
# Cray X1 Software

Don Mason

May 15, 2003

# Topics


Cray Software Roadmap

Cray X1 Software Components

- Programming Models
- Programming Environment
- Operating system

Release Plans

Cray Software Documentation

# Software Roadmap


Cray
Programming Environments

2010
2009
2008
2007
2006

Sustained Pflops

Black Widow ++

Black Widow +

Black Widow

Cray X1e

Cray SV1ex:

Cray T3E-1350

Cray X1
(2h02)

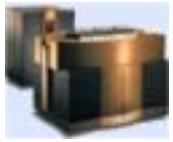

Cray T90

UNICOS/mp

UNICOS & UNICOS/mk

# Cray X1 Software Components

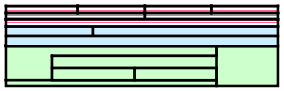


**Cray Programming Models**

**Cray Programming Environment**

**Cray UNICOS/mp O.S.**

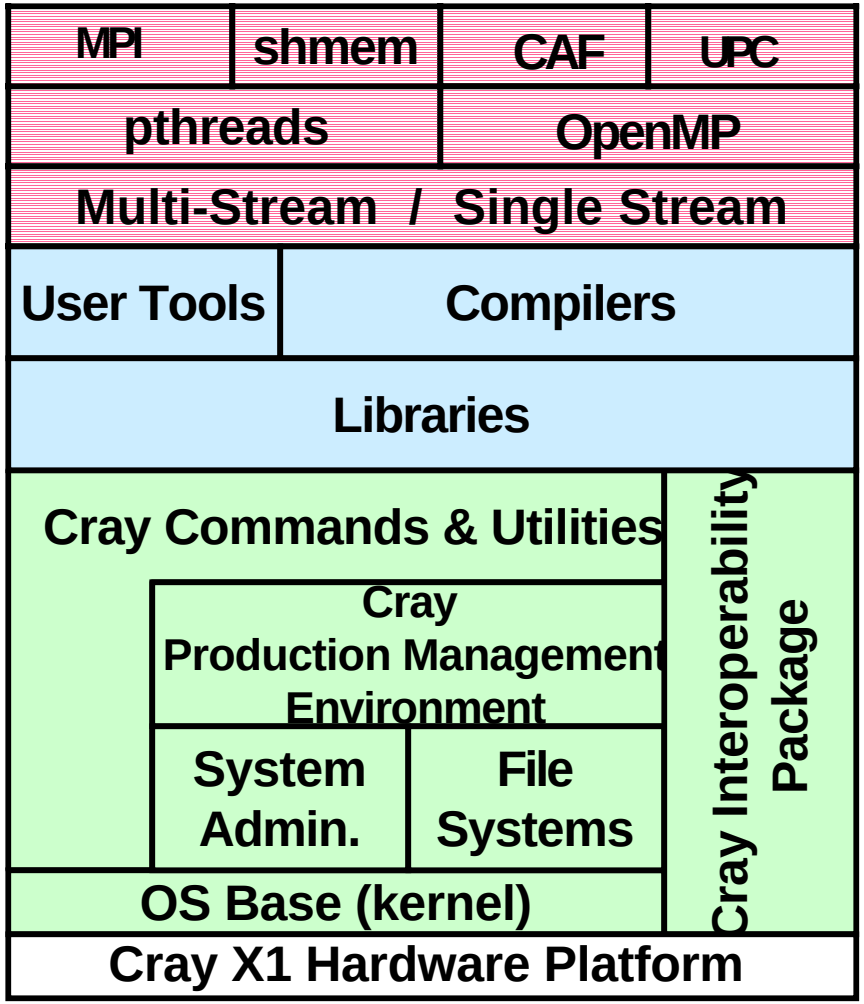

# Programming Model Basics

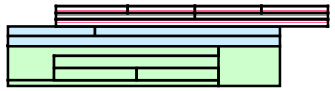


- MPI is our most important programming model.
  - Our focus is on performance; the X1 with high bandwidth, global read/write access is a good match for MPI codes.
  - We track industry standards (MPI-2)
- Co-array Fortran (CAF) and UPC usage is growing.
  - Programmers may optimize MPI applications by selectively introducing Co-Array Fortran and UPC syntax.
  - The U.S. Department of Energy is sponsoring portable implementations of these models (www.pmodels.org).
- OpenMP is supported for Fortran and C on a Cray X1 node; OpenMP can be used with distributed models such as MPI or Co-Array Fortran.
- The Cray X1 Multi-Streaming Processor allows automatic, parallel execution of tightly coupled execution streams.
- Cray compilers support both Multi-Streamed and Single Stream execution modes. Both modes are supported for all supported programming models.

# UPC & CAF Programming Model Advantages

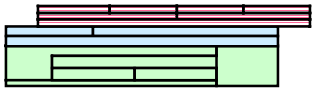


Expressability (maintainability)

- Communication expressed as memory operations rather than subroutine calls with numerous, “faceless” arguments.

Non-redundancy (maintainability)

- Communication expressed as one memory operation rather than two “mirror-image” send/receive calls.
- Size, type of communication inherent in code rather than redundantly stated as arguments to send/receive calls.

Performance

- Eliminates subroutine overhead
- Allows compiler to automatically overlap communication and computation

# UPC/Co-array Fortran Example

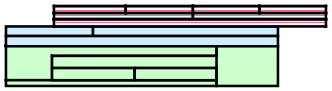

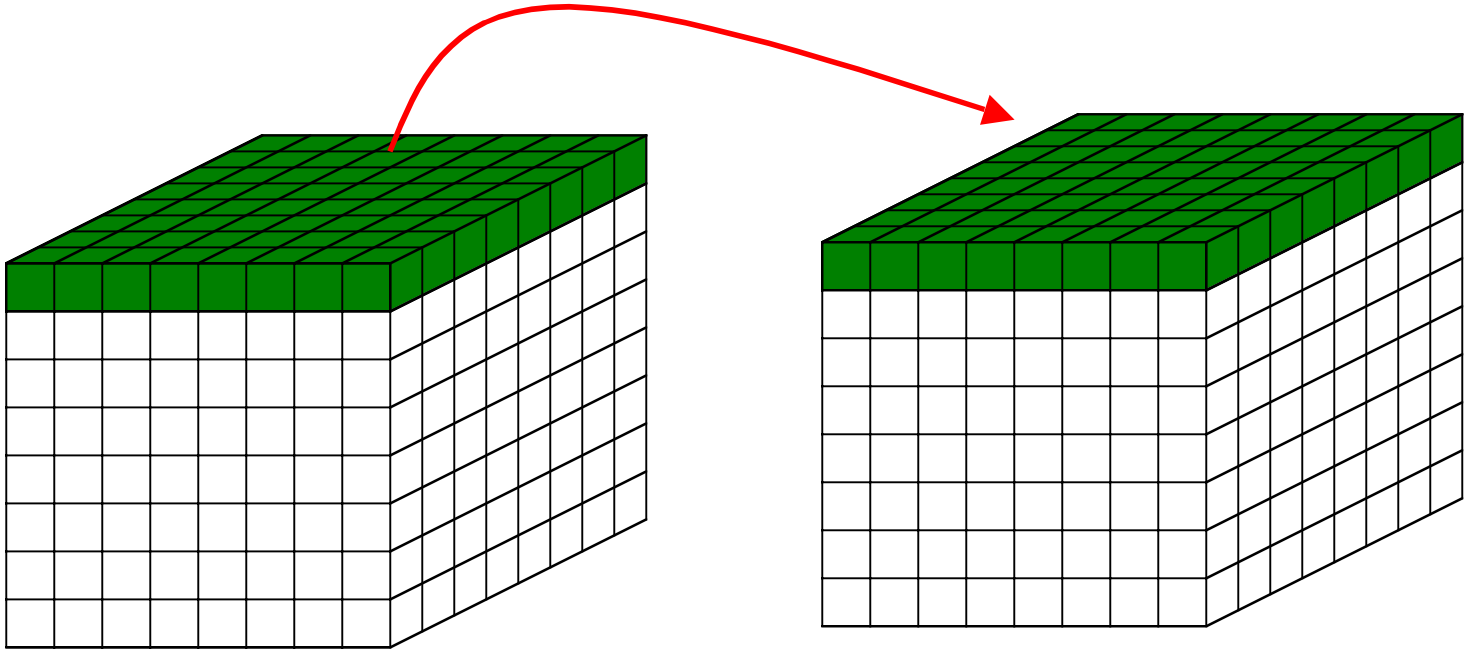

# MPI and Co-array Fortran Comparison

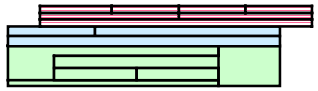


**MPI:**

```
call mpi_send_init(ai(1,1,1),   ihp*ip, mpi_real, imgi(nrpoc+1), &
   9905, mpi_comm_world, mpireq(1), mpierr)
call mpi_send_init(ai(1,1,2), 2*ihp*ip, mpi_real, imgi(nrpoc-1), &
   9906, mpi_comm_world, mpireq(2), mpierr)
call mpi_recv_init(ai(1,1,4),  ihp*ip, mpi_real, imgi(nrpoc-1), &
   9905, mpi_comm_world, mpireq(3), mpierr)
call mpi_recv_init(ai(1,1,5), 2*ihp*ip, mpi_real, imgi(nrpoc+1), &
   9906, mpi_comm_world, mpireq(4), mpierr)
call mpi_startall(4, mpireq, mpierr)
call mpi_waitall (4, mpireq, mpistat, mpierr)
```

**Co-Array Fortran (similar syntax for UPC):**

```
ai(:,:,4:4) = ai(:,:,1:1)[imgi(nproc-1)]
ai(:,:,5:6) = ai(:,:,2:3)[imgi(nproc+1)]
call sync_all(imgi(nproc-1:nproc+1))
```

# Programming Environment Basics

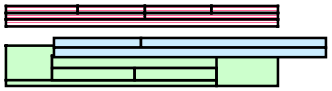


- Cray provides Fortran90, C and C++ compilers, and the CAL assembler with support for both 32 bit and 64 bit execution.
  - New for the Cray X1, default data types follow industry standards for ease of application and open source product porting.
- Cray provides several libraries that deliver a high percentage of Cray X1 peak performance to applications:
  - Math and scientific libraries include BLAS, LAPACK, SCALAPACK.
  - Visualization libraries including X11R6 and Motif
  - IO libraries including ffio and global io
- Cray provides the Cray Performance Analysis Tool (CPAT) and the Etnus TotalView debugger to support program development. Both tools support line mode and GUI interfaces.

# Programming Environment Components

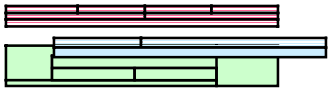


## User Development Tools

### Cray Tools

- Cray Performance Analysis Tool (CPAT).
  - combines previous Cray perf. tools incl. profiling and tracing
  - accesses Cray X1 HPM
  - no recompile to use
- C and Fortran loopmark

### 3rd Party Tools

- Etnus TotalView
  - Cray support, tracks Etnus releases

## Compilers

- Common vector code generation enhanced to provide support for:
  - large cache management
  - large register set
  - multi-streaming processor
- Fortran 90 standard support with select Fortran 2000 features
- Standard C99 support
- Standard C++ support with Dinkumware class libraries
- Industry standard container sizes; 'default 64' option for Fortran

# UNICOS/mp Operating System Basics

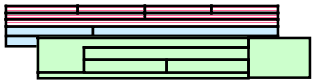


- UNICOS/mp is a full-featured, POSIX compliant OS with UNICOS/mk architectural features.
  - Supports Cray X1 scalability to hundreds of multi-stream (or thousands of single stream processors) with a single system image.
  - Extensive, high performance IO including direct IO and asynchronous IO
- UNICOS/mp includes UNICOS & UNICOS/mk production management features.
  - Application scheduling based on the psched mechanism
  - New for the Cray X1 is support for partitioning a Cray X1 system
- UNICOS/mp provides a high degree of interoperability for Cray X1 systems in a customer's computing environment.
  - A Cray Network Subsystem supporting high performance GBE
  - New for the Cray X1 is support for ADIC's StorNext SAN File System and access to a broad suite of data management capabilities.

# Operating System Components

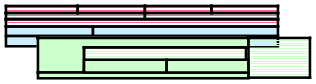


**Cray Commands & Utilities**

- POSIX commands & utilities
- UNICOS & UNICOS/mk commands (e.g. 'assign')
- csh, ksh, sh, tcsh

Cray Open Source

- perl
- bash
- tcl
- autoconf
- cvs
- others

**Cray Production Management Environment**

**System Admin.**

- Error logging
- System-wide dump
- System-wide SAR

**File Systems**

- XFS journaled
- RAID
- auto failover

**OS Base (kernel)**

**Cray Interoperability Package**

# UNICOS/mp Production Management Components

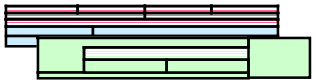


## Cray Production Management Environment

| | | |
|---|---|---|
| **Cray Application Placement (psched)**<br>• Application Migration<br>• Gang Scheduling<br>• Pre-emption<br>• Allocation Control | **Cray Checkpoint/Restart**<br>• Application, system or user initiated<br>• Restart in different placement<br>• Distributed applications | **Cray Accounting**<br>• Standard SVR4<br>• Project accounting (acids)<br>• Process limits |
| **Cray Resiliency Management**<br>• Module addition<br>• Down CPUs<br>• Impact only apps with resources assigned | **Cray System Partitioning**<br>• Run multiple OS versions<br>• Operate from CWS | **Cray Security**<br>• National Information Assurance Partnership (NIAP) EAL2 security assessment. |

# UNICOS/mp Interoperability Components

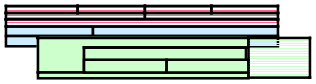


## Cray Interoperability Package

| | |
|---|---|
| NTP | Sockets |
| NIS | TCP<br>IPv4 |
| StorNext<br>SAN File System | NFS V3 |
| KERBEROS | Globus Support<br>(under evaluation) |
| Open SSH | PBS PRO |

# Cray T90/SV1/T3E Software Releases Ongoing


## UNICOS 10.0.1. 2 (May – Cray T90/SV1)

No new features

Approximately 100 bugfixes since 10.0.1.1

After May, additional fixes as needed via Field Notice mechanism

No plans for further UNICOS releases

## UNICOS/mk 2.0.6 (Cray T3E)

Archive fix packages as needed throughout the year

No plans for further UNICOS/mk releases

## CF90 3.6.0.2 / C++ 3.6.0.2

February fix release, fixes packages ongoing

# Cray X1 Software Release Plans


## UNICOS/mp 2.1 (Released in March)

Initial 64 MSP support

psched notification of downed CPUs

checkpoint/restart of MSP pthread applications

Access Control Lists (ACLs)

## PE 4.3 (Released in March)

UPC libraries

'Ordered' directive added to Cray Streaming Directives

Reorganize loop restructuring

## X11 6.3 (Released in March)

Several libraries: libX11, libXt, libXaw

Commands: xterm, imake

# Software Release Plans (cont.)


**Coming UNICOS/mp features:**

## UNICOS/mp 2.2 (Planned for July 2003)

- Initial 128 MSP support (general 128 MSP support in August 2.2.x update)
- Psched support for PBS Pro application limits

## UNICOS/mp 2.3 (Planned for Sep/Oct 2003)

- Initial 256 MSP support
- ADIC SAN File System

# Software Release Plans (cont.)


## Coming PE features:

### PE 5.0 (Planned for Jun 2003)

- SSP mode support
- Sun hosted cross-compiler
- LibsSci enhancements:
- - Adding ScaLAPACK and BLACS
- - Performance improvements for BLAS, FFTs, LAPACK
- Loopmark for C

### PE 5.1 (Planned for Sep 2003)

- OpenMP

# Software Release Plans (cont.)


**Totalview 6.x.x (Planned for May)**

gui interface

**Cray Open Software (COS) 3.4 (Planned for May)**

tk, gnuplot

**Motif 2.1 (Planned for June)**

Libraries

UIL not yet supported

**Message Passing Toolkit (MPT) 2.2 (Planned for June)**

MPI-2 extensions

SSP mode support

**Totalview 6.x.x (Planned for September)**

Distributed programming model support

# Cray X1 Software Documentation


## Cray X1 System Overview

Optimizing Applications on Cray X1 Systems

## Cray X1 User Environment Differences

Cray Programming Environments Release Overview and Installation Guide
Cray C and C++ Reference Manual
Cray Fortran Compiler Commands and Directives Ref. Manual
Fortran Language Reference Manual, Volume 1
Fortran Language Reference Manual, Volume 2
Fortran Language Reference Manual, Volume 3
Cray Fortran Co-array Programming Manual
Cray X1 Application Programmer's I/O Reference Manual
Cray Message Passing Toolkit Release Overview

## Migrating Applications to Cray X1 Systems

Dinkum C++ Library Documentation
Etnus TotalView Getting Started
Etnus TotalView New Features
Etnus TotalView Reference Guide
Etnus TotalView Users Guide
TotalView Release Overview, Installation Guide, and User's Guide Addendum for Cray X1 Systems
FLEXlm End Users Guide
Cray Assembly Language (CAL) for Cray X1 Systems Reference Manual

Cray Open Software Release Overview/Installation Guide

## Cray X1 System Administration Differences

UNICOS/mp General Administration Guide
UNICOS/mp Networking Facilities Administration
Cray X1 System Configuration and Operations Guide
UNICOS/mp Disks and File Systems Administration
UNICOS/mp Resource Administration
UNICOS/mp Release Overview
UNICOS/mp Installation Guide
Cray Network Subsystem (CNS) Overview/Installation Guide
Cray Workstation (CWS) Installation Guide
Cray Workstation (CWS) Release Overview
Cray Programming Environment Server (CPES) Release Overview/Installation Guide
Cray Programming Environment Server (CPES) Administration Guide
PBS Pro 5.3 Administrator Guide
PBS Pro 5.3 User Guide
PBS Pro Release Overview, Installation Guide, and Administration Addendum for Cray Systems
PBS Pro 5.2 External Reference Specification, PBS-3BE01
Common Installation Tool (CIT) Reference Card
CrayDoc Installation and Administration Guide