SF 298

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	31 March 1998	Final Progress Report, 1 Jan 1997–31 Dec 1997

4. TITLE AND SUBTITLE
Smart Photonic Networks and Computer Security for Image Data

5. FUNDING NUMBERS
DAAH04-95-1-0123

6. AUTHOR(S)
M. J. Flynn, M. Morf, J. Gill

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Stanford University
Sponsored Projects Office
Godzilla Modular
Stanford, CA 94305-3027

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
The views, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words*)

Smart Photonic multistage Interconnect Networks (SPINs) were investigated for secure and robust high speed image data communication. To cope with the plethora of user and mission requirements flexible architectures with intelligent and reconfigurable subsystems are required.

New coding algorithms for image and other types of data were developed, that provide or extend simultaneous security (privacy, integrity, reliability, availability, and protection from reverse engineering) based on space-time code diversity. They exploit parallelism and scalable hierarchical multiplexing schemes, by flexible combinations of space-, time- and frequency-multiplexing of photonic interconnect and computing based architectures.

Security and robustness can be supported in the hardware via photonic FPGA's, software reconfigurable gate arrays based on novel photonic quantum devices. The function of such devices are well matched to photonic communication (transmission, routing and conversion) and high speed processing (coding, security, data-fusion, etc.)

VLIW type architectures were shown to be ideal for very high-speed switching/routing and very high performance processor architectures of the type required for communications and multi-media processing, or more generally on latency tolerant applications. The latter are ideally matched to regular or photonic FPGA based processing. As a baseline, we focused on network and multi-media applications.

14. SUBJECT TERMS
Integrated data security, space-time code diversity, network architecture and processing, bandwidth and latency of networks and switching devices, photonic technologies, fundamental limits of optical switching and memory.

15. NUMBER OF PAGES
7

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Contents

A Scientific Personnel Supported and Honors/Awards/Degrees

Prof. Michael Flynn. Recent honors: Tesla Medal, 1998, International Tesla Society (Belgrade). Honorary D.Sc., University of Dublin, 1998.

Co-investigators Martin Morf, John Gill.

Graduate students Kevin Rudd, Alice Yu, Andrew Zimmerman. (PhDs expected between August 1998 and December 1988.)

B Report of Inventions

Recent patent applications:

S94-160 "Vertical Cavity Optical X-Modulator," J. Trezza, J. Powell, M. Morf, J. Harris. Patent Application 08/778,817 filed 1/3/97.

S96-175 "Systems Applications of X-modulators (Transmission/Reflection Modulators)," J. S. Powell, M. Morf, J. Harris.

C Scientific Progress and Accomplishments

1. Coding and Security.

   We continued our investigation into new space-time code diversity based algorithms for image and other types of data, with the goal of providing or extending simultaneous security (privacy, integrity, reliability, availability, and protection from reverse engineering) by exploiting parallelism and scalable hierarchical multiplexing schemes to build photonic network architectures.

2. Basic Technologies.

   Various photonic technologies were studied in terms of their suitability for Smart Photonic Networks. The work on quantum well devices is summarized in two theses [2,3]. Extensions of this work to quantum dots leads to very high density (10^{11} gates per chip), and quantum computing. We also explored the concepts of agregation (creating reliable functions from unreliable elements) and super-resolution (non-diffraction limited effects that may have applications to communications and computing, including all optical gates, small efficient antennas and interconnects.) The architectural implications of these new basic technologies are revolutionary. Our most recent work on Architectural Issues and Nano-Technology is to be reported at a special session on Photonics and Nano-Technology at the Optical Society of America meeting in the Fall of 1998.

3. Switching and Processor Architectures.

   We continued our studies of very high-speed switching/routing architectures and very high performance processor architectures and simulation tools. As a baseline, we focus on network and multi-media architectures that take advantage of the very robust and distributed TCP/IP protocol.

Thesis work on VLIW architectures [4] is coming to completion, that has the potential of a significant impact on very high performance processor architectures for communications and multi-media, or more generally on latency tolerant applications. The latter seem to be ideally matched to FPGA based processing. We are studying how to support such high-throughput proccessing with very high-speed photonic or very high density quantum dot based technologies.

Other thesis work on efficient multi-media type processing [5] is complementing our architectural studies. Applications to very low latency photonic interconnects between processors were pursued, in part in discussions with DEC WRL.

C.1 Coding and Security

Understanding the interactions of Coding and Security, Switching Architecture Simulations, and Basic Technologies is our basic goal.

A list of the topics studied includes: network security, primary security requirements, standard approaches to security requirements, security via secret sharing, use of space-time-code Diversity, reliability using multiple paths, integrity using multiple paths, privacy using multiple paths, Shamir's polynomial interpolation method, unified reliability, integrity, and privacy, network and routing issues, and ramp schemes.

Summary and review of the coding and security work

Primary security requirements include:

Reliability—data communications protocols should be robust enough to withstand link and node failures and misrouting;

Integrity—data is not modified accidentally or deliberately tampered with, by replacement, insertion, or deletion;

Privacy—confidentiality of the data in the network should be maintained even if an eavesdropper can tap one or more links in the network.

There are a number of standard approaches to deal with security requirements; we pursued alternative methods to achieve such requirements that are better matched to photonic technologies.

Secret sharing can be used in a unified approach to satisfy these security requirements. Secret sharing takes advantage of the possibility of transmitting shares of a message over different paths, at different times, and perhaps using diverse coding schemes. We call this capability of multiple transmissions *space-time-code diversity*.

Reliability can be achieved using multiple paths, i.e., a message may be transmitted more than once, either in time or space. For example, two copies of messages can be sent over two different paths (e.g., one direct link and one path through an intermediary). Communication is successful if either copy of the message arrives. Security is limited to link-level encryption. If any link is compromised, the message will be intercepted.

Integrity can be achieved using multiple paths, i.e., if a message is transmitted over three different paths, an altered copy can be detected and corrected using the two correct copies. If

$n = 2k+1$ distinct paths are used, we can protect against coordinated tampering with k message copies. (Achieving integrity using such a repetition code for error correction is inefficient.)

Privacy can be achieved using multiple paths, i.e., assume two disjoint paths from transmitter X to receiver Y. If an eavesdropper can tap only one link, then one bit of information m can be sent from X to Y with perfect security. X generates a random bit r and sends the "shares" m XOR r and r over the two disjoint paths. Y reconstructs the message m from the two shares XOR-ing the two received shares, thereby canceling r, i.e., getting back the message m. Neither of the shares in transit gives any information about m.

Secret sharing schemes were introduced by Blakley (1979) and Shamir (1979) as a solution to the preceding key management problem. The key (the "secret") is broken into n pieces called "shares" in such a way that a) any subset of k or more shares can recover the secret; b) no subset of $k-1$ keys contains any information about the secret. Such a secret sharing method is called a $(k;n)$ threshold scheme.

Secret sharing schemes can function in "courier mode". The messages themselves, rather than the cryptographic keys, are broken into shares that are given to different couriers—that is, they are transmitted over diverse paths in space or time. Secret sharing can provide integrity and reliability in addition to privacy. If we can guarantee that a) at most l shares are lost or are delivered with detectable errors, and b) at most t shares are tampered with or arrive with not obviously detectable errors, then a $(k;n)$-threshold scheme can be used to achieve perfect security (privacy, integrity, reliability) as long as $n - k >= 2t + l$. For example, 6 disjoint paths can resist tampering with one share, loss of one share, and interception of two shares.

For a suffciently rich network, there will be many disjoint paths from transmitter X to receiver Y. New routing algorithms will be needed. Instead of simply finding a good route from X to Y, we must now find n reasonably good paths that do not intersect. The number of logical paths may be increased by use of transmissions at different times. This will improve reliability in the presence of transient failures, but does not in general improve privacy or integrity if a link or node has been compromised.

Practical implementations will most likely take advantage of "ramp schemes". The bandwidth expansion required by secret sharing can be reduced if the requirement that $k-1$ shares provide absolutely no information about the secret s is relaxed. If instead we use a modified scheme: a) Given up to j shares, no information is gained. b) From $j+1$ to k, the remaining uncertainty about the secret vector decreases linearly down to 0 with each additional compromised share. This allows for $m = k - j$ secrets to be communicated using n shares, improving the communications rate. In this case, we trade improved communication rate/efficiency with reduced security; however, it is important to note that the information obtained by the shares in a ramp scheme is on the set of secrets as a whole and not on any individual secret. If each secret is further encrypted, an eavesdropper will have great diffculty in taking advantage of any partial information obtained.

In summary, our Space-Time-Code Diversity has several advantages. It a) provides a mechanism for enabling a three-way trade-off between the aggregate bandwidth, computational overhead, and the level of security; b) offers a means of communicating with very high levels of security; c) is a "distributed solution" to security problems; c) has low complexity implementations well matched to high speed photonic networks. Communication over multiple paths, at different times, and with different coding methods is a unified approach to security requirements.

C.2 Basic Technologies, Photonics, and Nano-Technologies

Our basic technology studies, including Photonics and Nano-Technologies, are in part based on interactions with projects that include quantum well devices (with 3, 2, 1, and 0-D quantum confinements), quantum dot laser diodes, amplifiers, and modulators in free space, wafer plane, and fiber/optical transmission line forms.

More recently, the Harris group's work was extended to include very high density quantum dots (10^11 gates, lasers, or modulators per chip), and the use of quantum dots for quantum computing—an integrated alternative to their approach using NMR.

There are a number of common architectural themes that thread through our joint work. In one theme, we are exploring the concepts of *agregation*, i.e., creating reliable functions from unreliable elements. Our past work on using scattering theory for developing fast algorithms and high performance architectures is a very good example of such a theme; it goes far beyond the simple physical analogies that are evoked to justify, for instance "simulated anealing", "genetic algorithms", and "neural network" (NN) algorithms—all happen to be slow and inefficient! In contrast to NN, genetic regulatory circuits in evolutionary biology are different from NN; the former are very hierarchical, event driven, and have feedback loops. Scattering theory suggests aggregating "layers" or "grains" to achieve overall effects, such as implementing a reliable function from unreliable components. In the past, aggregation had been suggested using scalar functions, such as majority voting, or the serial composition of parallel switches, to improve reliability. In contrast, scattering theory suggests the aggregation of multi-input–multi-output (MIMO) functions to achieve overall reliability, avoiding single points or bottle-necks of failure. This is quite evident now in quantum computing, where probability density matrices are concatenated, e.g., in error correction operations, all of the MIMO type.

Another "common theme" is *super-resolution*, typified by non-diffraction limited effects, such as NMR and photon stimulated-emission and absorption; in both cases the wavelengths (λ) involved are much longer than the feature sizes (resolution in NMR, atom-diameter vs. λ-photon.) These studies are expected to have applications to communications and computing, including ultra-fast all optical gates, small efficient antennas and interconnects. The architectural implications of these new basic technologies will be revolutionary. We will report these recent developments on "Architectural Issues and Nano-Technology" at the 1998 Optical Society of America meeting, see [24].

C.3 Very High Performance Communication Processor Architectures

Switching and Processor Architectures

Our studies of very high-speed switching/routing architectures and very high performance processor architectures for photonic networks have identified two complementary architectural approaches on how to achieve the necessary high-throughput performance and low implementation complexity required for photonic or nano-technologies.

The thesis work in our group on VLIW architectures [4] has the potential of a significant impact on very high performance processor architectures required for very high speed smart photonic networks. VLIW architectures with their high degree of parallelism and the potential for a high degree of pipelining, as well as their potential for low hardware complexity (mostly statically compiled).

An alternative architectural approach that promises even higher degrees of parallelism and pipelining is based on FPGAs, under the heading of Adaptive Computing Systems. This is one of the subjects of our DARPA ACS program contracts, it calls for wireless network implementations of the security work carried out for this contract. We already demonstrated DES, and IDEA encryption algorithms running at up to .5Gbit/sec rates with FPGAs. Therefore, photonic or quantum dot technology based FPGAs predictably would have very high-speed and very high-density performance, especially in either very fine-grain (e.g. bit manipulations) or very latency tolerant high-throughput applications.

VLIW Approaches to Network Processors

Network processors require a significant amount of performance and one of the most effective approaches to achieving this level of performance is to exploit instruction-level parallelism (ILP). Many modern processors exploit ILP and are able to achieve performance that exceeds one operation per clock cycle.

A common approach to exploiting ILP, typical of superscalar processors, is to use hardware-based dynamic analyses to extract independent operations from the execution stream; these operations can then be executed in parallel out of their original order. By not restricting the execution order to that codified in the program, these processors are capable of working around Dynamic Events that arise during execution. These events have many causes and their effects are to delay or to interrupt the normal flow of execution. Although processors using dynamic analyses are highly tolerant of Dynamic Events, the complexity of these processors can limit their peak performance and scalability and often result in a long design time, difficult verification, and large die size.

An alternative approach, typical of VLIW (Very Long Instruction Word) processors, is to use compiler-based static analyses to produce an explicitly parallel execution stream; operations can then be executed in parallel without requiring any dynamic analyses. The simplicity of these processors results in high peak performance and scalability while providing a shorter design time, easier verification, and smaller die size than would be required if dynamic analysis hardware was used. However, the simplicity of these processors can limit their ability to effectively manage Dynamic Events and often requires that the compiler schedule code pessimistically to preclude all avoidable problems; unavoidable problems, such as exceptions and interrupts, must still be managed and can significantly reduce the achieved performance.

In spite of the difficulty that VLIW processors have managing Dynamic Events, the simplicity of these processors offers an intriguing base from which to develop an efficient high-performance ILP architecture. This research has considered two different approaches to VLIW processors: enhancing VLIW processors with Replay Buffers and extending VLIW processors with out-of-order execution.

C.4 Enhancing VLIW processors: Replay Buffers

Replay Buffers provide zero-penalty cycle Dynamic Event management for static VLIW processors. Their implementation requires minimal architectural impact and no software scheduling restrictions. Because they are effective and efficient they are an ideal solution to the problem of Dynamic Event management for low-complexity static VLIW processors. Figure C.4 shows how Replay Buffers fit into the normal pipeline to register-file (or other result consumer) data flow.

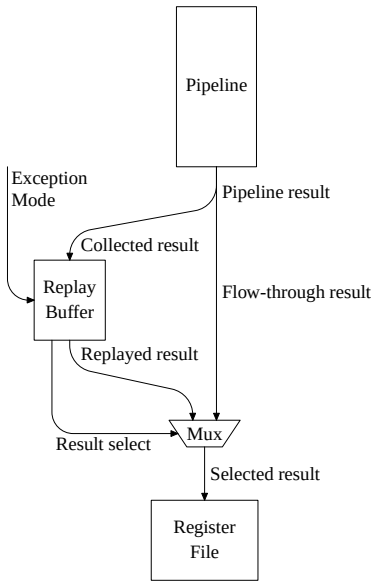

Figure 1: Replay Buffers easily integrate into normal pipeline to register-fi le data flow

During normal execution, results flow directly from the pipeline to the register-file. When an exception occurs normal results are collected into the Replay Buffers maintaining their original result ordering while exception results are delivered directly to the register file; following the completion of the exception, the collected normal results are then replayed in the original result ordering thus hiding the fact that the exception occurred from the execution stream.

Figure C.4 shows how an exception from operation x is managed by Replay Buffers. At the point that the exception is raised, the results from the pipelines are collected in the Replay Buffers and initiation of new operations (here, operation $x+3$) is suspended. Following servicing the exception, initiation resumes and replay begins. As can be seen in the figure, the collected results are replayed with the same timing as when they were collected. Thus the ordering of writes to $\mathtt{R}_x$ is preserved across the exception ensuring the correct execution of these operations independent of whether or not an exception occurs during their execution. We call this behavior *side-effect-precise* exception handling.

Because result collection begins immediately on raising the exception and normal execution and replay begin immediately on completion of servicing the exception it is clear that there is no penalty associatied with the exception and that the cost of the exception is simply the cost of servicing the exception. Note that this example shows the case when servicing the exception takes longer than draining the pipelines into the Replay Buffers—this is not a requirement and is shown only for simplicity. Replay Buffers can take advantage of this fact and manage delays as simply as managing exceptions by treating a delayed result as an exception that requires no servicing to be performed. Using Replay Buffers to manage delays as well as exceptions results in a simple unified treatment of all Dynamic Events.

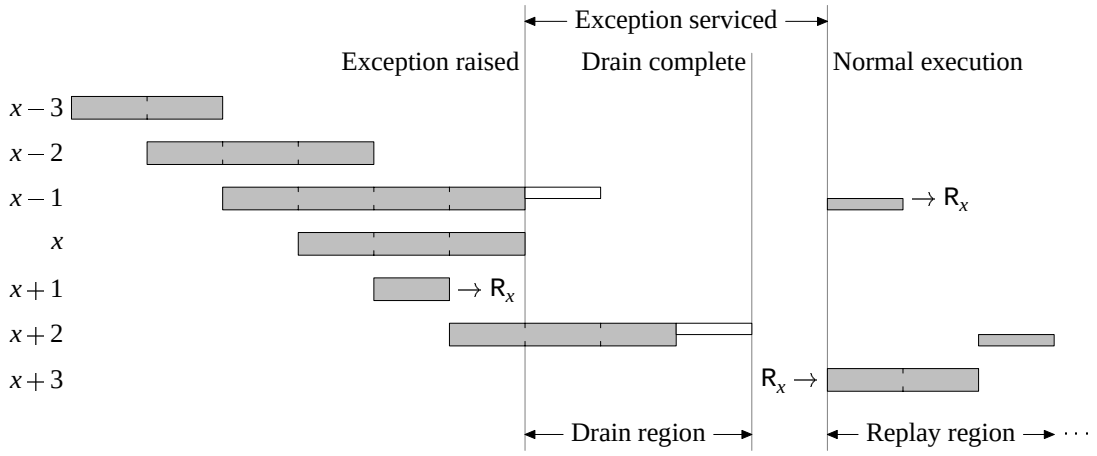

Figure 2: Management of exceptions with Replay Buffers

C.5 Extending VLIW processors: Out-of-order execution

While Replay Buffers provide zero penalty-cycle Dynamic Event management, they do not hide any of the cost of servicing the event itself. Out-of-order execution, however, provides the opportunity to reorder operation execution around delaying events and potentially to reduce the perceived cost of the event. Although out-of-order execution can effectively manages delaying events, it can significantly complicate interrupting events—these complications, resulting in increased hardware complexity, reduce the overall performance of the processor and offset the benefits of reducing the service cost. Nonetheless, out-of-order execution offers the opportunity to improve performance and is an important technique to consider.

Out-of-order execution poses a unique problem for VLIW processors. Since VLIW processors have their execution-stream carefully scheduled to take into account, changing the order of operation execution must be done in such a way that this scheduling is not affected. Our approach to out-of-order execution on VLIW processors is based on the three-phase split-issue (TPSI) execution paradigm which provides VLIW processors with a straightforward means to support any execution variation, including out-of-order execution, transparently.

TPSI execution uses the detailed specification of operation behavior, specifically the visible operation side-effects. Each operation is defined as a triple that specifies the timing and action of the read, compute, and write side-effects for an operation. For example, a simple increment operation $y \leftarrow y + 1$ that takes one cycle to execute would be specified as the triple

$$\begin{aligned}&\big((t_0 \leftarrow y@0),\\ &\ (v_0 \leftarrow t_0 + 1@0),\\ &\ (y \leftarrow v_0@0)\big)\end{aligned}$$

This triple describes that the source value y is read at cycle 0; the actual increment computation is then performed, also at cycle 0 although this is also dependent on the availability of the source value; finally, the write back to y is also performed at cycle 0. Temporary values t_0 and v_0 are used to buffer the source and result values in order to maintain the scheduled timing. A more complicated example of a traditional SAXPY operation $s \leftarrow a \times x + y$ that takes a total of

3 cycles could have the operation triple

$$\begin{aligned}&\big(\big(t_0 \leftarrow a@0, t_1 \leftarrow x@0, t_2 \leftarrow y@1\big),\\ &\ \big(u_0 \leftarrow t_0 \boxtimes t_1@0, v_0 \leftarrow u \boxplus t_2\big),\\ &\ \big(s \leftarrow v_0@2\big)\big)\end{aligned}$$

One feature of this SAXPY specification is that the sumanad y is not read until one cycle into the execution of the operation possibly allowing this operation to be scheduled one cycle early by the compiler. The partial computations $\boxtimes$ and $\boxplus$ may not produce the traditional results from $\times$ and $+$ operations but in combination result in the correct aggregate computation being performed.

This research developed the Nyfo VLIW simulator that is integrated into the University of Illinois, Urbana-Champaign, IMPACT Compiler suite. The Nyfo VLIW simulator (referred to as the NV simulator in previous reports) is an instruction-level parallel processor simulator that can simulate a wide range of superscalar and VLIW processor architectures and uses the TPSI model for both specification and simulation.

We used the Nyfo VLIW simulator to model an 8-wide VLIW processor with realistic function unit distribution and capabilities while varying both the compiler scheduling assumptions, degree of out-of-order execution supported, and memory system configurations. Our results have shown that out-of-order execution is quickly able recover from scheduling errors (micro-reordering) but has little capability to recover from major latency variations such as cache misses with long memory system latencies (macro-reordering). Based on these results, it is clear that out-of-order execution has significant benefit for maintaining compatability across a range of implementations and system configurations. However, since out-of-order execution has only limited ability to mitigate long memory latencies the degree or out-of-order execution should be limited to that required to perform any necessary micro-reordering and other techniques should be relied upon to reduce the actual or perceived memory latency.

Improvements in Video Compression for Multi-Media Applications

Our research in this area addresses a critical technological issue, one that affects a broad range of applications. Primarily, we focused on novel algorithms for video compression and developed more effective methods to transmit moving images across a communication system with very limited computational resources. The results of our research, reported in [5], are expected to have an impact on applications such as:

- Broadcast television in the adopted HDTV standard
- Transmission of image sequences across the internet for web browsing or correspondence
- Videotelephony, such as personal videophones or video conference calls
- Distribution of movies-on-demand from a centralized computer server.

We considered three key aspects of a video compression system. We first analyzed video fidelity metrics, to determine what kind of metric would be more effective at predicting the visual fidelity of the output from different video systems. Such a fidelity metric is important when multiple types of video systems need to be evaluated. Without such a metric, extensive

psychological viewing tests would need to be conducted, which can be both time consuming and expensive, since live subjects are needed. An effective video fidelity metric is consistent, efficient, and affordable.

Then we considered the other two key components of a video compression system, bit rate and complexity. The bit rate, or equivalently, compression ratio, refers to the amount of image data, transmitted in units called bits, that is needed to represent a certain amount of information. The bit rate should ideally be as compact as possible, which still preserving the same fidelity. Compact organization of the data allows the data to either be transmitted faster, or sent/stored using less resources. A superior video compression system should deliver one or more images using the most compact data representation possible.

The speed of execution for a video compression algorithm is determined by its computational complexity. An algorithm should carry out a given task within certain latency restrictions, while still computing the same, or virtually the same, result. A robust algorithm tries to minimize both bit rate and computational requirements without compromising video fidelity significantly.

We addressed all three aspects of the video compression system in more detail, see [6].

In this research, we assessed the capability of four metrics (average MSE, average SNR, ANSI parameters, and ITS metric) to determine the fidelity of video sequences. To do this, we first defined the ideal requirements for a video fidelity metric, in terms of monotonicity, degree of change, and consistent behavior. Then, we constructed a series of highly reproducible degraded sequences containing artifacts common to DCT-based transform coders, such as H.263, and evaluated the performance of each metric on those sequences. From the resulting data, we determined the accuracy and reliability of each of those metrics. Our analysis concluded that the more well-established average MSE and average SNR metrics exhibited superior performance over the ANSI parameters and ITS metric, see [7].

We used a standardized video compression system, known as the H.263 standard, which has optimal performance at low bit rates for video-telephony applications, to develop a more efficient algorithm for video encoding. To do this, we characterized H.263 bitstreams and observed that, often, data units, known as "inter-macroblocks," are reduced to zeros after a two-step process within the system called Discrete Cosine Transform (DCT) and Quantization. We took advantage of this observation and proposed a new H.263-compatible algorithm that predicted when there be all zeros in inter-macroblocks at that stage. This eliminated a significant portion of computation usually needed for those inter-macroblocks. Simulation results show that, relative to the base (original) H.263 encoder, our early-detection mechanism can reduce the computational requirements of the DCT/Quantization by as much as 81% while simultaneously reducing the bit rate, with virtually no visible change in video fidelity, see [8].

We extended the work done previously to improve the H.263 system even further. By optimizing the software used to generate H.263 bitstreams, we reduced the computation time of the encoder by four times its original value. Then, we used a zero-detection algorithm, similar to the one used for the DCT/Quantization, to improve the performance of another section of the H.263 encoder, known as the motion estimator. Our simulations demonstrate that the combined improvements to the H.263 encoder ultimately reduces the computational time of the original encoder by four to nine times its original value.

C.6 Packet Switching Photonic Network Switch Design and Routing Algorithms

Maturity of photonic technology makes it possible to construct all optical network switch to avoid optical-to-electrical signal conversion for routing. To realize all optical packet switching, our current network topology and routing algorithms have to be reexamined and modified to satisfy the necessities of all optical network switching such as a fast routing decision, consideration of hardware implementation, buffering etc. In this report, first, we present three methods to improve a parallel routing algorithm for a packet switching Benes network by taking advantage of input patterns and making a dynamic latency decision on routing. Secondly, optical implementation of a multiple output port network switch is presented. In many levels of networking from multiprocessor interconnection to wide area networking, multiple latencies resulting from this scheme could improve the overall performance when combined with smart routing schemes. Finally, we present an interpretation of multistage networks using symmetric groups. Cayley graphs for symmetric groups and their coset graphs suggest interesting alternative ways to construct new multistage interconnection networks.

Introduction

In the past decade, many researchers proposed various new network topologies and routing algorithms to meet ever-increasing bandwidth requirement in various levels of networking and make use of state-of-the-art technology to achieve better performance. For instance, in multiprocessor interconnection, efforts to improve overall performance by means of exploiting parallelism using multiprocessors require much larger communication bandwidth while interconnection is required to maintain low latency, good scalability and connectivity. In WAN and LAN, cheap costs of propagation medium such as optical fibers and advent of commercial products for optical switching such as Lithium Niobate (LiNbO3) Optic devices [9] and Self-Electrooptic-Effect Devices (SEED) [8] made it possible to build simple optical interconnection. Recent achievements such as Multiple Quantum Well (MQW) modulators [6] and Sagnac exchange/bypass gates [7] which we will discuss in section 2 provide complete 2x2 network switching capabilities that is essential to build a multistage interconnection network. After a brief introduction of these photonic devices in section 2, sections 3 and 4 will go over improved parallel routing algorithms for a packet switching Benes network and a multiple output port network in turn. Especially, in section 4, we will show the design of a multiple output port Batcher/Banyan network as a baseline model. Figure 1 and 2 show the comparison of two major network topologies in terms of latency and hardware cost assuming the use of $2x2$ switching elements. An optimal network line shows a theoretical bound for the required number of stages to realize a full permutation. From this figure, we can recognize two things. First, the Benes network is quite close to an optimal network. But, relatively large control latency, $O(\log_2(n))$ or $O(\log(n))$ depending on the assumption of pipelining, makes the Benes network unattractive for packet switching. Parallel routing algorithms allowing dynamic control latency would be a good solution to this problem from that aspect. Secondly, the Batcher/Banyan network provides simple self-routing capabilities but suffers from large latency, $O(log_2(n))$, and hardware costs, $O(n\log_2(n))$. However, this large latency problem can be relieved by allowing multiple latency in routing and letting early routed permutation leave a network as early as they can through multiple ports. These two ideas are the major motivation that we reexamine these networks as a candidate for an all optical network switch. Finally, in section 5, we will introduce the relation between symmetric group and multistage interconnection network (MIN) and explain how its Cayley graphs and

coset graphs can be used to design a MIN.

Figure 1

Figure 2

A brief introduction to photonic devices

Realizing switching in optical domain has been researched in many years [6-9]. Four major methods of realizing switching are amplitude modulation, phase modulation, polarization change, and directional change. Combination of these methods also provides switching. In this paper, we consider two photonic devices, multiple quantum well (MQW) modulator [6] and Sagnac gate [7]. The first device is in the category of the combination of the first and forth methods and the second device is using the first and second methods. From a functional point of view, both devices are realizing a Fredkin logic, figure 1 in Appendix, which basically connect two inputs to two outputs either in interchanged or non-interchanged way depending on the third control input. The structure of a MQW modulator consists of a slightly asymmetric Fabry-Perot cavity containing top and bottom mirrors composed of 10 and 12.5 period GaAs-AlAs quarter wave stacks surrounding an undoped cavity of quantum wells. The picture is shown in figure 2 of Appendix. From top to bottom, it is doped as p-i-n. Under no bias, light incident on the top mirror is propagated into two paths: reflection and transmission. The transmitted wave is again reflected and transmitted on the bottom mirror. The reflected wave on the bottom is propagated backward and finally combined with the initially reflected wave on the top. The cavity is designed such that these two waves are 180deg out of phase with same amplitude and cancel each other. As a result, we don't see any reflected light and a cavity is transparent, which is a cross mode in a Fredkin logic gate. Under bias, activated quantum wells in the intrinsic region absorb all the transmitted light from the top mirror so that we don't see any transmitted light coming from the bottom mirror and only observe the reflected portion of incident light on the top mirror due to no cancellation. This is a set-through mode in a Fredkin logic gate. While a MQW modulator is an electrically controlled optical datapath logic, a Sagnac gate, figure 3 in Appendix, is an optically controlled optical datapath logic. Injected optical signal at input A goes though 3dB coupler and enters a polarization maintaining 3dB coupler in the second stage. Inside this device, optical signal is divided into two and delivered to the upper and lower polarization beam split couplers. Under no control signal at input C, these two optical signal travel the same distance in clockwise and counterclockwise direction though optical fiber and interfere each other constructively at the upper input of the second stage. As a result, optical signal injected at input A is emerged at output A after some delay. Under existence of control signal at input C, optical signal propagating in the clockwise direction co-travels with a control input and experiences Kerr effect. That is, increase in intensity causes light to travel slowly in the non-linear medium, which in turn causes phase difference. The length of fiber is such designed that phase difference between two opposite traveling waves is (. This difference makes original optical signal to interfere itself destructively at the upper input of the second stage and constructively at the lower input. Consequently, optical signal injected at input A is emerged at output B. Considering this device is symmetric, the whole logic performs the Fredkin logic with another input at input B. Table 1 in Appendix shows the comparison of two devices.

New Parallel Algorithms for Packet Switched Photonic Benes Networks

The Benes network has received much attention in interconnection network literature because of its $O(n \log n)$ hardware cost and $O(\log n)$ depth [1–4]. Most known sequential route assignment algorithms, such as the looping algorithm for Benes (1962) networks, are designed for circuit switching systems where the switching configuration can be rearranged at a relatively low speed, $O(n \log n)$. To realize optical packet switching on a Benes network, it is natural to resort to a parallel algorithm that speeds up the routing decision by parallel computing on SIMD architecture. The best parallel algorithm reported [3,4] has a time complexity of $O(\log 2n)$, where n is the number of inputs to a network and we assume SIMD controllers are fully connected. With an additional assumption of fully pipelined stages of network, which requires the controller cost of $O(n \log n)$ [4] rather than $O(n)$, control timing cost decreases to $O(\log n)$. While this algorithm is relatively fast over any existing algorithms for a packet switching Benes network, total latency of the switching network is still $O(\log 2n)$. That is, from a latency point of view, there is no advantage to using the Benes network over a self-routed Batcher/Banyan network that has a latency of $O(\log 2n)$ and much simpler control algorithms. Here, we present three methods that can take advantage of many input patterns that can be routed less than $O(\log n)$ per stage by allowing multiple control latencies in each stage.

Permutations and Symmetric Groups, Cayley- & Cayley coset graphs

Given a set of generators for a finite group G, we can draw a graph, called a Cayley graph, in which vertices represent the elements of the group G and the edges represent the action of the generators. That is, there is an edge from an element a to an element b iff there is a generator g such that $ag = b$ in the group G. A symmetric group, usually denoted as S_n, is composed of $n!$ elements that correspond to all the permutations of n symbols. For instance, S3 is composed of 123, 132, 213, 231, 312, and 321. It can be easily seen that 132 and 213 can also represent generators and they are two generators enough to generate the whole group. In fact, there are infinitely many sets of generators that can generate the same group. Their difference is interpreted as different numbers of edges connecting each elements in a corresponding Cayley graph and their connectivity. In general, a coset graph can be used in several ways to build a multistage interconnection network. One is to find the small number of hops in which one can visit from identity to all the elements in the coset graph by choosing a shuffle pattern for each stage. In this process, a coset graph can even be decomposed into another coset graph recursively to find simple control algorithms or to simplify the coset graph more. Since the smallest number of hops doesn't give the simplest control algorithm, various configurations should be tested to find good balance between latency and control. Another way to use this coset graph is to find the number of hops as we did in the first method, but in an accumulative way. That is, after we construct several stages of the networks and visit some intermediate nodes, we only need to visit those nodes that were not visited yet. This method is actually constructing a multiple input port network similar to one we proposed in section 4. If we find a sequence of generators that have an identical inverse, this method can be used to build a multiple output port network as well. In fact, Benes network is the example.

In our future work we plan to: a) develop efficient non-overlapping segment detection algorithms for central processors in the method II of Benes network; b) carry out performance analysis for a modified Batcher/Banyan network (multiple output ports) and a radix sorter network (multiple output ports) under various loads with and without smart scheduling; d) design

of algorithms for a parallel compiler or a scheduler to find optimal scheduling of processes for multiple latency interconnection; e) design MINs using coset graphs.

Our studies on this topic are reported in more detail in [1].

D Summary

Our work on Smart Photonic Networks and Computer Security for Image Data has made significant progress in three major areas, switching and processor architectures, coding and security, and basic technologies. We have developed new very high-speed switching/routing architectures and very high performance processor architectures, and associated simulation tools. We focused on network and multi-media architectures that take advantage of—or extend—existing standards whenever possible. Our VLIW architectures and FPGA based work has the potential of a significant inpact on very high performance processor architectures for communications and Multi-Media, or more generally on latency tolerant applications. Our FPGA based work, started under this BMDO/ARO contract, has led to our current DARPA contract on reconfigurable wireless communication systems (Multi-Band, Multi-Service Wireless Infromation Transfer Systems.) We are studied how to support very high-throughput proccessing and very low latency interconnects with very high-speed photonic or very high density quantum dot based technologies (10^11 gates per chip.) In coding and security, our new space-time code diversity based algorithms for image and other types of data, are capable of providing or extending simultaneous security (privacy, integrity, reliability, availability, and protection from reverse engineering) for photonic network architectures. This was achieved by exploiting parallelism and scalable hierarchical multiplexing schemes, that are well matched to photonic and nanotechnology based network architectures.

E List of Manuscripts Submitted or Published

[1] H.-J. Lee, M. Morf, and M.J. Flynn, "Packet Switching Photonic Network Switch Design and Routing Algorithms," Stanford University, Computer Systems Lab Technical Report, CSL-TR-97-734, September 1997.

[2] J. A. Trezza, "Creation of Efficient Quantum Well Optoelectronic Switches," Stanford University, E.E. Dept. Ph.D. Thesis, February, 1995.

[3] J. Powell, "Systems Applications of Vertical Cavity Modulators," Stanford University, E.E. Dept. Ph.D. Thesis, February, 1997.

[4] K. W. Rudd, "VLIW Processors: Efficiently Exploiting Instruction-level Parallelism," Stanford University, E.E. Dept. Ph.D. Thesis, August, 1998.

[5] A. Yu, "Low Complexity Video-Encoding and Quality Metrics" Stanford University, E.E. Dept. Ph.D. Thesis, to be submitted December, 1998.

[6] Alice Yu, Ruby Lee, and Michael Flynn, "An Evaluation of Video Fidelity Metrics," COMPCON Digest of Papers, San Jose, California, pp. 49-60, February 23-26, 1997.

[7] Alice Yu, Ruby Lee, and Michael Flynn, "Early Detection of All-Zero Coefficients in H.263," Proceedings of the Picture Coding Symposium, Berlin, Germany, pp. 159-164, September

10-12, 1997.

[8] Alice Yu, Ruby Lee, and Michael Flynn, "Performance Enhancement of H.263 Encoder Based on Zero Coefficient Prediction," Proceedings of the Fifth ACM International Multimedia Conference, Seattle, Washington, pp.21-29, November 9-13, 1997.

[9] J. Campello, J. T. Gill, M. Morf, and M. J. Flynn, "Smart photonic networks and computer security for image data," SPIE International Symposium on Voice, Video, and Data Communications, Dallas Texas, November 1997.

[10] M. Morf, M. Flynn, "Architectural Issues and Nano-Technology," special session on Photonics and Nano-Technology, chair L.S. Lome, BMDO, Optical Society of America meeting, October 1998.

[11] W.H. Mangione-Smith, B. Hutchings, D. Andrews, A. DeHon, C. Ebeling, R. Hartenstein, O. Mencer, J. Morris, K. Palem, V. Prasanna, H. Spaanenburg, Seeking Solutions in Configurable Computing, IEEE Computer Magazine, Dec. 1997.

[12] Oskar Mencer, Martin Morf, Michael J. Flynn, "Hardware Software Tri-design of Encryption for Mobile Communication Units," International Conference on Application Specific Signal Processing, 1998.

[13] Oskar Mencer, Martin Morf, Michael J. Flynn, PAM-Blox: High Performance FPGA Design for Adaptive Computing IEEE Symposium on FPGAs for Custom Computing Machines, Napa Valley, 1998.

[14] Oskar Mencer, Martin Morf, Michael J. Flynn, Pipelined CORDICs for Reconfigurable Computing The Sixth FPGA/PLD Design Conference & Exhibit, Pacifico Yokohama, Yokohama, Japan, June 24-26, 1998.

[15] Oskar Mencer, Mark Shand, Michael J. Flynn, "FireLink: A High-Performance Adaptive Firewire Interface," IEEE Computer, special issue on High Performance Network Interfaces, Fall 1998. (Digital Systems Research Center TechNote 1998-012).

F Technology Transfer

We have a number of ongoing cooperations and technology transfer related industry-university interactions:

1. B. Pezeshki, SDL Inc., Stanford alumni and codeveloper of the Stanford photonic modulator technolopgy, is the SDL representative interacting with us in developing and applying photonic technologies.
2. J.A. Trezza and J.S. Powell, now at Sanders (Lockheed Martin), continue to interact with us on X-gate/modulator technology. They are realizing an integrated Mach-Zehnder version of Princeton's TOAD switch for Optivision. This enables technology cross-comparisons.
3. A. Nowatzyk, formerly at SUN Microsystems, now at Digital's Western Research Labs in Palo Alto, is now interacting with us in developing photonic very high speed low latency network technologies for multiprocessor and network switching and routing applications. Nowatzyk was the architect of Princeton's photonic networking effort that used the TOAD device. He is now interested in our X-gate/modulator and related technologies.

4. We are interacting with J. Massey from the ETH in Zurich and Cylink on security algorithms. The latest versions (FASTER) seem to be better adapted to the fine grain photonic and FPGA technologies.