Workshop for the Cell/B.E. Processor

The Cluster Workshop

Building Low-Cost Clustering Infrastructure with the Cell/B.E.

Chris Poulin (Patterns & Predictions)
July 10 2008

First Slide -> Why the Cell/PS3?

ps3cluster.org

Workshop for the Cell/B.E. Processor

The Cluster Workshop

- The Cell/PS3
- Gravity Grid
- Further Work
- Current Work
- Why Grid Focus
- So where to go/Hypothetical

Georgia Tech: Poulin (July 10 2008)

- Why the interest in PS3 clustering? A: The Cell Processor 'packs a punch'. One of the authors (Khanna) estimates that his MPI computations run much faster than on desktop workstation chipsets, and that a 8 node Cell cluster has comparable if not better performance to a 200 Node IBM Blue Gene system. And of course this computing power is available at a significantly lower price.

Next Slide -> The Gravity Grid

ps3cluster.org

Workshop for the Cell/B.E. Processor

The Cluster Workshop

- The Cell/PS3
- Gravity Grid
- Further Work
- Current Work
- Why Grid Focus
- So where to go/Hypothetical

Georgia Tech: Poulin (July 10 2008)

- This led to the Gravity Grid by Khanna(http://gravity.phy.umassd.edu/ps3.html)

Next Slide-> Further Work

ps3cluster.org

gravity.phy.umassd.edu

- Groups
- Users
- Webmail

PlayStation3 Gravity Grid

This section is dedicated to the ongoing research projects of our group that involve ***supercomputing*** in Physics.

The **Sony PlayStation 3** has a number of unique features that make it particularly suited for scientific computation. First, the PS3 is an open platform, which essentially means that one can run a different system software on it, for example PowerPC Linux. Next, it has a revolutionary processor called the **Cell processor** which was developed by Sony, IBM and Toshiba. This processor has a main CPU (called the PPU) and several (six for the PS3) special compute engines (called SPUs) available for raw computation. Moreover, each SPU performs vector operations, which implies that they can compute on multiple data, in a single step. Finally, its incredibly low cost make it very attractive as a scientific computing node i.e. part of a cluster. In fact, its highly plausible that the ***raw computing power per dollar*** that the PS3 offers, is significantly higher than anything else on the market today!

Thanks to a very generous, partial donation by **Sony**, we have a **sixteen PS3 cluster** in our department, which we call ***PS3 Gravity Grid***. Check out some pictures of the cluster here: **1) the PS3's arrive**; **2) the rack arrives**; **3) front view of the cluster**; **4) side view of the cluster**.

We are using "stock" PS3s for this cluster, with no hardware modifications. They are networked together using an inexpensive netgear gigabit switch. For Linux installation, there are several guides available on the internet. For YDL Linux, consider using the guide by Terrasoft Solutions. For Fedora 8, I found this guide particularly useful. For deploying a parallel job on this cluster, we use a code that implements a standard domain decomposition approach, based on message-passing (MPI). There are more details available on our code below. For compiling, we use GCC and also IBM's XL compilers for the Cell, that are available as part of IBM's Cell SDK. The MPI distribution that we are using is the recently released, OpenMPI distribution for PowerPC Linux.

Projects

- Binary Black Hole Coalescence using Perturbation Theory (GK)

  This project broadly deals with estimating properties of the gravitational waves produced by the merger of two black holes. Gravitational waves are "ripples" in space-time that travel at the speed of light. These were theoretically predicted by Einstein's general relativity, but have never been directly observed. Currently, there is an extensive search being performed for these waves by the newly constructed ***NSF LIGO*** laboratory and various other such observatories in Europe and Asia. The ***ESA*** and ***NASA*** also have a mission planned in the near future - the ***LISA*** mission - that will also be attempting to detect these waves. To learn more about these waves and the recent attempts to observe them, please visit the LISA mission website.

The evolution code for the extreme-mass-ratio limit of this problem (referred to as ***EMRI***) is essentially like an inhomogeneous wave-equation solver which includes a very complicated source-term. The source-term describes how the smaller black hole (or star) affects the space-time of the larger one. Because of the computational complexity of the source-term, it is often the most numerically intensive part of the whole evolution. On the PS3's Cell processor, it is precisely this part of the computation that is farmed out to six SPUs. This approach essentially eliminates the entire time spent on the source computation and yields a speed up of over a ***factor of five*** over a PPU-only computation. It should be noted that the context of this computation is ***double-precision floating point*** operations. In single-precision, the speed-up is significantly higher.

Overall, a single PS3 performs better than the highest-end desktops available and compares to as many as 25 nodes of an IBM Blue Gene supercomputer. And there is still tremendous scope left for extracting more performance through further optimization. More on that soon.

Furthermore, we distribute the entire computational domain across the sixteen PS3s using MPI (message passing) parallelization. This enables the entire cluster to run together in parallel, working on the computation in an efficient way. Each PS3 works on its part of the domain and communicates the appropriate data to the others, as needed.

Questions? Feel free to contact Gaurav Khanna about this research and the PS3 Gravity Grid.

People

Core Faculty

- Dana Fine, UMassD Math, MA
- J. P. Hsu, UMassD Physics, MA
- Jae-Hun Jung, UMassD Math, MA
- Gaurav Khanna, UMassD Physics, MA

Collaborative Faculty

- Martin Bojowald, Penn State, PA
- Lior Burko, UAH, AL
- Daniel Cartin, NAPS, RI
- Scott Hughes, MIT, MA
- Jorge Pullin, LSU, LA

Current Students

- Sarah McLeod, UMassD Physics, MA
- Emanuel Simon, UMassD Physics, MA
- Francis Boateng, UMassD Physics, MA
- Subir Sabharwal, UMassD Physics, MA
- Joshua Liberty, UMassD Physics, MA
- Ian Nagle, UMassD Physics, MA
- Pranesh Sundararajan, MIT Physics, MA

Past Students

- Vishnu Paruchuri, UMassD Physics, MA
- Jessica Rosen, UMassD Math, MA
- Peter Goetz, UMassD Physics, MA
- Seth Connors, UMassD Physics, MA
- Zhenhua Ning, UMassD Physics, MA
- Nobuhiro Suzuki, UMassD Physics, MA
- Mike O'Brien, UMassD Physics, MA
- Matt Strafuss, UMassD Physics, MA

PS3 Gravity Grid in the Media

- Wired Magazine & ABC News
- Computerworld
- Standard Times
- Herald News
- NZ Herald News
- DNA India
- GCN News
- eSchool News
- New Scientist Magazine
- Daily Telegraph UK
- The Age & Sydney M. Herald
- PSXExtreme Interview
- Other ..

Webmail Not Enabled

Webmail is not enabled for this server. Please contact your system administrator.

Object not found!

The requested URL was not found on this server. The link on the referring page seems to be wrong or outdated. Please

Mac OS X Server

- Groups
- Users
- Webmail

Workshop for the Cell/B.E. Processor

The Cluster Workshop

- The Cell/PS3
- Gravity Grid
- Further Work
- Current Work
- Why Grid Focus
- So where to go/Hypothetical

Georgia Tech: Poulin (July 10 2008)

- Further Activity: The Gravity Grid was expanded out to 16 PS3s (Khanna and Poulin). And external MPI based grids tested on non-PS3 hardware (Poulin). Furthermore, the PS3Cluster Guide written (Poulin and Khanna), first draft completed (July 2008). Beta release for Georgia Tech conference participants (username: "cluster" password: "workshop").

Next Slide -> Current Work

ps3cluster.org

Workshop for the Cell/B.E. Processor

The Cluster Workshop

- The Cell/PS3
- Gravity Grid
- Further Work
- Current Work
- Why Grid Focus
- So where to go/Hypothetical

Georgia Tech: Poulin (July 10 2008)

- Current Project Activity: Khanna has been awarded an NSF grant http://www.nsf.gov/awardsearch/showAward.do?AwardNumber=0831631 Please email him for more information as he conducts extensive testing of the Cell. Poulin is currently programming Grid based environments in Open MPI (on a PS3 Grid), Stackless Python, Condor , MapReduce/Hadoop, and the Google App Engine.

Next Slide -> Why Grid Focus?

ps3cluster.org

Workshop for the Cell/B.E. Processor

The Cluster Workshop

- The Cell/PS3
- Gravity Grid
- Further Work
- Current Work
- Why Grid Focus
- So where to go/Hypothetical

Georgia Tech: Poulin (July 10 2008)

- And why the focus on Grids?... These systems empower massive calculations such as those needed in science such as computational linguistics and analytics (Poulin) or computational physics (Khanna).

- But...these systems are hard to program. If not conceptually difficult, they are time intensive (man-hours). Once one has an implementation in one environment, it is heavily tied to that architecture. This is a 'non-trivial problem' as it is currently the challenge of one of the co-author's dev teams (Poulin). So...

Next Slide -> Where to go?

ps3cluster.org

Workshop for the Cell/B.E. Processor

The Cluster Workshop

- The Cell/PS3
- Gravity Grid
- Further Work
- Current Work
- Why Grid Focus
- So where to go/Hypothetical

Georgia Tech: Poulin (July 10 2008)

- Therefore...we are increasingly focused on, and as 'The Cluster Workshop', we are calling for Distributed Compile AND Runtime Environments to empower Grid applications (especially on the Cell). To date, MPI is one of the best raw systems, while the Google App Engine is one of the best known service (that we know). The Google App Engine has: 1) A widely accepted Python 2.5 / Django code interface. 2) Serious computing power behind the environment. 3) And of course it's Google, so we know about it. However it is propietary, in one language (Python), and may present other issues in terms of adoption (see this article as to why).

- A hypothetical scenario... what if we could build systems, such as games, that fully utilized peer-to-peer node calculations for communications, graphics rendering, and storage? Note that there are already proprietary (closed) systems out there that do this sort of thing (see here). But what if we could build strong data 'fabrics' so that these systems were more common and more accessible infrastructure for developers?

- Any ideas are welcome!

Thank you! -> Chris Poulin and Gaurav Khanna

ps3cluster.org

PS3Cluster Guide: Beta (Preview)

The Cluster Workshop

The Cluster Workshop

The Cluster Workshop is a pending distributed computation working group with The University of Massachusetts (Dartmouth Campus). As there are many challenges that face those that develop software in the building of truly 'distributed' applications, our project's goal is to make clustered computing easier. Our prior experience includes Gaurav Khanna's Gravity Grid, and Chris Poulin's work to embed systems in distributed AI. The Ps3cluster Guide website is our first attempt to bring clustering information to the programmer. In this case we hope to get others to build a low cost cluster infrastructure on commodity hardware. Our future ideas include the development of frameworks to enable easier programming in scientific, enterprise, and gaming applications.

Members

Chris Poulin - Principal Researcher: Gaurav Khanna - Principal Researcher

Andrei Kersha (Fedora 8 Testing) : Josh Gigantino (Image capture & Testing)

Contact Info

Chris Poulin
chris@poulinholdings.com
+1 617.755.9049

Site CSS Template derived from 'Andreas08.css' by Andreas Viklund