

# CRAY MTA



## *The Promise of Parallelism Realized*

There are three fundamental factors that limit the scaling of conventional parallel computers:

- Overhead resulting from communication and synchronization
- Imbalance in processor workloads
- Inability to exploit multiple levels of parallelism in real-world applications

As an integrated system of software and hardware, the CRAY MTA was designed to eliminate these limits to parallel performance. It begins with Cray's powerful compilers automatically parallelizing, as threads, every level of a program's hierarchy. These threads run on up to 128 RISC-like hardware streams per processor. Each processor issues an instruction from one of its ready threads at each cycle, switching between threads at no cost. Each stream has the same components as a conventional processor—instruction counter, register set, stream status word, and target and trap registers. While some streams wait for memory operations to complete, others use the processor's resources to move their threads along, enabling the processor to tolerate even long memory latencies while performing useful work. About 40 active streams per processor effectively overlap all memory latency with productive computation, achieving significant parallelism within a single processor.

The interconnection network is capable of delivering 8 bytes of memory to each processor every clock

period. This bandwidth scales with the number of processors, enabling a scalable, flat-shared memory. Synchronization between threads is done in memory, at negligible cost. Issues of data locality and the complexities of message passing protocols are eliminated, making programming far simpler than on distributed memory machines.

Cray's Multithreaded Architecture supercomputers represent a fundamental breakthrough, whose impact is just beginning to be felt—scalable, easy to program, parallel computing.

### *Hardware*

CRAY MTA systems are constructed from system modules. Each module contains:

- a computational processor
- an I/O processor
- memory units
- network routes

The network interconnection is capable of supporting data transfers to and from memory at full processor rate in both directions, as well as all of the connections between the network routing nodes themselves.

Just as CRAY MTA system bandwidth scales with the number of processors, so too does its latency tolerance. The current implementation can tolerate several hundreds of cycles of memory latency, representing a comfortable margin; future versions of the architecture will be able to extend this limit without changing the programming model as seen by either the compilers or the users.

### *Software*

A sophisticated easy-to-use parallel programming environment is provided with the CRAY MTA. Fortran 77, Fortran 90, C and C++ compilers offer a high level of automatic parallelization. Compiler analysis and performance programming tools are available. These tools, CANAL and TRACEVIEW, have a user-friendly graphical interface.

The CRAY MTA's scalable, uniform, shared memory allows fast prototyping of parallel code and high levels of programmer productivity. Scientific application programmers are freed to concentrate on solutions, not computer science.



## CRAY MTA

### Technical Specifications

- 64-bit data, addresses, and instructions
- Up to 128 threads per processor
- Up to 8 concurrent memory references per thread
- 1EEE 754 floating point arithmetic
- No data caches
- 8KB level 1 and 2MB level 2 instruction caches
- Fortran 77, Fortran 90, C, and C++ with customary extensions
- Automatic parallelization and vectorization
- Interprocedural analysis and optimization
- Symbolic debugging of optimized parallel code
- Graphical performance debugging tools
- Transparent, scalable parallel I/O
- 64-bit fast file system with variable block sizes
- Multi-user support for large and small tasks
- Checkpoint/restart capability
- Water cooled
- Automatic logic diagnosis via full scan

### CRAY MTA SYSTEMS CONFIGURATIONS

Model	Processors	Memory	Performance	Bisection Bandwidth
MTA 16	16 CP	64GB	>12 GFLOPS	125GB/s
MTA 32	32 CP	128GB	>24 GFLOPS	250GB/s
MTA 64	64 CP	256GB	>48 GFLOPS	500GB/s
MTA 128	128 CP	512GB	>96 GFLOPS	1,000GB/s
MTA 256	256 CP	1TB	>192 GFLOPS	2,000GB/s

**CRAY**

Corporate Headquarters  
411 First Avenue South, Suite 600  
Seattle, WA 98104-2860  
USA  
phone (206) 701-2000  
fax (206) 701-2500  
[www.cray.com](http://www.cray.com)