Cray X1 Review

Gary M. Johnson
OASCR/MICS

Advanced Scientific Computing Advisory Committee Meeting
April 5 and 6, 2004
Hilton Washington Embassy Row Hotel
Washington, DC

Charge

1. Review and analyze what is currently known about the performance of the Cray X1 on capability-limited science applications. Draw on all available information from all Cray X1 installations.
2. Based on this analysis, inform ASCR on the relative suitability of the Cray X1 to the science problem set of the Office of Science:
    - For which applications is it well-suited?
    - For which applications is it ill-suited?
3. Inform ASCR on the expected performance of larger X1 installations and their potential impact on capability-limited science applications.
4. Review the evaluation process currently being used by the CCS and recommend improvements, if necessary.

Review Committee

F. Ronald Bailey, Chair	AMTI/NASA Ames Research Center
Michael Colvin	University of California, Merced
James J. Hack	National Center for Atmospheric Research
Richard Matzner	University of Texas
James C. T. Pool	Center for Advanced Computing, Caltech
Daniel V. Pryor	IDA/Center for Computing Sciences

Cray X1 Review

Building 5700, Room F234, ORNL

Agenda

Tuesday, 10 February 2004

7:30 a.m.	Coffee, Pastries	
8:00 a.m.	Welcome & Charge to Reviewers	Gary Johnson
8:15 a.m.	Interagency Scene & Ultrascale Computing	Ed Oliver/Fred Johnson
8:45 a.m.	Cray X1 System Overview	Steve Scott
9:15 a.m.	Overview of X1 Software	Don Mason / John Levesque
10:15 a.m.	Break	
10:30 a.m.	Overview and Status of Cray X1 Evaluation at CCS	Thomas Zacharia / Pat Worley
12:00 noon	Working Lunch with Applications Presentations	
	Astrophysics	Tony Mezzacappa
	Climate	John Drake
	Fusion	Don Batchelor
	Materials	Thomas Schulthess (Stocks/Maier)
2:00 p.m.	Applications Experience on the X1	
	Arctic Region Supercomputing Center	Virginia Bedford
	Army High Performance Computing Research Center	Paul Muzio
	Boeing Company	Suresh Shukla
	Engineering Research and Development Center	John West
3:00 p.m.	Break	
3:15 p.m.	Questions	Review Team
4:00 p.m.	Open Issues Follow-up	Ron Bailey
5:00 p.m.	Adjourn for the day & tour of CCS	

Cray X1 Review

Building 5700, Room F234, ORNL

Agenda

Wednesday, 11 February 2004

7:30 a.m.	Coffee, Pastries	
8:00 a.m.	Open Issues Follow-up	Ron Bailey
9:00 a.m.	Writing Assignments	Ron Bailey
9:30 a.m.	Write Preliminary Report	Review Team
12:00 noon	Adjourn	

Findings

1. Review and analyze what is currently known about the performance of the Cray X1 on capability-limited science applications. Draw on all available information from all Cray X1 installations.

In summary, the committee feels there are a number of important scientific problems that can clearly demonstrate the benefits of the more powerful computational capabilities of the Cray X1 processors (when compared to commodity counterparts), the advantages of the high-performance interconnect (both in terms of latency and bandwidth) and the large memory. Therefore, based on the material presented, and what the Committee knows of similar scientific applications, it is our belief that the Cray X1 should be regarded as a one of the most powerful high-performance computer systems for enabling capability-limited science.

Findings

2. Based on this analysis, inform ASCR on the relative suitability of the Cray X1 to the science problem set of the Office of Science:
   - For which applications is it well-suited?
   - For which applications is it ill-suited?

In general, the machine does very well on codes that have been written as vector codes. Most codes that solve PDEs can be cast in this way, including important codes from climate modeling, and fluid flow. In the cases presented the bigger the problem the better in terms of demonstrating the capabilities of the machine. These conclusions are consistent with reported experiences using the Japanese Earth Simulator, which has reported extraordinary performance characteristics for climate, plasma simulation, and molecular dynamics application codes.

At the other end of the scale are large, heterogeneous codes from computational chemistry, whose performance was not good on the Cray X1. This is in part because the size and complexity of these codes has limited the amount of performance tuning done to date. … the performance problems with these codes can be attributed in large part to the designs of many of these codes, which do not take advantage of the vector nature of the simulations they perform.

Findings

3. Inform ASCR on the expected performance of larger X1 installations and their potential impact on capability-limited science applications.

   The committee believes increasing the size of the current Cray X1 (or X1-e) configuration would be a highly worthwhile investment given some of the scientific presentations that suggested they could take immediate advantage of a larger system and address problems currently out of reach using other computational platforms. … expanding this facility to more fully explore the capabilities of this architecture is regarded to be of strategic national importance.

Findings

4. Review the evaluation process currently being used by the CCS and recommend improvements, if necessary.

   The committee is impressed by the openness of the evaluation process and its progress to date; the strategy of having an expert staff, working closely with the vendor and application community; and the outreach component of the evaluation process. The committee feels the CCS evaluation of the Cray X1 is an extremely important component in mapping the future of computational science and engineering …

   The evidence of a mutually beneficial collaboration between ORNL and Cray was exemplary – a model that other vendors and sites should replicate! …

   The "half system" was delivered in July and the full system passed acceptance test on December 15, 2003, less than two months ago! Considering this timescale, the progress is remarkable.