BM/C^3 INFORMATION TECHNOLOGY DISTRIBUTED PROCESSING AND INFORMATION WARFARE

J. L. Hayes
US Army Space & Strategic Defense Command
Huntsville, AL 35807-3801

Ira W. Merritt
US Army Space & Strategic Defense Command
Huntsville, AL 35807-3801

James C. Hayes
COLSA Corporation
Huntsville, AL 35806

John K. McFee, Jr.
The MITRE Corporation
Huntsville, AL 35806

Jon Sauer
University of Colorado
Boulder, CO 80309-0525

Abstract

The US Army Space and Strategic Defense Command (USASSDC) Advanced Technology Directorate (ATD) currently manages several research programs that have the potential to significantly advance the current state of the art in information technology for future Battle Management/Command, Control, and Communication (BM/C^3) systems. These programs address some of the challenges associated with "full spectrum dominance" in information warfare by providing new and innovative technologies for advanced distributed processing. The definition of information technology as it applies to BM/C^3 is provided, as well as our vision for the future of distributed processing and its role in future BM/C^3 systems. We propose that the realization of more effective BM/C^3 systems utilizing megacomputer architectures to support the human-in-control will require continuing technological advances in high speed communications, architectural structures, automated decision support, modeling and simulation (M&S), and parallel processing algorithms. The current research in optimistic computing, and photonic interprocessor routing and switching, and the applicability of this research to distributed BM/C^3 is discussed. Finally, the future research plans, including the application to the BMDO's development of a Virtual Distributed Hardware-in-the-Loop (HWIL) Test Bed (VDHTB), are described.

1.0 Introduction

The modern warfighter is called upon to rapidly formulate concepts, plan operations, make decisions, and win battles on a dynamic, high tempo, extremely complex battlefield. To accomplish this mission effectively, the warfighter must rely on widely distributed (sometimes mobile) interconnected computer systems to acquire, manage, process, and display the vast amounts of information on which he must base time-critical decisions. BM/C^3 information technology must address the challenges associated with utilizing multi-sensor, multi-source, multi-media data, information, and knowledge to perform planning, control, coordination, monitoring, assessment, communication, and security in mission-oriented tactical and strategic environments. The USASSDC ATD's Innovative Science and Technology (IS&T) and Small Business Innovative Research (SBIR) programs have the potential to address these challenges and significantly advance the current state of the art in BM/C^3 information technology. The ATD's current information technology research programs address the challenges associated with distributed information systems by providing new and innovative concepts for advanced distributed processing, high speed communications, advanced computer security, automated decision support, and software engineering technologies.

2.0 Information Technology for Distributed BM/C^3

Achieving and maintaining situational dominance on the battlefield depends upon first achieving and sustaining information dominance.[1] Information dominance of the future battlefield will become increasingly difficult and will require that battlefield commanders at all levels be capable of readily visualizing the current tactical or strategic situation, and rapidly ascertaining the outcome of alternative courses of action (COA) within the enemies' "Observation, Orientation, Decision, and Action" (OODA) cycle.[1] Whether the objective is efficient troop movement, effective artillery placement, or ballistic missile defense, success ultimately depends on commanders evaluating information and making decisions better and faster than the enemy. LTG Wilson A. Shoffner, Commander, Combined Arms Command, Fort Leavenworth, KS, states in [2] that on the future battlefield "There will be less time to make decisions, and because the battlefield is more lethal, decisions will have a greater consequence."

The functionality and capabilities required to achieve information dominance on the battlefield will reside primarily within a BM/C^3 system. Battle management is the full-time automated process of analysis, planning, organizing, direction, coordination, and control over sensor devices and weapons; it reflects policy, rules of engagement and operational doctrine; and, is always subject to human control and override.[3] This final attribute ensures that in any BM/C^3 system, the ultimate decision maker will always be the human-in-control. Command and control in support of BM is the process by which the decision maker selects among competing options in order to achieve strategic and tactical objectives; communications is the means by which his command decisions are made and executed.[4]

The ultimate decision maker in any BM/C^3 system, regardless of the level of automation, will always be the human-in-control, and any technology effort directed at improving the quality (timeliness, correctness, clarity, appropriateness, comprehensibility) of the data on which he must base his command decisions will ultimately improve the quality of those decisions resulting in increased warfighter effectiveness. As a consequence, the real threat to future military success appears to lie in the possibility that the human-in-control may not be able to acquire, process, and utilize information as quickly and effectively as his enemy.

To address this potential threat, future BM/C^3 systems must be increasingly capable of faster acquisition and processing of widely distributed multi-sensor, multi-source, multi-media data, information, and knowledge to perform BM/C^3 functions in mission-oriented tactical and strategic environments. BM/C^3 information technology research must address the information processing challenges associated with distributed information systems by providing new and innovative concepts in advanced distributed processing, high speed communications, and automated decision support.

Based on past evolutionary trends, it is conceivable that current technological advances in the BM/C^3 system infrastructure areas of processors, communications, and software will be in a state of dramatic decline around the year 2012.[5] This belief carries with it the implication that the military of the future may not be able to achieve and maintain information dominance of the battlefield unless it begins now to move away from BM/C^3 systems based on autonomous computational resources and actively pursues alternative strategies. In response to this hypothesis, the ATD envisions future BM/C^3 systems being comprised of distributed virtual supercomputers or "megacomputers" [5] made up of thousands, or perhaps millions, of widely, possibly globally distributed, uncoordinated, collaborative processors. These megacomputers may ultimately be capable of PetaFLOPS performance, a million billion operations per second, which exceeds the combined computing power of all the computers now on the Internet.[6] The achievement of these goals, depicted notionally in Figure 1, will depend upon radical advances in device technology, architectural structures, and parallel processing algorithms. Assuming that sufficient technological advances are made, and those advances lead to system level program decisions, advanced BM/C^3 concepts based on these technologies could be a reality over the next two decades.

Future US information dominance of the battlefield is too important to rely solely on anticipated advances in commercial-off-the-shelf (COTS) technology, but instead requires the intelligent augmentation of COTS technology with government sponsored research. The ATD's goal is to sponsor information technology research that can form the basis for future distributed information systems, of the type discussed above, which will function effectively in dynamic battlefield environments. Continuing technological progress in advanced distributed processing, high speed communications, advanced computer security, automated decision support, M&S, and software engineering technologies are imperative, and will ultimately lead to more advanced warfighting capabilities.

The current research in three technology areas: optimistic computing, photonic interprocessor routing and switching, and information warfare, should ultimately provide essential infrastructure to support the realization of megacomputer technology and enable future BM/C^3 system effectiveness. The following sections describe proposed applicability of these advanced technologies to achieving information

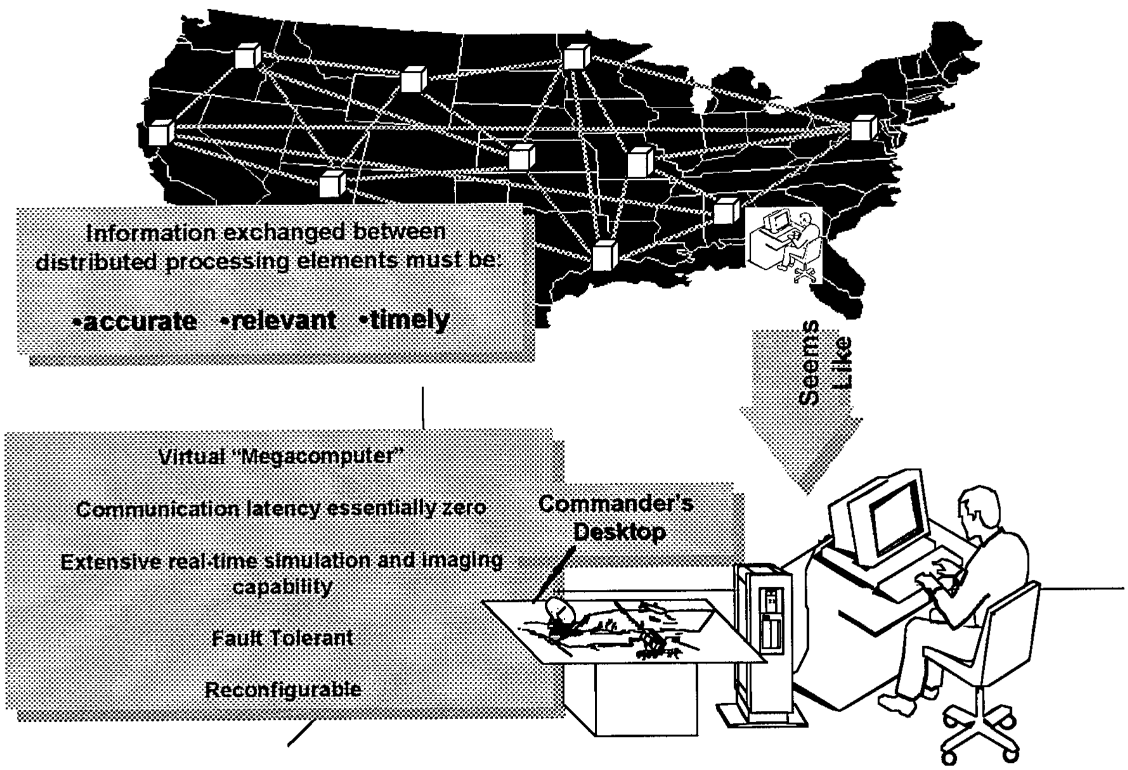

Figure 1, Widely Distributed Information Processing

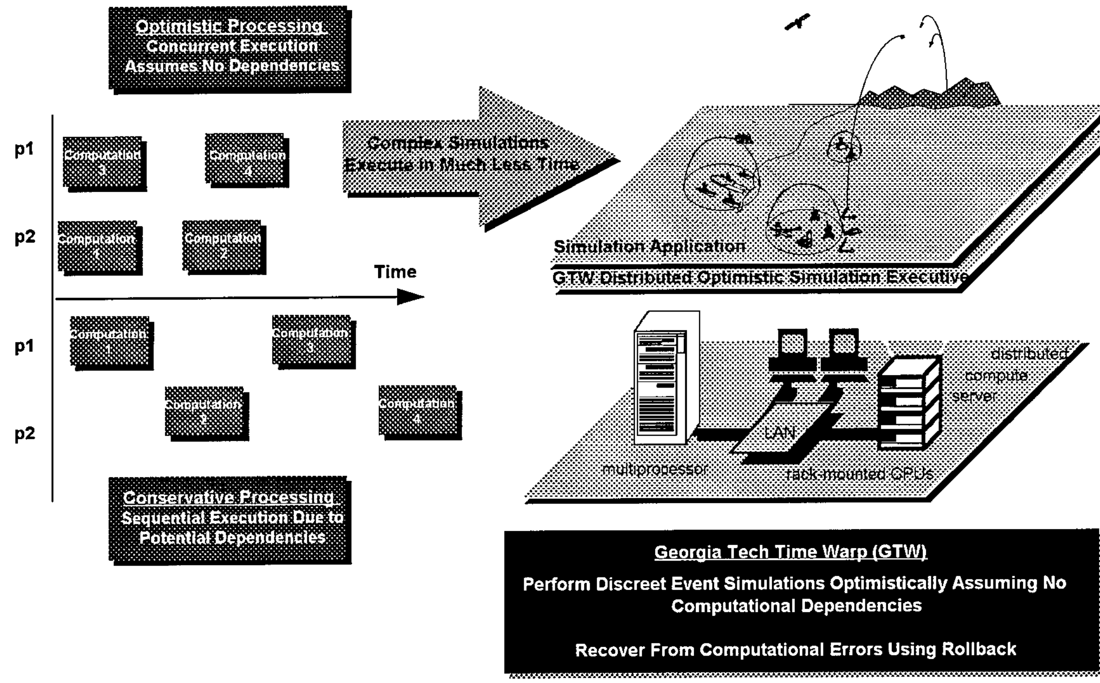

Figure 2, Georgia Tech Time Warp (GTW)

dominance through information technology for BM/C^3 systems.

3.0 Advanced Technology Directorate Research Efforts

3.1 Optimistic Computing

3.1.1 Overview of Research

As complex systems of the type described above become a reality, performance becomes a critical factor; as a result, efficient application algorithms and exploitation of parallel execution will play critical roles in determining the ultimate performance increases that can be achieved.[6] In this light, the ATD is currently managing a research program with Georgia Tech (sponsored by BMDO/IS&T under contract DASG60-95-C-0103) which is investigating the development of a high performance optimistic simulation executive to simplify the development and speed up the execution of discrete event simulations on parallel or distributed COTS hardware.

Research into the optimum mechanisms for synchronizing parallel/distributed discrete event simulations has to date progressed along two distinct paths, optimistic (rollback-based) and conservative (blocking-based).[7] It is believed that optimistic execution offers fundamental advantages over conservative approaches including greater concurrency, and better transparency of the synchronization mechanism, both of which reduce the effort associated with simulation development. Conservative synchronization mechanisms strictly avoid the possibility of causality errors by not allowing the processing of an event until there is no possibility that another event with an earlier time-stamp will be received. In contrast, optimistic synchronization allows events to be processed without regard to the temporal relationships of the processed events, correcting causality errors when they are detected.

Georgia Tech Time Warp (GTW) is an optimistic simulation executive which is an implementation of the "virtual time" paradigm originally proposed by David Jefferson in [8] called the Time Warp mechanism. A generalization of the GTW mechanism and its potential application are shown in Figure 2. The GTW realization of this mechanism seeks primarily to minimize the significant event processing overheads associated with the original Time Warp proposal.[9] Information regarding the specifics of the implementation and utilization of GTW is contained in [10]. A major application area for this technology has been, and continues to be, the synchronization of distributed/parallel discrete event simulations. To accomplish this, Time Warp systems, including GTW, rely on some implementation of "lookahead (peering into the future) - rollback" as the fundamental synchronization mechanism for distributed simulation objects, or as they are called in GTW, logical processes (LPs). This approach allows a distributed LP to execute based upon its own local virtual clock without regard to synchronization conflicts with other processes, i.e., optimistically. A typical GTW program is comprised of collections of these distributed LPs that communicate with one another by exchanging time-stamped events or messages. If an LP receives a message from another process in its past, i.e., the time-stamp of the message is less than the current local virtual time, a conflict arises. These unavoidable conflicts are resolved by "rolling back" the offending process(es) to the virtual time just before the conflict occurred and canceling all intermediate side effects. This all serves to dramatically simplify the development of parallel discrete event simulations by virtually decoupling the logical processes and freeing the developer from process synchronization concerns. Another advantage of optimistic synchronization techniques are potentially shorter simulation execution times than can be achieved, as compared to sequential or other parallel discrete event simulation approaches. Performance improvements, i.e., speedup, associated with allowing processes to execute optimistically will generally outweigh the costs associated with occasionally "rolling back". Assuming fixed rollback costs, Lipton and Mizell have shown the Time Warp can outperform the Chandy-Misra conservative approach to performing parallel discrete event simulation by an arbitrary amount. That is to say that given n processors, Time Warp can outperform Chandy-Misra by a factor of n. Lipton and Mizell further show that the opposite is "not" true: Chandy-Misra can only outperform Time Warp by a constant factor.[11]

It has been shown that the potential speedup of simulations built on the GTW executive is directly related to the number of processors, the number of events that must be eventually rolled back, processor idle time, and the overhead associated with parallel execution.[7][12] This relationship can be expressed as:

Speedup = $n(1-(R+I+O))$, where,

- n = number of processors,
- R = fraction of events "rolled back",
- I = fraction of time each processor is idle,
- O = additional overhead for parallel execution.

It has been shown in [7] that the fraction of events rolled back (R) and processor idle time (I) will decrease as the degree of parallelism and associated message density increase. As message density approaches infinity, the number of events rolled back will approach zero, and the potential speedup of GTW-based simulations will approach the number (n) of processors onto which the simulation is distributed. GTW innovations to the original Time Warp mechanism have included adaptive optimistic synchronization protocols, fast recovery from synchronization errors (direct message cancellation), efficient memory reclamation (on-the-fly fossil collection), optimistic I/O, and load balancing algorithms enabling background execution, all of which combine to reduce the time associated with rollback and message cancellation, further increasing the speed of execution. Speedups approaching the number of processors utilized, as high as 38 times faster using 42 processors, have been demonstrated experimentally and described in [9].

This research is currently exploring two advanced computing technologies related to GTW. The first is optimistic parallel processing techniques that entail executing computations based on possibly incomplete information. Current work is focused on developing methodologies to parallelize existing sequential simulation software. While an extensive body of work exists concerning the parallelization of scientific codes, new techniques are required for irregular discrete event simulations. Parallelization of sequential discrete event simulations has not been widely studied. The second technology being explored in this area is concerned with developing novel, interactive simulation techniques to provide analysts with much more sophisticated means of interacting with their simulation tools. Rather than the traditional run-analyze-run cycle, analysts will be able to interactively manipulate the simulations to rapidly evaluate the impact of key command and control decisions. Techniques that enable the analyst to (1) execute the simulation backwards (reverse execution) in order to evaluate the underlying causes of certain simulated events (e.g., to understand the reasons a simulated mission failed to achieve its objectives) and (2) dynamically "clone" [13] or replicate the simulation during its execution to simultaneously explore many different possible "futures" in order to compare and evaluate alternate choices are being developed. These interactive simulation techniques are being developed to fit "hand-in-glove" with optimistic parallel computing technologies.

3.1.2 Proposed Distributed BM/C^3 Applicability

A major function of any future BM/C^3 system will be the preparation of defense plans that are both reactive to the current battlefield situation, and proactive in developing contingencies that address the expected future situation. The ATD is planning to demonstrate the relevancy of this technology to facilitating the performance of faster-than-real-time defense evaluation. The requirement for this capability in future BM/C^3 systems is derived from the notion that a real-time tactical defense planning capability will require a faster-than-real-time defense evaluation capability. This capability would provide the mechanism for executing potentially hundreds of full-scale system simulations with various threats, configurations, and resources in a sufficiently short amount of time that the results could be evaluated and provided to the decision maker to enable more effective decision making on the battlefield. It is believed that this technology will eventually enable the performance of this function in a tactical distributed computing environment and provide the backbone for a real-time defense planning capability.

Optimistic simulation technology has demonstrated the potential to speed up simulation execution in proportion to the number of processors; near forty-fold speedups have been observed in the laboratory for benchmark simulations running on a Kendall Square Research KSR-2 multiprocessor [9]. We anticipate that such speedups will enable the real-time utilization of battle management simulations executing in faster-than-real-time as decision support tools within tactical and strategic BM/C^3 systems. In this context, faster-than-real-time implies the execution of a simulation in less wall clock time than the events being simulated would actually take to complete. The challenge associated with realizing such a capability is that missile defense simulations used to test battle management plans against anticipated threats are currently so large and time consuming that they cannot be performed on-line in "hot" situations. As a result, the simulations must now be executed prior to deployment, often using out-dated intelligence information. We believe that optimistic simulation technology, such as, GTW, will enable faster-than-real-time execution of missile defense scenarios and rapid, interactive "what if" analysis to evaluate and refine complex battle management plans, and develop alternative COA. We envision powerful, on-line decision aids, based on interactive parallel simulation technology, will become as ubiquitous and accessible to analysts in the field as they are to personnel at bases and command and control centers in the US.

A proof-of-principle (POP) demonstration is planned to show the applicability of the optimistic computing research. The proposed demonstration will show that optimistic computing technology, such as, GTW, can simplify parallelization and significantly speed up the execution of an existing sequential BM/C^3 discrete event system simulation. The simulation chosen for this demonstration is the Theater High Altitude Area Defense (THAAD) Integrated System Effectiveness Simulation (TISES). The TISES was chosen because it is a validated BM/C^3 simulation, is relevant to the THAAD program, and is readily available. The demonstration is intended to show that this technology provides a viable mechanism for realizing the Program Office's desire to eventually develop a faster-than-real-time Defense Planning and Evaluation capability. In addition to providing a relevant demonstration of the optimistic computing technology to a specific Theater Missile Defense (TMD) BM/C^3 data processing challenge, the speed-up of the TISES could provide a direct and immediate benefit to the program office in their ongoing TISES parallelization activities.

The TISES was developed to provide the THAAD Program Office with a system engineering analysis tool to support the end-to-end performance evaluation of a multi-battery THAAD demonstration/validation system in a many-on-many environment. In the TISES, each simulated battery consists of a battery tactical operations center (TOC), three to nine local launchers, one local radar, and connectivity to a battalion TOC. The release of the TISES software used for this research (version 3.1.1), is a sequential, non-real-time, discrete-event, engineering level simulation written primarily in Ada83, running on Silicon Graphics (SG) platforms under the IRIX Operating System (OS).

The TISES simulation architecture is highly modular and consists of several detailed system segment models, including a BM/C^3 model, which execute within a simulation framework. The large size of the TISES precludes parallelizing and speeding up the entire simulation, or even the entire BM/C^3 model. As a result, this research effort has focused on parallelizing and speeding up the engagement planning sub-module within the BM/C^3 model. The general approach to the parallelization of the TISES sub-module will be to covert TISES from Ada83 to Ada95, map TOCs to GTW LPs, configure GTW to save the proper TOC state variables, and map LPs to appropriate processors. The TISES/GTW research is being performed using a 4 processor SG Origin 200 high-performance workgroup server, and a single CPU SG O2 workstation serving as a front-end for the Origin 200. With this configuration, a speedup on the order of four times is anticipated for the parallelized TISES sub-module.

3.1.3 Future Directions

The natural progression of this research would be to extend it toward the development and demonstration of an optimistic simulation based real-time defense planning and evaluation capability to support National Missile Defense (NMD) BM/C^3 lookahead and predictive battle planning objectives. In addition, the ATD is investigating other potential applications to NMD data processing challenges including the application of this technology to radar resource utilization planning.

3.2 Photonic Interprocessor Routing and Switching Research Efforts

3.2.1 Overview of Technology

The University of Colorado research (sponsored by BMDO/IS&T contract DASG60-95-C-0112) is focused on the design, construction and demonstration of a prototype (8 to 24 node) multi-GFLOP supercomputing system based on next-generation COTS workstations. These workstations will be connected by a 10-20 Gbit/sec optical interconnector network having node separability of 10-40 Km. End-user systems will consist of highly integrated complexes of general-purpose heterogeneous COTS workstations that provide robust, distributed, fault-tolerant performance, as required for battlefield BM/C^3. The key elements of this effort are deflection routing and a ShuffleNet topology.

Deflection routing, was first proposed by Baran in 1964[21] and has been extensively studied by others[22][23][24]. A variation on deflection routing called "2Space-2Time" (2S2T) switching improves performance and is well suited to optical applications[25]. For this effort, Sauer[26] has incorporated 2S2T switching, ShuffleNet topology and multi-media (wire, fiber, and laser) network connector modules. Deflection routing eliminates the need for electronic buffering of message packets when there is contention for a switched output port. If two incoming packets need to use the same output port to minimize delay to the next destination, then one packet is granted the requested port and the other "deflected", or routed immediately to another port. The payload (data portion of a packet) remains in optical form and is sent to a routing control processor (RCP), which makes routing decisions and modifies the header. The objective is to optimize network latency for the case of a packet not being deflected, at the expense of increased

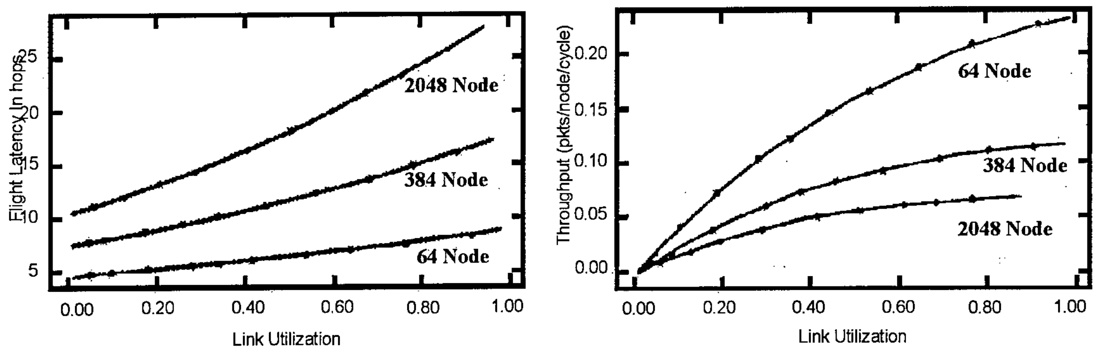

Figure 3, Latency and Throughput Versus Link Utilization

latency for a deflected packet. A requirement for deflection routing is that a deflected packet can still reach its destination, thus implying a network having at least one path between any pair of switching nodes. The "distance" between source and destination is measured in "hops" or links along the shortest (minimum link count) path connecting the nodes.

node. Wait buffer time is the time a packet waits at the sender for a free out port*. Network capacity is the sum of the throughputs of all users. Packets remaining in the network longer, due to deflections, waste an increased amount of bandwidth, and throughput decreases. Figure 3 shows how average flight latency and throughput depend on link utilization[27].

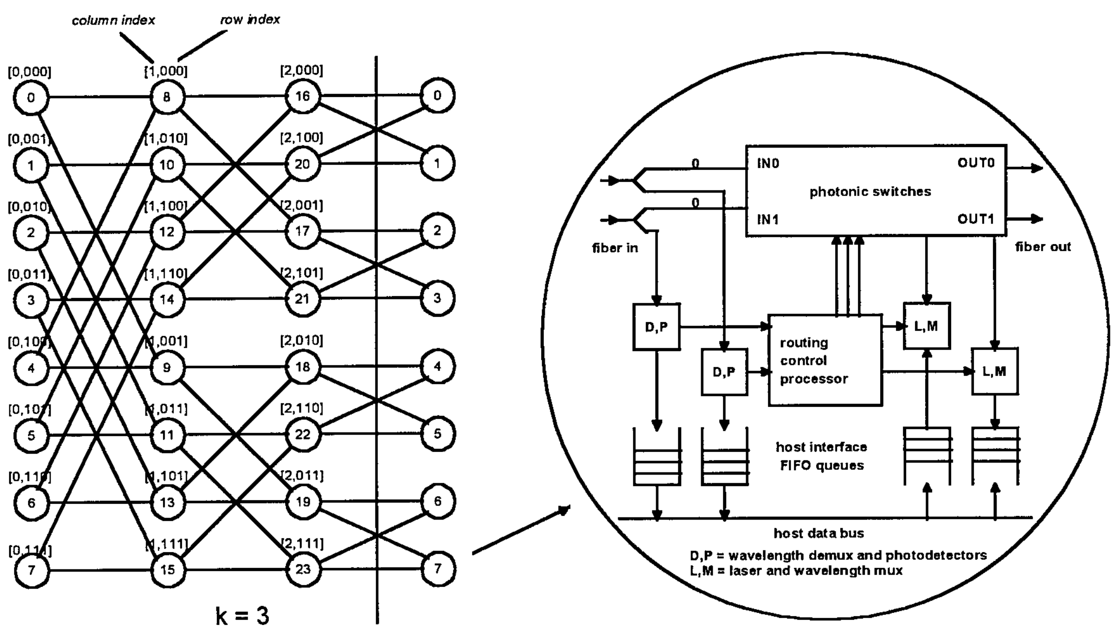

ShuffleNet

Node

Figure 4, 24-Node ShuffleNet and Diagram of Node

Latency has two components: "flight latency" and "wait buffer time". Flight latency measures the time between a packet's entry into the network and its reception at the destination

* There is an additional delay called transmission time which is bits per packet divided by link bit-rate.

In a 2S2T node, there is a space-time permuter which can exchange two incoming packets in time in an effort to reduce deflections. During every packet cycle, the 2S2T permuter considers 4 packet slots (2 slots which have just arrived on both input ports and 2 slots immediately ahead of them in time), permutes the slots as necessary to reduce deflections, then routes them to the node's output ports. The RCP at each node assigns each through-going or host-generated packet to an output port by generating settings for the photonic switches shown inside the node diagram in Figure 4. Each switch can be placed in a "cross " or "bar" state. A packet contains priority bits which get updated dynamically. The simplest priority scheme increments a packet's priority each time it is deflected. This associates an "age" with each packet; the older packet wins when two packets contend for an output port. If both have the same age, the winner is randomly chosen. Age priority reduces the variance of the flight latency probability distribution[26]. Only when an output port is not used by through-going packets can the host inject a packet. Because the RCP is pipelined, it can output new switch settings in each packet cycle, even though it may take multiple packet cycles to compute the settings. The benefit of this approach is illustrated in Figure 5[25]. 2S2T switching reduces the tail of the flight latency compared to spatial switching alone, at the expense of slightly higher delay per hop.

Networks with uniform link capacity can utilize deflection routing for contention resolution. Networks most suitable are those with low-degree nodes and multiple routes between any two nodes. Two commonly studied topologies with these characteristics are the ShuffleNet[28] and the Manhattan Street Network (MSN)[29]. This effort has focused on the ShuffleNet because, at least at uniform loads, the ShuffleNet provides lower latency than the MSN[26]. Unidirectional links, 2S2T nodes, and link utilization of 80% or less, to prevent throughput degradation that occurs in very highly utilized deflection networks, is incorporated in design. Routing decisions are simple and amenable to pipelined feed-forward logic due to a regular and static network topology.

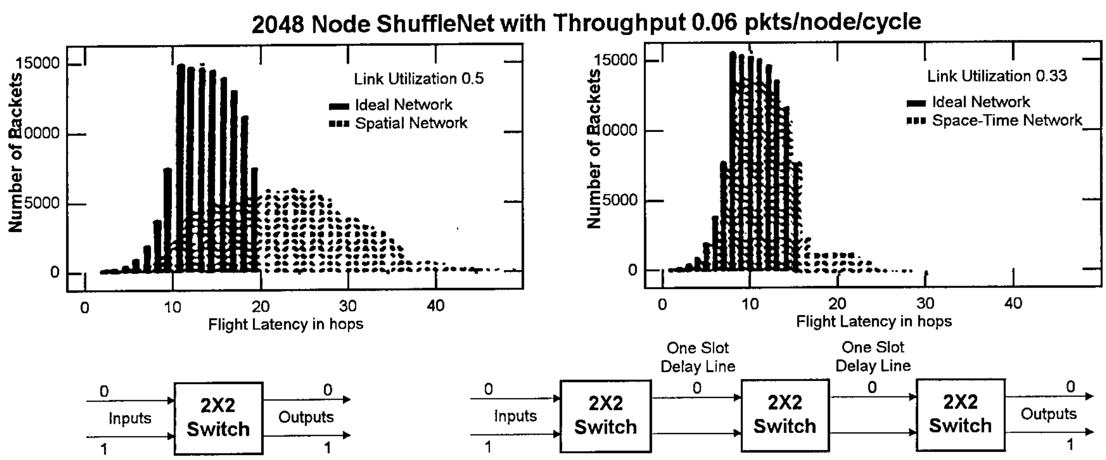

Figure 5, Flight Latency Histogram for 2048-Node ShuffleNet

The University of Colorado ShuffleNet[26] contains k columns of p^k nodes, with nodes in each column connected to nodes in the next by a p-way shuffle. The total number of nodes with p(in/out ports)=2 is $N=k2^k$. Figure 4 shows a 24-node network. Column and row indices are presented in brackets. The nodes form 2^k different rings, with row indices for nodes in a given ring being cyclic shifts of one another. $Node_i$ with column and row indices $[c_i,r_i]$ has ports connected to nodes with column index (c_i+1) mod k. When a packet is deflected, it is sent around the network again, i.e., it must visit each column again. The penalty for deflection grows at the same rate as the number of required columns, approximately $\log_2 N$ for the ShuffleNet per N nodes, versus N/4 for the dual-ring FIDDI.

A property of this ShuffleNet design is that for any source node, there exists many destination nodes having equal distance (hops) from either source output port. Thus, as a packet is routed, it may encounter intermediate path nodes from either output port without incurring additional hops. A packet that can be routed out either port is called a "don't-care" packet, while one requiring a specific output port in order to reach its destination in minimum hops is called a "care" packet for the node. If *d* is the distance between the packet's current node and its destination node, then the packet will care about its output port assignment only during the last minimum selected (*d,k*) hops. Since the ratio of "care" hops to "don't-care" hops decreases with growing network size, deflection routing is less costly for larger ShuffleNets [26].

Extensive simulations have shown the performance of different sized networks under varying link utilization[26][27][28]. The first working prototype of a four node network was demonstrated, Boulder, CO, March 97. This demonstration provided the first milestone for program integration through the hosting of a radar signature generation (RASIG) model. The hosting of this application software was accomplished with minimum rewrite and no breakage. This demonstration provided the benchmark for future hosting of Georgia Tech's GTW software (FY 98). An enhanced four node system will be demonstrated to the Program Element Offices at the USASSDC Advanced Research Center (ARC), Huntsville, AL, December 97. An eight node system capable of supporting parallel fiber, 4 color Bit-Per Wavelength (BPW) and laser links will be demonstrated, October 98.

3.2.2 Proposed Distributed BM/C^3 Applicability

For many military situations, a distributed and robust processing system would enhance operational performance. Attributes of the system would include:

- Architectural design supporting supercomputing capability with COTS processors,
- Standard hardware/software such as COTS heterogeneous PCI bus standards and distributed Message Passing Interface (MPI) for optimized network drivers and message passing,
- Processing scalability which approaches near linear capability per *n* processing elements,
- Geographic scalability with link formats that can accommodate emerging optical high-bandwidth long-haul (100 Km),
- Inherent battlefield fault-tolerant architecture possessing dynamic reconfigurability, the capability to maintain full functionality with few failed components and graceful degradation,
- Network protocol that supports chaotic data input and short, high priority "panic" messages that abruptly alert of component failure.

The University of Colorado technology incorporates the above attributes and is well suited for providing robust National Missile Defense (NMD) BM/C3. Other military applications include:

- HWIL testing,
- Faster-than-real-time TMD defense planning,
- Distributed global cooperative engagement,
- Interconnected TMD TOCs,
- Netted fleet of Unmanned Aerial Vehicles (UAVs).

3.2.3 Future Directions

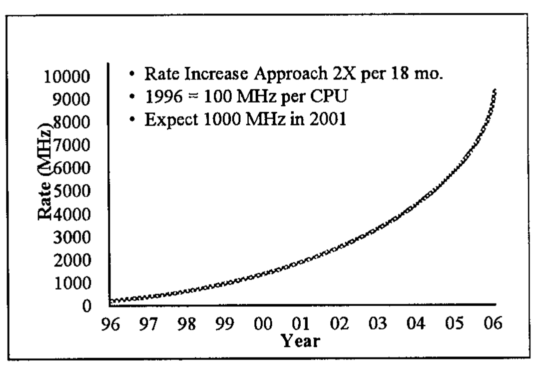

Figure 6, CPU Speeds (Projected) 1996-2006

Today, especially in systems designed to address larger computational problems, remote data access, and distributed HWIL test environments, performance is more often limited by low bandwidth and high data latency, than by processor speed. Although the physics of electromagnetic propagation sets limits on the bandwidth and latency ultimately attainable, practical constraints are more often set by legacy architectures and protocols. While once excellent solutions, they now present hindrances to exploiting the enormous and seemingly inexorable advances in processor speed, Figure 6.

This effort, unlike that of the telecommunications community, is directed at processors rather than humans as customers. With computers requiring 10,000 times faster reconfiguration than telecommunications, issues such as latency, bandwidth and message size are addressed by instructions, CPU and cache, rather than human voice/video, Figure 7.

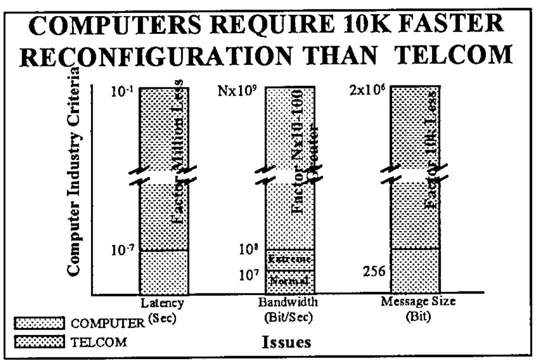

Figure 7, Computer Versus Telecom Requirements

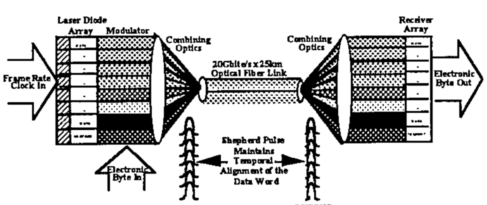

Figure 8, WDM Single Fiber Link with a Shepherd Pulse

University of Colorado network interface routers support multi-wavelength encoding. Bits of a packet can be transmitted through several wavelengths in parallel, using BPW encoding[33]. BPW makes it possible to use a single optical switch to transfer a parallel data word. This feature should allow this interconnect technology to be incorporated with Air Force Office of Scientific Research (AFOSR) soliton "Shepherd Pulse" initiative and the Jet Propulsion Laboratory (JPL) Wavelength Division Multiplexing (WDM) optical long-haul program, Figure 8. Additionally, University of Colorado interfaces allow near speed-of-light satellite interconnects in support of BMDO worldwide "Grid" concepts, Figure 9.

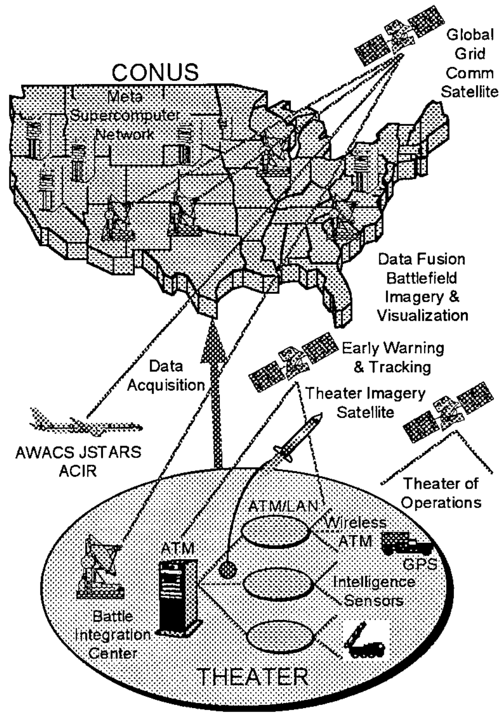

Figure 9, BMDO "Grid"

3.3 Information Warfare

According to the report, "Concept for Future Joint Operations: Expanding Joint Vision 2010", achieving and maintaining information superiority requires three distinct elements: first are the information systems which constitute the architecture — earlier portions of this paper describe our vision of one possible form of the processing portion of this architecture; second is that information which is relevant for military operations to achieve superiority; and, third is the conduct of offensive and defensive information operations.

3.3.1 Overview Of Technology And Applications

Successful exploitation of information superiority requires a redundant seamless network that links all aspects of military operations. The ShuffleNet architecture with photonic interconnects is capable of providing a high-speed seamless architecture. However, the system must also have other characteristics. It must be secure and have built in self protection capabilities against internal and external compromise and disruption. These capabilities can be achieved in several ways and will require multiple technologies. The very large number of processors that may be part of the ShuffleNet architecture provides excellent functional redundancy and this architecture may eliminate the risk of failure caused by loss of a number of nodes. This robustness can be achieved if the system can recognize failures and can reinitialize to recreate lost data or reconfigure to bypass degraded functions/machines. The GTW optimistic algorithms may provide a mechanism for recovering from the loss of distributed processing nodes. Future research should investigate utilizing GTW and the concept of "lookahead - rollback" to facilitate restoring the state of the system, "rolling back," to the point just prior to the loss of the processing element(s) (analogous to the concept of Global Virtual Time (GVT) in GTW) and restarting after redistributing the affected processes to other processing element(s).

Information superiority also requires multilevel security. Nothing is achieved in developing a fast system that meets all our processing needs, but, one that is easily compromised by an opponent. The enemy must not have access to our information. We must know the classification of our information and data; we must know who is authorized access to it; furthermore, we must be able to provide the information needed by the warfighter. ATD's SBIR contractors have developed NSA Class A1 firewalls to deny unauthorized entry to our computer networks. The Gemini Trusted Firewall/Guard permits a virtual private network over the Internet, but faster firewalls are needed. We are working with BMDO to develop a multilevel secure image OS that will automatically give access to classified multimedia data according to the security clearance of the user. And, we are working on physical and algorithmic means to reduce radio frequency interference and magnetic interference in our high speed computer and communication systems. One approach being developed uses low cost conductive thin films to absorb and reflect RF energy that may be self-generated in a nearby board or microcircuit or that may originate externally from either a friendly or hostile source. Other approaches to reduce external influences may use non-linear control algorithms, chaos theory or fractals to both encrypt and reject external influences.

3.3.2 Future Directions

We will continue to expand cooperation with other government agencies and with universities and industry to identify high payoff technologies that can ensure continued superiority of US information systems. Special emphasis is being placed upon developing technologies that allow COTS systems, with low cost changes, to be used effectively in the hostile environments that are expected on future battlefields. We carefully observe and manage SBIR developments and help focus contractor efforts into high payoff areas such as optical processors, nonlinear optics, RF and microwave sources and antennas, and low cost manufacturing techniques. We also participate in BMDO and State Department international programs for joint technology development to better understand the status of foreign R&D and to factor their implications into our development plans.

3.4 Other Activities

BMDO is currently pursuing the development of a Virtual Distributed HWIL Test Bed (VDHTB) which will link geographically distributed HWIL testing facilities to simulation facilities using COTS and enabling BMDO technologies to address the many challenges associated with BM/C^3 information management. In order to prove that the VDHTB is a viable concept, BMDO will utilize existing infrastructure, such as the DREN, High Performance Computing Modernization Office (HPCMO) hardware and software, and IS&T technology investments to perform a POP test program. The POP test program will demonstrate the feasibility of utilizing geographically distributed HWIL facilities, physics-based (i.e., high-fidelity) phenomenology simulations and system component emulations, and HPC centers to perform NMD and TMD-related information processing functions in real-time. The proposed approach will allow the use of expensive or unique hardware, software, or human resources at distributed, remote locations. The demonstration of this capability will be carried out in three experiments.

The first experiment will demonstrate the linking of a HWIL facility (JPL) containing the Quantum Well IR Photodetector (QWIP) whose output will be processed by the 3-dimensional artificial neural network (3DANN) to a remote physics-based system simulation of a cruise missile engagement at the Naval Research Laboratory (NRL). This test will demonstrate emerging hardware and software necessary for linking two sites for distributed computing and will prepare the QWIP for a flight against a Cessna aircraft (stand-in cruise missile) that begins in January 98.

The second experiment will demonstrate distributed access to an HPC center in real-time. The test will demonstrate real-time access by an NMD-type BM/C^3 Planner running on an emulated cluster of workstations (JPL) to a computationally complex sensor simulation hosted on a geographically remote workstation cluster located at the ARC.

The third experiment will demonstrate real-time assessment of high fidelity phenomenology (NRL) and system component emulations (Rome Labs) by a geographically distributed high-fidelity war game type simulation executing on a local cluster network (ARC). The purpose of this experiment is to demonstrate the feasibility of utilizing geographically distributed ground tests to produce results in real- to near-real-time that compare favorably with live flight test results.

One result of this research will be the establishment of the ATD IS&T BM/C^3 Laboratory at the ARC. This laboratory will provide both researchers and technology insertion agents a testbed in which fundamental principles regarding shared distributed BM/C^3 decision making may be formulated and metrics developed. The laboratory will allow for a synergistic insertion/transition of the above emerging ATD technologies to the warfighter, to provide increased capability and produce maximum effectiveness. Activities envisioned consist of laboratory connectivity with additional nodes at SSDC and MITRE, Huntsville, as well as, with NRL and JPL for NMD, and TMD BM/C^3 experiment insertions into global HWIL simulation exercises.

4.0 Summary

The technology in optimistic computing, photonic interprocessor routing and switching, and information warfare is clearly being advanced. The challenge now is to identify relevant applications and conduct convincing demonstrations so that system level decision makers will be in a position to implement the alternate BM/C^3 constructs that the new technology supports.

MITRE, under contract DAAB07-97-C-E601, provides a unique combination of technical support, experience, and cross-service army-wide perspectives on all relevant aspects of BM/C^3 technology, policy, and implementation. COLSA is the integration contractor for this effort, under contract DASG60-89-C-0092.

5.0 References

[1] Osterholz, J. L., "Merging C4I With Modeling and Simulation", Proceedings of MILCOM '95, pp. 1129-1133, November 1995.

[2] Rebbapragada, Dr. V. L., "Distributed Battle Management for Command and Control", IEEE.

[3] Strategic Defense System Glossary, October 1988.

[4] Andriole, S. J., Halpin, S. M., Editors, "Information Technology for Command and Control", IEEE Press, 1991.

[5] Lewis, T., "The Next $10{,}000_2$ (16) Years: Part 1", Computer, pp. 64-70, April 1996.

[6] Thomas, Dr. S., "PetaFLOPS Scale Computing Systems, Opportunities, and

Challenges", Proceedings of the IEEE Aerospace Applications Conference, Vol. 3, pp. 31-49, February 1996.

[7] Gupta, A., Akyildiz, F., Fujimoto, R., 'Performance Analysis of Time Warp With Multiple Homogeneous Processors", IEEE Transactions on Software Engineering, Vol. 17, No. 10, pp. 1013-1027, October 1991.

[8] Jefferson, D., "Virtual Time", ACM Transactions on Programming Language and Systems, Vol. 7, No. 3, pp. 404-425, July 1985.

[9] Das, S., Fujimoto, R., Panesar, K., Allison, D., Hybinette, M., 'GTW: A Time Warp System for Shared Memory Multiprocessors."

[10] Fujimoto, R., Das, S., Panesar, K., 'Georgia Tech Time Warp (GTW Version 2.3) Programmer's Manual", April 1995.

[11] Lipton, R. J., and Mizell, D. W. "Time Warp vs. Chandy-Misra: A Worst-case Analysis", Proceedings of the SCS Multiconference on Distributed Simulation, Vol. 22, pp. 137-143, January 1990.

[12] Carothers, C., Fujimoto, R., Lin, Y., England, P., 'Distributed Simulation of Large-Scale PCS Networks", Proceedings of the MASCOTS Conference, January 1994.

[13] Hybinette, H., Fujimoto, R., 'Cloning - A Novel Method for Interactive Parallel Simulation", Proceedings of the Winter Simulation Conference, December 1997.

[14] Bell, T. E., Riezenman, M. J., 'Communications [technology analysis and forecast]", IEEE Spectrum, pp. 27-37, January 1997.

[15] Hollenbach, Capt. J. W., Misch , G. L., 'Linking C4I and Modeling and Simulation Systems", Proceedings of MILCOM '95, pp. 1126-1128, November 1995.

[16] Meadows, S. I., 'In 2025, Fast, Compact Units Will Range Over Vast Arenas", National DEFENSE, April 1997.

[17] Reddy, R., Garrett, R., 'Future Technology Challenges in Distributed Interactive Simulation", Proceedings of the IEEE, pp. 1188-1195, August 1995.

[18] Sherman, J., "Welcome To The Future: The Army Tests Its Information Age Brigade In The Desert", Armed Forces Journal International, pp. 12-13, May 1997.

[19] Shoffner, LtGen W. A., 'Future Battlefield Dynamics and Complexities Require Timely and Relevant Information", PHALANX, The Bulletin of Military Operations Research, March 1993.

[20] Theater High Altitude Area Defense (THAAD) Simulation Support Plan, February 1996.

[21] Baran, P., "On Distributed Communication Networks", IEEE Transactions on Communications Systems, Vol. 12, pp. 1-9, 1964.

[22] Borgonovo, F., Fratta, L., and Tonelli, F., "Circuit Service in Deflection Networks", Proceedings of IEEE INFOCOM, Vol. 1, pp. 69-75, 1991.

[23] Krishna, A., and Hajek, B., "Performance of Shuffle-like Switching Networks with Deflections", IEEE INFOCOM, pp. 473-480, 1990.

[24] Acampora, A., and Shah, S., "Multihop Lightwave Networks: A Comparison of Store-and-Forward and Hot-Potato Routing", IEEE Transactions on Communications, Vol. 40, No. 6, p. 1082, 1992.

[25] Ramanan, A., Jordan, H., Sauer, J., and Blumenthal, D., "An Extended Fiber-optic Backplane for Multiprocessors", Proceedings of Hawaii International Conference on System Sciences, 1994.

[26] Sauer, J., "A Photonic, Distributed, High-Performance Computer Interconnect", Final Report, US Army Contract DASG60-93-C-0148, December 1995.

[27] Ramanan, A. V., "Ultrafast Space-Time Networks for Multiprocessors", Ph.D. thesis, University of Colorado at Boulder, Department of Electrical and Computer Engineering, 1993.

[28] Hluchyj, M. G., and Karol, M., "ShuffleNet: An Application of Generalized Perfect Shuffles to Multihop Lightwave Networks", IEEE INFOCOM, pp. 379-390, March 1988.

[29] Maxemchuk, N., "Regular Mesh Topologies in Local and Metropolitan Area Networks", AT&T Technical Journal, Vol. 65, pp. 1659-1685, September 1995.

[30] Feehrer, J., Ramfelt, L., and Sauer, J., "An Optical Deflection-Routed Multiprocessor Interconnect", 6th European Research Consortium for Informatics and Mathematics Workshop, pp. 103-115, June 1994.

[31] Feehrer, J., Ramfelt, L., and Sauer, J., "Design and Implementation of a Prototype Optical Deflection Network", Computer Communications Review, Vol. 24, p. 191, October 1994.

[32] Feehrer, J. R., Ramfelt, L. H., and Straub, D., "Implementation Details for Optical Deflection-Routed Multiprocessor Interconnect Prototype", Packet Network Laboratory Technical Report, University of Colorado Optoelectronic Computing Systems Center, March 1995.

[33] Blumenthal, D., Feuerstein, R., and Sauer, J., "First Demonstration of Multihop All-Optical Packet Switching", IEEE Photonics Technology Letters, March 1994.

Acknowledgments

Richard Fujimoto and Chris Carothers, Georgia Institute of Technology, provided input for the overview of optimistic computing research.