Computer Architecture Research and the Real World

Arvind, Anton T. Dahbura, Alejandro Caro

Computation Structures Group
Memo 397

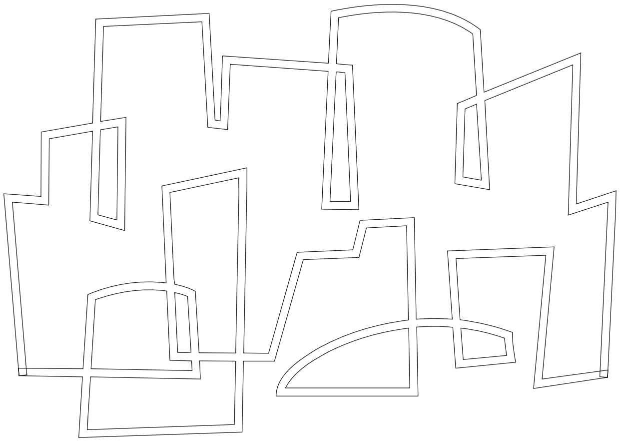

LABORATORY FOR COMPUTER SCIENCE

MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

Computer Architecture Research and the Real World[1]

Computation Structures Group Memo 397
April 23, 1997

Arvind

Anton T. Dahbura

Alejandro Caro

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Computer Architecture Research and the Real World[†]

Arvind[‡] Anton T. Dahbura[§] Alejandro Caro[¶]

April 23, 1997

Abstract

In the mid 1980s, the U.S. Defense Advanced Research Projects Agency (DARPA) decided to explore new ways to increase the speed and extent of technology transfer from academia to industry, especially in the area of computer architecture. As a result, DARPA organized a number of large computer architecture projects as partnerships between universities and companies. This article describes one of the earliest of these collaborations between the Massachusetts Institute of Technology's Laboratory for Computer Science and Motorola, Inc.'s Computer Group. This research effort demonstrated that university-industry collaborations can produce excellent results, but it also showed that such partnerships can be quite risky. The goal of this article is to share the authors' experiences and to make recommendations that will improve the chances of success of similar future projects.

Keywords: University-industry collaboration, computer architecture, parallel computers

In the mid 1980s, the U.S. Defense Advanced Research Projects Agency (DARPA)[1] funded a number of large computer architecture projects to be carried out by partnerships of universities and industrial organizations. DARPA had precise goals for this new type of research collaboration: it hoped to expedite the process of technology transfer from academia to industry, and it wanted to improve university access to the design and manufacturing expertise available in industry. In addition, DARPA expected industrial partners to pay for at least half of the cost of each research program.

This article focuses on the history of one of the earliest of these research collaborations conducted by the Massachusetts Institute of Technology's Laboratory for Computer Science (LCS) and Motorola, Inc.'s Computer Group (MCG). MIT and Motorola worked together on two parallel computing projects: the Monsoon Project [1] and the StarT Project [6], from 1989 until November 1995. These two projects demonstrated that a university-industry collaboration can be an excellent vehicle for performing systems research, but it also showed that such partnerships can incur substantial risk.

We describe the obstacles encountered during the Monsoon and StarT projects, and we summarize our experiences in a set of recommendations to fellow researchers. We hope these recommendations will improve the chances of success of future collaborations between industry and academia.

[†] This article is based on a presentation given by Arvind at the DARPA Principal Investigators meeting held in San Antonio, Texas, on February 15, 1996.

[‡] Arvind is Johnson Professor of Computer Science and Engineering at MIT. He was the head of the team at the MIT Laboratory for Computer Science that collaborated with Motorola in the Monsoon and StarT projects.

[§] Anton T. Dahbura was the director of the Motorola Cambridge Research Center. He is currently a consultant for Digital Equipment Corporation.

[¶] Alejandro Caro is a graduate student at the MIT Laboratory for Computer Science. He has been a member of the Monsoon and StarT project teams since 1989.

[1] DARPA has changed its name back and forth to ARPA several times. Throughout this article, we will use DARPA to refer to this agency.

The article also examines the issue of whether or not it is actually useful to build machines in an academic setting. We think this activity is more than useful; it is crucial to maintaining the preeminence of U.S. computer engineering research and education.

1 DARPA: An Agency with a Vision

The MIT/Motorola DARPA-funded projects were successors to several earlier DARPA-funded projects at MIT on implicitly parallel programming and dataflow architectures [2]. Until 1987, the research was conducted using an emulation environment consisting of a network of Lisp machines. Unfortunately, the capabilities of the emulators were eventually pushed to the limit, and important research problems, like resource management algorithms for large parallel applications, became impossible to study. The limitations of the emulation system also made it difficult to motivate students and research staff to carry out even more ambitious experiments.

Early dataflow results were promising, but credibility in the external research community began to suffer due to the lack of real hardware. It became obvious that systems had to be built, but MIT could not do it alone. The scope of the project and the expertise required to complete it dictated the need for an industrial collaborator. Partnering with a company also increased the chances that dataflow and implicitly parallel programming would become available to the computing mainstream.

At DARPA's request, MIT invited all major U.S. computer manufacturers for a day-long meeting in March 1988. Over 45 executives and engineers from 20 companies, as well as two DARPA program managers, attended. MIT's greatest challenge at this meeting was to convince the companies to accept the fact that there would be no commercial products at the end of the first three-year phase of the project. An industrial partner would have to commit to a long-term collaboration in order to reap the benefits of the research effort.

Motorola was an active participant at the March 1988 meeting. At that time, the company was emerging as a technological innovator and world-class manufacturer, and it was enjoying tremendous growth in the cellular phone and pager markets. Motorola managers were being urged to seek out "distinctive competences," and the company was eager to explore alternative computer architectures, like dataflow, perhaps due to its failure to anticipate the impact of RISC microprocessors. After six months of negotiations with three companies, Motorola became MIT's industrial partner.

2 The Monsoon Project

From 1988 to 1991, MIT and Motorola collaborated on the design and fabrication of research prototypes of the Monsoon dataflow machine. Monsoon was to demonstrate the feasibility of a dataflow machine and to provide a vehicle for executing large Id programs [5]. The decision to build Monsoon was taken at just the right time in the evolution of dataflow architectures and languages. Dataflow machines had been built in Japan (SIGMA-1) and Britain (Manchester Dataflow), so Monsoon was certainly not the first implementation effort. However, the early machines relied on extremely complex hardware which Monsoon did not require. So, Monsoon was a clean, simple design that enabled advanced research in parallel programming. It had to be built.

Monsoon required a significant amount of engineering (see box), but the project was exposed to little technological risk. The machine was built using well-understood manufacturing processes and conservative circuit technology. Most of the hardware design, as well as the Id compiler and run-time system implementation, was performed by MIT. Fabrication and low-level system software implementation were done by Motorola. In total, sixteen two-node Monsoon systems—each consisting of one processing element (PE) and one I-Structure memory module (IS)—were

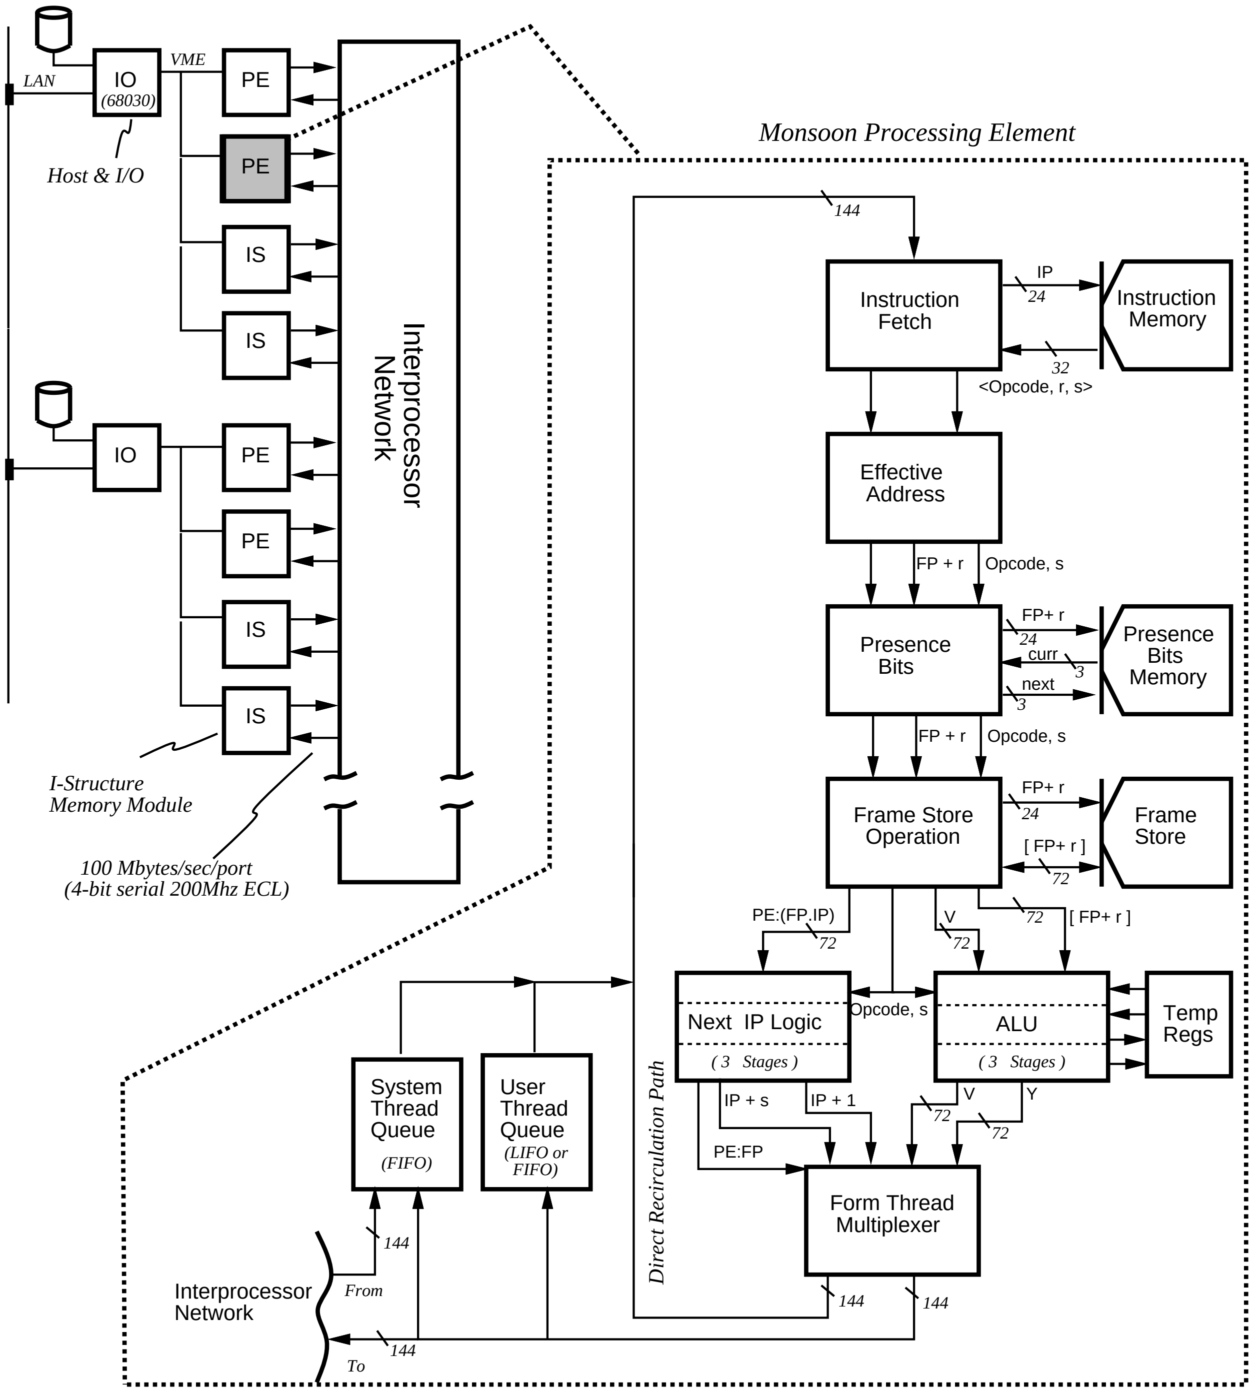

Figure 1: Monsoon Architecture

constructed and delivered to universities across the country. Two 16-node systems—with eight PE and eight IS modules—were delivered in 1991 to MIT and Los Alamos National Laboratories (see Figure 2).

A Monsoon parallel machine consisted of a collection of 64-bit, pipelined PEs connected to each other and to a set of IS memory modules via a multistage packet switch network [1] (see Figure 1). Each PE was implemented on a single 9Ux400mm surface-mounted printed circuit board that supported a VME port for diagnostics and input/output and two unidirectional 800Mb/second network links. The 10MHz processor core was byte-sliced into eight 10,000 gate 1.5 micron CMOS arrays. The floating point unit was fully pipelined, yielding up to 10 million double precision floating point operations per second. The PE contained a 256 KWord (32 bit) instruction memory, a 256 KWord (72 bit) data memory, and two 64 KWord (144 bit) thread queues.

Like the PE, the IS was implemented on a single 9Ux400mm surface-mounted circuit card with a VME port and two unidirectional network links. Each IS supported four million words (72 bits) and processed four million requests per second. The high-speed packet network consisted of 4x4 Packet-switched Routing Chips (PaRC), designed by MIT, located on each Switch board (SW) that provided a bandwidth to each port of 800Mb/second. Data Link Chips (DLC), also designed by MIT, and high-performance flex cables reliably transferred data between boards.

3 An Encouraging Success

The Monsoon Project was an overwhelmingly positive experience for all involved. The project finished on time and within budget, Id programmers wrote and executed substantial applications, and the team reported encouraging speed-up results [4]. In fact, several standard Id applications like GAMTEB—a Monte-Carlo photon transport code—frequently ran on huge data sets for hours at a time, something that would have been impossible on an emulated architecture. The behavior of these large parallel applications—studied in great detail using the extensive hardware monitoring capabilities of Monsoon [3]—motivated much of the work on resource management and scheduling that carried over to the StarT project.

The Motorola Cambridge Research Center (MCRC) near LCS was an additional by-product of the Monsoon project. This facility was staffed with a group of world-class researchers who proved invaluable in bridging the vast geographical and cultural distance between the Monsoon group at MIT and the MCG engineers in Phoenix, Arizona.

The partnership with Motorola made it possible for MIT to build a substantial number of high-quality prototypes. As interest in dataflow and implicitly parallel programming spread, MIT established several important and enduring collaborations with other computational science groups. MIT and Motorola also received numerous broad patents in dataflow technology, giving Motorola a strategic advantage in the field.

Once Monsoon systems were in the field and researchers had a chance to experiment, the limitations of the machines became apparent:

- Though good speed-up results were reported for a variety of important applications, the absolute performance of the machines was relatively low, as is typical for research prototypes; furthermore, the largest configuration was limited to only eight processors.

- Monsoon was unable to execute code written in standard languages such as C or Fortran due to the lack of compilers.

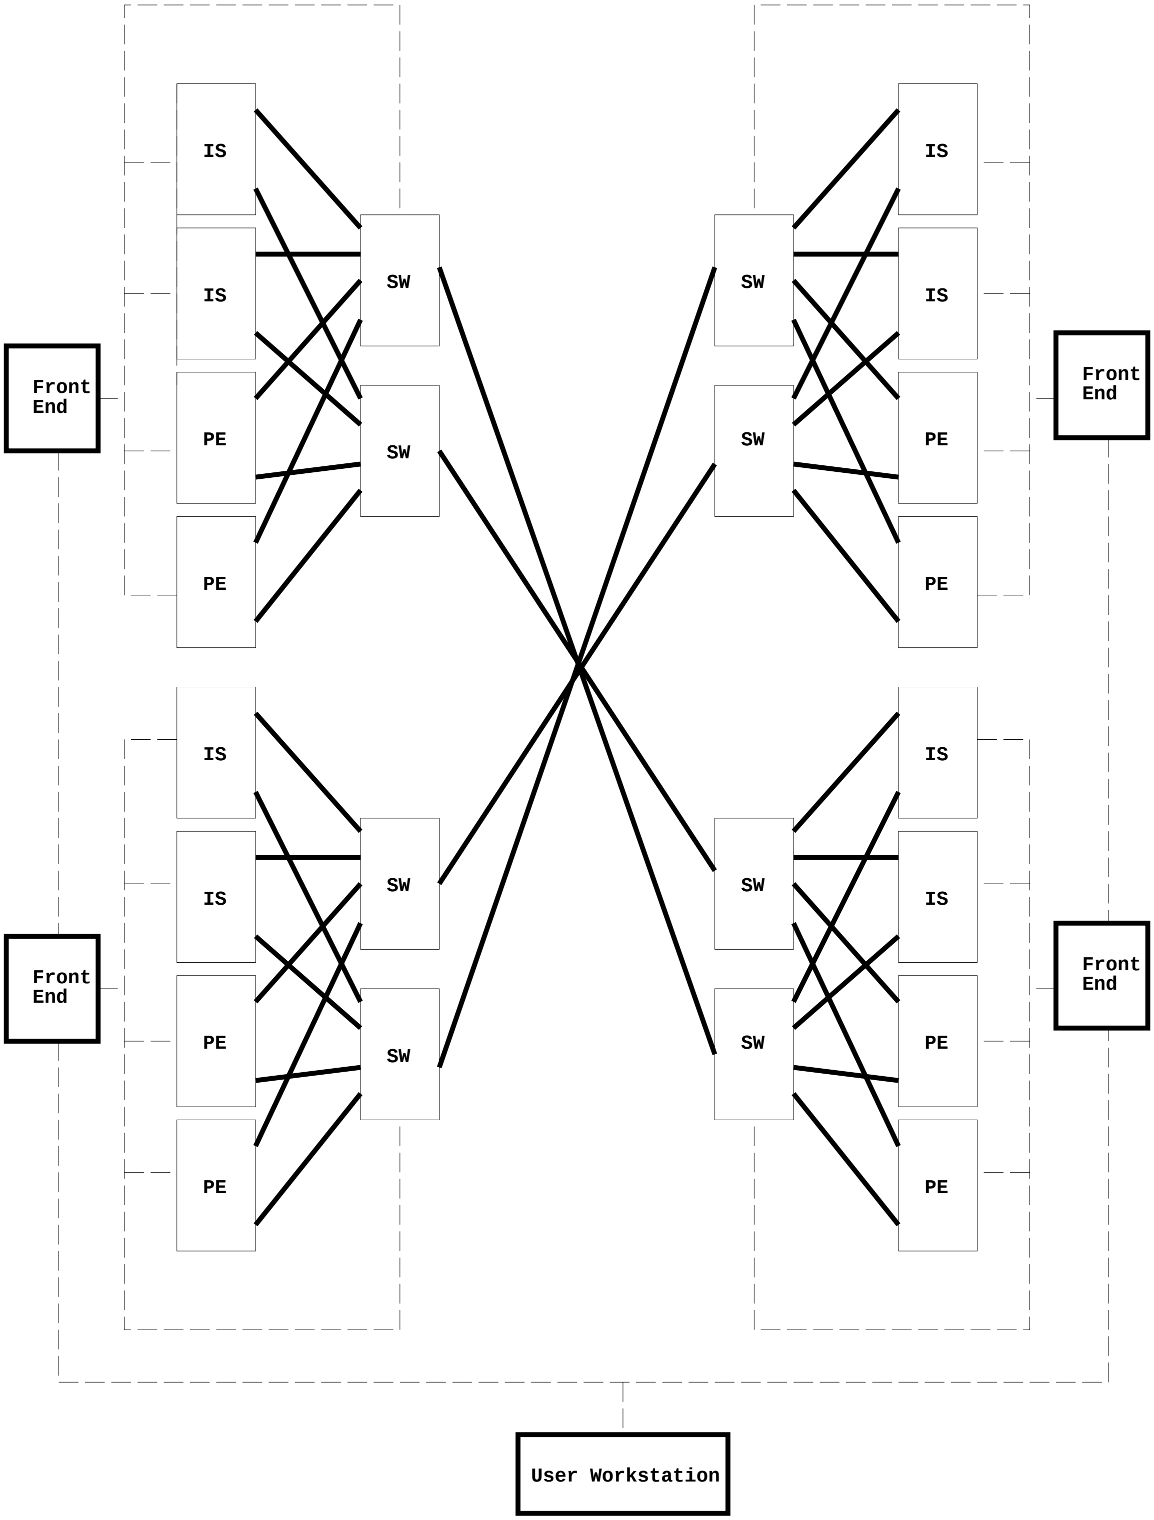

Figure 2: Sixteen Node Monsoon (to be replaced by photo?)

- The machines did not include adequate I/O subsystems, greatly limiting the codes that could be executed.
- The bulk of the hardware was based on custom and semi-custom circuitry instead of commodity components.
- The lack of a full operating system made program development difficult.
- It was a burden to provide maintenance to Monson machines outside of MIT.

With these successes and limitations in mind, the team faced a critical decision for the next phase of the research: should the next-generation system be a refinement of Monsoon to address its shortcomings, or should something more "commercially-viable" be produced? The decision was made to marry commercial RISC processors and dataflow technology in an attempt to produce a state-of-the-art parallel machine that would support fine-grained parallelism while maximizing the use of off-the-shelf hardware and software. These goals led to the StarT Project.

4 StarT: RISC and Dataflow Join Forces

The StarT Project was launched by MIT and Motorola in mid-1991. The original architectural plan [6] called for a collection of processing nodes connected in a fat-tree topology. Each node was to employ two modified RISC processors for executing computations, handling remote memory requests, and performing scheduling and synchronization operations. The architecture was designed to support most programming models and to obtain the greatest possible leverage from the use of existing hardware and software. The processor chosen for StarT was the Motorola 88110, a 32-bit RISC architecture that was the first superscalar microprocessor produced by Motorola.

The fat-tree network for StarT was based on the MIT Arctic packet routing chip [7]. Arctic was twice as fast as Monsoon's PaRC and was expected to drive the interconnection network at 1600Mb/second/link in each direction, while accommodating packet sizes ranging from 16 to 96 bytes. From the beginning, the project leaders understood that Arctic would be a complex piece of custom hardware, yet the risks of such an undertaking were outweighed by the fact that Arctic represented a significant advance in switching technology.

The interface between the processor and the network was perhaps the most ambitious part of the StarT design. Experience with Monsoon indicated that communication *latency* should be kept to a minimum if a machine is to exploit fine-grain parallelism in application codes. Consequently, the team designed a number of logic modules for the 88110 to act as a tightly-coupled network interface. These modules were to be integrated directly on the processor chip since the 88110 was designed from the beginning to incorporate additional functional units, such as graphics co-processors (see Figure 3). The research team viewed these modifications as a unique opportunity to add important features to a commerical microprocessor.

A subtle but important shift took place between the Monsoon and StarT projects. The bulk of the Monsoon design work was performed at MIT. In contrast, it was agreed that Motorola would assume a wider role in StarT. As Motorola proceeded with the circuit-level designs of the modified 88110 (called the 88110MP, for MultiProcessor), it committed significant resources to the project. Over thirty hardware engineers were working on the 88110MP design at its peak, many more than had ever worked on Monsoon.

The shift in responsibilities made MIT team leaders uneasy. MIT had little control over a major portion of the project and could not participate in the detailed design of the hardware. These were

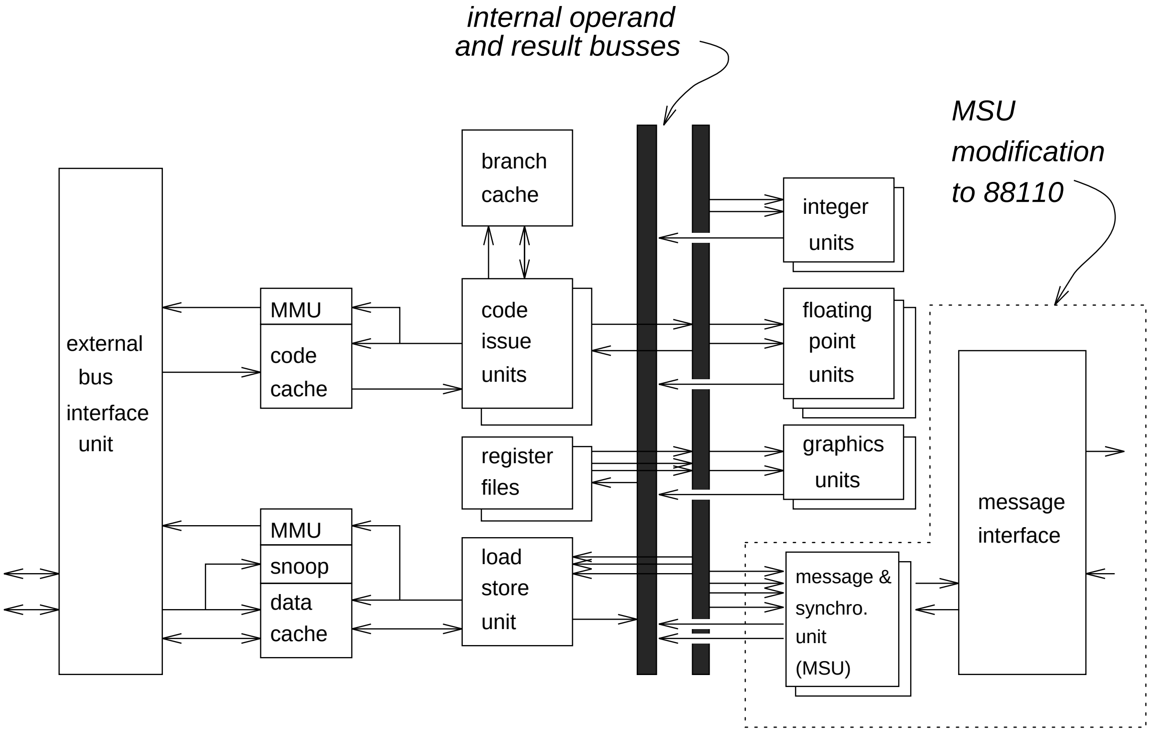

Figure 3: The 88110MP Architecture

the realities of collaboration with industry. MIT had to accept that StarT could not be built unless important design tasks were delegated to Motorola.

5 The Real World

Shortly after the StarT Project began, Motorola, IBM, and Apple Computer formed an alliance to design and manufacture the PowerPC family of RISC microprocessors [10]. This was a bad omen for the StarT project. All evidence indicated that the 88110 would, in the best case, become Motorola's secondary microprocessor priority. Yet, MIT was reassured by Motorola management that the PowerPC would not affect the project, and Motorola stated in no uncertain terms that the company remained committed to the 88110-based StarT architecture. But predictably, the intensity of the Motorola design effort began to wane and reached a low-point in June 1993. At that time, it was clear that completing the 88110MP was not a realistic goal given Motorola's new focus on PowerPC.

Even without the PowerPC muddying the waters, the StarT team learned a valuable lesson about the perils of developing a derivative microprocessor. The plans for the 88110MP called for it to be completed roughly at the same time as the mainline 88110 product, but by a team of engineers assigned to the StarT project. This meant that the 88110MP was based on an early version of the 88110 design, a design which was likely to be buggy. Since both development efforts operated concurrently, it became impossible for the 88110MP team to track the bug fixes implemented by 88110 team. After the initial 88110 design snapshot was accepted, the 88110MP team was essentially on its own.

The summer of 1993 was traumatic for the StarT team, both at MIT and Motorola. For technical reasons explained below, MIT seriously considered cancelling the project rather than switching to PowerPC. The Motorola Computer Group was also in turmoil as it attempted to assimilate the PowerPC family into its product plans. Finally, in a somber meeting of the entire technical staff, MIT decided to go back to the drawing board and to start afresh on a PowerPC-based StarT machine. The work of the past two years would largely be thrown away.

It took several additional months for the StarT team to formulate a redesign of the machine based on the much anticipated 64-bit PowerPC 620 processor [8]. Sadly, most of this delay was due to organizational barriers within Motorola that prevented information from flowing freely between its Microprocessor and Memory Technologies Group (MMTG), the group responsible for PowerPC design and manufacturing, and MCG. The MIT team slowly realized that MCG and MMTG operated as independent companies and that MCG was actually a minor customer of MMTG when compared to others like Apple Computer.

The switch to the PowerPC architecture forced significant compromises in the goals of the StarT project. Foremost, the PowerPC 620 made an on-chip network interface impossible. Unlike the 88110, PowerPC 620 dies lacked space for additional functional units. In addition, the development team for the 620 was driven by an aggressive time-to-market schedule. It could not afford delays to incorporate the research contributions of the StarT team.

The lack of an on-chip network interface changed the model of computation for the StarT machine. The overhead for sending and receiving data between the processor and the network increased from an estimated 10–20 machine cycles with the 88110MP to perhaps 200 cycles with the 620. The PowerPC-based StarT architecture gradually evolved into a cluster of shared-memory multiprocessor (SMP) sites. The StarT project maintained its research appeal, nonetheless, since the network was to be connected to the high-bandwidth second-level cache interface of the 620. Furthermore, by dedicating one of the processors from each SMP site to protocol processing, the design made it possible to build global cache-coherent memory across sites.

In early 1994, the bright future of the StarT project was clouded once again. Motorola announced that it was purchasing a significant stake in Companie de Machines Bull, S.A. The Motorola Computer Group forged a strategic agreement to obtain high-end PowerPC-based systems from Bull to resell in the U.S. in return for Bull reselling Motorola's low-end and mid-range systems in the European market. This agreement led to an immediate decline in Motorola's interest in developing its own high-end systems (like StarT) since Bull had already invested in its own in-house technology for clustering SMPs.

In late 1994 and early 1995, the frustration within the MIT group grew. Several key designers at Motorola left the project. To make matters worse, it became apparent in July 1995 that the PowerPC 620 was critically delayed. Motorola eventually advised MIT to modify the StarT design to incorporate a different microprocessor of the PowerPC family.

The architecture was redesigned once again—this time around a 32-bit PowerPC 604 SMP [9]. Each redesign of the machine consumed tremendous physical and emotional resources. Each also weakened the morale of the members of the entire technical team. Eventually in November 1995, a final blow was struck when the chief StarT architect at Motorola left the company. Soon after, Motorola informed MIT that it did not see any merit in continuing the project. Even if Motorola had wanted to go on, MCG now lacked the proper technical personnel.

After considerable deliberation, MIT decided to bring the whole project in-house. MIT and Motorola dissolved their research partnership amicably, and MIT entered into a joint-study agreement with IBM Research. The decision to complete the StarT project at MIT proved to be a blessing in disguise. Originally, team leaders worried that it would overwhelm the staff with too much work, but in truth, the decision rejuvenated the entire research effort. The joy of in-house control—gone since the days of Monsoon—returned, and there was a dramatic improvement in productivity and morale. The quality of the research also improved. Students and staff were not encumbered by the requirements of producing a commercially viable product; they were free to explore a much wider range of implementation ideas.

The latest version of the StarT architecture has been named StarT-Voyager (see Figure 4). It supports global cache-coherent memory and fast message passing for all sizes of messages. The new architecture has great research potential, due to its flexibility and the speed of the network, yet it bears little resemblance to the original StarT architecture and no similarities to Monsoon. The StarT-Voyager project is expected to produce two 32-site machines. Each site is to contain two PowerPC 604 processors, memory, and a Network Endpoint Subsystem (NES) to connect to the Arctic switch fabric. One of the machines will stay at MIT and one will be shipped to IBM Research.

6 The End of System-Building in Academia?

Should academics even try to build a machine like StarT? Some see the project as worthwhile but plagued with bad luck. Other critics, more disturbingly, take the problems of the project as proof that the era of building computers in academia has come to an end.

The critics point out that other mature engineering faculties do not undertake real-world projects: civil engineering professors and students do not go out in the jungle to build dams, and their aerospace counterparts are not building space shuttles. There is little reason, they argue, for building computers in an academic setting. It's obvious that industry can do so more cheaply and more efficiently. Academics should stick to small scale investigations of "paper machines."

We believe these arguments are short-sighted. Computer architecture is not a mature field. If

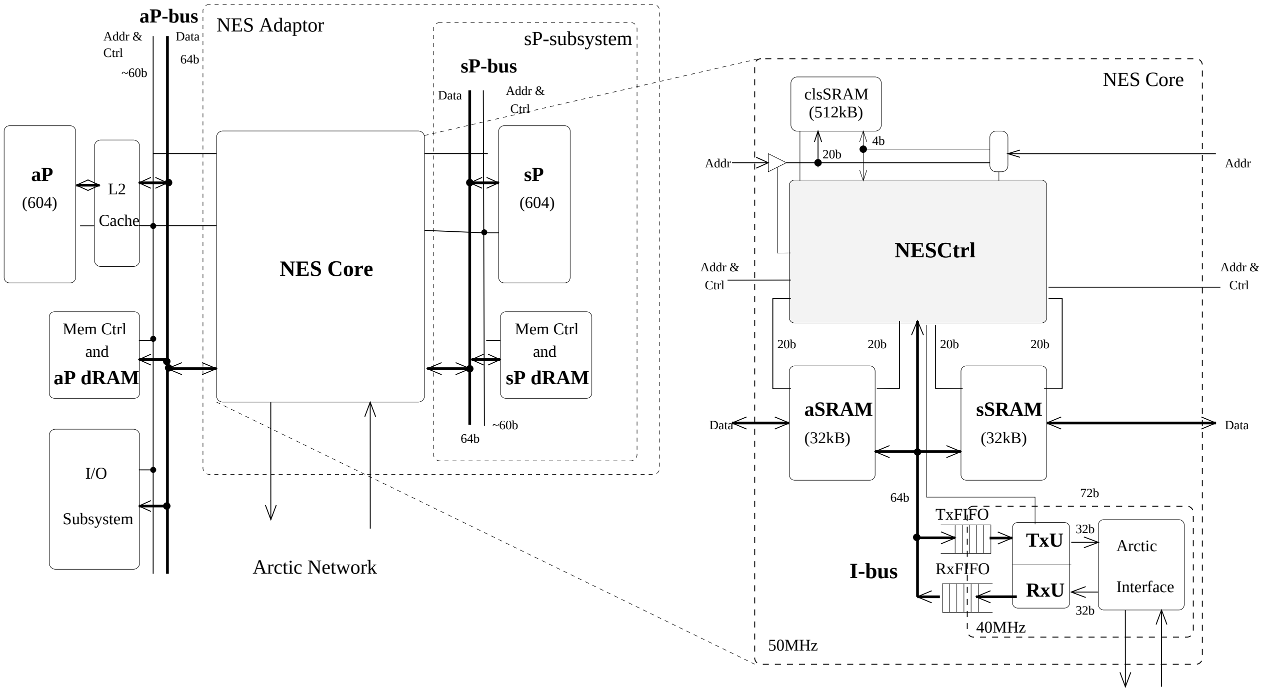

Figure 4: A Site in the StarT-Voyager Architecture (to be simplified)

it were, companies like Intel would have little incentive to design radically new architectures every two years. In the area of parallel computers, most commercial companies are still taking their first steps. Computer architecture is actually one of the fastest-changing fields of scientific endeavor.

Our experiences with Monsoon and StarT have shown that the challenges of building a real machine have a positive impact on the quality of research. In both cases, the research team members tended to “go the extra mile” because they knew their efforts would be rewarded with tangible results. Designs had to be specified in great detail and without any loose ends; all the interactions between the architectural mechanisms had to be understood completely for the machine to function properly. Nothing could be left to be patched-up later. Indeed, we believe that large systems projects cannot be designed properly without the expectation that they will be implemented.

In contrast, a paper machine project tends to emphasize simulation measurements over the implementation of mechanism. Unfortunately, these measurements are often of questionable quality since they rarely result from executing real-world workloads. It is hard enough to convince application developers to write large programs for prototype machines that are built; it is almost impossible to convince them to do so for simulated architectures. Application developers expect a new machine to offer a better way to solve a problem. Simulated machines do not provide this opportunity.

What has to be rethought in academia is the way that systems-building projects are planned and completed. The need for industrial cooperation must be taken as a given. The challenge, then, is twofold: how to innovate within the contraints of such a partnership, and how to organize partnerships so that they will succeed.

7 The Secrets of a Successful Collaboration

The Monsoon and StarT projects provided completely opposite experiences for the researchers involved. From the researchers' perspective, the Monsoon project was a great success and achieved all of its goals. The StarT project, though not a technical failure, suffered serious setbacks due to the weaknesses of university-industry collaborations. In this section, we discuss the most important factors that influenced the outcome of the MIT/Motorola partnership.

7.1 Beware of Market Forces

In the StarT project, market forces overwhelmed long-term research plans. Although no one could have anticipated the variety of strategic shifts and events that occurred, the project should have been designed to minimize the impact of any *reasonable* set of events. In addition, team leaders should have been prepared from the outset to redirect the project as obstacles were encountered. As it happened, redesigns consumed an inordinate amount of time because StarT team leaders had made few provisions for changes in the original project plans.

7.2 Don't Lose Sight of the Big Picture

The MIT team was very diligent during project redesign phases. Team members were eager to overcome setbacks and quickly immersed themselves in the technical details of each new design. In retrospect, they were perhaps too diligent. The MIT team became too reactive and tended to ignore the big picture, often overlooking the development of long-term negative trends. MIT researchers and engineers felt productive as they completed and refined new versions of StarT designs, but they became frustrated as new non-technical hurdles—which should have been spotted—continued to emerge.

7.3 Hold Realistic Expectations and Minimize Risks

Expectations for the StarT project were unrealistic on both sides. The project was viewed by Motorola primarily as advanced product development rather than as basic technology development. These are quite different undertakings. Product development requires a constant and often rapid reordering of priorities according to the demands of the market. Computer architecture research, while somewhat constrained by the availability of implementation technology, has a more general, long-term focus that is not affected by short-term commercial trends. Once a research project is complete, the resulting technology can be transferred to the commercial domain by the industrial partner via an advanced products group. The research program should not undertake both tasks at once.

The strong commercial bent of the StarT Project was due to two factors: first, MIT team members were intent on proving themselves to their industrial counterparts; and second, MIT wanted to overcome the perceived skepticism in the technical community regarding the commercial viability of dataflow architectures. Looking back, it is clear that the expectation that StarT would lead directly to products increased its dependence on anticipated, but non-existent, commercial technology. The degree of technological risk exposure was too high for what was supposed to be a long-term research project. Simply put, research projects must be resilient to changes in implementation technology. Commercial projects, on the other hand, can shoulder substantially more technological risk since the effects of cancelling a product are not as long-lasting as the effects of cancelling an entire university research program.

7.4 Project Plans Must be Resilient

StarT showed that once an industrial partner loses interest in a project, there are few incentives—including additional government funding—that can restore its commitment. Motorola involvement in StarT withered as the company gradually redirected personnel and resources to other programs. Even the Motorola Cambridge Research Center became isolated and largely irrelevant to the goals of MCG.

Personnel changes in particular had a strong negative impact on the health of the partnership. MIT/Motorola team leaders at first did not realize the extent of these effects. Over the seven years spanned by the collaboration on Monsoon and StarT, highly supportive managers, executives, system architects, and engineers at Motorola, MIT, and DARPA were replaced. Looking back, it is clear that the vision that existed at the beginning of the project was not inherited by subsequent generations of project management. Personnel changes have a lasting, profound effect on the success of a research project and should not be taken lightly.

The Monsoon project was a gratifying experience for all involved partly because the objectives of the project remained constant during its lifetime. The StarT project, on the other hand, encountered obstacles that forced several major redesigns. Eventually, one of the greatest challenges to the management team was to motivate the technical staff to persevere in the face of constant setbacks. Even the most dedicated team members became demoralized when they had little to show for their efforts after several years of work.

7.5 Information Flow is Crucial

Cutting-edge systems projects such as StarT often require that academics access confidential information of the industrial partner, such as proprietary source code, documentation, and product plans. Universities need to accept that their staff and students may be required to enter into non-disclosure agreements (NDAs) with companies. In turn, industrial partners need to be flexible in crafting NDAs tailored to the needs of academics rather than corporations. During the StarT Project, MIT students were prevented from participating in the hardware part of the project unless they signed individual NDAs with Motorola. MIT always refused to enter into an institutional version of such an agreement on behalf of the entire project team.

The MIT students who signed NDAs benefitted professionally from having access to the confidential information. They were able to learn a great deal about how processors, architectures, and systems are built in industry. Still, a few felt uncomfortable by the whole process and thought it had no place in an academic setting. From a social point of view, the NDAs also created two classes of students; the "haves" with access to information and the "have-nots" without access. This division had a detrimental effect on research dialogues at MIT.

The NDAs also made it difficult to give technical talks about StarT in an open forum. At the time, one of the most closely guarded parts of the PowerPC 620 design was its ability to connect to memory-mapped peripherals through a fast second-level cache interface. MIT presentations at academic conferences had to be worded very carefully to avoid direct references to this feature, yet this feature was the very linchpin of the entire design.

7.6 Lessons Learned

The following list summarizes our experiences with Monsoon and StarT. We believe these recommendations, had we followed them, would have dramatically improved the chances of survival of the MIT/Motorola partnership. We hope fellow researchers can put these to good use in future collaborations.

- Project expectations—of the industrial partner, the university, *and* the government funding agency—should be set appropriately from the beginning and must be reviewed periodically to ensure that the high-level goals of each party are realistic, that warning signs are not overlooked, and that the vision and level of commitment are maintained in spite of personnel re-assignments.

- The project should be viewed as technology development, not product development. The industrial partner should identify an internal applied research or advanced development group to be the interface with the university partner.

- Short-term market or customer requirements should not be allowed to affect the long-term research goals.

- Both parties should staff the project with first-class technical personnel familiar with their organization's culture. The StarT and Monsoon projects were extremely fortunate in this respect.

- Team management should have a good understanding of their partner's organizational culture and decision-making processes.

- The technical should be in geographical proximity, or the collaboration is likely to fail. Despite the presence of advanced communications facilities, face-to-face unscheduled meetings tended to yield the best results in the StarT and Monsoon projects.

- An efficient mechanism needs to be created to make sensitive company information available to the university partner (including students). In turn, universities need to ensure that their policies allow for students, faculty, and staff to sign non-disclosure agreements.

- An exit strategy should be established from the outset and should be invoked when the project cannot be completed. The project should be designed, however, so that the loss of interest by the industrial partner will not necessarily lead to failure. The StarT project had an exit clause in its legal contract, but the clause was written to be used when one of the parties defaulted on its obligations. It was not appropriate for the circumstances that existed when the MIT/Motorola collaboration was dissolved.

8 Questions for the Future

The biggest dilemma that planners of systems research projects face is the amount of technological risk that is acceptable. Some argue that technology itself has driven, and will continue to drive, innovative machine architecture. Yet if this is the case, is a project that implements a new architecture using older technology irrelevant? Is a proof-of-concept (conservative) implementation sufficient for an architectural innovation to find its way into the commercial domain? Or must systems research projects use the latest, most aggressive implementation technology? The answers to these questions will help determine the level at which systems projects will be funded in the future. Indeed, is it possible to build innovative systems less expensively? Or is it inevitable that only a small number of high-cost systems projects are possible?

One of the paradoxes of industrial-university partnerships is that companies that understand the importance of research already possess internal research organizations and are less inclined to support similar efforts in academia. On the other hand, companies that lack a legacy of research experience may be eager to collaborate, but find it tempting to view research as an immediate business opportunity. How can we encourage industrial participation in research projects while maintaining appropriate expectations for the outcome of these projects?

University research must remain compelling to industry. Unfortunately, the impact of many research projects is difficult to measure. It is rare that a project leads to a radical result that can be commercialized quickly. More commonly, the benefits of research projects are the continuous incorporation of new techniques, processes, and design philosophies into commercial products. Given this gradual process, how can we measure the value of the technology transferred by systems research projects, especially since many of these benefits are indirect and unexpected? Specifically, how can investigators maintain the interest of the industry in long-term research when it increasingly values short-term results?

The Monsoon and StarT projects demonstrated that university-industrial partnerships can excel at conducting state-of-the-art computer architecture and systems research. The resources contributed by an industrial partner, in conjunction with the design talent in academia, make it feasible to build and study real machines, the ultimate test of any systems discipline. But as a research community—one that includes academic leaders, industrial innovators, and funding agencies—we must answer a number of crucial questions that will make future collaborations viable, for neither industry nor academia alone can sustain the current rate of innovation in systems research.

Acknowledgments

We would like to thank Jamey Hicks, Chris Joerg, Boon Ang, Derek Chiou, and Larry Rudolph for their comments and suggestions on early drafts of this article.

References

[1] K. R. Traub, G. M. Papadopoulos, M. J. Beckerle, J. E. Hicks and J. Young, "Overview of the Monsoon Project," *Proceedings of the 1991 IEEE International Conference on Computer Design,* Cambridge, MA, October 1991. †

[2] Arvind and R.S. Nikhil, "Executing a Program on the MIT Tagged-Token Dataflow Architecture," *IEEE Transactions on Computers*, Vol. 39, No. 9, 1989. †

[3] G.M.Papdopoulos, "Program Development and Performance Monitoring on the Monsoon Dataflow Multiprocessor," in *Instrumentation for Future Parallel Computing Systems*, ACM Press Frontier Series, October 1989, pp.91-110.

[4] J. Hicks, D. Chiou, B.S. Ang and Arvind, "Performance Studies of the Monsoon Dataflow Processor," *Journal of Parallel and Distributed Computing,* July, 1993. †

[5] R.S. Nikhil, "Id Language Programming Manual," Computation Structures Group Memo 284-2, MIT Laboratory for Computer Science, Cambridge, MA, July 1991. †

[6] G.M. Papadopoulos, G.A. Boughton, R. Greiner and M.J. Beckerle, "*T: Integrated Building Blocks for Parallel Computing," *Proceedings of Supercomputing '93,* Portland, OR, November 15-19, 1993.

[7] G.A. Boughton, "Arctic Routing Chip," *Proceedings of the 1st International Workshop on Parallel Computer Routing and Communication,* Seattle, WA, May 16-18, 1994. †

[8] D. Chiou, B.S. Ang, Arvind, M.J. Beckerle, G.A. Boughton, R. Greiner, J.E. Hicks and J.C. Hoe, “StarT-NG: Delivering Seamless Parallel Computing,” *Proceedings of the EURO-PAR '95 Conference,* August 1995. †

[9] Boon S. Ang, Derek Chiou, Larry Rudolph, and Arvind, “Message Passing Support on StarT-Voyager,” Computation Structures Group Memo 387, MIT Laboratory for Computer Science, July 1996. †

[10] C. May, E. Silha, R. Simpson, H. Warren, eds., *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Morgan Kaufmann, San Francisco, CA, 1994.

†Available from `www.csg.lcs.mit.edu`.