COMPUTER SYSTEMS

Computer Architecture

Rice University; Sarita V Adve; ***CAREER: An Integrated Approach for Improving the Performance, Programmability, and Portability of Shared Memory Multiprocessors*****; (CCR-9502500 A000); $96,991; 36 months.**

While shared-memory multiprocessors simplify many aspects of programming, it is widely believed that naive shared-momory programs are not likely to result in high performance on most shared-memory systems. The broad objective is to develop and evaluate techniques for improving the performance of shared- memory systems while simultaneously enhancing programmability and portability. A key problem that inhibited performance in shared-memory systems is that, unlike message passing programs, the exact communication requirements of shared-memory programs are not obvious. Therefore, compilers, runtime systems, and hardware typically need to make conservative assumptions, resulting in excessive communication and synchronization primitives provided on a particular machine (e.g., locks and barriers), the intended sharing behavior is obscured by these details. Furthermore, the low-level primitives that must be used to obtain good performance vary from platform to platform, thus inhibiting program portability. This research takes an integrated approach to improving the performance, and portability of systems. The approach broadly is: (1) to identify common implementation-independent communication patterns along with primitive to express those patterns in a programming language, (2) to develop heuristics that can be used by the compiler, runtime systems, hardware or some combination thereof to map these high-level primitives to the appropriate low-level primitives supported by the particular target system, and (3) to evaluate the resulting performance benefits on the system. By exposing the intended sharing behavior, many instances of unnecessary communication should be avoided, thus achieving higher performance. Furthermore, despite the larger variety of synchronization patterns that a programmer must choose from his or her task is simplified because details of implementing the patterns are left to the programming support system (note tha t the sharing patterns themselves must be understood by the programmer anyway). Finally, in combination with other projects within the overall research program, this research should demonstrate that these common, high-level synchronization patterns can be efficiently supported on a variety of platforms, thus enhancing the portability of the program. This part of the research program is focused on investigating a specific kind of target system - a hardware distributed shared- memory multiprocessor that provides software control for protocols and efficient active message support (e.g., Stanford Flash and Wisconsin Typhoon)- using execution-driven simulation. Other projects in this research program are already exploring similar issues on other types of target systems. Education: The educational objectives fall within three broad areas: curriculm development, teaching methods, and education outside the classroom. The objectives of the curriculm development part are (1) to introduce a course in parallel computing involving computational scientists and engineers along with computer scientists (both in terms of students attending and faculty teaching the course), (2) to introduce hands-on parallel programming experience as part of currently existing undergraduate computer architecture course, (3) to impart a sound education on performance analysis techniques, and (4) to impart a research experience to undergraduates by having a modest research project as part of the senior level architecture course and by actively recruiting undergraduates in the proposed research program through the Rice senior honor's thesis program. The teaching methods part of the education plan address growth as an effective teacher by exposure to research in teaching methods. The final part seeks to (1) encourage women and other under- represented groups to pursue computer science by actively interacting with such groups informally and through formal workshops, and (2) encourage an ex change between industry and academia through tutorials and short courses.

University of California – Berkeley; Alexander Aiken; ***Program Analysis: Logics, Algorithms, and Application*****; (CCR-9416973 A000); $33,492; 12 months. (Joint support with the Software Engineering Program, the Programming Languages and Compilers Program, and the Theory of Computing Program)**

This project is undertaking a systematic investigation of a general class of structural program analysis problems. The term structural program analysis refers to compile-time program analyses that deal with the types and structure of values computed by programs. This class of program analyses is broad, including within it the classical dataflow analyses of C and FORTRAN compilers, type inference for languages such as ML, and analysis of class membership of values in object-oriented languages. The primary goal of this work is the synthesis of set constraint techniques for structural program analysis in a common framework; in particular, an implementation of set constraints suitable for use in a general class of program analysis problems will be developed and distributed for use in research and teaching. There has been considerable activity directed towards using set constraints for structural analyses in recent years. Most of the work has been directed either towards studying the theory of set constraints itself, or at applications of set constraints to particular problems. Set-based analysis is the exception; set- based analysis provides a high-level characterization of a class of program analyses in terms of set constraints. Currently, however, there is no general-purpose implementation of set constraints available for experimentation with set-based (or any other class of) analyses. This system is expected to be expressive enough to encompass all set-based analyses and many analyses that are not set-based. Aside from the primary goal of developing an "off-the-shelf" implementation of set constraints that can be used in many settings, the project also pursues three secondary goals. The first goal is to develop a common technical language for structural analyses. The second goal is to describe the structure of the design space for structural program analyses using this common technical framework. The tradeoffs between accuracy and speed in structural analyses are poorly understood. By studying these problems in an abstract setting, it should be possible to make general statements about the relative efficiency and precision of various program analysis techniques that apply across a wide spectrum of applications. The third goal is to design algorithms using the above results. General algorithms are sought that would provide the best precision-speed tradeoffs for various classes of analysis problems.

University of Cincinnati; Fred S Annexstein; ***Research Initiation Award (RIA): Network Emulations for Improved Parallel Computation*****; (CCR-9309470 A002); $27,056; 12 months.**

This research concerns theoretical aspects of parallel algorithms and architectures that relate directly to the performance of “real” parallel computers. The parallel architectures of interest in this study are multipurpose arrays of identical processing elements. To fully exploit a parallel system, algorithms must be translated into efficient code that can be mapped to a target architecture. The research builds upon earlier work by formulating enhanced versions of mapping problems for these parallel architectures. The particular focus in this proposal is the study of such mapping problems with an eye toward improving the realism of the models, in order to more accurately predict and optimize computational performance. The mapping problems investigated are based on the modeling of computational structures with graphs. The graph-theoretic approach provides fundamental structural information that drives the research program. The framework of graph-embeddings is used to study the problem of programming network emulations. The merit of particular emulation-solutions are quantified in terms of measures of efficiency, including slowdown factors, processor and memory utilization, load balancing, and implementation overhead.

Rutgers University – New Brunswick; B. R Badrinath and Tomasz Imielinski; ***Designing Communication Support for Mobile Hosts in Distributed Systems*****; (CCR-9509620 A000); $45,000; 12 months.**

The rapidly expanding technology of cellular communications, wireless LAN, and satellite services will make it possible for users to access information anywhere and at anytime. In this environment, users will be mobile, use low powered computers, and communicate with the network over narrow bandwidth links. Such an environment introduces new challenges in the design of communication mechanisms due to constraints such as bandwidth, energy, cost, and the availability of computing resources due to host mobility. Hence, communication protocols should be designed for adaptability to changing environments brought by mobility. The research project will design and build support for mobile communications applications that explicitly take into account bandwidth limitations, energy limitations, and mobility using the following approaches: (1) The communication between the client and the server will be via an intermediary. Indirect interaction allows for optimization and development of specialized protocols for the wireless link and mobile host without having to change any protocol that is running on the fixed hosts. (2) Design environment aware communication primitives that will have access to the current environment state and also be notified about any changes. Thus, the type of communication to use can be tailored to the prevailing conditions.

University of Washington; Jean-Loup Baer; ***Reducing Memory Latency in High-Performance Systems*****; (CCR-9401689 A001); $95,237; 12 months.**

The focus of the research is the design and evaluation of hardware assets and software techniques to improve the performance of the memory hierarchy of high-performance systems. Processor speed has steadily improved during the last decade but memory latency and bandwith have been progressing at a slower pace. Several techniques to reduce or tolerate high memory latencies have recently surfaced. Some of these techniques are evaluated using trace-driven and instruction-level simulations of benchmarks. A major thrust of the effort is to devise improvements in the simulation techniques, e.g., by using selective sampling and parellel discrete event simulation, so that performance results can be obtained faster. The experiments in parallel simulation serve as a test bed to customize programs for parallel execution. Additional projects continuing present reserch include the design of cache coherence protocols with various block sizes for transfer and coherence. These protocols are geared towards a reduction in contention in the interconnect of shared-memory multiprocessors and a decrease in the amount of traffic caused by sharing. Similar goals are the subject of software techniques such as compiler directives for bypassing the cache and placement of data in the right levels of the memory hierchy.

University of Wisconsin – Madison; Suresh Chalasani; *Research Initiation Award (RIA): Parallel Algorithms for Scalable Multicomputers*; (CCR-9308966 A002); $30,000; 12 months.

This research explores three aspects of algorithmic techniques, partitioning schemes and fault-tolerance schemes of parallel algorithms by designing and implementing several classes of algorithms on the CM-5 parallel machine. The motivation for this project is in developing a methodology for the design and implementation of fast, efficient, and robust parallel algorithms for scalable multicomputers. In the context of algorithmic techniques, the results of the research enable the designer to choose between synchronous and asynchronous algorithms, and allow the designer to use different techniques, such as iterative and divide-and-conquer, at different levels of a parallel algorithm. The research also identifies the tradeoffs among various partitioning schemes for a given application. In the context of fault-tolerance techniques, the research shows how to exploit the fault-tolerance inherent in iterative algorithms, and devise suitable fault-tolerance schemes for applications such as solving partial differential equations. Preliminary results, obtained from implementations on the CM-5 machine, indicate that asynchronous iterative algorithms achieve speedups over their synchronous counterparts, and that several divide-and-conquer algorithms can be implemented efficiently on the CM-5 parallel machine. As part of this project, many subproblems that are specific to the CM-5 system are formulated and solved.

University of Texas – Austin; Craig M Chase; *Research Initiation Award (RIA): Efficient Algorithms for Automatic Parallel Program Decomposition*; (CCR-9409736 A001); $20,383; 12 months.

Parallel program decomposition is the process by which a single application is decomposed into multiple processes for execution on parallel hardware. For real programs, the execution speed of a good decomposition can exceed that of a poor decomposition by factors of ten or more. Unfortunately, optimal program decomposition is known to be an NP-hard problem, and efficient exact solutions are unlikely to exist. Hence, the key to finding automatic approaches to the program decomposition lies in discovering efficient algorithms that produce approximate solutions. A novel decomposition algorithm with nearly linear time complexity is explored. The algorithm initially decomposes the program into a sequence of small, parallelizable sections of code, known as partitionables. A locally optimal decomposition is selected for each partitionable based on information contained in an empirically derived database. Dynamic data redistribution is combined with a static technique called decomposition smoothing to produce a good global decomposition. The principal contributions of the research are: (1) to define new methods for program decomposition that go beyond simply minimizing communication and consider aggregate parallel execution performance, (2) to establish the extent to which fast decomposition strategies can produce acceptable solutions for practical applications, and (3) to collect and disseminate a database of performance results from kernels of scientific applications.

Princeton University; Douglas W Clark and Kai Li; *Architectural and Organizational Issues for High-Performance Uniprocessors*; (CCR-9423123 A000); $83,412; 12 months.

Caches and write buffers are becoming increasingly important to memory hierarchy designs since they are the key components for bridging the widening performance gap between microprocessors and DRAMs. This research studies the details of memory-referencing behavior in contemporary uniprocessors such as multiple referencing behavior . In particular, the research investigates the fine distinctions that characterize data references with respect to their cache and write-buffer behavior, and also studies possible architectural and organizational enhancements that will benefit writes.

Dartmouth College; Thomas H Cormen; *Research Initiation Award (RIA): Virtual Memory for Data-Parallel Computing*; (CCR-9308667 A002); $29,997; 12 months.

The project implements a virtual-memory system for data-parallel computing by taking advantage of the technical advances in algorithms for parallel disk systems. Such a system will benefit parallel computing in many scientific and other application areas by making it simpler and faster to run code on objects that are several gigabytes or even larger. Typically, when the memory requirements of an application exceed the available random-access memory on parallel machines, programmers keep their data on a parallel disk system and perform disk accesses in their source code. The resulting codes are difficult to develop and often run slowly. The proposed system will help avoid the additional software development time and complexity associated with such large problems. Moreover, because the system employs recent advances in algorithms for planning parallel disk accesses, it runs more efficiently. The project entails the design and implementation of an appropriate set of primitive data-parallel operations in a virtual-memory context. The implementation may include further development of efficient algorithms for parallel disk systems. In addition, the project includes the development of one or more compilers that compile languages such as Fortran and C into these primitives.

University of Southern California; Michel Dubois and George P Papavassilopoulos; *Asynchronous Algorithms: Scalable Algorithms for Multiprocessors*; (CCR-9222734 A002); $70,777; 12 months.

This project is concerned with the analysis, implementation and simulation of asynchronous algorithms. Asynchronous algorithms do not require synchronization and thus are particularly suitable for large-scale multiprocessors or massively parallel systems. Moreover, asynchronous algorithms are tolerant to processor and link failures and easily adapt to changes in parameters or sensor data in real-time embedded systems. The theoretical analysis of convergence considers models with stochastic delays. Asynchronous algorithms running on a multiprocessor are processes with random communication delays. With this approach complex problems such as the effects of random message delays due to variable load conditions in the interconnection network and the reliability issues related to probabilistic link or processor failures are easily analyzed. The approach is also applied to the analysis of problems in which the solution changes with time and of algorithms with good time-adaptation. Finally, the convergence of algorithms which are trapped in periodic orbits can be improved through randomization techniques. Six asynchronous algorithms with the above features are implemented on the Intel Touchstone DELTA machine accessible at Caltech. Extensive evaluation of each algorithm is performed in order to assess the effectiveness of asynchronous algorithms in the context of large-scale multiprocessor systems.

University of Wisconsin – Madison; James R Goodman; *A Study of Long-Term Architectural Limits to Computing: Galileo*; (CCR-9509589 A000); $84,682; 12 months.

Galileo is an investigation into the implications of long range technology advances across the computer field. While projections of narrow scope are difficult and risky because of the rapid progress over a very broad range of technology, it is possible to extrapolate current trends and make intelligent guesses about future computer building blocks. Extrapolating current rates of progress to as little as a decade hence suggesting that incredibly powerful computer systems might not be only possible, but even affordable. Current technology trends will undoubtedly be affected by new discoveries and by limitations long before such performance levels are reached. It is the goal of this research to identify those development directions that are likely to come into conflict or otherwise be retarded, and to investigate the implications of these effects. Future machines will undoubtedly much lower costs of all operations, but some operations will become much cheaper than others. A major goal is to identify the relative improvements on different kinds of operations, and to understand how these changes will affect architecture, the programming environment, and the algorithms used. The processor has long been the heart of a computer system, and though much has been written about other limitations, the CPU still plays a central role in today's architectures. Increasingly the processor is becoming less effective, limited not by its peak processing power, but by the ability of the system to provide work to keep it busy. While processing power is rapidly becoming cheaper, the effective use of memory and communication is becoming the critical challenge in achieving higher performance. Memory costs have declined rapidly and steadily for more than twenty years, and will continue to decline in cost per bit. Access time has not improved dramatically, however, and high-bandwidth, low latency-memory will always be relatively expensive. Limitations in memory access tile and bandwidth are already posing difficult design challenges, and this will only be aggravated by technology trends. This research is investigating how a system can get higher performance form its memory system by applying cheap processing power wherever it can be used effectively. In some sense, the system becomes a collection of memory modules with processing power distributed in ways to maximize the effective use of memory by minimizing communications between memory modules. Memory access time and memory latency are intimately related. Both are critical for a high-performance system, and either can be optimized at the expense of the other. Today memory latency seems to be the more urgent problem, and much current research is devoted to reducing or tolerating memory latency. Assuming that some of these techniques are successful, the more fundamental problem of memory bandwidth emerges. While it may be possible to build systems of arbitrarily high bandwidth, it appears that one of the most important emerging challenges is to make the most effective use of whatever bandwidth is available between modules. By today's standards, future systems will have enormous processing power, huge amounts of memory, and incredible communications bandwidth. Computer systems will be limited by memory modules of insufficient capacity and communication paths of limited bandwidth, with processor distributed generously wherever needed to optimize the effectiveness of memory and communication bandwidth. Galileo is investigating how the architecture of such systems should differ from systems of today. Galileo is investigating how the architecture of such systems should differ from systems of today, and how such systems should be programmed.

University of Washington; Edward D Lazowska, Henry Levy, and John Zahorjan; *System Support for High Performance Computing*; (CCR-9200832 A003); $408,796; 12 months. (Joint support with the New Technologies Program Program, the CISE Institutional Infrastructure Program, the Software Engineering Program, and the Operating Systems and Systems Software Program)

The fundamental issue addressed in this work is the structuring of operating systems to take advantage of multiprocessors so as to yield the best possible application performance. The division of labor between the kernel and the user level should be different on a multiprocessor than on a uniprocessor. On a uniprocessor, the kernel is responsible for address spaces, communication, resource allocation, threads, and thread scheduling. On a shared memory multiprocessor, the kernel is left in charge of address spaces and resource allocation, but communication, threads, and thread scheduling are implemented at user level through runtime support in each address space. Furthermore, each address space has direct control over certain key aspects of virtual memory management and file management. It is important to understand that while "micro-kernel" designs such as Mach 3.0 take the important step of moving the file system, the virtual memory system, etc., out of the kernel to user-level servers, they leave cross-address space communication, threads, and thread scheduling in the kernel. There are structural differences between these monolithic kernels, micro-kernel designs and the "nano-kernel" advocated here. In experimental work over the past several years, researcher constructed proof-of-concept prototypes in most of the essential areas that were addressed: locking protocols, user-level threads, cross-address space communication, the kernel interface, processor allocation policies, and user-level virtual memory management. The objectives in this work are of two kinds, direct and indirect. Direct objectives include: (1) integrating these proof-of-concept prototypes into a single system, convincingly demonstrating the viability of the overall approach; (2) solving the problems that stand in the way of this integration, some of which are known and some of which can only be discovered by undertaking the task. Indirect objectives include: (1) use the "core" of the resulting system to study a number of important policy questions (thread scheduling, processor allocation, virtual memory management) that can only be explored given such an experimental testbed; (2) use this core as a base for exploring high performance system support for more advanced architectures, specifically large scale machines with distributed memory, distributed/parallel programming using high-bandwidth low-latency networks, and wide-address machines

Cornell University; Miriam E Leeser; ***NSF Young Investigator (NYI) Award: Toolkit Development for the Design of Floating-Point Arithmetic Hardware and Software*****; (CCR-9257280 A004); $62,500; 12 months.**

The goal of this research is the development of a toolkit based on formal methods for aiding designers in developing floating point arithmetic hardware and software. A novel feature of this research is the introduction of theorem proving techniques directly into the design process. The toolkit includes a formalized version of the IEEE floating point specification and theorems that reason about that specification.

University of Minnesota – Twin Cities; Zhiyuan Li; ***CAREER: A Compiler-Assisted Approach to High Performance Memory Architectures*****; (CCR-9502541 A000); $45,000; 12 months.**

The widening gap between processor speed and memory latency will continue to be a critical issue to the performance of computer systems. The latency problem is even more acute on multiprocessors that share the same space, in which a main memory module is typically remote to all or all but one of the processors. It is therefore important to reduce unnecessary remote memory references and unnecessary network activities, and to achieve high hit ratio at the private, primary caches. The project explores coherent designs of compilers and memory architectures to improve the average memory latency, and the network traffic. The main thrust of this work is to combine compilers' ability to predict program behavior with harrdware's ability to record memory reference patterns at run time. The research focuses on three major issues: (1) The use of a compiler-assisted and hardware controlled coherence mechanism on multiprocessors with private cache memories; (2) The use of sophisticated interprocedural analysis to support the proposed coherence mechanism, via array data flow analysis and static task scheduling; and (3) The effects of program transformations, false sharing and limited directory pointers on the proposed coherence mechanism. This work covers a variety of multiprocessors that share a common address space, including bus-based systems, the Non-Uniform Access (NUMA) multiprocessors and the Cache-Only Memory Access (COMA) multiprocessors. The main objectives of the PI's education plan in this funding period are (1) to redesign the compiler and language design course sequence for the Computer Science Department of Minnesota and to solidify the interfaces of these sources with other language design sources; (2) to augment the main compiler and language design sources by seminar courses which address new issues regarding the interfaces of the compilers with the rest of the computer systems, and (3) to achieve teaching excellence in these c ourses so as to provide students solid understanding of the principles of programming language design and implementation and compiler techniques which will be important to their scientific and engineering career. The PI will conduct two experiments in this funding period. One is to develop a language design source to be taught after the compiler courses and the other is to add an operating system implementation course to the curriculum. Additional educational activities of the PI include exploration of new avenues for creating industry co-op programs for students in the computer systems area and exploration of new avenues or providing professionals with excellent access to both the basic courses and the seminar courses, e.g. through the university's Distributed Multimedia Learning Center.

Michigan State University; Philip K McKinley; ***Collective Communication in Wormhole-Routed Massively Parallel***

Computers; **(CCR-9503838 A000); $66,218; 12 months.**

The supercomputing industry is now dominated by parallel architectures. In many such systems, memory is physically distributed among processors in order to achieve better scalability, requiring nodes to communicate through a network. In order to reduce communication latency, many systems have the wormhole-routing switching strategy, in which messages are pipelined through and a single destination, or collective. Communication operations may be either point-to-point, which involves a single source and a single destination, or collective, in which more than two processes participate. Collective communication operations are particularly important in scientific computing, where large data arrays are distributed over the local memories of the nodes, but are manipulated through single operations. This research project addresses the design of collective communication operations for wormhole-routed networks. The research is intended to explore architectural and algorithmic alternatives in the design of communications subsystems for next generation wormhole-routed networks. The project can be divided approximately into four subareas. First, both models will be used to develop new algorithms for collective operations in addition to those already under investigation. Target platforms will include both direct networks, as well as wormhole-routed switch-based interconnects. Second, the resulting algorithms will be applied in order to support higher-level parallel programming constructs, such as virtual process topologies, parallel I/O operations and distributed shared memory. Third, the resulting communication services will be used to improve the performance of parallel applications, in particular parallel numerical algorithms. Finally, the use of the new models will be investigated in wormhole-routed systems that serve purposes other than parallel computing, such as high speed local area networks. The methods used will be analysis, simulation, and experimental studies on commercial wormhole-routed systems. The significance of this research is that it will provide new solutions to the communication needs of parallel applications. Besides providing a better understanding of the issues and problems that must be addressed in design and use of collective operations in wormhole-routing systems, this project will produce tangible results in the form of new protocols, prototype implementations of new algorithms, and new applications software, all of which will be made publicly available for use by other researchers.

University of Pittsburgh; Rami Melhem; *A Travel Support Grant for the Conference of Massively Parallel Processing Using Optical Interconnections (MPPOI); San Antonio, Texas; October 23-24, 1995*; **(CCR-9521486 A000); $7,680; 12 months.**

This award supported travel of nine students and two minority and women faculty members to the Second Conference on Massively Parallel Processing Using Optical Interconnection in San Antonio, Texas, October 23-24, 1995.

SRI International; Jose Meseguer; *System Level Issues for Multiparadigm Computing and SIMD and MIMD/SIMD Architectures*; **(CCR-9505960 A000); $76,497; 12 months.**

Parallel computing applications differ widely in the type of computational paradigm that is best suited to efficiently exploit their parallelism. Current software systems for parallel hardware support only single-paradigm computing. The objective of this research project is to investigate system-level support for efficient multiparadigm parallel computation. Previous work on coordination languages and multiparadigm sequential programming languages form a solid foundation for this research, although significant new ideas are required. Although current high-performance computer architectures are moving toward a combined MIMD/SIMD organization, very little software and system support exists for such hybrid computer architectures. This research will thus also focus on aspects of multiparadigm computing for massively parallel SIMD and MIMD/SIMD machines, including the Rewrite-Rule Machine as a testbed architecture. The research results will enable future systems designers to better appreciate the engineering trade-offs to efficiently support multiparadigm computing in a parallel environment. The results of this research are expected to enhance the future programming ease of high-performance massively parallel computing systems.

Southern Methodist University; Dan Moldovan; *Marker-Propagation Networks: A Quantitative Analysis*; **(CCR-9406998 A000); $75,000; 12 months.**

Marker-propagation networks (MPNs) represent a powerful computational model especially suitable for artificial intelligence and other applications where information is often uncertain and incompletely specified. MPNs refer to networks of active nodes and labeled links. The power and novelty of this model derives from the fact that the processing instruments are programmable data patterns of flexible length called markers. Node processing is activated by the arrival of markers. Marker propagations through the network are not driven by destination addresses, instead, they depend on local conditions in the nodes, network topology and the information programmed in each marker. This is a major difference between MPNs and other parallel processing models, and it is the reason why MPNs are capable of handling weakly or incompletely specified applications. Previous work in the area of marker–propagation processing assumed much simpler models. The model studied in this project operates asynchronously, and is highly programmable. Some of the critical issues related to the implementation of MPNs on existing parallel computers are possible programming paradigms, sources of parallelism, load balancing, design tradeoffs, and others. The approach taken in this project is a quantitative one, meaning that the model is implemented on several parallel computers, performance is measured and the results compared. The MPN model offers scalability and has the potential of being a preferred architecture for future supercomputers dedicated to knowledge processing applications.

University of Virginia ; Paul F Reynolds; ***Delta Cache Coherence Protocols*****; (CCR-9503143 A000); $63,000; 12 months.**

Finding a good solution to the cache coherence problem is critical to the effective use of medium and large scale multiprocessors. The high cost of remote memory access in such systems makes it important that processes be able to access data locally, but allowing processes to hold local copies of shared writable data requires finding a good way to keep the copies consistent. This project will investigate a new family of cache coherence protocols called delta cache protocols that maintain cache coherence with little or no synchronization. Delta cache protocols are directory, as opposed to snoopy, protocols. The delta cache protocols are based on isotach time systems, time systems implemented by a class of networks called isotach networks. Isotach networks are characterized not by topology - they can be imp lemented on arbitrary topologies - but by the guarantees they provide about the relative order in which operations appear to be delivered. Isotach systems require custom switching elements, but can be built with off- the-shelf processors and in a variety of configurations ranging from tightly -coupled multiprocessors to clusters of workstations. A simulation study of isotach systems shows that they outperform comparable conventional systems under realistic workloads, in many cases by an order of magnitude or more. In extreme cases, they outperform conventional systems by an arbitrary amount. The project is a three year study of the cache coherence problem with the goals of exploring the design space for cache coherence protocols opened by isotach networks; and evaluating the performance of delta cache protocols. The exploratory work includes protocol and cache design; extending the approach to a variety of topologies and to other levels of the memory hierarchy; and devising ways to reduce the network bandwidth requirements of the protocols and to improve the efficiency of the protocols through data migration. The performance study will simulate isotach and conventional systems with caches and compare their performance under a variety of workloads. The simulation will build on the already completed simulation of isotach systems without caches.

University of Massachusetts – Amherst; Arnold L Rosenberg; ***Algorithmic Support for Parallel Architectures*****; (CCR-9221785 A002); $90,000; 12 months.**

This research is based on the theme of using algorithmic techniques to enhance the performance and flexibility of parallel architectures. Specific projects span a spectrum from traditional topics, such as task scheduling and load balancing, to less familiar ones, such as using emulations to achieve architectural flexibility (e.g., reconfigurability); from foundational questions concerning the potential and the limitations of various modalities of computation and communication (e.g., the limitations of bufferless communication), to more immediately applicable studies of algorithmic techniques that endow a parallel architecture with facilities that would be prohibitively expensive to provide in hardware (e.g., a multigauging capability). The program is motivated by three observations: (1) the speed of processors in parallel architectures increases faster than the speed of the communication subsystem; (2) the potential of parallel computation is realized only when programmers are liberated from having to manage the detailed physical resources (number of processors, interconnection topology, etc.) of their machines; (3) technological developments are continually creating new modalities of parallel computation. All three of these phenomena increase the importance of developing algorithms that enhance the computing power of parallel architectures. Observation (1) suggests that we can now relegate to software chores that previously demanded the speed and concomitant cost of hardware implementation. Observation (2) points out the increasing need for algorithms that enhance the portability of parallel programs. Observation (3) suggests the need to understand emerging modalities of parallel computation, in order to exploit them effectively.

University of Washington; David H Salesin; ***Presidential Faculty Fellowships*****; (CCR-9553199 A000); $37,500; 12 months.**

This Presidential Faculty Fellow award is made in recognition of significant achievements of this principal investigator as a researcher and teacher, with potential for future contributions to this country's scientific and engineering enterprise. The research supported by this award is in the general area of image synthesis in computer graphics. In particular, the following topics will be investigated: (1) Real-time walkthroughs of complex environments (2) Controlled-precision photorealistic image synthesis, in which reliable error metrics are used to improve both speed and accuracy and (3) Computer-generated pen-and-ink illustrations.

Oregon State University; Vikram A Saletore; ***Machine Independent Parallel Execution of Numeric and Symbolic Computations*****; (CCR-9210568 A004); $5,000; 0 months.**

This research deals with scheduling algorithms for parallel execution of tasks in parallel computations. The project develops innovative scheduling schemes for parallel loops in numeric computations and dynamically spawned work in AND

and OR parallel symbolic computations. The work focuses on the design and implementation of scheduling algorithms for shared and distributed memory machines in the framework of the Chare-Kernel parallel programming system. The schemes is evaluated with a comprehensive set of benchmark programs for both numeric and AND/OR computations

University of California – Santa Barbara; Klaus Schauser; *CAREER: Efficient Communication Primitives and Realistic Models for Modern Multiprocessors*; (CCR-9502661 A000); $45,000; 12 months.

Efficient communication is central to parallel computing. Modern multiprocessors obtain high bandwidth, low latency user-level communication through special hardware support, often implemented as a dedicated communications co-processor. Future parallel machines, including high performance networks of workstations, will be built using communications co-processors. The goal of this research is to study how the presence of a dedicated co-processor affects the performance of interprocesor communication in modern parallel machines. In particular, the project focuses on two aspects. First, it studies the impact of communications co-processors on the performance of various communication primitives and determines the architectural requirements for implementing efficiently many-to-one communication such as active messages. Second, it studies how the communications so-processor can be captured by a theoretical performance model. The goal is to derive a performance model that is realistic enough to characterize machines and accurately predict performance, but still simple enough to be used in designing and analyzing parallel algorithms. To study these questions, the project takes a vertical approach, spanning theory, applications, languages, systems software, and architecture. The project constructs a complete system which allows taking an in-depth look at how efficient communication affects all of these layers. The project leverages off the work on active messages, the parallel languages Split-C, and the LogP performance model, developed at UC Berkeley. The main experimental platform is a 64-processor Meiko CS-2, which the UCSB Computer Science Department acquired through an NSF CISE infrastructure grant. The CS-2 reflects, a modern high- performance network of workstation enthronement equipped with communications co-pocessors, making it extremely well suited for the study. The principal investigator has already implemented active mess ages and Split-C on the CS-2. The project develops a realistic model for parallel computation which captures the capabilities of communication co-processors. To create a framework to characterize communication and measure its performance, the project implements a simulator based on this model as well as a Split-C profiler. This framework is used to fine-tune applications and for architectural studies. Besides answering these important research questions, the work also provides infrastructure which is used for teaching purposes as well as by other researchers. The principal investigator ismstriving for a well balanced career as a professor in computer science, combining teaching and research. His areas of interest are parallel computing, computer architecture, languages, and compilers. His educational objective is to teach students in these fields, to get them excited about computer science and engineering, and to prepare them for a successful professional career.

Northwestern University; Eric Schwabe; *Research Initiation Award (RIA): Graph Embeddings, Network Emulations, and Parallel Computer Architecture*; (CCR-9309111 A002); $25,695; 12 months. (Joint support with the Theory of Computing Program)

This project studies graph embeddings and network emulations, two areas which proved quite useful in the theoretical analysis of parallel architectures. The goals of the research include increasing understanding of the relative computational strengths of different interconnection network structures by proving new embedding and emulation results, and studying more closely the relationship between graph embeddings and other general notions of network emulation. In addition, problems in the related area of efficient mapping of parallel algorithms to parallel architectures, are studied.

Harvard University; Margo I Seltzer; *CAREER: High Performance Storage Systems*; (CCR-9502156 A000); $41,975; 12 months.

The demand for network services is growing at an exceptional rate. This growing demand requires that the necessary infrastructure be developed to facilitate rapid and easy access to vast, distributed information stores. The focus of this research is the essential base technologies required to make distributed information systems generally useful. Specifically, this research includes: high performance file system design and analysis, rotationally optimal disk scheduling, and distributed caching data. The research will result in a high-performance, distributed information server, utilizing new file system techniques, advanced disk scheduling algorithms, and caching based upon network topological location (geographical push-caching). The file system design borrows the essential characteristics of today's log-structured and journaling file systems, combining them with adaptive, rotationally optimized disk scheduling techniques. New caching techniques designed for scalability enhance the file system's feasibility by reducing user-perceived latency for distributed requests. The teaching component emphasizes curriculum development for both computer science concentrators and non-concentrators and training programs for undergraduate and graduate teaching assistants. For concentrators, the focus on undergraduate curriculum complements

Harvard's commitment to the establishment of a strong computer systems research group. Undergraduate courses have been and will continue to be modified to incorporate the fundamental concepts in an area, the classic literature, and research activities. This combination of class activities is the cornerstone for producing undergraduates who can fully participate in ongoing research projects. For non-concentrators, the introductory computer science sequence has been modified to emphasize design and problem solving skills over programming mechanics. Future developments will extend this idea still further with the development of an introductory design course. The course is intended to serve both as a common ground for computer science and engineering concentrators as well as a broad introduction for non-concentrators. The focus of the graduate curriculum development is on the application of the scientific method to systems research. Students are asked to critique published work, design experimental evaluation platforms, analyze and critique data sets, and produce quality research. The last aspect of the teaching proposal is the development of teaching assistant training seminars. Professor Seltzer has been working closely with the Derek Bok Center for Teaching and Learning, participating in their workshops and presenting panels and seminars with her teaching assistants. The orientation procedure she has developed will be enhanced so that it will be suitable for teaching assistants in all compute science classes, and potentially other fields.

Santa Clara University; Weijia Shang; *CAREER: Fundamental Problems in Parallel Programming*; (CCR-9502889 A000); $45,000; 12 months.

Parallel computers are difficult to program. Often, users have to specify the parallel execution, i.e., to specify data distributions to different local memory modules (as in High Performance Fortran), data communications between processors, and processor synchronizations. The insufficient programming support limits the applicability and further development of parallel systems. The objectives of this research are (1) to search and find solutions (optimal or feasible) for those fundamental problems in programming parallel systems such as how to distribute/replicate input data to different local memories in the system to minimize inter-processor communications, how to schedule and partition the computations to minimize the overall running time, how many processors should be used and how to choose the right message size for a message vector; (2) to develop new techniques for data alignment, scheduling partitioning and estimating the number of processors to be used for optimal execution of algorithms with nested loops; and (3) to implement these techniques and add them to some existing software tools such as PVM and MPI. Education: The education plans are as follows: (1) Undergraduate computer architecture lab at Santa Clara University has not been changed for six years and is obsolete. It is planned to update the lab and syllabi and redesign the students' projects to include recent advances in computer design, technology and debugging software to have students more motivated and prepared when they graduate. (2) Develop new parallel processing lab projects. One project is to have students run parallel programs in PVM environment so students can see and think those fundamental problems in parallel programming. (3) Be involved in mentoring women and underrepresented undergraduates and graduate students. (4) Do research projects with several senior undergraduate students, especially women engineering students to motivate them to pursue graduate studies in computer engineering at SCU

New York University; Alan R Siegel; *Managing the Memory Hierarchy: From Theory to Implementation, Application, and Practice*; (CCR-9504175 A000); $93,685; 24 months.

This project describes work in portable, high performance algorithm development, with experimental system support for I/O intensive programs. For these programs the data access time is the predominant component of the running time. The work has three thrusts. (1) Multiprocessor Multi-disk File Systems. In ongoing research, an experimental multiprocessor multi-disk file system has been built, which runs on Intel Paragon. The file system was built to support a prototype for a standard application programmer interface; the interface has been designed to facilitate the production of portable programs. A major portion of this project is to evaluate the efficiency of theoretically optimal data access algorithms that have already been designed, and to validate the paradigms used to design these algorithms for a two-level memory hierarchy. (2) Additional Two-Level Memory Model Algorithms. LAPACK is a natural (and very appropriate) package for upgrading to accommodate for out-of-core data sets. A portion of these algorithms will be redesigned and implemented as prototypes for testing and validation. Completed programs will be released for public use and evaluation. (3) Three-level Memory Model. A significant research goal is to develop a memory which captures critical performance and latency behavior between main memory, disk systems, and tape systems, for many of the contemporary computer architectures. Such a model is intended to support the development of I/O- efficient algorithms.

University of Massachusetts – Amherst; Ramesh Sitaraman; *Research Initiation Award (RIA): Communication and Fault Tolerance in Parallel Computers*; (CCR-9410077 A001); $29,987; 12 months. (Joint support with the Theory of Computing Program)

The project focuses on two important inter-related issues in the design of large scale parallel computers: (1) inter-processor communication and (2) fault-tolerance. The nature of the research is one of synthesizing theoretical and practical approaches to parallel computing. Efficient inter-processor communication is at the heart of efficient parallel computing. With some exceptions, much of the theoretical work focuses on provably optimal algorithms that are often too complex to implement. On the other hand, much of the implemented routing hardware have simpler designs but lack theoretical guarantees as to their performance. This project attempts to bridge this gap and focus on simple routing schemes that are provably good (though not always optimal). Fault tolerance is becoming ever-increasingly important in the world of parallel computers with a large number of fallible components. The study of fault-tolerant routing schemes focuses on algorithms that can route packets efficiently even in the presence of faulty switches. The two research directions represent very important but largely unchartered territory. The first concerns the question of intrinsic limitations on the fault tolerance of networks. The second direction concerns reconfiguring an increasingly important class of parallel computers known as heterogeneous parallel computers.

Harvard University; Michael D Smith; *NSF Young Investigator (NYI) Award: A System Architecture for High Performance Networking*; (CCR-9457779 A001, A002); $67,000; 12 months.

Experimental research is performed to explore and develop sophisticated compiler algorithms and novel computer architectures that meet the application-to-network needs of such data-intensive applications as video-on-demand and real-time image retrieval.

University of Wisconsin – Madison; Gurindar S Sohi; *Studies in Instruction-Level Parallel Processing*; (CCR-9303030 A003); $128,519; 12 months.

This research investigates instruction-level parallel (ILP) processors. The goal is to develop processors that can sustain over 10 instructions per cycle when executing arbitrary non-numeric applications such as compilers, simulators, CAD, text processing, and database programs. The research builds on the previous studies of ILP to expand the foundation for ILP processors. A novel paradigm for ILP processors, the multiscalar processing paradigm, continues to be developed, with the goal of bringing the multiscalar paradigm into the mainstream of processor design. Issues relevant to the operation of a multiscalar processor such as data memory systems, instruction memory systems, inter-operation communication mechanisms, and task specification and execution, are studied using detailed simulation models. Also investigated is software support for multiscalar processors.

University of Wisconsin – Madison; Gurindar S Sohi and James E Smith; *Prototyping Multiscalar Processors*; (MIP-9505853 A000); $261,670; 12 months. (Joint support with the Experimental Systems Program Program)

This project is evaluating a new architectural paradigm that can extract and exploit the parallelism in sequential code. This new approach uses both software scheduling in the compiler, as in VLIW, and hardware scheduling at run-time, as in superscalar architectures. The compiler segments code into large blocks of instructions that form subgraphs of the control-flow graph, though not necessarily basic blocks. The compiler appends synchronization information to each block that describes which registers must be shared with blocks. The compiler appends synchronization information to each block that describes which registers must be shared with other blocks. The blocks or tasks are passed to separate identical parallel execution units, each of which executes its task sequentially. Tasks are scheduled optimistically, so that some tasks may be executed by a unit when they would not be executed in a sequential system; in such cases, the tasks are "squashed," which means that their results are not written to memory or registers. At the end of each task, it blocks until the system determines that the task will not be squashed; at that time results are committed. The collection of execution units appears logically to be one unit, with a single register file. Shared register values are passed on an inter-execution-unit network, and tasks that need shared values block until the values are produced. During the first two years of this award is devoted to compiler design for the new architecture, to comparison with other architectures using simulation, and to conceptual refinement of the architecture. During later years, work is expected to focus on implementation.

Louisiana State University – Baton Rouge; Jerry Trahan and Ramachandran Vaidyanathan; *Algorithmic Scalability in Reconfigurable Bus-Based Models*; (CCR-9503882 A000); $45,231; 12 months.

For a model to be a workable component of a parallel processing system, it must be scalable. Researchers often design algorithms so that the assumed machine size for running the algorithm grows with the problem size. (This is particularly the case with algorithms for reconfigurable busbased models.) A physical machine size is, of course, fixed. For a scalable model, however, the arbitrary sized machine for running the algorithm can be mapped to the smaller physical machine that actually solves the problem without loss of efficiency. A compiler can perform this mapping, so algorithm designers and programmers can design algorithms for arbitrary size machines and need not design different algorithms or write programs for different sized problems. No reconfigurable bus-based model is known to be scalable, however, unless its computational power is significantly reduced. To better understand the extent to which these models are scalable, the investigator is less restrictive than traditional scalability. This research will investigate scalability of two main types of reconfigurable bus-based models, namely, the Reconfigurable Mesh (R-Mesh) and the Reconfigurable Multiple Bus Machine (RMBM). The RMBM comprises four versions: the B-RMBM, S-RMBM,F-RMBM, and E-RMBM. The investigators have previously established relative scalability results for the S-RMBM and for "separable" algorithms of the E-RMBM, the most powerful version. The research will seek to investigate scalability of general E- RMBM algorithms and study different scalability approaches for different ranges in the relationship of simulated to available machine sizes. Scalability questions for the R- Mesh will also be investigated. These will include relative scalability and scalability of certain classes of algorithms, topics not yet examined for the RMesh, and extensive attempts at complete scalability of the RMesh, a topic that has received only to narrow attention in the past. Scalability questions will also be studied for directional versions of the R-Mesh and RMBM, which have proven particularly useful for directed graph algorithms.

State University of New York – Stony Brook; Amitabh Varshney; *CAREER: Improvement, Simplification, and Clustering in Polygonal Datasets for Virtual Reality Applications*; (CCR-9502239 A000); $42,148; 12 months.

The objective of the research being done if to devise a set of techniques and tools that will allow improvement, simplification, and clustering of polygonal datasets within a unified framework. This is an important first step towards the long-term objective of formulating techniques to enable interactive visualization of complex virtual environments. The techniques being used to solve the above problems are al based on the notion of "rolling a sphere", formalized by the concept of a-hulls. For improving the existing polygonal datasets that have cracks and small missing polygons the intuitive idea is to roll a sphere that will fill-up these minor holes.Such flaws are common in present-day datasets and no general techniques are currently known that will handle them in a unified manner. Simplification of polygonal objects will be achieved through simplifying the genus of the objects. This provides a natural extension to, and is orthogonal to, the existing techniques for creating topology-preserving multiresolution hierarchies. The research being done will remove topological features that are less than a user-specificable size. Clustering of polygonal datasets is of considerable use in devising efficient methods for solving global illumination and visibility preprocessing for interactive walkthroughs of synthetic environments. This research will accomplish clustering of polygons in a dataset by using sufficiently large rolling spheres. The goal of the educational activities being done is to expand the current entertainment-oriented view of virtual reality by presenting it as a useful tool for serious tasks in engineering, science, and medicine. The activities include attempts to reach out to the high-school students, undergraduates, and the community at large through slide-and-video shows and hands-on demonstrations of the scientific applications of virtual environments. The activities also include the creation and development of a new graduate-level course on virtual environments. The tools and techniques developed for the research component of this project will be used to impart a hands-on education to inculcate a feel for application-oriented research and development work in virtual environments.

University of Washington; Mary K Vernon; *Parallel System Resource Management*; (HRD-9550429 A000); $155,118; 12 months. (Joint support with the Visiting Professorships for Women Program)

Effective resource management algorithms are critical to the performance of a computer system. This project will address open questions related to effective resource management in parallel computer systems, using a combination of experimental techniques, simulation, and analytical performance models. In particular the project will investigate questions concerning high performance parallel processor allocation policies, the interplay between parallel processor allocation and memory management, the interplay between parallel programming models and appropriate workloads for evaluating parallel resource management algorithms, and the impact of workload characteristics on the performance of particular parallel system resource management algorithms. The interactive component of this project involves: (1) teaching two courses, one a graduate seminar and one a core graduate or undergraduate course, (2) hosting a Northwest Workshop on Academic Careers for Women, modeled after the recent Computing Research Association (CRA) workshops of the same title, (3) participation in the mentoring and outreach efforts of the Women in Engineering program at the University Washington, and (4) joint research with Computer Science and Engineering faculty and students at the University of Washington.

University of California – Davis; Kent D Wilken; *Compiler and Architecture Support for Enhanced Program Profiling*; (CCR-9312880 A001); $5,000; 0 months.

The initial objective of this research is to develop techniques that significantly reduce the performance cost incurred by program profiling. Compiler techniques will be developed that exploit program semantics and features found in existing architectures to reduce the cost of profiling. New architecture techniques will be developed that support low-cost profiling. The high cost of existing profiling approaches precludes on-line profiling. However, on-line profiling is highly desirable because it could provide more accurate characterization of a program's behavior. A specific goal of this work is to reduce the cost of program profiling so that on-line profiling becomes feasible. On- line profiling is the enabling technology for adaptive code- optimizations, which will also be studied in this project. Code optimizations that adapt to changing patterns in program behavior can potentially provide significant performance improvement compared with the traditional static compile-time optimizations. An additional objective of this project is to develop new low- cost program profiling approaches that can more completely characterize program behavior, including profiling of multiple- basic-block traces and accurate characterization of data cache behavior.

University of Wisconsin – Madison; David A Wood; ***PYI: Memory Systems for Cost-Effective Computer Architectures*****; (CCR-9157366 A004); $75,000; 12 months.**

The goal of this research is to develop cost-effective computer architectures that take advantage of rapidly changing technologies. The thrust of this work is in memory systems, which must supply information to increasingly-faster processors. As microprocessors approach the cycle time of supercomputers, the memory system (including I/O) becomes the main difference between a personal computer and a mainframe. Current research focuses on developing new techniques for analyzing memory system performance. One new approach uses a renewal-theoretic model to improve the accuracy of time-sampled trace-driven simulations. This model promises to substantially reduce simulation time (10-fold) with only a small decrease in accuracy (less that 10%). Another new technique is cache profiling. Similar to traditional execution time profilers, a cache profiler relates a program's cache performance to its source listing. Programmers and compiler writers can use this information to optimize the cache performance of their programs, restructuring code and changing algorithms. Computer architects use this tool to better understand cache access patterns, leading to improved cache designs, synchronization mechanisms, and prefetch strategies.

New Mexico Institute of Mining and Technology; Victor Yodaiken; ***Research Initiation Award (RIA): Verification and Analysis of Operating Systems and Computer Architectures Using Moore Machines*****; (CCR-9409454 A000); $45,760; 12 months. (Joint support with the Software Engineering Program)**

This research spans two areas of computer engineering that are currently considered quite far apart. The research is aimed at problems in the design and implementation of operating systems, computer architecture and digital circuits. And the research is based on the algebraic theory of Moore machines (state machines with output). A new method for defining, composing, and reasoning about very large-scale Moore machines and their products was developed in prior work. The research will use this method to analyze and validate (or invalidate) some interesting systems and design principles - including bus protocols, fault tolerance and scheduling mechanisms, and asynchronous circuit design rules. The research goals are to: clarify understanding of the examples. Clarify or develop design principles for reliable systems, and investigate the basic mathematical properties of real-time sensitive and concurrent state transition systems.

Fault Tolerance

Yale University; Richard Beigel and James Aspnes; ***Parallel Fault Diagnosis*****; (CCR-9415410 A000); $60,000; 12 months. (Joint support with the Theory of Computing Program)**

Computer systems of the future may consist of thousands or even millions of processors. Inevitably some of the processors will fail. Faulty processors need to be repaired so that the system will continue functioning.. This project will investigate improvements to the existing algorithms for parallel fault diagnosis. Second, and more importantly, the project will develop dynamic models for fault diagnosis that are more realistic than previous models.

University of California – Irvine; Douglas M Blough; ***Design and Evaluation of System-Level Fault Tolerance Mechanisms for Parallel Computers*****; (CCR-9318495 A001); $70,736; 12 months.**

Recent developments have dramatically increased the importance of ultra-reliable parallel computing. First, parallel processing has moved from a specialized scientific computation tool to an instrument for general-purpose computing. At the same time, computers have begun to control virtually every aspect of human life and are, therefore, increasingly relied upon to provide continuous service. These trends make the development of general but practical techniques for fault-tolerant multiprocessing a necessity. This project focuses on loosely-coupled parallel computers. These systems are ideally suited for ultra-reliable applications due to their built-in hardware redundancy that can be used to achieve fault tolerance. The project has both basic and applied research components. The basic research component studies new models and algorithms for the problems of multiprocessor system fault diagnosis and fault-tolerant routing. The applied research component involves the development of an experimental testbed for fault-tolerant multicomputer systems. A Transputer-based MIMD multicomputer system provides the testbed hardware and low-level operating system. Special-purpose software is developed to provide a system-level fault tolerance framework, a fault simulator, and a data collection tool for the testbed. The testbed allows experimental evaluation of the system-level fault tolerance mechanisms developed in the basic research component and is also made available to other research groups working on multiprocessor system fault tolerance.

California Institute of Technology; Jehoshua Bruck; ***NSF Young Investigator (NYI) Award: Efficient Fault-Tolerant Parallel and Distributed Computing*****; (CCR-9457811 A001); $100,000; 12 months.**

This research program addresses a number of fundamental issues related to the design of novel algorithms, protocols and architectures that enable efficient fault-tolerant parallel and distributed computing for scientific and commercial applications. The program is a blend of basic research and experimental systems activities. It includes work on the understanding of new computing paradigms inspired by neural networks as well as designing and prototyping of novel communication schemes for parallel and distributed systems.

Arizona State University; Partha Dasgupta; ***Collaborative Research: High Performance Parallel Processing: Fault Tolerant Computing on a Network of Workstations*****; (CCR-9505519 A000); $65,888; 12 months.**

This project deals with both theoretical and prototyping research in parallel processing to harness the power of workstation clusters. The primary techniques will based on aggressive scheduling, evasive memory and dispersed data management. The formal work expands the previous results of the PIs in asynchronous computing in several directions: (1) Extend the models to models of fine-grained programs. (2) Permit the inclusion of architecture specific characteristics into theoretical frameworks. (3) Techniques to distinguish inconsistencies in memory states as viewed by different processors resulting from time-outs, interrupts and failures (or indefinite postponement of a computation). (4) Techniques to allow non-deterministic executions. (5) Model input/output operations of fault-tolerant programs. The experimental part of the project will build a scaled version to execute parallel programs on clusters of workstations. The prototype will address issues related to the eager scheduling of threads, the implementation of evasive memory and data management. The prototype will attempt to evaluate the effectiveness of the environment in terms of efficiency. This project is conducted jointly by P. Dasgupta of Arizona State University and Z. Kedem, of New York University.

Carnegie Mellon University; Ronald Bianchini, Jr.; ***NSF Young Investigator (NYI) Award: Fault Tolerance in Distributed Systems*****; (CCR-9257973 A003); $62,500; 12 months.**

The first area addressed in this project is fault-tolerance in distributed systems with primary focus on continuous systems operation and collaboration. Collaboration is achieved when multiple distributed entities contribute toward a common goal. For collaboration to be successful in a distributed system, the goal must be achieved in a potentially unstable network environment. The second area is the investigation of telecommunication switching architectures. Specifically, the design and deployment of high-speed multimedia networks.

New York University; Zvi M Kedem; ***Collaborative Research: High Performance Parallel Processing: Fault-Tolerant Computing on a Network of Workstations*****; (CCR-9411590 A000); $53,469; 12 months. (Joint support with the Experimental Systems Program Program)**

This project deals with both theoretical and prototyping research in parallel processing to harness the power of workstation clusters. The primary techniques will based on aggressive scheduling, evasive memory and dispersed data management. The formal work expands the previous results of the PIS in asynchronous computing in several directions: (1) Extend the models to models to fine-grained programs. (2) Permit the inclusion of architecture specific characteristics into theoretical frameworks. (3) Techniques to distinguish inconsistencies in memory states as viewed by different processors resulting from time-outs, interrupts and failures (or indefinite postponement of a computation). (4) Techniques to allow non-deterministic executions. (5) Model input/output operations of fault-tolerant programs. Thee experimental part of the project will build a scaled version to execute parallel programs on clusters of workstations. The prototype will address issues related to the eager scheduling of threads, the implementation of evasive memory and data management. The prototype will attempt to evaluate the effectiveness of the environment in terms of efficiency. This project is conducted jointly by P. Dasgupta of Arizona State University and Z. Kedem, of New York University.

Carnegie Mellon University; Bruce M Maggs; ***NSF Young Investigator (NYI) Award: Routing Protocols and Mechanisms for Tolerating Failures in High-Speed Networks*****; (CCR-9457766 A001); $47,762; 12 months. (Joint support with the Theory of Computing Program)**

This project studies protocols for admitting and routing calls in distributed networks, and for tolerating failures in these networks. Its goal is to understand how to best use high-speed networking technology, such as ATM, to support parallel supercomputing, and to advance the NII (National Information Infrastructure) initiative. In the supercomputing domain, it seeks to develop network structures, routing algorithms, and mechanisms for tolerating faults that are both practical and provably effective. In the realm of the NII, it is analyzing the behavior of simple algorithms, including greedy algorithms, for admitting calls in wide-area integrated-service networks.

Cornell University; Sam Toueg; ***Unreliable Failure Detectors for Reliable Distributed Systems*****; (CCR-9402896 A000); $90,000; 12 months. (Joint support with the Operating Systems and Systems Software Program and the Theory of Computing Program)**

The starting point for this research is a fundamental problem in fault-tolerant distributed computing: reaching agreement among processes in a system that is subject to failures. It is well- known that this problem, called Consensus, has no deterministic solution in asynchronous systems, even it is assumed that communication is reliable and no more than one process may fail. The impossibility of solving Consensus (and other related problems such as Atomic Broadcast) is one of the most severe obstacles to implementing fault-tolerant applications in asynchronous systems. In recent work the PI has introduced a novel approach to circumvent such impossibility results: he showed that unreliable failure detectors can be used to solve Consensus (and Atomic Broadcast), even if the information that they provide about failures is highly inaccurate, e.g., even if they make an infinite number of mistakes. Since such failure detectors can be implemented in realistic distributed systems, and since various considerations make the asynchronous models especially attractive, this work suggests an approach to fault-tolerance that is viable in practice. The objectives of this research are to broaden the applicability of this approach by removing the limitations of the earlier work, and to explore in more concrete terms its practicability. Specific goals include: tolerating communication failures, including network partitions (the earlier work assumed reliable links); tolerating process failures of various types (the earlier work dealt with crash failures only); considering shared-memory systems (the earlier work dealt with message-passing systems); and solving other problems that are central to fault-tolerant distributed computing, including Group Membership and Group Multicasts (the earlier work solved Consensus and Atomic Broadcast). In order to assess the cost and benefit of using unreliable failure detectors complexity questions will also be explored. Finally, the practicality if this approach will be validated by implementation on an experimental platform.

Performance Evaluation

Georgia Institute of Technology; Ian F Akyildiz; ***End to End Modeling, Analysis and Parallel Simulation of Integrated Network Architectures*****; (CCR-9404726 A000); $145,834; 24 months.**

The research focuses on an important, so-far, unresolved issue of end-to-end modeling, analysis, and parallel simulation of Broadband Integrated Service Digital Network (B-ISDN) architectures. ATM (Asynchronous Transfer Mode) is the multiplexing and switching techniques for the B- ISDN architecture. The end-to-end models to predict performance measures such as delay and throughput in order to design an accurate admission control protocol for ATM networks are unsatisfactory. They are either computationally too expensive (And thus impossible to perform on-line in real- time) or too simplistic and inaccurate. There is considerable room for improvement in accuracy, without prohibitive complexity. Also the concept of "Quality Service" (QoS) has been trivialized in most analytical results to be just the cell loss probability. Consequently, results, e.g., in call admission and call policing, are primarily judged by their effectiveness concerning cell loss probability. The introduction of additional measures (e.g., delay and delay jitter bounds) will further restrict the set of possible admissible traffic mixtures, and as a result, the often-used linear call admission method (and the related equivalent bandwidth definition) becomes insufficient. Hence, other call admission techniques may very well be needed. The effectiveness of the call admission algorithms can be investigated while utilizing results from end-to-end models. There is a need to exploit the special control features that can be provided by an ATM network, such as the allocation of buffers to connections in intermediate nodes, or the re- negotiation of the traffic descriptor or of the QoS during the lifetime of a connection. In this project an analytical method based on Pade Approximation is used to analyze the end-to-end model for ATM networks. Also a fast parallel simulation approach is used as a complementary technique. Note that the methods can also be used as general purpose performance evaluation tools and can be applied to a larger set of modeling problems.

Fisk University; Amy W Apon; ***Research Planning Grant (RPG): Performance Measurement and Modeling of Parallel Computer Systems*****; (CCR-9503116 A000); $0; 18 months.**

The purpose of this Research Planning Grant is to perform preliminary research and to develop a full proposal which will be submitted to NSF in the fall of 1995. The primary objective of the research is a measurement and modeling study of parallel computer systems. The platforms to be used include the Kendall Square KSR1 and the Intel Paragon multiprocessor systems located at Oak Ridge National Laboratory, and a network of Sun Solaris 2.3 workstations located at Fisk University, an HBCU/MI. The method of evaluation is to measure and evaluate various system design alternatives using alternative queueing network models of parallel computer systems. Queueing network models are chosen because, unlike other methods such as Markov models and Petri net models, queueing network models do not require the enumeration of all possible states that the system may enter, but rather, use computationally efficient techniques to determine mean performance metrics of the system. The use of parallel and distributed computer and communications systems is becoming commonplace. However, the computer systems which are the focus of this research have characteristics which are difficult to measure and model. The systems may have forking and joining in the task graph, variations in the multiprogramming level and use various synchronization techniques. System issues such as thread scheduling and placement are difficult to model accurately. Thus, the development and application of computationally efficient techniques for evaluating the performance of these types of systems is an important and significant area of research. The project includes a component for undergraduate research experiences at Fisk University. Given appropriate guidance, undergraduate have the opportunity to alter parameters, run experiments, and participate in the measuring and modeling process.

University of Arizona; Mary L Bailey; ***Efficient Simulation of Multicomputer Programs*****; (CCR-9212018 A002); $3,750; 0 months.**

Given a multicomputer program, its input and target architecture, we want to predict the synchronization strategy resulting in the fastest simulation of the program. Simulators are often present in environments facilitating parallel program design. In addition, they are used to compare the performance of different parallel architectures on a given program. In both instances, speeding up the simulator's execution decreases development time of parallel programs. The major emphasis of this research is on comparing different synchronization strategies for multicomputer simulations. The project will proceed in two stages. First, a simple multicomputer simulation paradigm having a few communication primitives and a simple timing model will be used. Methods for predicting the execution times for this paradigm using both analytical and empirical techniques will be developed. Together with sequential simulation, a variety of parallel synchronization strategies, such as synchronous, time-driven, asynchronous, and hybrid strategies will be considered. Second, more complex multicomputer simulators, including additional communication primitives and more complex timing models will be investigated.

Towson State University; Darush Davani; ***Optimization of Discrete Event Simulation Models: Single-Run-Based Approaches*****; (CCR-9505731 A000); $19,448; 12 months. (Joint support with the Numeric, Symbolic & Geometric Computation Program and the Theory of Computing Program)**

The aim of this project is to obtain new results which will help reduce the time and computational effort required when using the computer simulation models for optimization found in many discrete-event systems. The focus is on the enhancement of the Score Function and Likelihood Ratio methods to provide accurate estimates for derivative and perturbation, respectively. The goals are improvements the existing variance reduction techniques, theoretical unification, some extensions, and advancements in the practical implementation of these "single run" methods. The project will develop procedures and algorithms, investigate their accuracy versus execution time behavior, and present numerical results for a wide variety of queuing, inventory, and reliability models.

University of Massachusetts – Amherst; Wei-Bo Gong; ***Rational Approximation for Some Computationally Complex Performance Analysis Problems*****; (CCR-9506764 A000); $51,197; 12 months.**

Performance index functions with integer parameters are important in the evaluation and analysis of computer systems, communication networks and general distributed and parallel processing systems. These functions are often difficult to evaluate for large integers. On the other hand, most of these performance functions have nice properties such as monotonicity, convexity, analyticity, as well as obtainable asymptotic behavior. This research uses such properties to extrapolate the performance functions based on their values at small integers. The objective of this research is to further establish the theoretical foundation and to develop guidelines, algorithms and their software implementations that will be applicable to a wide array of systems. The results of this research will make it possible to express the performance measure function in terms of the system size of dimension as a low-degree rational function for some general classes of systems. The successful completion of this research will provide a framework in which various difficult modeling, control and optimization issues can be attacked.

North Dakota State University – Fargo; Rajendra S Katti; ***Research Initiation Award (RIA): A Tool for the Performance Evaluation of Parallel Computers*****; (CCR-9308099 A002); $29,427; 12 months.**

A tool to aid in the design of parallel computers by computing their performance is developed. Currently queueing networks or stochastic Petri nets are the main methods used for computing the performance indices of parallel systems. A method that uses probabilistic automata to model parallel computers will make the tool easy to use, comprehensive and efficient. Probabilistic automata are a modular way of modeling parallel systems. The modeler is relieved from dealing with a global model of a parallel system as the system is divided into modules that are dealt with separately. Problems such as the design of architectures for certain problem domains and the mapping of algorithms to architectures are studied with the aid of such a tool.

University of California – Davis; Charles Martel; *Designing Parallel Algorithms Using Modeling, Simulation and Experiments*; (CCR-9403651 A001); $62,165; 12 months.

Parallel computers have the potential to significantly outperform current computer systems. In order to exploit this potential, techniques for creating and evaluating efficient parallel programs are needed. Current parallel computers have significant overheads and latencies for messages and this will continue for at least several years. Thus it is essential that programs for parallel computers be designed with these overheads in mind. Parallel computers are complex systems, so it is often difficult to predict the performance of programs based only on analytic models. Thus, it is important to evaluate parallel programs using both simulations and implementations. Networks of workstations are already widely available and they are likely to become even more prevalent. They provide an easy way to integrate the latest workstation technology into a parallel system, and they provide a cost effective parallel computer. Thus programs developed for a network of workstations have the potential to have a significant impact. This research focuses on developing and evaluating programs for networks of workstations. It results in effective algorithms for both basic problems such as load balancing, median finding, and simple graph algorithms, as well as for more complex applications such as modeling air or fluid flows. It also results in improved network support for algorithms on networks of workstations. As indicated, the performance of these programs is evaluated using simulations, so that the performance can be predicted for a range of settings, and by measurements of actual implications. A second focus of this research is to model cache-coherent shared memory systems, and to develop algorithms with good locality to help improve performance. Since algorithms for networks of workstations often exploit data locality, some of the algorithms developed for the network of workstations settings will also be effective on cache-coherent shared memory machines. Simulations and tests help determine which of the network algorithms perform well in this setting.

Princeton University; Margaret R Martonosi; *CAREER: An Integrated Hardware and Software Performance Monitoring System*; (CCR-9502516 A000); $45,000; 12 months.

Over the past 15 years, processor speeds in modern computers have increased at a much faster rate than main memory speeds. Thus, relative memory latencies in both sequential and parallel computers have become quite large, and continue to increase rapidly. Such trends form the motivation for this research, which attacks the issues of how to monitor, analyze and improve computer performance in the face of a large processor-memory performance gap. Because of this performance gap, it becomes increasingly important to develop performance monitoring systems that allow programmers, compiler writers, and system designers to identify and tune portions of their design where memory overhead is limiting performance. To accomplish these goals, this project focuses on several area: (i) exploring hybrids of compile-time and run-time software-based techniques for application memory performance analysis, (ii) combining these hybrid monitoring techniques with dynamic compilation to implement flexible, low-overhead, detailed performance tools, (iii) designing and implementing novel, software-aware, hardware performance monitoring techniques, and (iv) designing software tools based on this novel hardware support. By studying both hardware and software solutions concurrently, one gains insights into the strong and weak points of each approach. The end result of this research is a suite of performance tools that span a range in the level of monitoring detail provided and the level of hardware support required.

Carnegie Mellon University; John P Shen; *Analytical Model of Superscalar Processor Performance*; (CCR-9423272 A000); $239,999; 36 months. (Joint support with the Microelectronic Systems Architecture Program Program)

Trace-driven simulators (TDS) are widely used by processor architects to accurately assess the expected performance of a processor design. While necessary and useful, TDS have some shortcomings. Traces of realistic programs are usually very large and expensive to store, and simulation of traces can very time consuming. Furthermore, simulation does not yield a characterization of the factors that determine performance, and hence can not recommend specific changes to enhance performance. The two-fold objectives of this research are (1) to investigate alternatives to TDS for accurately quantifying processor performance; and (2) to seek to characterize factors that determine superscalar processor performance. The effort will involve developing an analytical model of machine behavior, based on which processor performance can be computed without requiring tedious trace-driven simulation. It is hoped that the same analytical model, in addition to providing a useful tool, will also provide a characterization of superscalar processor performance. This characterization can potentially reveal fundamental laws governing processor performance that can be mathematically formulated. In this research project, processor performance is modeled as a composition of program parallelism and machine parallelism. Program parallelism is represented by two distribution functions that are derived from the execution trace of benchmark and captures the control and data dependencies in the program. This derivation is analogous to performing a type of transform on the execution trace from the time domain to the transformed domain. Consequently the two distribution functions can be viewed as concise spectral representation of the program parallelism. Machine parallelism represents the resources available in the processor for the concurrent processing of instructions. A machine is modelled as a network, or circuit, of primitive components, e.g. buffers and scalar pipelines. Macro models of the components and the model for their serial and parallel interconnections will be developed. With such a circuit of model microarchitectures, complex superscalar processors can be effectively and rigorously modeled. While the research appears somewhat theoretical in nature, a great deal of experimentation is required and expected. The plan is to work with realistic programs with extremely long traces as well as a wide range of state-of-the-art implementations of real instruction set architectures. The Alpha 21164 and the Power PC 620 designs and the SPEC92 benchmark suites are used as the primary experimentation vehicles.

Northwestern University; Valerie Taylor; ***NSF Young Investigator (NYI) Award: A Tool for the Improvement of Uniprocessor Performance*****; (CCR-9357781 A002); $67,500; 12 months.**

The goal of this research is to investigate and develop techniques needed to produce a tool that allows the applications programmer and computer designer to work to increase uniprocessor performance. Computational units, such as the floating-point units, are becoming faster and readily available because of the rapid advances in technology and computer-aided design (CAD) tools. The performance, however, is limited by the rate at which data is applied to the units. To achieve good uniprocessor performance, the computational units, memory, and memory ports must be efficiently utilized. This may require that the applications programmer redesign the algorithms to better exploit the unique features of the architecture or the computer designer utilize CAD tools to design special-purpose processors that exploit the unique features of the algorithms or for best results, both. This research investigates and develops: (1) techniques for extracting statistics on the utilization of the aforementioned processor resources, and (2) methods for displaying the performance statistics to benefit both the applications programmer and the computer designer. The results of this research permit the applications programmer and computer designer to interact closely to address the problem of efficiently solving challenging problems such as climate modeling, vehicle dynamics, semiconductor modeling, and quantum chromodynamics, which are beyond the scope of existing computers.

Computer Graphics, Visualization and Computer-Aided Geometric Design

Purdue University; Chandrajit Bajaj; ***Geometric and Solid Modeling with Algebraic Surfaces*****; (CCR-9222467 A003); $78,101; 12 months. (Joint support with the Numeric, Symbolic & Geometric Computation Program)**

The approaches in computational geometry to date have been largely theoretical and deal with discrete and combinatorial structures such as point sets, lines, polygons, polyhedra and at most general arrangement of hyperplanes. This project considers the design and implementation of efficient algorithms for an arbitrary arrangement of algebraic curves and surfaces and specifically compact solids bounded by curved algebraic surface patches. The class of algebraic surfaces, a subset of arbitrary analytic surfaces, provides enough generality to accurately model almost all complex rigid objects. Additionally, algebraic curves and surfaces have compact storage representation and form a class which is complete under all common operations required by a geometric modeling system. Finally, representing physical objects with algebraic surfaces encompasses all prior modeling approaches. The main research thrusts of this project are: (1) the design of efficient algorithms for basic operations on algebraic surfaces (these manipulations include intersection, parameterization, Ck-interpolation, and approximation), and (2) the design of efficient algorithms for basic manipulations on solids bounded by algebraic surface patches (these manipulations include: boolean set operations, convex hull, decomposition, convolution, offset, blending, and smoothing operations).

Carnegie Mellon University; David Baraff; ***Research Initiation Award (RIA): Non-Interpenetrating Rigid Body Simulation*****; (CCR-9308353 A002); $22,514; 12 months.**

The ability to realistically simulate physical objects' motion behaviors has become an area of great importance in science, engineering, training and planning. Realistic simulation requires accurately modeling the behavior and interactions of contacting objects. In particular, it is necessary to be able to simulate the same non-interpenetrating motion that real- world objects undergo. An important subset of physical simulation is simulations where all objects are modeled as being perfectly rigid. The abstraction of a perfectly rigid body, though physically unobtainable, is an excellent and reasonable approximation for many real-world applications in robotics and computer graphics. Even with the rigid-body assumption, simulation is made difficult by the need to deal with complex surface geometries and contact phenomena such as friction. This project aims at extending the capabilities of current physical simulation systems that model contact and collision between systems of rigid bodies and explore practical uses for simulation systems. Specifically, the investigation centers on the development of new geometrical algorithms that allow for simulations with greater levels of geometric complexity, and the extension of computational complexity bounds and algorithms for the problem of contact with friction. This project will also explore using non-interpenetration constraints to aid in interactive computer-aided design and modeling applications.

University of California – Berkeley; Brian A Barsky; ***SGER: Developing New Geometric Modeling and Scientific Visualization Techniques for Curved Optical Surfaces*****; (CCR-9520703 A000); $50,000; 24 months.**

The proposed research is suitable for an SGER, due to its interdisciplinary nature, involving applications of geometric modeling and scientific visualization to provide improved representation and display techniques for the cornea. The application has the potential of revolutionizing the field of optometry. Instead of the current trial-and-error method of fitting contact lenses a much faster and precise procedure can be developed. Furthermore, it will enable fitting contact lenses to people whose corneas have complex and/or irregular shapes, something which is impossible with the current state of the art. A successful completion of the project may induce the manufacturers of contact lenses to embark on a development effort, but at present they don't have the financial incentive to take the risk. Thus the small investment in this project can have significant impact. The PI is highly qualified and is strongly motivated, due to a personal vision problem. I recommend funding.

University of Washington; Anthony DeRose; ***Subdivision Methods for Computer-Aided Design and Engineering*****; (CCR-9502694 A000); $47,250; 12 months.**

Representing curved shape is a fundamental problem in geometric modeling and design. A variety of approaches currently exist for modeling curved shape, by far the most popular of which is a boundary representation (B-rep) whose surfaces are trimmed NURBS (non-uniform rational B-splines). Unfortunately, building a modeler based on trimmed NURBS is exceedingly difficult for at least two reasons. First, trimming two NURBS patches to agree along a common edge usually involves complex surface/surface intersection algorithms and substantial approximation. Second, even in the absence of sharp edges, patching is often required since a single NURBS patch is capable of modeling only surfaces topologically equivalent to planes, cylinders, and tori. Modeling surfaces of other topological types (even those as simple as a sphere) require the surface be decomposed into a quiltwork of individual NURBS patches. Complicated continuity conditions must then be imposed to guarantee that adjacent patches join smoothly. Subdivision provides a promising alternative to NURBS. Using subdivision, a polyhedron is mapped to a sequence of denser polyhedra whose limit is a curved object. Each new polyhedron is derived from the previous polyhedron by splitting each of its faces into four new faces and positioning its new vertices using a fixed averaging mask. Different masks yield surface features such as corners, curved edges, or smooth faces in the limit object. The method is simple, requiring only polyhedral modeling. The method is general, producing objects with smooth faces and sharp curved edges without trimming or patching. Finally, subdivision provides a hierarchical representation that allows the design of efficient multi-resolution processing algorithms. This research addresses a combination of theoretical and applied problems whose solution is essential if subdivision is to be of use as a practical method for modeling curved shape. Specifically it develops: Curvature continuous subdivision methods; Subdivision methods for local refinement and introduction of shapes and edges; Functional subdivision methods over irregular triangulations (necessary for FEA) and algorithms for: Conversion to and from other representations such as NURBS, Editing and display of subdividion surfaces, Finite element analysis over subdivision surfaces, Computing the intersection, union, or difference of subdivision surfaces.

Stanford University; Patrick M Hanrahan; ***Efficient Light Transport Algorithms for Computer Graphics*****; (CCR-9508579 A000); $85,962; 12 months.**

Many realistic image synthesis algorithms have been devised, but unfortunately they tend to be very inefficient, often requiring hours or days of calculation. As a result, the visual cues provided by these algorithms (shading, shadows, transparency, reflections, etc.) still have not been used routinely by applications requiring visualization, primarily because they are not supported in real-time on current generation graphics workstations. However, both increased performance and greater functionality can be expected in future generation workstation. This research project develops efficient high quality image synthesis algorithms. The project will attack several specific areas. Continuing to develop the theory behind rendering so that it has a sound basis in physics and mathematics continuing the development of fast hierarchial algorithms based on wavelets for solving the rendering equation. Investigating various hybrid algorithms that combine the best features of the finite element and Monte Carlo approaches. Finally investigating various issues that result from the need to handle very large, complex environements. The algorithms will be implemented in an experimental image synthesis system that is being developed and they will be tested on a variety of different scenes and environments.

Washington State University; John C Hart; ***Research Initiation Award (RIA): Modeling, Rendering and Animation of Implicit Surfaces*****; (CCR-9309210 A002, A003); $35,266; 12 months.**

Implicit surfaces are shapes described by the points where a function is zero. Sphere tracing is a recently developed method for rendering implicit surfaces that operates on implicit functions that indicate the distance to the implicit surface. Such implicit functions are developed for common modeling primitives, CSG operations, some blends, and some domain transformations which deform objects, but many other surfaces, such as offset surfaces, and other operations, blends and deformation remain. Sphere tracing is a robust efficient implicit surface rendering algorithm. It ray traces implicit surfaces by tracking along each ray by the distance to the implicit surface, as reported by the implicit function. This iteration converges on the first ray-surface intersection. Sphere tracing also provides an area sampling method for antialiasing. Enhancements to sphere tracing include the sphere buffer which allows sphere tracing to capitalize on spatial coherence, spatial partition, and a method for eliminating all iteration from implicit surface rendering. Also addressed are the problems of animating implicit surfaces, particularly the consistent texturing of an implicit surface as it deforms, possibly changing genus.

Carnegie Mellon University; Paul S Heckbert; *NSF Young Investigator (NYI) Award: Simulating Global Illumination*; (CCR-9357763 A002); $25,000; 12 months.

Synthesizing realistic images is a fundamental challenge for computer graphics. One of the most difficult aspects of this task is the simulation of global illumination: the interreflection of light between surfaces. This research applies finite element methods to the simulation of global illumination. The work is divided into three subprojects: (1) error metrics for global illumination, (2) adaptive meshing for radiosity, and (3) generalization to scenes with arbitrary reflectance. Rather than attempting to measure "realism", which is a rather nebulous perceptual quality, simpler, quantitative accuracy measures are being developed. Using these error measures, existing and future image synthesis algorithms can be evaluated and quantitatively ranked. Adaptive meshes for radiosity mesh generation is only partially automated in many radiosity systems, and little is known about the errors introduced by various meshes. Recent experiments on simplified, two-dimensional scenes show that the use of adaptive meshes and higher degree approximations can lead to much faster and/or more accurate solutions than those of existing algorithms. More accurate, better looking pictures result when the radiosity mesh follows discontinuities in the solution function such as those at shadow boundaries. These "discontinuity meshing" techniques have been embraced by researchers. In this research, meshes for radiosity are compared and evaluated using error metrics as an objective criterion. This rigorous analysis of radiosity functions suggests new mesh generation techniques allowing more accurate simulations. Generalization to scenes with arbitrary reflectance, the final sub-project, is the generalization of these methods to scenes with realistic light scattering. The reflectance of real surfaces is not purely diffuse or specular, but is described by a complex bidirectional reflectance distribution function (BRDF). Recent work suggests that the use of spherical harmonics or other appropriate basis functions greatly facilitates the approximation of direction-dependent light functions. Such representations for functions of both position and direction are being analyzed and developed. The ultimate goal is the development of accurate, efficient algorithms for simulating global illumination on complex scenes with arbitrary reflectance and transmittance.

State University of New York – Stony Brook; Arie E Kaufman; *Volume Synthesis Theory and Algorithms*; (CCR-9205047 A003); $106,914; 12 months.

This project deals with volume synthesis, its theoretical framework, its underlying voxelization algorithms for synthesizing volumetric (voxel-based) objects from their geometric description, and its discrete rendering algorithms for visualizing these objects. The three main objectives of this research are: (1) to study the theoretical framework in 3D discrete topology for the voxelization and discrete rendering algorithms in order to guarantee the generation of a "well-behaved" discrete representation in terms of connectivity, thickness, separability, minimization, coverage, and accuracy; (2) to develop and implement efficient voxelization algorithms for 3D lines, 3D circles, surfaces, and solids, using a variety of connectivities, thicknesses, and spatial occupancies, with a specific focus on both binary and non- binary ("fuzzy") voxelization; and (3) to develop and implement discrete algorithms that capitalize on the discrete representation of the volume buffer for fast viewing of the scene, using approaches such as volumetric ray tracing and template-based projection.

Stanford University; Marc S Levoy; *PYI: Computer Graphics to Visualize Scientific and Medical Data*; (CCR-9157767 A004); $62,500; 12 months.

This project focuses on techniques for photorealistic visualization of sampled scalar, vector, and tensor fields of three spatial dimensions — i.e. volume data. A key requirement of the photorealistic approach is real-time image generation. The goal is to provide tools of sufficient power and polish that scientists and physicians, not their assistants, will use them. Since no commercial workstation is capable of rendering usefully large volume datasets in real-time, the research plan includes a significant architectural component.

University of Washington; David H Salesin; *NSF Young Investigator (NYI) Award: Studies in Image Synthesis*; (CCR-9357790 A002); $62,500; 12 months.

Photorealistic image synthesis is an important problem for such diverse fields as architecture, lighting and industrial design, remote sensing from satellites and robots, scientific visualization, simulation systems, and entertainment and advertising. It also requires enormous computation. Carefully chosen approximations are therefore essential for making the computation efficient. Unfortunately, few such approximations admit good error metrics, which makes controlling image fidelity difficult at best. In many applications this accuracy is vital. This research develops a controlled-precision rendering approach for general systems of light transport both to provide reliable bounds on accuracy, and also to allow an algorithm to concentrate its effort where the error reduction is largest. In addition to photorealistic image synthesis, research begins in non-photorealistic rendering, an area with great commercial potential that has received relatively little attention to date. Mechanisms are developed to support the generation of effective illustrations from three-dimensional continuous-tone images. Specific research goals include creating a taxonomy of traditional illustration techniques; developing new modeling paradigms, imaging models, and user-interaction methods for specifying and computing illustrations; establishing methods for simplifying geometric models from models with a higher level of detail; and developing techniques for simulating traditional artist tools, such as pencil, chalk, air brush, and watercolor.

Indiana University – Bloomington; Peter Shirley; ***Realistic Computer Rendering of Complex Scenes Using Quantitative Perceptual Error Metrics*****; (CCR-9401961 A000); $43,520; 12 months.**

This research is concerned with the design of computer graphics that produces realistic images of environments that are complex with respect to geometry, reflection behavior, and light transport. This design will be pursued with the additional goals of real-time performance on general purpose hardware that will exist within one to two decades, and compatibility with display algorithms based on human perception.

University of Washington; Werner Stuetzle, Anthony DeRose, Thomas E Duchamp, and John A McDonald; ***3D Scanning: From Physical Objects to Electronic Models*****; (DMS-9402734 A001); $120,017; 12 months. (Joint support with the Integration Engineering Program, the Computational Mathematics Program, and the Statistics Program)**

The goal of this project is to develop methods that will enable rapid, automatic, and inexpensive "3D scanning"; that is, the use of optical scanning technology to create electronic models from physical objects such as a clay model of a car, or a turbine blade. 3D scanning is at present an important technology, used in a number of critical US industries, such as automobile and aerospace production. However, the expense of current scanner hardware and the need for extensive human intervention has prevented 3D scanning from achieving its full potential. The methods developed by our interdisciplinary research team will make 3D scanning far less expensive, faster, and more automatic. The goal of 3D scanning is the inverse of computer aided manufacturing: given a physical object, such as a clay model of a car, a turbine blade, a chair, or a house, create an electronic model that captures its shape and color. 3D scanning is at present an important technology in a number of critical US industries, such as aerospace and automotive production. However, the expense of current scanner hardware and the need for extensive human intervention in both data collection and modeling has prevented 3D scanning from achieving its full potential. This grant will result in new algorithms and software that will enable new scanner systems that are faster, more automatic, smaller, less expensive, model other visual properties in addition to shape, and produce higher quality and more usable output.

Rice University; Joe Warren; ***Subdivision Methods for Computer-Aided Design and Engineering*****; (CCR-9500572 A000); $47,725; 12 months.**

Representing curved shape is a fundamental problem in geometric modeling and design. A variety of approaches currently exist for modeling curved shape, by far the most popular of which is a boundary representation (B-rep) whose surfaces are trimmed NURBS (non-uniform rational B-splines). Unfortunately, building a modeler based on trimmed NURBS is exceedingly difficult for at least two reasons. First, trimming two NURBS patches to agree along a common edge usually involves complex surface/surface intersection algorithms and substantial approximation. Second, even in the absence of sharp edges, patching is often required since a single NURBS patch is capable of modeling only surfaces topologically equivalent to planes, cylinders, and tori. Modeling surfaces of other topological types (even those as simple as a sphere) require the surface be decomposed into a quiltwork of individual NURBS patches. Complicated continually conditions must then be imposed to guarantee that adjacent patches join smoothly. Subdivision provides a promising alternative to NURBS. Using subdivision, a polyhedron is mapped to a sequence of denser polyhedra whose limit is a curved object. Each new polyhedron is derived from the previous polyhedron by splitting each of its faces into four new faces and positioning its new vertices using a fixed averaging mask. Different masks yield surface features such as corners, curved edges, or smooth faces in the limit object. The method is simple, requiring only polyhedral modeling. The method is general, producing objects with smooth faces and sharp curved edges without trimming or patching. Finally, subdivision provides a hierarchical representation that allows the design of efficient multi-resolution processing algorithms. This research addresses a combination of theoretical and applied problems whose solution is essential if subdivision is to be of use as a practical method for modeling curved shape. Specifically it develops: Curvature continuous subdivision methods; Subdivision methods for local refinement and introduction of shapes and edges; Functional subdivision methods over irregular triangulations (necessary for FEA) and algorithms for: Conversion to and from other representations such as NURBS, Editing and display of subdividion surfaces, Finite element analysis over subdivision surfaces, Computing the intersection, union, or difference of subdivision surfaces.

University of California – Santa Cruz; Jane Wilhelms and Allen Van Gelder; *Hierarchical Exploration and Visualization of Large and Complex Multi-Dimensional Data Sets*; (CCR-9503829 A000); $110,897; 12 months.

This project investigates the use of hierarchical methods for accessing and visualizing large and complex multi- dimensional data sets. Visualization can provide an efficient and powerful means of exploring such data However, most visualization methods are developed for moderately sized non-hierarchial data sets that can be imaged acceptably fast on graphics workstations. In practice, there are many data sets that cannot be effectively explored using these methods. The project is developing and implementing visualization approches designed for very large, multi-dimensional data sets over a variety of grid types. Such data sets present significant problems in the management of time and space resources, which will be addressed by this research. The unifying theme will be the use of hierarchical data structures to manage both resources effectively. Addressing time efficiency, techniques are being developed for the selective traversal and rapid interactive imaging of the data. For space efficiency, various data models, such as wavelets, are investigated to achieve compression, while being directly amendable to visualization. These hierarchies are being integrated with existing visualization techniques for isosurface extraction, direct volume rendering, and flow visualization, many of which have already been developed at UC Santa Cruz and elsewhere for non-hierarchial data. Techniques for rectilinear, structured, and unstructured grids are studied. Parallelization of these methods of SIMD and MIMD machines is also is being explored. This research is developing new algorithms for visualization of large, complex data sets, as well as making software systems available for use by scientists. Software is modular to facilitate integration with existing visualization packages. While the main objective is visualization, the hierarchical methods developed will also be of use for any data analysis hampered by access problems in dealing with very large data sets.

Miscellaneous

State University of New York – Buffalo; Jin-yi Cai; *Uncheatable Benchmarks*; (CCR-9319393 A002); $64,346; 12 months.

Benchmarks have been used to test everything from the speed of a processor to the access time of a memory system. The computing community relies on them heavily to assess how fast a given hardware or software system operates. They are of fundamental importance in everyday computing. However, up until now, the study of the art of designing a good benchmark has focused on making the benchmark "realistic" in predicting how well it will perform for the intended applications; the issue of making benchmark results trustworthy has been relegated to "trusted" or third party agents, and little attention has been paid to the question of making benchmarks themselves "uncheatable." The project studies the problem of how to make benchmarks resistant to tampering and hence more trustworthy. The plan is to use modern cryptography and complexity theory to make this possible. Basically the trust in individuals and organizations is replaced by trust in the impossibility of breaking certain computational problems. The advantage of this new approach is clear: the new schemes will be scientifically trustworthy. There is a whole array of problems that need to be formulated and solved, which are important in making uncheatable benchmarks; some are practical issues, some are theoretical.

Princeton University; Richard J Lipton; *Uncheatable Benchmarks*; (CCR-9304718 A002); $96,968; 12 months.

Benchmarks have been used to test everything from the speed of a processor to the access time of a memory system. The computing community relies on them to assess how fast a given hardware or software system operates. They are of importance in everyday computing. However, until now, the study of the art of designing a good benchmark has focused on making the benchmark "realistic" in predicting how well it will perform for the intended applications; the issue of making benchmark results trustworthy has been relegated to "trusted" or third party agents, and little attention has been paid to the question of making benchmarks themselves "uncheatable." The project studies the problem of how to make benchmarks resistant to tampering and hence more trustworthy. The plan is to use modern cryptography and complexity theory to make this possible. Basically the trust in individuals and organizations will be replaced by trust in the impossibility of breaking certain computational problems. The advantage of this approach is clear: the schemes will be scientifically trustworthy. There is an array of problems that need to be formulated and solved, which are important in making uncheatable benchmarks; some are practical issues, some are theoretical.

California Institute of Technology; Paul Messina; ***Parrallel I/O Methodologies for I/O-Intensive Grand Challenge Applications*****; (CDA-9318145 A002); $608,000; 12 months. (Joint support with the Networking and Communications Research Program Program, the New Technologies Program Program, the CISE Institutional Infrastructure Program, and the Computational Engineering Program)**

This multi-disciplinary project investigates and develops strategies for efficient implementation of I/O intensive applications in computational science and engineering. Scalable parallel I/O approaches are pursued by a team of computer scientists and applications scientists who work together to: (1) characterize the I/O behavior of specific application programs running on large massively parallel computers, (2) abstract and define I/O models (templates), and (3) implement and test application-level I/O tools on large-scale computations. The Pablo performance analysis environment provides the foundation for the performance instrumentation and analysis. The application programs are already fully operational on advanced architecture systems. The principal computers used are the Intel Touchstone Delta and Paragon systems at Caltech, each with over 500 computational nodes. Five application areas are included: fluid dynamics, chemistry, astronomy, neuroscience, and modelling of materials-processing plasmas. The parallel programs for these applications cover a range of I/O patterns and volume, and the techniques that are developed in this project are of relevance to a broad spectrum of engineering and science applications. In addition, by overcoming their current I/O limitations, the specific applications targeted in this award achieve significant new science and engineering results. The project devises and implements general methods for scalable I/O using today's advanced computers, immediately applies those methods to carry out unprecedented applications in several fields, and uses the methods and experience gained in the first half of the project to tackle the I/O issues in future sustained teraflops computers. This project is carried out by a team of researchers who have many years of experience in using parallel computers for large-scale applications, in measuring and characterizing the behavior of applications on such computers, and in creating methods and tools that facilitate the use of such systems. Building on this considerable experience, the project concentrates initially on exploring the I/O behavior of existing applications and on devising ways to get adequate and scalable I/O performance. The five application areas included in this award cover a wide range of I/O needs and behavior, including the use of out-of-core methods, reading data from high-speed instruments, interactive and irregular data flows in modelling, and pipelining of data from distributed computing resources on high-speed networks. A part of the approach is an integration effort that concentrates on ensuring close communication and synergy among the components of the project and on integrating the results. The objective of the project is to develop methods and insights that are as generally applicable as possible. The award is supported by the Advanced Projects Research Agency as well as NSF programs in engineering and computer science.

University of Southern California; Ulrich Neumann; ***CAREER: Closed-Loop Augmented Environment Systems*****; (CCR-9502830 A000); $41,222; 12 months.**

Research Component: Closing the loop in AE systems. The project addresses augmented environment (AE) and virtual environment (VE) tracking system inadequacies through a new "unified-system" approach to AE system design. Even with the best head-mounted displays (HMDs) and graphic systems available, the shortcomings of tracking systems hinder or preclude the usefulness of many AE/VE applications. Such applications, including on-line surgical visualization, assembly or maintenance assistance, and interactive training simulations, require robust solutions to the following problems. 1) Achieving millimeter (position) and 0.1 degree (orientation) accuracy. 2) Readily achieving precise calibration of tracking systems to the real world. 3) Eliminating (or minimizing) latency in tracking (and graphics) systems to prevent gross misalignment during head motion. This project aims to develop new technology and improved methods for making video-based AE systems usable and productive for commercial and scientific applications. The long term research goal is the accurate alignment of real and virtual objects in an augmented environment under all conditions of user activity including rapid head motion. Achieving such performance required the integration of new image and tracking function components into an AE system. A 2D alignment error metric forms a basis for the compensation of a variety of 3D tracking and calibration errors. Corrections are computed and applied at every frame creating a closed loop system which compensated for both static and dynamic (latency-induced) misalignment. The project will produce increased understanding of closed loop AR systems by 1) implementing a software simulator to facilitate rapid exploration of closed loop AE system design options for feature detection error measurement, and error compensation, 2) designing a new AE system architecture with components which have simple, well-defined interfaces to other co mpenents in the system, and 3) developing a prototype of a complete closed-loop AE system and quantifying its performance characteristics. Education Component: Richer understanding through experience. Teaching responsibilities include at least one course during each Fall and Spring semester. A two-course Computer Graphics sequence for graduate students is under development. These courses build on the introductory undergraduate course and emphasize project creation. Hands on system-building learning opportunities, at the appropriate levels for students, are a major focus of the educational activities. In addition to syllabus design and class teaching, directed-research and advising promotes student-faculty interactions. A CD-ROM with computer graphics simulations of graphics concepts will be developed for both introductory and advanced levels.

Harvard University; Michael O Rabin; ***Asynchronous Parallel Computing, Fault Tolerance, and Randomization*****; (CCR-9313775 A001); $106,340; 12 months.**

An important class of practical parallel computers consists of machines comprising many large grained processors. These processors execute asynchronously and memory is assumed to be distributed. Trying to execute a general n-thread parallel program P on such a parallel system poses theoretical and practical challenges. One must provide mechanisms for implementing the synchronization barriers, for load sharing and balancing amongst the processors, and for program executing that respects memory locality. The approach used in this research is to define algorithms and a program transformation (compilation) T so that T(P) executes correctly and efficiently on any P. This simulated execution does not require any specialized hardware or operating system mechanisms in the target machine. Of particular interest is the case that the parallel program is large-grained, i.e. that in each parallel step a substantial block of instructions is executed on some or most of the parallel threads. It is assumed that the parallel program is already optimized so as to make the blocks of code within the parallel step quite large. This project studies the most general case of large grained parallel programs, paying particular attention to memory locality and to memory fault-tolerance, employing randomization. A further goal is to achieve Las Vegas style results, i.e. small expected running time and no errors.

Duke University; Kishor S Trivedi; ***Industry/University Cooperative Research Center for Advanced Computing and Communication*****; (EEC-9418765 A001); $209,364; 12 months. (Joint support with the Industry/University Cooperative Research Centers Program, the Human Resources Development Program, the NASA-CACC Program, and the NSA-CCSP Program)**

The North Carolina State University Industry/University Cooperative Research Center for Advanced Computing and Communication is a multi-university Center with the addition of a research site at Duke University. Four companies are part of the Center and Duke University is a research site. The research agenda of the multi-site Center addresses high-speed networking, reliable and fault-tolerant systems, digital communication systems, distributed systems, and algorithms and digital image processing algorithms. Additional funding to augment the ongoing research program and to obtain a membership in the Center is being provided by the U.S. Department of Army and National Security Agency.