

# The Aries Decentralized Abstract Machine (ADAM)

Andrew Huang

Artificial Intelligence Laboratory  
 Massachusetts Institute of Technology  
 Cambridge, Massachusetts 02139

<http://www.ai.mit.edu>



**The Problem:** Large systems are challenging to manufacture and to operate with high availability. The problem extends from the manufacturing yields and reliability of individual chips to the integration of cabinets and cables, and it must be addressed as an integral part of any massively parallel architecture. This problem is confounded by the fact that the average service life of a massively parallel machine is fairly long, and maintaining a base of exact replacement parts becomes difficult as manufacturing processes evolve. Ideally, one would like to be able to replace worn-out parts with newer, faster parts without having to recompile or even restart the running applications.

**Motivation:** The high cost of ownership and diminishing performance per dollar of full-custom, massively parallel computers has stymied research and development in this area; it seems that the classic big-iron machines has succumbed to the attack of the killer micros. However, here in the future of state-of-the-art pure-play silicon foundries, there is less of an economic tie between an architecture and a process technology. In addition, silicon is becoming cheap enough relative to the cost of packaging and testing that providing redundancy at the chip-level to increase yield and reliability is an attractive possibility.

**Previous Work:** There is an extensive body of previous work in the arena of high performance computing. The specific seminal works that the architecture described here stems from include decoupled access/execute machines [7], dataflow machines [3], processor-in-memory/chip multi-processor machines [5], named-state techniques [4], and architectures, languages and runtime environments that feature dynamic latency reduction and/or load balancing, such as Cilk [6], Active Threads [8], Threaded Abstract Machine [2], and COOL [1].

**Approach:** The Aries Decentralized Abstract Machine (ADAM) provides an abstract programming model (figure 1) that is independent of implementation details. It is in many ways similar to the java virtual machine concept, except that it is tuned for high performance in a massively multithreaded environment with distributed shared memory. This layer of abstraction allows the underlying machine to be incomplete, imperfect, or inexact in many important ways.

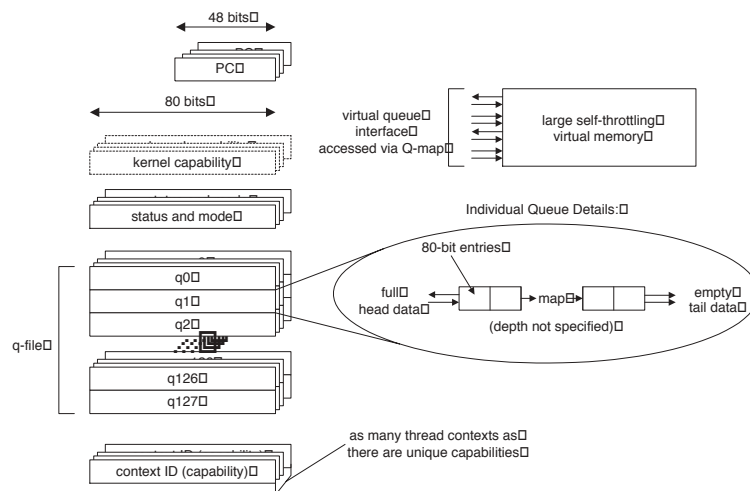


Figure 1: ADAM programming model

Like a JVM, an ADAM is started for each application that the user desires to run on the target machine. Thus, partitioning a job between several user jobs is a matter of distributing the nodes between the ADAMs. Since the ADAM run-time model supports efficient dynamic data migration, load balancing between jobs can be done by forcing a job's objects to migrate out of some nodes and moving those nodes into other job's allocation pools.

Static failures due to manufacturing defects can be avoided by coloring the node allocation map. Since the basic structure of ADAM presumes a PIM-style chip-multiprocessor, manufacturing yields can be improved since any die with at least one functioning processor could be considered salable die. Defects at the board and cabinet level can also be isolated at system integration time, allowing the machine to be come available before it is entirely assembled and debugged.

Since ADAM is not tied to any specific implementation, a machine that implements ADAM can be evolutionary. In its early stages, it may use emulation or emulation-assisted hardware; as the design is debugged and validated and as semiconductor processes improve, nodes can be switched over to higher-performance dedicated hardware. This can be done on the fly by forcing all data to migrate out of nodes to be serviced and bringing them off-line.

Dynamic failures can also be compensated for using the ADAM concept. Users can start multiple ADAMs running the same program on different nodes of a machine; correct results can be determined on the basis of a vote. Early-warning hardware monitors that track power supply health, bit error rates, and temperature can activate software that forces data to migrate out of a failing node. Also, ADAM's feature of isolating the entirety of a thread's state into a single capability simplifies the implementation of transactional rollback. The state of a thread before an operation can be copied to a backup location, which can be used as a rollback image in case the operation is aborted or failed.

**Impact:** The decoupling of a programming environment from the details of the hardware implementation gives the system architect a powerful tool for dealing with cost of ownership issues and development issues. The use of a system design similar to the ADAM vision would hopefully reduce time-to-solution and keep development costs at a level where the reward is worth the investment. Software teams can develop compilers that target the abstract machine concurrently with hardware teams that are developing the custom silicon high-performance solution. Significantly, machines with performance and cost points that are well below the ultimate machine's can be used for development purposes. Thus, application developers do not have to waste valuable machine runtime on the "big iron" to debug code; they can debug at their desk on a scaled-down implementation.

**Future Work:** The system design concept enabled by ADAM is currently being validated by the concurrent An-sible hardware implementation effort and the Couatl compiler development effort. The ADAM specification is currently defined by its gross parameters, but many of the important details such as the exception handling mechanism and the data migration mechanism are still under development.

**Research Support:** Support for this research was provided by the Air Force Research Laboratory, agreement number F30602-98-1-0172, "Active Database Technology".

#### References:

- [1] Rohit Chandra, Anoop Gupta, and John L. Hennessy. Data locality and load balancing in cool. In *Proceedings of PPPP IV*, pages 249–259. ACM, 1993.
- [2] David E. Culler, Anurag Sah, Klaus Erik Schauer, Thorsten von Eicken, and John Wawrzynek. Fine-grain parallelism with minimal hardware support: A compiler-controlled threaded abstract machine. In *Proceedings of ASPLOS IV*, pages 164–175. ACM, 1991.
- [3] Ben Lee and A.R. Hurson. Dataflow architectures and multithreading. *IEEE Computer*, 1994.
- [4] Peter R. Nuth and William J. Dally. The named-state register file: Implementation and performance. In *Proceeding of the First IEEE Symposium on High-Performance Computer Architecture*, pages 4–13, 1995.
- [5] Kunle Olukotun, Basem A. Nayfeh, Lance Hammond, Ken Wilson, and Kunyung Chang. The case for a single-chip multiprocessor. In *Proceedings of ASPLOS-VII*, Cambridge, MA, 1996. ACM.
- [6] Keith H. Randall. *Cilk: Efficient Multithreaded Computing*. PhD thesis, MIT, Department of Electrical Engineering and Computer Science, 1998.

- [7] James E. Smith. Decoupled access/execute computer architectures. In *Proceedings of the 9th Annual International Symposium on Computer Architecture*, Austin, Texas, 1982.
- [8] B. Weissman, B. Gomes, J.W. Quittek, and M. Holtkamp. Efficient fine-grain thread migration with active threads. *Proceedings of the First Merged International Parallel Processing Symposium and Symposium on Parallel and Distributed Processing*, pages 410–414, 1998.