

Investigating the Impact of the Cielo Cray XE6 Architecture on Scientific Application Codes

Courtenay Vaughan, Mahesh Rajan, Richard Barrett, Doug Doerfler, and Kevin Pedretti

Center for Computing Research, Sandia National Laboratories

P.O. Box 5800, MS 1319, Albuquerque, NM, 87185-1319, USA

Email: ctvaugh, mrajan, rfbarre, dwdoerf, ktpedre@sandia.gov

Abstract—Cielo, a Cray XE6, is the Department of Energy NNSA Advanced Simulation and Computing (ASC) campaign's newest capability machine. Rated at 1.37 PFLOPS, it consists of 8,944 dual-socket oct-core AMD Magny-Cours compute nodes, linked using Cray's Gemini interconnect. Its primary mission objective is to enable a suite of the ASC applications implemented using MPI to scale to tens of thousands of cores. Cielo is an evolutionary improvement to a successful architecture previously available to many of our codes, thus enabling a basis for understanding the capabilities of this new architecture. Using three codes strategically important to the ASC campaign, and supplemented with some micro-benchmarks that expose the fundamental capabilities of the XE6, we report on the performance characteristics and capabilities of Cielo.

Index Terms—High performance computing; parallel architectures; message passing communication; performance evaluation; scientific applications.

I. INTRODUCTION

Cielo, a Cray XE6, is the Advanced Simulation and Computing (ASC¹) Campaign's program's newest capability machine. Rated at 1.37 PFLOPS, with dual-socket oct-core AMD Magny-Cours nodes linked using Cray's Gemini interconnect, Cielo represents the latest evolution in multicore-based HPC architectures. The initial system, delivered toward the end of 2010, consists of 6,654 compute nodes, for a total of 106,464 processor core elements, capable of 1.02 PFLOPS. The final system, to be delivered in the spring of 2011 will consist of 8,944 compute nodes, for a total of 143,104 cores.

The two major programmatic requirements for Cielo are ease of migration of existing ASC multi-physics applications and strong performance of these applications at capability scale. Migration of applications is enabled by Cray's mature software development environment that includes a set of compilers (Cray, GNU, and PGI), Cray message passing library toolkit (MPT), job launch and management through ALPS, performance analysis (CrayPAT), and a debugging tool (TotalView). Performance and scaling are facilitated by the node processor and memory architecture, the new Gemini node interconnect, and the software and hardware components for the I/O sub-system. More important is that Cielo is an improvement to a successful architecture previously available to many of our codes[11], providing evolutionary and possibly revolutionary capabilities that may be important in future code

configuration issues, especially as we progress to even larger computing scales.

In this report we present results on performance analysis focused three codes from the ASC Applications Acceptance Test suite, representing a broad set of requirements of the ASC Tri-lab organizations (Lawrence Livermore, Los Alamos, and Sandia National Laboratories). Our goal is to understand the capabilities of the Cielo XE6 architecture, which is significantly aided by comparing with data from its XT-series evolutionary ancestors. Preliminary results were summarized in [12]. Since then, we have more directly explored the effects of these characteristics, including at much larger scale, stronger profiling, and supplemented with a set of micro-benchmarks.

To facilitate understanding of the measured performance, we break it down into three components: the impact of the processor core, the impact of the node memory architecture, and the impact of the node interconnection network. Although it is difficult to attribute performance effects clearly in terms of these individual characteristics, our interpretation of the results lead us to some strong conclusions. First, the Magny-Cours-based node architecture, with four NUMA regions each with independent memory controllers and dual-channel DDR3 memory configuration, is a key design feature that clearly benefits the bandwidth requirements of our applications, given the inevitable push for large core counts in each compute node. We also find that the Gemini interconnect provides an evolutionary performance improvement to codes that send large messages. More importantly, we find that codes that send many small messages realize significantly stronger performance, an issue critical to effective use of very high processor counts, which expected to be critical at the exascale[2], [6].

This report is organized as follows: We begin with a description of the Cielo architecture, with an emphasis on its evolution from the Cray XT-series. Next we describe the applications used in this study, focusing on the issues required to achieve strong performance at large scale. We include micro-benchmark results that help us understand the processor, node, and interconnect performance. Then we describe our experiments involving the full application codes and their results, followed by a summary of this work and our future plans.

¹<http://nnsa.energy.gov/aboutus/ourprograms/defenseprograms/futurescienceandtechnologyprograms/advancedsimulationandcomputin>

II. ARCHITECTURE OVERVIEW

Cielo is the result of an evolutionary path, beginning with the XT3, Cray's third-generation massively parallel processing system, which built upon on the T3D and T3E systems. The XT series is based on commodity AMD Opteron processors, and for the machines procured by the ASC program, a Cray custom interconnect named SeaStar[5], and a light-weight kernel (LWK) operating system[8]. Nodes are connected by a SeaStar router through HyperTransport in a 3-dimensional torus topology.

The XT series evolved with the development of multi-core processors, progressing from dual- to quad- to hex-, and now with the XE6, to oct-core processors (see Table I). During this progression, a node went from a single processor to two processors, referred to as dual socket. Memory also progressed, from DDR-1 at 333 MHz to the XT6 DDR-3 at 1.33 GHz. Although progress, the per core memory capabilities have not kept pace with the computational capabilities of the node, and we've seen a progressive degradation in per-core performance[1]. The SeaStar interconnect also evolved, now at version 2.2.

The XT6 brought a significant change to the node architecture, and then the XE6 introduced a new node inter-connect. Each of these are described and examined below. The Cielo system configuration, the XE6 instantiation examined herein, is shown in Figure 1.




Fig. 1. Cielo XE6 architecture. Image courtesy of Cray Inc.

A. Node architecture

Beginning with the XT6 and continuing on to the XE6, two AMD Opteron 8-core Magny-Cours processors share a compute node, illustrated in Figure 2. Each Magny-Cours processor is divided into two memory regions, which AMD calls NUMA nodes, each consisting of four processor cores. Thus each compute node consists of 16 processor cores, evenly divided among four NUMA nodes, which are connected using Hyper Transport version. All links run at 6.4 GigaTransfers per second (GT/sec). So the 24-bit links between die in a processor




Fig. 2. XE6 Magny-Cours based node. Image courtesy of Cray, Inc.




Fig. 3. XE6 Gemini (bottom) and SeaStar (top) interconnects. Image courtesy of Cray, Inc.

run at 19.2 GigaBytes per second (GB/sec), the 16-bit links between processors run at 12.8 GB/sec per direction, and the "cross" 8-bit links between processors run at 6.4 GB/sec. The impact of this NUMA memory organization is investigated using the STREAM benchmark, taking care to set processor and memory affinity of all the MPI tasks running on a NUMA node using the numactl utility. For the XT4 dual core system, we were seeing about 2.12 GBytes per second per core, and for the XT4 quad cores, we were seeing about 2.00 GB/sec per core. Table II shows node level performance of the XT5 and XT6.

Magny-Cours processors are also available with 12 cores divided into 6-core NUMA nodes, which form the basis of the new Hopper II computer at NERSC². A discussion of the choice to use 8-core processors in Cielo is included in [3].

B. Gemini interconnect

The Gemini interconnect is the most significant architectural difference between the XE6 and prior-generation Cray XT-series supercomputers, which use the SeaStar interconnect. Gemini and SeaStar are both custom system-on-a-chip ASICs developed by Cray that implement a high performance 3-D torus interconnect where each node is connected to its six nearest neighbors, as shown in Figure 3. Gemini achieves

²<http://www.nersc.gov/nusers/systems/hopper2/>

Arch	Year	Node			Processor			Memory		Network			
		Number of compute nodes	cores/ socket	sockets/ node	Number of cores	GHz	Ops / clock	DDR	GHz	Type	Latency μ sec	Injection BW GB/sec	Peak Link BW GB/sec
XT3	2005	10,368	1	1	10,368	2.0	2	1	0.33	SeaStar 1.2	5.5	1.2	4.8
XT4	2006	12,960	2	1	25,920	2.4	2	1	0.4	SeaStar 2.1	5.5	2.2	4.8
XT4	2007	6,720 6,240	2 4	1	38,400	2.4 2.2	2 4	1 2	0.4 0.8	SeaStar 2.1	5.5	2.2	4.8
XT5	2010	160	6	2	1,920	2.4	4	2	0.8	SeaStar 2.2	5.5	2.2	4.8
XT6		20	8		320			3	1.33				
XE6	2011	6,654 8,944	8	2	107,264 143,104	2.4	4	3	1.33	Gem 1.2	1.3	6.1	4.7, 12-bit links 9.4, 24-bit links

TABLE I

CRAY XT SERIES EVOLUTION TO XE6. PHYSICAL CONFIGURATIONS ARE THOSE INSTALLED AT SANDIA NATIONAL LABORATORIES, AND NOW THE CIELO XE6, ACQUIRED BY THE ACES PROGRAM[3] AND LOCATED AT LOS ALAMOS NATIONAL LABORATORY.

XT5		
memory / node	0	1
0	10.3	7.8
1	7.8	10.4

XT6				
memory / node	0	1	2	3
0	13.4	6.9	6.8	5.6
1	7.0	13.8	5.6	6.8
2	6.9	5.6	12.39	6.8
3	5.7	6.7	6.8	13.8

TABLE II

XT5 AND XE6 LOCAL AND REMOTE NUMA NODE BANDWIDTH USING ALL CORES IN A NUMA DOMAIN, IN GB/SEC ; MEASURED USING THE STREAM TRIAD BENCHMARK[10].




Fig. 4. Message Injection Rate: Gemini vs. SeaStar

higher packaging density than SeaStar by supporting two physical nodes per Gemini chip, but logically each direction (X, Y, and Z) has the same number of network links. Every other hop in the Y dimension takes place within the Gemini ASIC.

Gemini has been architected to provide high performance support for fine grained remote load-store-style messaging, as is typical of partitioned global address space (PGAS) languages. This also results in significantly improved MPI messaging rates compared to SeaStar, as shown in the SMB message rate micro-benchmark [4] results shown in Figure 4. The Gemini achieves over an order of magnitude higher messaging rate than SeaStar for small messages. This translates to a significant performance boost for MPI applications that send many small messages in rapid succession, as is observed in the xNOBEL application results presented in Section IV-C.

Gemini also provides an evolutionary improvement to the achievable asymptotic bandwidth for point-to-point communication. As with SeaStar, there are two potential bottlenecks to consider: injection bandwidth and link bandwidth. Injection bandwidth is limited by the speed of the Opteron to Gemini HyperTransport link, which runs at 4.4 GT/s. Link bandwidth is determined by the signaling rate and the width of the link. Due to Gemini's double-density packaging, links in the




Fig. 5. Point-to-Point Bandwidth: Gemini vs. SeaStar

X and Z dimensions are twice the width of links in the Y dimensions (24-bits vs. 12-bits wide). The uni-directional streaming bandwidth micro-benchmark results shown in Figure 5 illustrate this difference clearly. For the configuration tested, communication in the Y-dimension is limited by link bandwidth while communication in the Z-dimension is limited

by injection bandwidth. SeaStar is always limited by injection bandwidth due to its slower 1.6 GT/s Oporton to SeaStar HyperTransport link.

III. OVERVIEW OF THE APPLICATIONS

Acceptance of the Cielo architecture was based on the performance of a suite of six application programs[12], called the ASC Applications Acceptance Test, representing a broad set of requirements of the ASC Tri-lab organizations. Three of these applications, CTH, SAGE, and xNOBEL, form the basis of the performance analysis reported herein. They have been selected for their ability to illustrate the impact of the architectural characteristics that differentiate Cielo from its ancestors. Each application is briefly described below.

A. CTH

CTH is a multi-material, large deformation, strong shock wave, solid mechanics code developed at Sandia National Laboratories[7]. CTH has models for multi-phase, elastic viscoplastic, porous and explosive materials, using second-order accurate numerical methods to reduce dispersion and dissipation and produce accurate, efficient results. For these tests, we used the shaped charge problem, in three dimensions on a rectangular mesh, illustrated in Figure 6.




Fig. 6. CTH shaped charge simulation

Computation is characterized by regular memory accesses, is fairly cache friendly, with operations focusing on two dimensional planes. Inter-process communication aggregates cell data into MByte size messages.

B. SAGE

SAGE (SAIC Adaptive Grid Eulerian) is a multi-dimensional multi-material Eulerian hydrodynamics with cell-based AMR[13]. Although the problem set used here does not include refinement, the code does not take advantage of the static nature of the memory layout, and thus the runtime profile is representative in the important sense. For example, inter-process communication is through a bulk-synchronous gather/scatter abstraction, which collects off-process data and inserts it into doubly indexed arrays; the receiver unpacks the message also using a doubly indexed array.

C. xNOBEL

xNOBEL is a one, two, three dimensional, multi-material Eulerian hydrodynamics code developed for solving a variety of high deformation flow of materials problems, with the ability to model high explosives[9]. Runtime is communication intensive, requiring the transmission of many relatively small messages.

IV. EXPERIMENTS

Each of the applications in the ASC Applications Acceptance Test measures performance based on an application-relevant “Figure of Merit” (FOM). The general intent of this metric is to capture Cielo’s capabilities with regard to supporting increased fidelity and computational efficiency of the particular application. Toward that end, the FOM are carefully chosen for each application to be representative of the performance characteristic of interest, and are intended to measure Cielo’s ability to scale across tens of thousands of cores. That is, the goal is to capture the runtime characteristics of the application code rather than its algorithmic performance. For SAGE, and xNOBEL, the FOM measures the number of mesh cells that are processed per time step per processing element per unit of time. The CTH FOM is wall clock time for 100 time steps.

The problems were selected by the application teams to be representative of the sorts of computations and communication requirements that are common across a breadth of uses. Of course no one problem set can fully capture all of the ways the application might exercise the machine, but again, our intent is more focused on investigating the machine capabilities, and these three applications and problems sets stress important, distinct areas.

All experiments were run in weak scaling mode, in an MPI-everywhere configuration, whereby each MPI rank is assigned to a distinct processor core. Placement of the MPI processes onto the system is explicitly managed using executable launch command line options that enforce processor-memory affinity.

Point-to-point message traffic for these experiments is shown in Figure 7. Some general runtime profiling information




Fig. 7. Per core point-to-point message traffic for the three applications. Note that the total volume for xNOBEL continues well off the view of this graph, up to 55,000 messages per minute.

is shown in Table III. (Note that MPI time is exclusive of MPI_SYNC time.) This information motivates our choice of applications for this study. Each of the codes is based

code	Percent time spent in		
	compute	MPI	MPI_SYNC
CTH	52.2	38.4	9.5
SAGE	84.3	11.8	3.9
xNOBEL	33.3	22.8	44.0

TABLE III
COMPUTATION AND COMMUNICATION PROFILE, AS MEASURED ON XT5
FOR 256 CORE RUN.

on Eulerian hydrodynamics. Each is configured for the bulk synchronous programming (BSP) model. The major distinction is how the computations and inter-process communication are implemented. CTH operates on a structured three dimensional mesh, aggregating two dimensional “face” data into very large messages within its BSP model. Thus although it spends a large percentage of time in inter-process communication, few messages are exchanged, and relatively little time is spent synchronizing among the parallel processes. SAGE employs a nearest neighbor gather/scatter communication model, spending the least amount of time of the three applications in MPI operations. However, it spends a significant amount of time preparing the data involved in inter-process communication, which will be strongly impacted by memory bandwidth. xNOBEL sends a relatively huge number of small messages, which will expose the message injection rate capabilities of Gemini within the context of a full application. Access to an XT6 and XE6 distinguishes the effects of the node architecture and interconnect.

Problem sets are configured for power-of-two node counts, with full node utilization of cores. This presents an issue with regard to the hex-core XT5, since we are not using one full node for each job. However, that is only meaningful for small core counts, which is not the focus of this work (and neither is the XT5). We begin with a single socket of the Magny-Cours-based node in order to help us understand dual-socket XT6 and XE6 performance, but then maintain power of two node counts.

A. CTH

As seen in Figure 8, CTH performance degraded with the XT4 change from dual- to quad-core, attributed to intra-node contention as well as the decrease in processor clock speed (2.4GHz to 2.2 GHz). Performance improved with the change to XT5 with the increased clock speed (back to 2.4 GHz) and the move to dual-socket nodes. The change to the Magny-Cours-based XT6 node significantly improved performance for relatively small core counts (up to 32 cores on 2 nodes), but this resulted in increased contention for the interconnect at higher cores counts. (A core can have a maximum of six communication neighbors; for this problem that number is reached for some cores once 128 cores are used.) The XE6 move to Gemini addressed the network contention, maintaining the small core scaling performance throughout higher core counts.

As already observed, the CTH message aggregation implementation would benefit most strongly from increased




Fig. 8. Performance of CTH

interconnect bandwidth, although the node and interconnect architecture have provided meaningful performance enhancements.

B. SAGE

Performance degraded throughout the XT-series multicore evolution, attributable to on-node memory contention. The XT6 change to the dual-socket Magny-Cours node significantly reversed this trend (illustrated in Figure 9), due to its




Fig. 9. Performance of SAGE

improved support for the SAGE bandwidth requirements. The addition of the Gemini interconnect for the XE6 provides little improvement due to SAGE’s relatively small, well-managed MPI requirements.

C. xNOBEL

Results are shown in Figures 10. The move from dual- to quad-core XT4 degraded performance only slightly, attributable to the reduced clock speed. Performance improved dramatically with the XT5 change to dual-socket (and the increased clock speed), but this advantage was mostly eliminated as increased core counts highlighted the network contention.




Fig. 10. Performance of xNOBEL

The move to the XT6 node architecture again improved performance at small core counts, but even more quickly saturated the network. The increased message injection rate of the Gemini interconnect successfully addressed this issue for the XE6, better handling xNOBEL’s use of many small messages.

D. Large scale results

Experiments described thus far were necessarily constrained to less-than capability scale computing. This allowed for multiple runs, with multiple profiling views, within the resource availability constraints inherent in a new system in the midst of acceptance testing. However, some large scale runs were executed, with results aggregated in Figure 11, presented in




Fig. 11. Weak scaling parallel efficiency.

terms of weak scaling parallel efficiency. Base performance is 16 processor cores, up to 65,536 cores (2^{16}) for CTH and SAGE, and 32,768 cores (2^{15}) for xNOBEL.

At very large scale, CTH maintains its scaling profile. This is due to its “bursty” bandwidth requirements, whereby its very large messages are well managed by the node and interconnect architecture. Its scaling curve would likely move

up with an increase in interconnect bandwidth capabilities, but its message traffic is otherwise well-managed.

The trends seen at smaller processor scales continue for SAGE and xNOBEL until we exceed 8,192 cores, where we see performance degrade. However, the problem sizes were necessary constrained in order to avoid an issue with 32-bit integers. Therefore, even though this is a weak scaling problem, there is relatively little computation relative to communication. Additionally, MPI collectives have not been highly tuned for this architecture at this scale, which may have some effect. Our analysis at smaller scale suggests that with an appropriate amount of work, scaling trends will be maintained.

We will further investigate more closely as the machine becomes more readily available at this scale. Intermediate processor counts should provide some insight to drive deeper profiling as we regain stronger access to the machine. For SAGE our expectation is that additional performance improvements would be driven by a reduction in message latency costs and that xNOBEL could be improved by an increase in the message injection rate and perhaps a scheduling algorithm for its large number of messages requirement.

V. SUMMARY AND FUTURE WORK

In order to study the performance characteristics of the ASC campaigns’s new Cielo Cray XE6 capability computing platform, we configured and ran experiments using three application codes that have been identified as critical to this computer’s success. Some micro-benchmarks supplemented our understanding of the runtime characteristics of the new node and interconnect architecture. Although these applications are focused on problems of interest to the ASC campaign, in some important ways their implementations and runtime characteristics are representative of a much broader set of scientific computation codes. In particular, the on-node bandwidth requirements, particularly using indirect addressing of memory, the bulk-synchronous parallel programming model, and message aggregation techniques are commonly found throughout many areas and implementations of scientific computation. More importantly, the results described herein provide insight into the effects of the architectural characteristics employed by Cielo in order to address issues critical to large scale computers based on multi-core processors.

Porting applications throughout the XT series evolution, and now to the XE6, has been straightforward. However, throughout this evolution, performance of applications typically degraded, or at least did not improve appreciably, most often attributable to bandwidth or network contention. The addition of NUMA nodes to the XT6 node architecture, combined with the faster DDR-3 memory, has reversed this trend, putting application performance above that of any of the previous XT generations, despite the lack of increase in processor clock speeds. This information is motivating explorations into alternative code configurations, specifically with regard to on-node threading (e.g. via OpenMP) and Partitioned Global Address Space (PGAS) languages.

The addition of the Gemini interconnect to the XE6 provides an evolutionary improvement to most of our applications, that is, those configured to reduce the number of messages they transmit by increasing their size through message aggregation and other means and those that are latency constrained. More notably, Gemini's significantly increased message injection rate can provide significant performance improvements to codes that send relatively many smaller messages, as demonstrated here by xNOBEL. One implication of this is the potential for an even greater impact on Partitioned Global Address Space (PGAS) languages, which we are also investigating in the context of important computations.

We will continue to investigate and report on the performance characteristics and capabilities of the Cielo architecture, focusing on the issues described above. In particular, we are configuring experiments that we hope will more fully expose the casual relationships between the processor, node architecture, and node interconnect. Further, we look forward to comparing the effects of the Cielo node architecture with that of the Hopper II Cray XE6 recently installed at NERSC.

ACKNOWLEDGMENTS

The Cray Application Performance Analysis Group provided invaluable support, guidance, and advice throughout this work.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energys National Nuclear Security Administration under contract DE-AC04-94AL85000.

REFERENCES

- [1] S.R. Alam, J.A. Kuehn, R.F. Barrett, J.M. Larkin, M.R. Fahey, R. Sankaran, and P.H. Worley. Cray XT4: an early evaluation for petascale scientific simulation. In *Proceedings of the 2007 ACM/IEEE conference on Supercomputing*, SC '07, pages 39:1–39:12, New York, NY, USA, 2007. ACM.
- [2] K. Alvin, R. Brightwell, and S. Dosanjh et al. On the Path to Exascale. *International Journal of Distributed Systems and Technologies*, 1(2), 2010.
- [3] James Ang, Doug Doerfler, Sudip Dosanjh, Scott Hemmert, Ken Koch, John Morrison, , and Manuel Vigil. The Alliance for Computing at the Extreme Scale. In *Proceedings of the 52nd Cray User Group*, 2010.
- [4] Brian W. Barrett and K. Scott Hemmert. An Application Based MPI Message Throughput Benchmark. In *Proceedings of the 2009 IEEE International Conference on Cluster Computint (Cluster'09)*, 2009.
- [5] Ron Brightwell, Kevin T. Pedretti, Keith D. Underwood, and Trammell Hudson. Seastar interconnect: Balanced bandwidth for scalable performance. *IEEE Micro*, 26:41–57, 2006.
- [6] A. Geist and S. Dosanjh. IESP Exascale Challenge: Co-Design of Architectures and Algorithms. *International Journal of High Performance Computing*, 2009.
- [7] E. S. Hertel, Jr., R. L. Bell, M. G. Elrick, A. V. Farnsworth, G. I. Kerley, J. M. McGlaun, S. V. Petney, S. A. Silling, P. A. Taylor, and L. Yarrington. CTH: A Software Family for Multi-Dimensional Shock Physics Analysis. 1993.
- [8] S. Kelly and R. Brightwell. Software architecture of the lightweight kernel, catamount. In *Proceedings of the 47th Cray User Group*, 2005.
- [9] Los Alamos National Laboratory. xNobel benchmark, 2009. LA-CC 09-61.
- [10] J.D. McCalpin. STREAM: Sustainable memory bandwidth in high performance computers. www.cs.virginia.edu/stream.
- [11] James L. Tomkins, Ron Brightwell, William J. Camp, Sudip Dosanjh, Suzanne M. Kelly, Paul T. Lin, Courtenay T. Vaughan, John Levesque, and Vinod Tipparaju. The Red Storm Architecture and Early Experiences with Multi-core Processors. *International Journal of Distributed Systems and Technologies (IJ DST)*, 1(2), April – June 2010.
- [12] C.T. Vaughan, M. Rajan, D. Doerfler, and R.F. Barrett. Poster: From Red Storm to Cielo: Performance Analysis of ASC Simulation Programs Across an Evolution of Multicore Architectures. In *SC '10: Proceedings of the 2010 ACM/IEEE conference on Supercomputing*, 2010.
- [13] R.P. Weaver and M. L. Gittings. Massively parallel simulations with DOE's ASCI supercomputers: An overview of the Los Alamos Crestone project. In *Proceedings of the Chicago Adaptive Mesh Refinement Workshop*, 2004. LA-UR-04-0012.