

APPLICATION-SPECIFIC OPTICAL INTERCONNECTS FOR EMBEDDED MULTIPROCESSORS

Neal K. Bambha
US Army Research Laboratory
Adelphi, MD, 20783

Shuvra S. Bhattacharyya
Department of Electrical and Computer Engineering
University of Maryland
College Park, MD, 20742

ABSTRACT

Embedded systems are distinguished from general-purpose computers in that they consist of special-purpose hardware and software optimized for a specific task. They are pervasive in Army systems, appearing in soldier radios, sensor systems, vehicle control, communication systems, and many other applications. This paper focuses on multiprocessor embedded systems targeted towards signal, image, and video processing applications requiring large computing power and having real-time performance requirements. As transistor sizes shrink, interconnects represent a significant bottleneck for embedded systems designers. Several groups are researching optical interconnects to cope with this trend. Optical interconnects enable new system architectures. These new architectures in turn require new methods for high-level application mapping and hardware/software co-design. In this presentation, we discuss high-level scheduling and interconnect topology synthesis techniques for embedded multiprocessors. We focus on designs that are streamlined for one or more digital signal processing (DSP) applications. That is, we seek to synthesize an *application-specific interconnect topology* for a multiprocessor DSP design. We show that flexible interconnect topologies that allow single-hop communication between processors offer advantages for reduced power and latency.

We have previously shown that multiprocessor scheduling algorithms can deadlock in the general case of a topology graph that is not strongly connected, or if communication is limited to be single hop. We have also demonstrated an efficient algorithm that can be used in conjunction with existing scheduling algorithms for avoiding this deadlock [1]. In this presentation we discuss the advantages of performing application scheduling and interconnect synthesis jointly, and present a probabilistic scheduling/interconnect algorithm utilizing graph isomorphism to pare the design space. We demonstrate the performance advantages that an application-specific interconnect topology can produce for several DSP bench-

marks.

1. OPTICAL INTERCONNECTS

In recent years, optics have played an increasing role in multiprocessor systems. Commercial high-performance computers now use fiber ribbons to connect multiple processing nodes. Other examples include storage area networks using fiberchannel, and optical clock distribution to reduce clock skew across a chip. Programs such as the DARPA VLSI Photonics [4] program are pushing to integrate photonics technology on a single chip. Intel is currently backing an effort to bring “fiber-to-the-processor” [3]. The idea is to break the processor to cache bottleneck by using an optical waveguide integrated on the processor chip.

2. CONNECTION TOPOLOGIES

Electrically connected multiprocessor systems generally have a regular interconnection pattern, due to the physical constraints imposed by two-dimensional circuit board layout. Some examples include ring, mesh, bus, and hypercube interconnect topologies. Using these topologies, communication between remote processors requires multiple hops, which increases both latency and power, and increases contention throughout the network.

In contrast, optically connected multiprocessors, particularly those utilizing free space optics and three dimensions, are free to utilize arbitrarily irregular interconnection networks. Once the signal is in the optical domain, there is very little attenuation, so the energy required to transmit a unit of data is essentially independent of distance. The required energy instead is a function of the number of electrical-to-optical conversions that must be performed [2], which in turn is determined by the number of hops. With single-hop schedules the overhead associated with routing data

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 00 DEC 2004		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Application-Specific Optical Interconnects For Embedded Multiprocessors				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army Research Laboratory Adelphi, MD, 20783; Department of Electrical and Computer Engineering University of Maryland College Park, MD, 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM001736, Proceedings for the Army Science Conference (24th) Held on 29 November - 2 December 2005 in Orlando, Florida. , The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Application	N	$\Delta(E)(\%)$	$\Delta(M)(\%)$
FFT1	7	16	8
Karp10	6	24	4
Irr	8	16	(2)
Qmf4	7	32	3
NN16-3-4	8	58	2
Sum1	6	1	4
Laplace	7	4	(3)
FFT2	7	12	2

Table 1: Reduction in communication energy ($\Delta(E)$) and makespan increase ($\Delta(M)$) of single hop schedule over three-hop schedule.

through intermediate processors is eliminated. Furthermore, due to the flexibility of the communication medium, it is generally possible to avoid multi-hop communication operations by simply activating direct communication channels between the source and destination processors. Together, these properties make it desirable to limit the number of hops per communication operation when exploring configurations (interconnection patterns and task graph mappings) for an optically connected, embedded multiprocessor.

3. SCHEDULING AND INTERCONNECT SYNTHESIS ALGORITHMS

In order to quantify this effect, we scheduled several DSP benchmark applications using our modified scheduling technique, which takes the number of hops as an input parameter. We scheduled the benchmarks with hop constraints of one hop and three hops, and compared the communication energy required. For our purposes, we assumed all communication tasks transferred the same number of bits, so the energy cost of all IPC actors was equal. Table 1 shows the reduction in the required communication energy for single-hop schedules over three-hop schedules for the benchmark applications. For these benchmarks, we found that any undesirable effect on the makespan of the additional constraint for single-hop schedules was very small, as can be seen in Table 1. In two of the benchmarks (Irr and Laplace), the makespan was in fact better (lower) when we limited the scheduler to single hops.

We present a genetic algorithm for synthesizing efficient interconnection networks for embedded multiprocessors. The algorithm works in conjunction with a list scheduling algorithm to jointly optimize both the schedule and the interconnect topology. The algorithm is able to account for different distributions of local vs. global (long) interconnect routing tracks via a processor *fanout*




Figure 1: Comparison of number of graphs (top) with $|P| = 5$ nodes to those that are isomorphically unique (bottom).

constraint.

Figure 1 compares the number of possible graph labelings (for a graph with 5 nodes and varying numbers of edges) with the number of isomorphically unique graphs. By searching only isomorphically unique topologies, our interconnect synthesis algorithm pares the design space significantly and searches more efficiently.

We evaluated our interconnect synthesis algorithm on several DSP benchmark application graphs. We calculated how the makespan improves as the maximum fanout constraint is increased. This amounts to an area/performance tradeoff in the system. We also compared the performance of systems with topologies available with electrical interconnects vs. optical interconnects. These topics will be described in the presentation.

REFERENCES

- [1] N. K. Bambha, S. S. Bhattacharyya, and G. Euliss. Design considerations for optically connected systems on chip. In *Proceedings of the International Workshop on System on Chip for Real Time Processing*, pages 299–303, Calgary, Canada, June 2003.
- [2] R.K. Kostuk, J.W. Goodman, and L. Hesselink. Optical imaging applied to microelectronic chip-to-chip interconnections. *Applied Optics*, pages 2851–2858, 1985.
- [3] B. Lemoff. Fiber-to-the-processor: REALLY high-density interconnects for computers. In *SPRC Annual Meeting*, September 2003.
- [4] <http://www.darpa.mil/mto/vlsi>.