

Tentative Agenda

High Level Issues

Monarch History

Supplementary Data on Monarch/MOSIS Interactions

Supplementary Data on Project Results and Innovations

128-processor Prototype

BBN's Understanding of DARPA Concerns Relating to Monarch

Brief BBN Thoth Project Overview

BBN Understanding of Ft. Meade Concerns about Thoth

Our Proposal for Incremental Funding for Monarch

Software Issues for Monarch

Intangible Issues

Other Interactions Potentially Involving Monarch

Appendix A: Other Interactions Potentially Involving Monarch

Appendix B: BBN Viewpoint on MOSIS

Discussion of DARPA, Ft. Meade, and BBN Lines of Communication

Detailed Review of Project Tasks and Expenditures to Date

Detailed Review of Proposed Tasks and Expenditures

Discussion of an Appropriate Replacement for John Goodhue

Monarch History

The initial contract estimate was too low -- because we mis-estimated, because DARPA encouraged us to be even more aggressive than our natural inclination led us to be, and we transitioned BBN's parallel processing engineering staff to BBN ACI

We were reluctant to pursue the project with total commitment (tools, staffing) during the first year when we were on risk funding and there was serious worry on our part that DARPA might lose the funding before the contract was signed

It took us longer than we expected to become efficient at VLSI design

There was a mutual decision between DARPA and BBN to depend on MOSIS and there has been considerable trouble with MOSIS (see supplementary data on the Monarch History with MOSIS)

We worked on the complete system design and implementation rather than producing interim milestones

We proposed a system with 1 MIP processors with 32-bit words; we have designed a system with 6 MIP processors with 64-bit words, a factor of 12 increase in millions of bits of data processed per second -- this performance increase has been the result of several important and generally useful engineering innovations (see the supplementary data on project results and innovations)

Better communication between DARPA and BBN would have produced a better controlled project and DARPA would have been more aware of the programmatic trade-offs which were being made and the consequent technical results and innovations

Supplementary Data on Monarch/MOSIS Interactions

MOSIS is a wonderful idea; its goal was to allow small VLSI groups to try new kinds of system architectures

Losleben, Squires, and Kahn all agreed that it was appropriate for BBN to be fodder for MOSIS, i.e., to push its capabilities -- to provide a substantial demonstration that MOSIS works

We were happy with this agreement, and still think it was the right thing to do

However, the DARPA/MOSIS plan for 1.25 micron technology was overly optimistic-- in retrospect they should have aimed at 2 micron.

Also, MOSIS was squeezed for funding (a mistake we believe), and was encouraged to do innovation to get funding rather than to focus on service; DARPA support in the community tools for MOSIS vanished, which hurt its acceptance by the community (although we don't claim it hurt us much)

We still believe in MOSIS; without MOSIS what we will do for under \$10M could cost five or ten times as much -- this is the point that Mead and Conway were trying to make

[We have previously provided DARPA with our recommendations for how MOSIS should be operated to be the national resource it should be -- see Appendix.]

Supplementary Data on Project Results and Innovations (including related results from Butterfly and Thoth)

In the course of our parallel processing work on the Butterfly, Monarch, and Thoth, we have produced many valuable results. These include new and novel ideas, techniques, and artifacts. Some are valuable new electrical engineering, some are new forms of algorithms, and some are new approaches to machine design. We have listed some of them here with a brief description of each. We propose that many of these results be published as quickly as possible to permit others in the DARPA parallel processing community to make use of the ideas.

Monarch System Design - The Monarch illustrates how a parallel computing system can be designed around the interconnection mechanism rather than around a processor or a memory. It illustrates the value of coordinating chip and system design.

Most computers are complex in their architecture because computer architect's tend to include special facilities for particular aspects of a problem, (e.g. auto incrementing addressing modes). The Monarch has an exceptionally straightforward structure which may be an example for other designs.

The Monarch is easy to describe — thousands of identical processors access a single multiported memory through a switching network. We expect that such a uniform architecture will let the programmer attack the parallelism in his program more effectively than other designs, removing machine imposed barriers such as declaring shared variables, local memory limits, and the coordination of message traffic. If a programmer can see parallelism in his application, he should be able to put it on this machine.

Because we have done a complete system design, the overall architecture of Monarch and the trade-offs that went into it can be published at once.

Algorithms for Large-Scale Parallelism - Our experience with Butterfly and the Thoth problems have led to several new ways to approach problems on a parallel machine. A few of these techniques are listed below.

Depth first search - We have developed a parallel algorithm for walking a search tree in a depth first manner using a large number of processors. In the past, parallel searches have been limited to breadth first searches, but depth first is much more appropriate for problems where the search space is large but there are many equivalent solutions.

Recursion Breaking - An algorithm which requires a number of steps, where each step depends on the previous result, is referred to as a recursion. We have developed several techniques that split the string of results into pieces so that the parallelism is enhanced.

Large Memory Techniques - In addition to a large number of processors, the Monarch has a large amount of memory. We have found many ways to precompute tables that aid in problem solution. Surprisingly, this often helps improve the parallelism in an application. These techniques may be useful for large memory uniprocessors (like the Princeton Massive Memory Machine).

Data Synchronization - The steal primitive provides efficient synchronization between parallel tasks with much finer gain than has previously been attempted in a shared memory machine.

Thoth Results - The Thoth analysis has produced very exciting theoretical results. We have now begun to run the Thoth benchmarks on the simulator (16-processors). The first benchmark simulated gives results better than the theoretical estimate. A thousand processor version of the simulator should be running soon.

Algorithms - In addition, we have examined numerical algorithms showing 10,000 way parallelism for dynamic programming, sorting, searching, pattern matching, bit operations and many others.

The above techniques can be published at once.

Electrical Engineering Results - The Monarch is designed around a high-performance interconnection network. We have developed techniques that allow operation at high data rates using normal CMOS technology. Some of the particular results are listed below:

Dynamic Delay Adjustment - In a normal system design, signalling between components of the machine is either synchronous or asynchronous. In a synchronous design, a central clock establishes when signals should be sent and when they should be received. Normally, the difference in delay between signals and the skew in clock phase as seen by different components limits the speed at which a machine can operate. This forces very special design techniques in the Cray and other machines where circuits are tuned to control skew and drives Cray to physically smaller machines.

The alternative is asynchronous design where the timing of signals is passed with the signals. Unfortunately, this requires arbitration when two events occur at nearly the same time. The uncertainty in which came first gives rise to the "classic synchronizer problem" and introduces additional delay.

With the help of Professor Lance Glasser, we have developed a third approach which relevant to systems where custom or semicustom LSI is being used. We distribute a central clock to establish an accurate frequency of operation, but introduce an artificial delay into every signal in the machine. This delay is automatically adjusted to correct for any signal skew. We have just recently received a patent on this technique.

This approach has greatly eased the design of the Monarch and has enabled us to consider a Thoth-sized machine. The technique will be of great value to many other system design projects.

High-Performance I/O Pads - The signalling rate of the Monarch will be at least 100 megabits per second, and theoretical analysis, timing simulations, and fragmentary empirical results from test chips indicate the signalling rate may be as much as 400 Mbps. This is unusual in an age where most CMOS is running at 16 to 33 MHz. We have developed special high performance pads in our Monarch design frame that can drive transmission lines at these high speeds. The techniques used here are probably of general interest in the DARPA VLSI community.

On-Chip Termination Resistors - We have integrated the line termination resistors into the Monarch chips so that we can turn them off when they are not needed. This provides better terminations, significant space savings, and this could dramatically reduce power consumption and thus requirements for heat dissipation. This work is similar to the recent work by Tom Knight (MIT and Symbolics).

High Speed CMOS design - We have been using MOSIS design techniques to make very high-performance CMOS designs for register files, switching networks, and other circuit components. These designs or techniques are at least 10 times faster than most other CMOS designs. The idea that CMOS can't run at these speeds is simply wrong, and our demonstration of high speed CMOS design will add to the list of successful high speed CMOS designs in the DARPA community.

We have theoretical results for each of the above at present and expect to have empirical results after the switch chip comes back from MOSIS (expected in February). We can publish theoretical studies now and empirical studies by next summer.

Switching Network Design - The Monarch switching network is unique. We are using switch nodes that have more outputs than inputs (e.g. 6 inputs routed to 8 destinations) to reduce the load on the outputs and reduce contention. The switches have two paths in parallel to each destination so that the probability of conflict is substantially reduced. This changes the network statistics so that it can run at close to full capacity without over congestion rather than having to run at quarter capacity to avoid congestion. It's a technique others should use.

Since we have switches with more outputs than inputs, there are a lot of wires between columns of switch nodes. In order to reduce this, we provide a statistical concentrator which reduces the number of wires between boards and reduces the number of switch nodes required in the second and third columns.

The Monarch design also provides for redundancy throughout. The system can be reconfigured around the failure of any processor, switch, concentrator, or memory subsystem in the machine.

This switching network design is well integrated into the Monarch computer but could be used in other systems as well. Thought of as a switch, it supports 8 Megabytes per second of memory reference on each port and can be built with from 4- to 64000-ports for a total bandwidth of from 32 Million to 0.5 Trillion bytes per second. Imagine using the Monarch switch as a very high performance very local area network, e.g. to connect the LISP machines in the MIT dataflow simulator.

We can publish the design, theoretical and simulation studies now; we can publish empirical studies by summer.

Processor Design - Our goal has not been to innovate in the design of the Monarch processor, but none-the-less, several innovations have occurred. The processor performs 64 bit operations. However, it is implemented in a serial fashion. At times, the data paths are two bits wide, and at others four bits wide, as appropriate for the situation. This has allowed us to hide the latency of the switching network.

The Monarch processor is mostly compatible with the MIPS core instruction set. However, in its implementation, there are aspects of a very large-word

processor. Instructions are matched to the pattern of periodic communication with memory. The processor may execute a memory reference instruction, then execute several register-register instructions or control instructions.

This is a nice enhancement of the reduced instruction set principles, and should be usable to increase several fold the speed of such instruction execution, on our own and other machines. We are able to publish this design now.

Memory System - The memory system of a computer is often an underrated component. In the Monarch, the memory system has been designed with as much care as the processor. The memory consists of a large number (half the number of processors) of independent memory modules. Each module is optimized for memory bandwidth through clever use of serial access modes.

The design also includes time multiplexing for better utilization, two-port memory system design, and tagged memory at little extra cost.

Our memory system design ideas can be published now; empirical results can be available in late 1988.

Control Network - In the Monarch design, every custom chip is connected to a monitoring and control network. (This network is in no way associated with the primary inter-processor connection network of the machine.) Initially, this was to allow us to parameterize a few chips and monitor a few statistics. We have been surprised at how powerful this notion has become. For example, using the control network, it is possible to test electrical continuity throughout the machine. Similarly, it is possible to run diagnostics on all of the chips in the machine from the control net directly. This is a level of diagnostic and maintenance support that may be revolutionary, and makes it possible to build reliable massively parallel machines.

The control network also allows control over the configuration of the machine. Switch outputs can be enabled or disabled, error detections recorded, and switch load measured directly. Even minor pieces of the machine can be controlled. For example, each switch has its own random number generator. The control network sets all of these generators to different values at initialization.

This technique should become a standard in DARPA VLSI projects. The control network ideas and design can be published now. Empirical results can be available in late 1988.

Packaging - Our board interconnection ideas have received some informal publicity. These should be described formally. Our ideas for controlled-impedance connectors could be used by others.

Tools - Other groups may also be interested in some of the tools we have developed.

We have developed a simulator for the Monarch which we are using internally and to sent to Ft. Meade. We would also be interested in providing this simulator to others working in the DARPA parallel processor community. It will be developed further under Thoth phase 2.

We have developed several VLSI tools to aid our design effort and a library of high-speed logic cells that may be of interest to other groups.

Services - We have built a small but very strong staff of high speed MOSIS CMOS designers (5 designers who have been with us for a while, 2 recent new hires). We have an extraordinarily creative team of parallel processing system architects and very extensive experience deploying parallel processors (50 Pluribus systems in mid-70's, 85 Butterflies in mid-80's, 228 parallel processor based Computer Image Generation systems in the last two years) and programming them for numerous applications (several communications systems, sonar signal processing, graphics processing, speech recognition, vehicle dynamics simulation, fluid mechanics, finite elements, AI, etc.). We could collaborate with other members of the DARPA community on new system architectures and applications.

128-Processor Prototype

We have been asked to provide a schedule and price to produce a 128-processor prototype.

A natural **phase 1** is to produce a pre-prototype testbed consisting of the following:

- o switch chip
- o memory controller chip
- o a rudimentary processor chip

Such a pre-prototype testbed will provide the following:

- o confidence that BBN adequately understands VLSI design
- o some data on the MOSIS yields for BBN's novel approach to a system architecture
- o many empirical results regarding power, cooling, signaling times, etc.

Based on a pre-prototype testbed milestone, DARPA can decide what it wants to do next:

- o decide that too much needs to be fixed and stop the project and publish the results
- o decide that the pre-prototype testbed shows that our design concepts and engineering to date mostly work and the project should continue

Phase 2 would include:

- o iterate on the switch design fixing any problems discovered in phase 1
- o iterate on the processor design fixing any problems discovered in phase 1 and giving it full capability
- o iterate on the memory design fixing any problems discovered in phase 1
- o implement the concentrator
- o iterate on our packaging and cooling designs based on the results of phase 1
- o construct a real 128-processor prototype

In the area of packaging there are two options for how the 128-processor prototype would be built -- a 128-processor prototype which is an end in itself, or a 128-processor prototype which is a step to a 1024-processor prototype. The latter is a little more expensive.

A 128-processor prototype will require MOSIS to be able to produce an adequate number of chips with adequate yields; it will also require additional system integration and engineering, but without high risk

A 1024-processor prototype will require MOSIS to produce thousands of production chips (or we will have to find an alternative supplier, which would require relearning some of our techniques); and like the 128-processor prototype, will require much additional system integration and engineering, but without high risk

BBN Understanding of DARPA Concerns Relating to Monarch

- a. Will the switching network work, e.g., problems with hotspots, etc.
- b. Is the architecture programmable by an experienced but normal programmer
- c. Will the architecture be cost effective
- d. Can BBN build it for an acceptable and predictable amount of money

There is also a question which DARPA has not focused on but about which other members of the parallel processing community have expressed concerns

- e. Are the high speed signaling technique and other architectural and electrical design details realistically implementable and useful

Ways of Reducing Concerns

- a. Run test cases on the simulator;

Explain in detail the technique of having messages back out of the switch and the reason this makes the Butterfly and Monarch switch work, and publish the switch simulation results;

Have an independent evaluator (e.g., Gottlieb from NYU) understand what we are doing and why it eliminates the worry about hotspots, etc.

- b. Run test programs on the simulator;

Where possible, ask Ft. Meade to report on the results of running their test programs on the simulator.

Make the simulator available to other parallel processor researchers in the community and let them try their programs on it

- c. Run test programs on the simulator;

Get empirical results from a pre-prototype testbed;

Present the results of discussions with a semiconductor manufacturer for implementation of Thoth

d. We can't build a machine for \$4.8M;

Normally a chip being built by a big team of people at a semiconductor company may cost as much as \$20M;

We can't prove we can build it;

We believe this is the biggest test of the DARPA/MOSIS concept that a smaller team can do novel system architectures using big VLSI for a modest amount of money, and we think we are succeeding;

We are attempting to do for under \$10M what might cost five or ten times that if done in a traditional manner;

We can only explain carefully and candidly what we have done (good and bad), how good the project team is, and BBN's commitment to the project, and hope DARPA will believe that it is a good gamble;

We should take the next sensible step at which we can succeed and have that be a milestone;

We should communicate partial results as we go along;

DARPA and BBN should be determined to remove any obstacles which arise;

We should involve other institutions as necessary to help the project succeed

If we succeed, it will be a breakthrough in both computer architecture and a demonstration of the DARPA/MOSIS concept.

e. Publish the key ideas;

Hold a seminar of the parallel processing community at which we present our ideas and techniques in detail (we believe that most well-known parallel processing experts spending sufficient time studying our ideas will become convinced of their value).

Listen to what other researchers say and watch what they do;

We bet that a number of the ideas and techniques from the Monarch design will find their way into other machine once other deigners understand them

Brief Overview of BBN's Thoth Project

Phase 1 -- Concept Study -- summer '86 to summer '87

- o CDC, ETA, BBN, TM, FPS, Goodyear, Gould, (IBM no bid)

Phase 2 -- Design Study -- Oct. 1, 1987 to Sept. 30, 1988

- o BBN, CDC, Loral/Goodyear/Convex

Phase 3 -- Jan. 1, 1989 to Dec. 30, 1991

- o One winner will build a machine

Evaluation criteria:

- o cost/performance/risk tradeoff
- o performance expectations significantly better than off-the-shelf from Cray, IBM, etc. in 1991
- o strong delivery date expectation
- o operational research tool -- replaces current system

Phase 2 schedule (excerpts):

- o preliminarily design review -- Feb. 1-2
 - hardware packaging design
 - hardware semiconductor teaming plan
 - software plan
 - any improvements in software benchmarks and analysis
- o simulator delivered -- Feb. 28
- o final preliminary design reports -- March 1
- o critical design review -- Aug. 1-2
 - details of our design
 - plan to be able to build it, e.g., teaming with vendor
- o [RFP received -- Aug.]
- o final critical design reports -- Sept. 1
- o performance analysis due -- Oct. 1
- o phase 2 final report due -- Oct. 31
- o [phase 3 proposal due -- Oct. 31]

BBN phase 2 funding:

- o \$1.287M

Phase 2 staffing:

- o **project management, administration, and documentation -- 12 man-months**
- o **hardware design -- 22 man-months (does not include time of semiconductor manufacturer)**
- o **software design -- 14 man-months**
- o **simulator work -- 6 man-months**
- o **performance analysis -- 10 man-months**
- o **integration plan -- 22 man-months**

BBN Understanding of Ft. Meade Concerns about Thoth

Ft. Meade thinks the Monarch architecture is very important because of its programming potential

Ft. Meade has the following worries:

- a. the signalling rate can't be achieved
- b. the dynamic delay adjustment technique won't work
- c. there is not a viable cooling and packaging concept
- d. it will cost as much as three times what BBN estimates to produce the machine, particularly for custom CMOS chips.
- e. BBN's ability to construct such a large machine

The above worries must be adequately satisfied by October 1988 in order for a contract to build the machine to be awarded; a finished prototype is not needed by then since Ft. has a three year schedule to do detailed machine design and implementation.

Ways of Reducing Concerns

- a. A pre-prototype testbed will produce critical empirical results
- b. Same as a.
- c. Same as a.
- d. Same as a;

Get data from a teaming partner, e.g., semiconductor manufacturer and/or computer vendor

- e. Team with an appropriate partner, e.g., semiconductor manufacturer and/or computer vendor.

Scale the machine down, e.g., to 16,000 processors and build several of those.

Our Proposal for Incremental Funding to DARPA

128-processor end-in-itself machine

- o phase 1 -- pre-prototype test bed -- gives DARPA what it needs to decide to go ahead with phase 2 or definitive negative research results -- gives us what we need for Ft. Meade
- o phase 2 -- build 128-processor prototype -- a 750 MIPS, 250 Mflops machine with fully sharable memory, which fits on several boards, is an interesting machine
- o optional phase 3 -- in additional or instead of phase 2, build a 128-processor machine aimed at producing a 1024-processor machine

Detailed task, schedule and cost estimates for phase 1 and phase 2 will be presented.

Software Issues

We should assume Mach and retargeting C to Monarch

Mach permits most Unix software to be used and will enable much Butterfly parallel processing operating system and utility software (e.g, performance analysis) and application software (e.g., math libraries) to be used; the language compilers (Fortran, LISP, and ADA) will have to be retargeted

No doubt software conversion from Butterfly to Monarch will require some reoptimization of the software to take advantage of the unique parts of the Monarch architecture

During the conversion, we should watch for opportunities for new algorithm innovation, e.g., the techniques required to use thousands of processors will no doubt be somewhat different than those required for hundreds of processors

In particular, we should take advantage of having a fully shared memory and steal instruction, e.g., massive use of a concept somewhat like future

Intangible Issues

This is one of the most experience and best parallel processing groups anywhere.

They have a beautiful machine design with enormous potential

They have a number of important engineering results and detailed engineering ideas and initial empirical results (from the switch chip) will be available in a few months.

They have some exceptional programming results -- see the Thoth study

Despite delays and increased costs, we feel that this project is one of the best at BBN, at DARPA, and in the U.S.; and it is important to BBN, DARPA, and the country for it to continue to a successful conclusion.

Other Interactions Potentially Involving Monarch

(see Appendix A for details)

BDM Optoelectronic Interconnect Study
USC/ISI Packaging Collaboration
Brandeis Lisp Study
Optoelectronic Technology Developers
Hybrid on Stacked Silicon (HYMOS)
DARPA Software and Architecture Contractors (several)
Neural Network Simulation
DEC (for Thoth)
National Semiconductor (for Thoth)

Appendix A: Other Interactions Potentially Involving Monarch

During the past year, we have had many interactions with other groups working with DARPA on research that is related to Monarch in one way or another. In three cases, we have proposed or actually started collaborative activities. In several other cases, collaborative activities were suggested to us but we did not have time to pursue them. This note summarizes some of the contacts that we have made.

BDM Optoelectronic Interconnect Study

The Defense Sciences office at DARPA has been sponsoring research into optoelectronic interconnection technologies for several years. During the fall and summer of this year, we put together a joint proposal with the BDM Corporation to study the possible application of the technologies under development in the context of the Monarch Architecture. The premise is that interaction between system builders in the architecture program and the materials scientists in the optoelectronics program will increase the probability that technology coming out of the optoelectronics program in three to five years will be useful in real computer systems. BDM is the proposed prime contractor they will be studying the applicability of optoelectronic interconnect technology to several (possibly all) of the architectures being developed under the Strategic Computing Program, and they will draw general conclusions about the technology. BBN will evaluate these technologies in the context of the Monarch.

The proposed procurement is sole-sourced to BDM, and BDM assigns >90% probability that a contract will be awarded in May. A copy of our proposal to BDM can be provided. Our proposal was incorporated directly into the BDM proposal. The DARPA program manager is John Neff. Our contact at BDM is:

Carl Friedlander
The BDM Corporation
7915 Jones Branch Drive
McLean, Virginia 22102-3396
(703) 848-7812
(703) 247-0360

USC/ISI Packaging Collaboration

The Advanced Production Technology Project at USC/ISI has been funded by DARPA to study packaging and interconnection technologies of potential use to members of the DARPA VLSI community building high performance systems.

We have been collaborating with ISI in two areas: Evaluation of Integrated Circuit Packages, and Evaluation of a polyimide multi-chip module technology.

In the first of the two collaborative efforts, we supplied ISI with a dozen or so test chips containing our I/O driver/receiver circuit, and they have collected samples of Integrated Circuit packages from several different vendors. They plan to bond our test chips into their packages and test the ability of each package to support high-speed signalling. ISI is using our test chips because they are the fastest circuits available in the VLSI community. The benefit to us is that it helps us find better packages for Monarch chips (we are afraid that our current packages will not work very well when we start using 1.2 micron CMOS technology). Our chips are at ISI. We do not know when they will start using them.

In the second of the two efforts, we gave ISI schematics for a subsection of the Monarch switch, which they used to design a multi-chip module. The module uses a high-density packaging technology that is being supplied by Microtec, a subsidiary of Augat. ISI is using our design as a demonstration vehicle because it stresses the key advantages of the Microtec technology: high wiring density, ability to dissipate lots of heat, and good impedance control. The benefit to us is that ISI is trying out a technology that is directly applicable to the Monarch but too new to be using in any of our prototypes. A summary of the project can be provided. The DARPA program manager is John Toole, of ISTO. Our contact is USC/ISI is Bob Parker. Ken Sedgwick has been the primary contact at our end.

The collaboration started late last spring, and ISI is nearly ready to send the module design to Augat for prototype fabrication. We have discussed the Monarch schedule slip with ISI, and ISI has discussed it with DARPA. DARPA and ISI have decided to put the project on hold until the results of our December 10 review are known.

Brandeis Lisp Study

Jim Miller, formerly a member of the Scheme development team at MIT, is now a professor at Brandeis. This fall, we put together an arrangement under which Jim has committed one of his graduate students for one semester to study Monarch Lisp implementation issues. In exchange, BBN ACI gave three Butterfly processor nodes to Jim's department at Brandeis. In the first stage of the project, they plan to code a set of basic Lisp operations (garbage collector inner loop, assq, memq, etc) to see how well the Monarch instruction set supports Lisp. So far, Jim's graduate student has spent a week or so on the project, and has produced code for one basic Lisp primitive.

Optoelectronic Technology Developers

Several developers of optoelectronic interconnect technology have expressed an interest in promoting their wares to us. It appears that DARPA has been pushing them to seek out contractors in the architecture program and explore possibilities for cooperative efforts. In all cases, the DARPA program manager is John Neff. Ken Sedgwick has been helping entertain and keep track of these contacts. Brief summaries follow:

Dean Tsang
MIT Lincoln Laboratories
863-5500 . X4427

Dean's group is developing free-space optoelectronic interconnection technology. They are interested in sending high-speed (1-10 gigabit/second) signals through air over moderate distances (less than an inch to 20 feet or so). They visited BBN twice during the fall of this year: once for a brainstorming session to see whether there might be a use for their technology in future Monarchs (we found several), and once to see whether they could convince us to submit a joint proposal to DARPA (we had no proposal-writing time available).

Don Alvarez
Naval Ocean Systems Center
San Diego
(619) 225-6221

Don's group is working on a technique for multiplexing high-bandwidth signals onto an even higher-bandwidth channel. He visited BBN last summer to discuss possible applications within the Monarch, and we came up with several. He called during the fall to say that he would be using the Monarch as an application example in his new proposal to DARPA. There have been no further interactions since then. We have a standing invitation to visit his lab at NOSC.

Kevin Kilcoyne
Rockwell International
Thousand Oaks, California
(805) 373-4207

Kevin's group is working on compact, low-cost high-bandwidth (multi-gigabit) fiber optic links. A basic theme of his work is that most fiber-optic technology has been optimized for expensive long-haul links for telecommunications applications (\$20K-\$30K for several kilometers). He is optimizing for computer systems, where the distances are much smaller and the number of links is much

higher. We have a standing invitation to visit his lab (Thousand Oaks is north of Los Angeles), but there have been no further interactions.

Allen Johnston
Jet Propulsion Laboratory
Pasadena, California
(818) 354-4054

We spoke with him once and put his telephone number in our optoelectronics file.

Hybrid on Stacked Silicon (HYMOS)

Myles Suer at Irvine Sensor Inc. (714-549-8211) contacted us. His company has developed a technique for stacking and interconnecting silicon chips in extremely dense modules. He told us that DARPA is funding him to build a memory module that packs 4 gigabytes into 100 cubic centimeters (smaller prototype version expected in early 1988). He plans to use 128 256-kbit static RAM chips, stacked vertically. The technology provides for 80 wiring channels on each side of the stack, and allows two of the four sides to be used (unless you want to be aggressive). His interest was similar to that of the Optoelectronics contractors. DARPA had suggested that he contact people in the architecture community to let them know about the technology and to discuss potential applications. Myles offered to stop by BBN on the afternoon of December 1. We declined politely. We have a standing invitation to visit his company in Irvine, California.

DARPA Software and Architecture Contractors

Steve Squires suggested that we communicate with the following fellow Strategic Computing contractors:

Trusted Information Systems: on the tail end of some Cronus business, Steven Crocker took some time to tell us about the work that he is doing on security issues for the Mach operating system. He has expressed an interest in collaborating with BBN on a study of security issues for Monarch.

Mach Development Group at CMU: BBN ACI has been working closely with this group on Butterfly Mach. We have had a few discussions with Rashid and others on Monarch.

NYU Ultracomputer Group: We made a get-acquainted visit to NYU; they are doing lots of interesting things, i.e., many of their ideas are close to ours. Pat Teller (of NYU) is visiting BBN in December.

Massive Memory Project: Dick Lipton at Princeton University is attaching a large memory (hundreds of megabytes?) to a Sun workstation so he can study the speed-space tradeoffs that occur when a large physical memory is available (reminiscent of our THOTH experience). We have talked with him about visiting in one direction or another, but have not followed through.

Len Kleinrock: His group is studying the statistical properties of multiprocessor interconnection networks. We have not contacted him.

IBM RP3 Project: Randy received an invitation from Pfister to visit. We have not followed up.

Arcadia: This is programming environment development project at UC Irvine. We have not followed up.

Camelot: Al Spector (one-time BBN summer employee) is now a professor at CMU working on a reliable distributed transaction management system called Camelot. We spoke briefly with Al after a talk that he gave at MIT, and again at a recent PI meeting. Camelot has several interesting properties with respect to Monarch: it uses the Mach kernel only, and might therefore run on a Monarch with a minimal Mach system; since it is a parallel application and assumes a very large (48-bit) virtual address space, it should run well on a Monarch.

Neural Network Simulation

The Tactical Technology Office at DARPA (program manager: Jasper Lupo) has commissioned a neural network study that covers five different areas: (1) comparison of standard processing techniques versus neural networks; (2) mathematical foundations; (3) simulation tools and technologies; (4) implementation techniques; (5) applications. The simulation tools and techniques part of the study is being run by Dr. Paul Kolodzy at MIT Lincoln Laboratories (271-0285). They are surveying the state of the art in computer architectures for simulation of neural networks in the range of 100-100,000,000 interconnections. Kolodzy visited BBN with four other members of his study committee on Tuesday, November 24. We met with them for a couple of hours. They will be including the Monarch in their report to DARPA. We offered to review the Monarch section before the report is submitted, and Kolodzy said he would send

us a draft. Incidentally, members of the committee had different opinions as to whether they wanted to hear about Butterfly, ranging from "absolutely not" to "why didn't you include the Butterfly in your presentation?".

Appendix B: BBN Viewpoint on MOSIS

**MOSIS STATUS, EXAMPLES, AND OPERATION:
AN ASSESSMENT**

**D. Walden
Bolt Beranek and Newman Inc.**

12/9/87

MOSIS Status

- o Semi-conductor technology is the central technology of the electronics hardware industry
- o MOSIS is a national strategic resource; it has facilitated a revolutionary advance in quick semi-conductor turn-around
- o At a time when the U.S. is worried about its industrial and defense competitiveness:
- o System architects now have an opportunity to realize systems which were not previously possible
- o Anticipated and needed military and commercial electronics developments depend on shrinking component size, especially the widespread use of general- and special-purpose parallel architectures
- o MOSIS permits rapid turn-around for designs, an important competitive advantage (e.g., the major Japanese advantage in consumer electronics has been said to be rapid turn-around)
- o MOSIS serves as a broker between users and many fabrication lines:
- o MOSIS provides exceptionally high leverage

How Should MOSIS be Operated?

- o MOSIS should remain independent of the semi-conductor industry
- o The MOSIS mission should be to remain state-of-the-art, not the lowest common commercial denominator
- o MOSIS facilities should be expanded
- o MOSIS should accept government and commercial work
- o ISI is an ideal manager for MOSIS

How Should MOSIS Be Funded?

- o MOSIS should have a two-tier pricing policy
- o Continue to facilitate multi-project chips to enable costs to be shared across users
- o Greater funding should be applied
- o Adequate funding of MOSIS could put the U.S. in a much better position to compete with the Japanese in micro-electronics
- o MOSIS is an extremely important capability which it is critical for the government to provide for electronics in the U.S.
- o Government and DoD policy setters should declare provision of MOSIS to be a keystone of the DARPA charter

Appendix: A Few Examples of MOSIS Use

- o Lincoln Laboratory wafer-scale integration work
- o Development of the pioneering RISC instruction computers at Stanford and Berkeley
- o North Carolina graphics pixel processor
- o Complex signal processor at MITRE
- o Cal Tech mosaic cube and other DARPA advanced computer architectures (e.g., BBN Butterfly and Monarch chips)
- o Intelligence community decrease in turn-around time for frequent routine designs
- o Dozens of DARPA contractors and universities are graduating experienced VLSI designers