Cray X1 Early Experience

Brad Blasing
Network Computing Services / AHPCRC
May 14, 2003

Agenda

- Intro – Why Cray X1
- Give a view of Early Production Systems experience vs. later LC Production Systems experience
- Unicos/mp comments
- Networking comments
- Porting code to Unicos/mp
- Memory management
- Job scheduling
- Future
- Performance

Army High Performance Computing Research Center

Mission

- Advance defense computational science through Army-university collaborative research
- Provide for university-Army technology exchange
- Support Army's societal, educational and outreach objectives
- Provide a state-of-the-art computing infrastructure
- Network Computing Services, Inc. is the prime contractor

Past Systems – AHPCRC Program

- 1990 – TMC CM-2
- 1991 – TMC CM-200, s/n 1
- 1991 – TMC CM-5, s/n 1
- 1998 – Cray T3E-1200 LC272, s/n 1 of T3E-1200 series
- 1998 – IBM SP
- 2000 – Cray T3E-1200E LC1088 upgrade
- 2001 – Cray SV1ex 16/64

Cray X1

- AHPCRC codes require high-bandwidth, low-latency memory access
- Run well with vector and shared memory model
- Gather/scatter capability also a plus
- Cray X1 chosen
- More on decision criteria in Paul Muzio / Richard Walsh paper later today at 2:00
- The AHPCRC had been interested in Cray X1 as a T3E follow-on for 5 years

Cray X1

- AHPCRC procured 2 Cray X1 AC early production systems in Sep. 2002 (s/n 1 & 3)
- In Feb. 2003, the AHPCRC took delivery of Cray X1 LC 32 MSPs, s/n 3
- The X1 LC system was upgraded to 64 MSPs on 5/2/03 to be the first full size production Cray X1 LC system.

Cray X1 AC Early Production Systems

- 2 (of the 5 built) early production X1 AC systems installed on 27 Sept 2002
- First Cray X1 not installed at a classified site
- Hardware Installation
  - Almost “Plug in the wall” installation; air cooling
  - Typical Cray high quality shipping and installation
- Software Installation
  - Easy, came preinstalled
  - Fair number of components to hook to networks and set up names, IP numbers, and passwords

Cray X1 AC Early Production Systems (cont.)

- Hardware limitations
  - 400 MHz clock speed (1/2 speed)
  - Internode bandwidth performance limitation
  - Self-sync hang with power of 2 stride
  - Pass 1 parts
- Software limitations – Unicos/mp 1.0
  - _" document describing limitations
  - No multi-volume (i.e. no volume manager) file systems (largest file system was 585 GB)
  - No DNS (back to /etc/hosts tables!)

Cray X1 AC Early Production Systems (cont.)

- Software limitations – continued
  - Totalview debugger not available
  - Checkpoint/restart pretty much not working
  - No ntpd time sync
  - No email (not even local delivery from PBS!)
- SV2 appeared everywhere
  - Boot scripts, dump scripts
  - Environment variables to control memory usage

Cray X1 AC Early Production Systems (cont.)

- I/O
  - What is now the CPES, was called the ION
  - ION handled network traffic in addition to programming environment
  - Large (physically) Sun, small networking capability
  - Replaced later with Dell/Linux CNS
  - Older generation LSI Logic disk 4600 product

- Programming Environment
  - Started with PE 4.1
  - Some bugs
  - Medium vectorization / streaming

Cray X1 AC Early Production Systems (cont.)

- Hardware Reliability
  - Great, no failures
  - High tech spray evaporative cooling – very reliable, no issues
- Software Reliability
  - Software MTBF “a day or two”
  - With hardware limitations that caused hangs it was sometimes hard to sort out hardware vs. software crashes
- Stream Triad: 371 MFLOPS, 23% of peak, 4.46 GB/sec (SSP)

Cray X1 AC Early Production Systems (cont.)

Cray X1 LC Production Systems

- Cray X1 LC s/n 7803 production system installed with 32 MSPs on 20 Feb 2003.
- Upgraded to 64 MSPs on 2 May 2003
- To be upgraded (2nd cabinet) to 128 MSPs on 23 May 2003
- Single system image 128 MSPs in August 2003

- Hardware Installation
  - Liquid cooling installation, low GPM, smaller pipes
  - Typical Cray high quality shipping and installation

Cray X1 LC Production Systems (cont.)

- Software Installation
  - Easy, came preinstalled
  - Less issues
  - Started with Unicos/mp 2.0
- Hardware limitations
  - None
  - 800 MHz clock speed (full speed)
  - Pass 2 parts
- Software limitations – Unicos/mp 2.0
  - Few, mostly items known to be not ready or in progress
  - X11 clients, full UPC, Totalview GUI

Cray X1 LC Production Systems (cont.)

- SV2 references are down to a few now
- I/O
  - No more ION – smaller Sun that runs PE only (CPES)
  - Dell/Linux CNS handles networking (much better)
  - Newer generation LSI Logic disk 5600 product (RS200)
- Programming Environment
  - Started with PE 4.3
  - Much better vectorization & streaming
  - About to install 4.3.0.1

Cray X1 LC Production Systems (cont.)

- Hardware Reliability
  - Great, few failures
  - High tech spray evaporative cooling – very reliable, no issues
- Software Reliability / Issues
  - Software MTBF “a week or more”
  - General MTBF feel is positive
  - We had some software problems that were worked through by Cray
  - Errant kernel thread caused negative scaling
- Stream Triad: 787 MFLOPS, 25% of peak, 9.45 GB/sec (SSP)

Cray X1 LC Production System

Unicos/mp Miscellaneous Comments

- It looks a lot like Irix
- Administer it similar to Irix
- You build kernel and commands on Irix
- You debug a crash dump on a separate Irix box
- Simulator runs on Irix
- The #ifdef IRIX versions of system programs compile better
- Imported a file system XFS bug from Irix

Unicos/mp Miscellaneous Comments

- Simple user queries are a little complicated
  - How much memory is my program using?
  - What processors (MSPs) is my job using?
- Boot script can power up/down modules to assist clearing out hung nodes
- CWS upgrade script with bug powered-down entire system at one point!
- Boots fast (crash, dump, boot < 10 min.)

Unicos/mp Miscellaneous Comments

- What is an CPU?
  - Continuing confusion about SSP's vs. MSP's
  - Some system tools deal with one, some with the other
  - Programs get built for one vs. the other
  - MSP building is default
  - Leads to common mistake of building tools as MSP
    - Think "gmake"
  - For accounting, 1 MSP CPU sec. is the "same" as 1 SSP CPU sec., but one is more capable than the other
  - When the capability to run parallel SSP work comes (PE 5.0), even more confusing (now my job seems to take 4x the cpu time)

Unicos/mp Miscellaneous Comments

- Variety of software components to keep track of:
  1. Unicos/mp 2.1
  2. PE 4.3
  3. CWS
  4. CNS
  5. Trigger 2.1
  6. CPES
  7. MPT
  8. PBS 5.x
  9. Totalview 5.x
  10. X11 6.3

Unicos/mp LC Configuration

- 25 Tbyte of disk
- Disk configured pretty stock (4+1 & 8+1 RAID)
- Experimented with 12+1 RAID
- 3 CNS (2 GigE, 1 HIPPI).
  - 1 GigE to outside DREN, 1 GigE for computer room use
- X1, CWS, CPES all Kerberized
- Kerberizing CWS interferes with Cray install tool

Unicos/mp Networking

- Early production systems used Sun ION
- GigE network performance was terrible
- Production systems use Dell/Linux CNS
- Seeing 50+% bandwidth of GigE (expect better with jumbo frames)
- Network diagram is complicated
- Our full X1 LC cabinet has 28 Linux L1's, 3 Linux CNS's, 2 Sun's, and 3 Storage controllers to telnet to, administer.

2 AC & 1 LC X1 Network Diagram

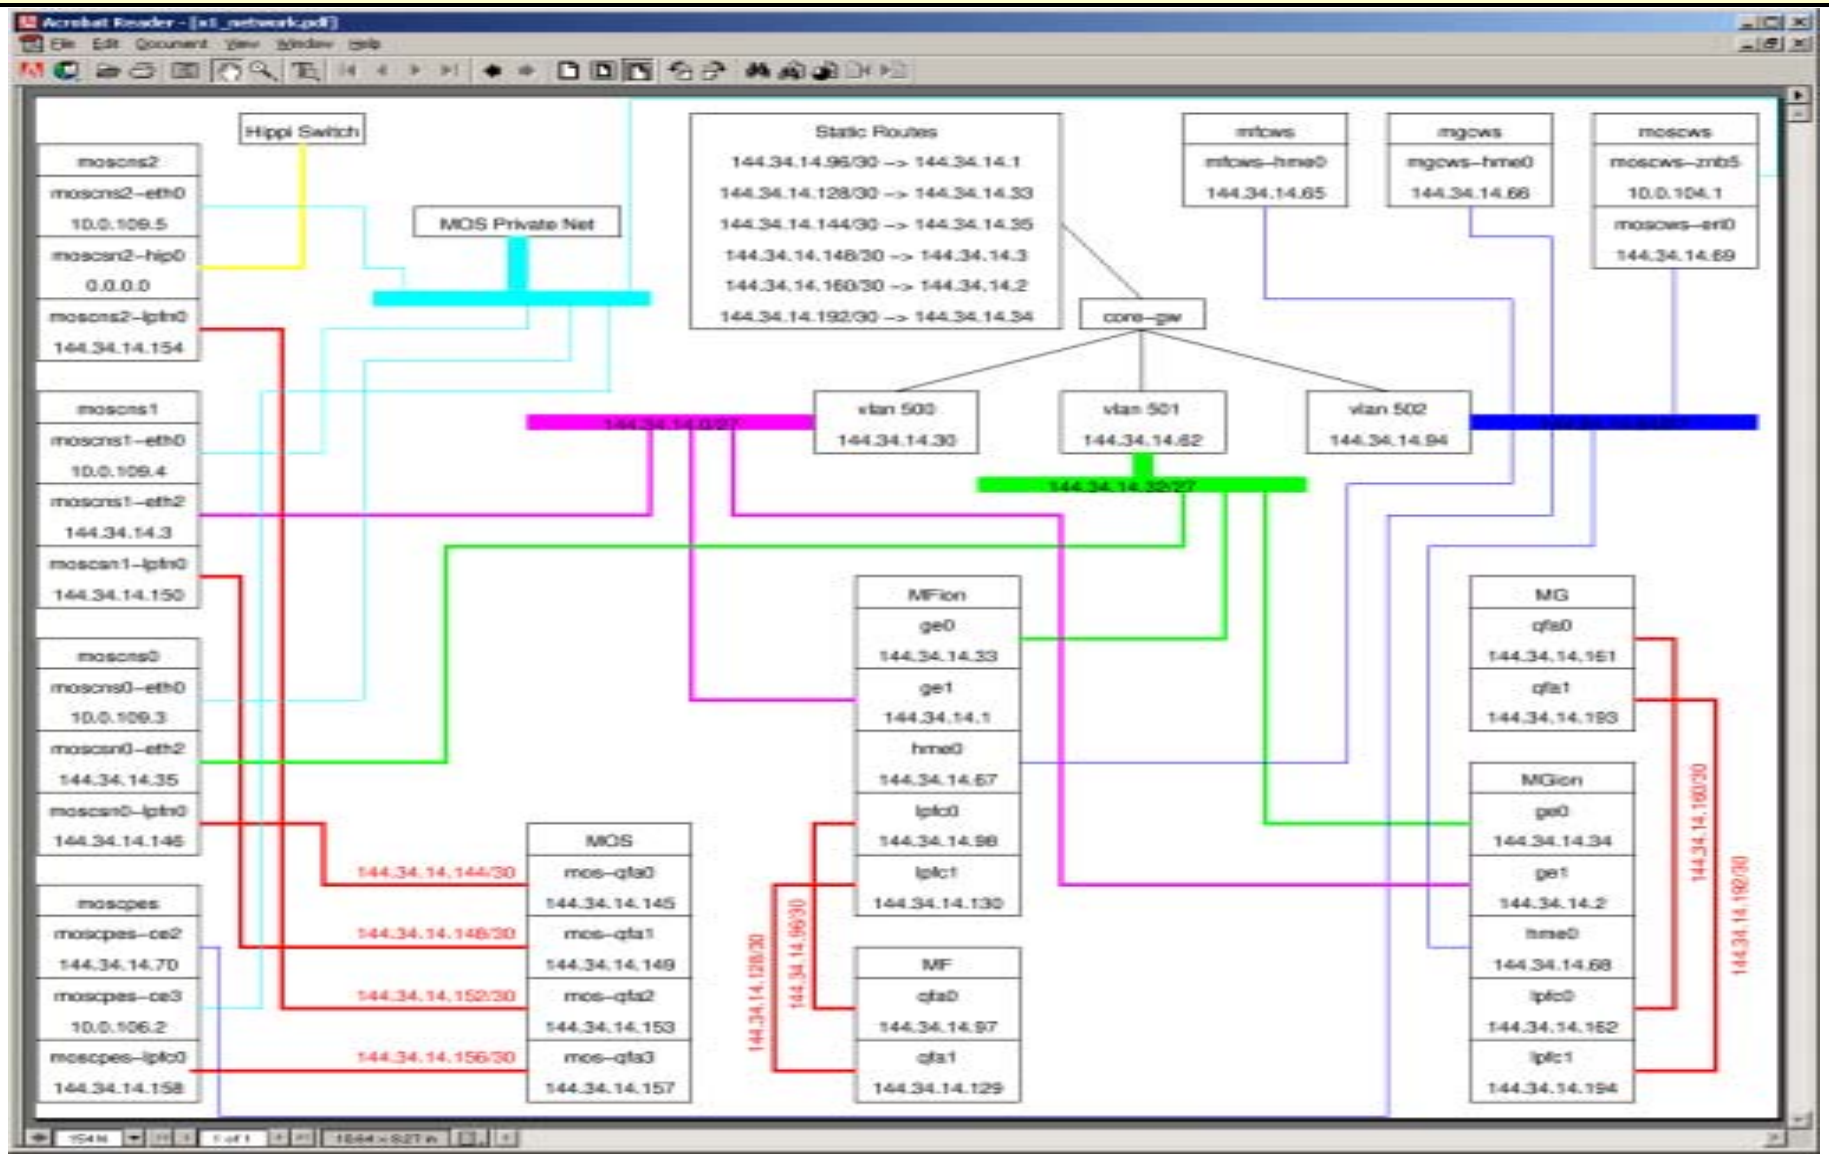

Porting Experience – Kerberos

- Had to discriminate between CRAY and CRAY X1
- #if defined(CRAY) or #ifdef CRAY isn't the right thing
- Had to create _UNICOSMP tag to discriminate X1 code
- Data type sizes are not the same as on traditional Cray vector systems (e.g. int is 32 bits), and not the same as Irix systems
- Long and int are not the same size (as they are on older Crays)

Data Type Size Comparison

	Linux Intel-IA32	Traditional Cray vector	Cray T3E	IRIX	Cray X1
short int	2	8	4	2	2
int	4	8	8	4	4
unsigned int	4	8	8	4	4
long int	4	8	8	4	8
unsigned long	4	8	8	4	8
size_t	4	8	8	4	8
float	4	8	4	4	4
double	8	8	8	8	8
char	1	1	1	1	1

Porting Experience – Kerberos

- There were several utilities that needed to be built to allow us to configure and compile Kerberos
  - autoconf - produces configure scripts
  - gmake - generates executables
  - bison - parser similar to yacc
  - m4 - macro processor(need version 1.4)
- Compile times were long, link times long also
- Compiling as commands different, from parallel MSP
- Compiling for parallel MSP is default
- Cannot install new binaries over existing running binaries

Porting Experience – Accounting

- NetworkCS has accounting code that runs on all systems
- C code issues
  - int/float are again 32 bits; change to long and double
  - Ripple effects (%d changes to %ld in many printfs)
  - #ifdef CRAY isn't right - #ifdef IRIX frequently is better
  - Need to specify "-h command" to compile SSP
  - Causes non-standard makefile's

Porting Experience – Accounting

- Unicos/mp issues
  - Some commands in Unicos location, some in Irix spot
  - Unicos has /bin, & /usr/bin, Irix has /sbin & /usr/bin
  - Unicos/mp has /bin, /usr/bin, & /sbin
  - sbrk() works differently due to X1 memory management
  - pacct file is different than Irix and Unicos
  - pacct file issues multiple records for parallel work non-atomically (can wrap over day boundaries or be delayed)
  - Memory integrals complicated and not working well yet

Porting Experience – Kernel Code (HFS2)

- HFS2 uses vnode file system interface
- Kernel code written for Irix ports very directly to Unicos/mp
- Building Unicos/mp from source straightforward
- Kernel data types and sizes same as from IRIX64
- Booting / debugging is fast
- Dumping is relatively fast
- Issues
  - No native kernel compiler yet
  - Kernel compiler error messages wrong & line #'s wrong
  - Crashmp does not run on live system yet

Compiling Environment Issues

- Non-native PE is slow
  - First kerberos compile: 5 hours; now 2 hours. 15 min on Linux
- Security policy bars compiling as root due to NFS
- Complicated GNU builds tax the current trigger PE and sometimes don't build correctly; due to trigger usually
- File systems between CPES and X1 must be mounted with specific mount options
- CPES and X1 must be time sync'ed
- Had some compiling trigger bugs
- Would rather have a native compiler

Memory Management

- Combination shared memory / massively parallel system is X1's advantage
- With advantage there is complexity
- Memory is a full virtual memory system
- Unicos/mp not reacting well to over-committing memory
- Processes get access to more memory by allocating more nodes (modules)
- The 4 MSPs on a node module share memory, so system admin must be prepared for swapping/paging.
- User must use environment variables to access all memory
- User needs to choose page size to avoid "no forward progress" abort

Job Scheduling

- Similar issues to Cray T3E (Unicos/mk)
- Several schedulers (psched, PBS) not working in together
- More complicated scheduling
  - MSP vs. SSP scheduling (SSP coming in PE 5.0)
  - Accelerated mode vs. flexible mode
  - Number of MSPs per node (aprun -N)
- No preemptive scheduling (aside from prime job)
- No backfill scheduling

Job Scheduling

- NetworkCS created “jsched” for T3E to solve these problems
- jsched raised T3E utilization from 80’s% to 90’s%
- jsched allowed large work to run without manual intervention
- jsched ported to SV1ex
- Intend to port to X1 Unicos/mp

Other Future Enhancements by AHPCRC

- File systems
  - Porting our SAN shared file system HFS2 to Unicos/mp
- Tapes
  - Porting TCP/IP based tape client to Unicos/mp to allow labeled tape access to X1
- Kerberos
  - Continuing Kerberos maintenance and enhancements

General Performance Being Seen

- Sustained performance is our key measure
- General numbers (specifics in other talks):
- Andrew Johnson CFD code: ~25-33% of peak
- MM5 (weather): ~30% of peak
- Stream Triad: 25% of peak
- Generally pleased with these early numbers
- Accomplished with modest optimization efforts

Finally

"The research reported in this presentation was performed in connection with contract DAAD19-03-D-0001 with the U.S. Army Research Laboratory. The views and conclusions contained in this presentation are those of the authors and should not be interpreted as presenting the official policies or positions, either expressed or implied, of the U.S. Army Research Laboratory or the U.S. Government unless so designated by other authorized documents. Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon."

Questions

Brad Blasing
Network Computing Services / AHPCRC
brad@ahpcrc.org
brad@networkcs.com

www.ahpcrc.org