VAN GELDER A., TARANTINO P., GIBBS J.: Hierarchical and parallelizable direct volume rendering for irregular and multiple grids. In *Proc. IEEE Visualization '96* (1996).
- [WWB97] WOOD J., WRIGHT H., BRODLIE K.: Collaborative visualization. In *Proc. IEEE Visualization '97* (1997), pp. 253–259.
- [ZCW02] ZHOU H., CHEN M., WEBSTER M.: Comparative evaluation of visualization and experimental results using image comparison metrics. In *Proc. IEEE Visualization 2002* (Boston, MA, 2002), pp. 315–322.
- [ZQ91] ZHANG X., QIN X.: Performance prediction and evaluation of parallel processing on a NUMA multiprocessor. *IEEE Transactions on Software Engineering* 17, 10 (1991), 1059.