



COMING SOON: RESEARCH IN A CLOUD

By Pam Frost Gorder

A trend is taking shape in the computing industry that could significantly change the way academic research is done. A few years from now, researchers who work with massive data sets might stop processing their data locally and find themselves outsourcing the job to massive commercial data clusters. The US National Science Foundation (NSF) is working with Google, IBM, HP, Intel, and Yahoo to promote the development of technologies that will make these super-sized clusters—called computing clouds—amenable to research.

Whither the Cloud?

The computing cloud concept is so new that it has no set definition yet, but the term likely grew out of the way engineers represent computer networks in diagrams (networks are often drawn as clouds). As more people use Web-based software applications for everyday activities such as emailing, blogging, and photo sharing, most of the action on the Internet is happening on networks far from users and their personal computers. It's happening on the multi-acre compounds of service providers such as Google, Yahoo, and Microsoft, where tens of thousands of servers share terabytes of data. The sites are called data centers or server farms, and they're large enough to effectively function as Webs of their own.

The average Internet user seems comfortable with the idea of server farms; most bloggers don't know the physical location of their blog data, for instance. They just know that the blog is out there, somewhere, in a "cloud." Could researchers be as comfortable with their own data being "out there, somewhere?" They might, given that the NSF's Cluster Exploratory (CluE) initiative's goal is to develop resources that will make cloud computing easy and reliable.

To Jeannette Wing, assistant director for Computer & Information Science and Engineering at the NSF, the agency's first cloud partnership is a case in point. Since February 2008, the NSF has been working with Google and IBM to turn an academic computing cluster into a resource for researchers.

"Here's the beauty of the situation: we don't even know where that cluster is. And it doesn't matter," Wing says. What

does matter is that the cluster offers software and services that the scientific community would benefit from using.

A second NSF cluster project was announced in July, and it brought several other commercial giants into the mix: the Cloud Computing Testbed (CCT) at the University of Illinois, Urbana-Champaign, will be partially funded by HP, Intel, and Yahoo and will be one of the six "centers of excellence" comprising the larger international HP, Intel, Yahoo CCT.

In a sense, the research community has already unconsciously moved toward solutions for data-intensive computing, says Indranil Gupta, one of the CCT's principal investigators. The testbed will focus on developing system software to make those solutions work for applications that draw on thousands of gigabytes or even hundreds of terabytes of data. "There are people who talk of petabytes and exabytes as well," he adds. "When you have this huge amount of data, it completely changes the cluster management problem from what it was before."

NSF's CluE initiative has awarded three research grants so far—to Carnegie Mellon University, Florida International University, and the University of Maryland, College Park. These early projects are intended to initiate cluster research activities, as well as work out any logistical problems for accessing the equipment. As many as 10 more grants might follow, up to US\$500,000 each, and the agency is currently reviewing new proposals it received this summer.

The CluE program officer, James French, says that the NSF is "looking for projects that can extend the data-intensive computing environment, or projects that are really only feasible using this platform. This is an enabling technology for projects that have previously been impractical or infeasible because of the volume of data needing to be analyzed."

Wing echoes this sentiment, saying that good CluE projects would demonstrate new ways of using the cloud for scientific research or address new problems not already adequately served by other architectures.

A Cloud of Their Own

There's no shortage of data-intensive research today. Most

notably, astronomy, climate science, microbiology, and particle physics suffer from an overabundance of data. The scientific community has long been pining for a powerful and well-maintained computing infrastructure, such as a worldwide computer grid, to support research. Experts have suggested that the grid currently supporting the Large Hadron Collider could be a potential candidate (“Physics Experiment Could Spawn Permanent Computing Grid,” *CiSE*, vol. 9, no. 6, 2007, pp. 5–8). But if CluE succeeds in making commercial computing clouds work for research, it could build a new kind of bridge between academia and the private sector.

How would such an arrangement work? Theoretically—in daily research practice, anyway—cloud computing wouldn’t feel all that different from more familiar forms of data-intensive computing. For example, when NASA releases new satellite data today, a climate scientist might download a subset of that data to a supercomputer center where her university has an account. She’ll run applications to process the data remotely, and she’ll know where her data are, insofar as she knows where the supercomputer center is. Or she might process the data on a grid, where—again—applications will process the data remotely. And she knows that her data are somehow distributed among university facilities that are part of that particular grid.

With cloud computing, she would still process her data remotely, but she wouldn’t know where the data were stored or where the processing was actually happening. Wing says the experience would feel nearly identical, except that the data processing will go much faster, for two reasons. First, clouds store data very close to the processors, so there’s no transfer delay. Second, clouds are made to analyze massive amounts of data in parallel.

She offered an example: “Rather than have one instance of a huge data set that you access periodically from a single computer, you can partition this data set into 100,000 chunks, and store each partition right next to its own processor. Then you can do a computation in parallel over 100,000 machines, instead of sequentially on one machine.”

The researcher wouldn’t be aware of what was happening, explained the CCT’s Gupta. Middleware and software will manage the experience. “Data can be located on multiple machines, but to the end user who is trying to query the data, it will all look like one,” he says.

Each type of computer, be it a supercomputer, grid, or cluster computer, is natural for some kinds of applica-

tions. Wing suggests that the scientific community can take a step back and ask, “Which architecture is best suited to my application?” Because of all the data-intensive computational needs that scientists currently have, a cluster computer might very well be the best architectural choice. It’s a great opportunity for the computer science community to come together with the other science disciplines to get the most from this new kind of computing configuration, she says.

A Matter of Trust

One issue that will have to be addressed is data provenance. Researchers will want to be sure that the data they’re accessing are the right data, and when necessary, they’ll want to protect sensitive or proprietary information.

Government policy makers are just beginning to catch up to the issue of cloud data security for the general public. In September 2008, the Pew Internet and American

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE

GRADUATE FELLOWSHIPS



Benefits Include:

- A competitive yearly stipend of \$32,400
- Payment of tuition and required fees
- Research practicum at a DOE laboratory
- Yearly fellows conference
- Renewable for up to four years

Application Deadline:
January 14, 2009.

For more information, visit:
www.krellinst.org/csgf

Contact: The Krell Institute
1609 Golden Aspen Drive,
Suite 101, Ames, IA 50010
515.956.3696
csgf@krellinst.org
www.krellinst.org/csgf



Sponsored by the U.S. Department of Energy Office of Science and
MNSA Programs. Administered for USDOE by the Krell
Institute under contract DE-FG02-97ER25308. This
is an equal opportunity program that is open to all
qualified persons without regard to race, sex, creed,
age, physical disability or national origin.





For more information visit: www.krellinst.org/csgf

MOLECULAR DYNAMICS ON GRAPHIC PROCESSING UNITS: HOOMD TO THE RESCUE

By Joshua A. Anderson and Alex Traveset

The insatiable sophistication of video games over the past decade has fueled a demand for powerful and fast graphic processing units (GPUs). Although GPUs provided huge opportunities for scientific applications early on, the considerable difficulties in programming them discouraged broad implementation. This landscape dramatically changed in Spring 2007 when Nvidia released the compute unified device architecture (CUDA; <http://developer.nvidia.com/object/cuda.html>), which provides low-level hardware access to GPUs via the C language, unleashing the possibilities of GPU computing to the scientific community.

What makes GPUs so powerful? In a typical CPU architecture, transistors are specialized to perform data processing as well as data caching and flow control. A GPU is designed to perform a large number of highly parallel instructions as quickly as possible. Therefore, the GPU architecture must maximize the number of transistors devoted to performing data processing at the expense of those that perform data caching and flow control. In simple terms, the GPU's traditional job is to render images represented as 2D arrays of independently calculated pixels. These arrays require relatively modest storage and little manipulation but must be processed as fast and efficiently as possible. GPU computing with CUDA uses the same parallel hardware, but instead of drawing pixels, performs integer and floating-point operations in millions of independent threads.

For these reasons, scientific applications—such as molecular dynamics (MD) simulations,¹ which can be executed as data-parallel computations with a high arithmetic

density—are particularly well suited for the GPU. In MD simulations, N particles are evolved through time according to Newton's equations of motion. Positions and velocities at the next time step depend only on those from the previous step. Therefore, MD can be cast as a data-parallel calculation with at least N parallel threads.

However, implementing MD on a GPU isn't as straightforward as it might appear. In any realistic MD simulation, each particle interacts with the remaining $N - 1$ particles. This $O(N^2)$ cost in computational time is prohibitively large in most applications, and methods exist to avoid it. If the interaction potential between particles is sufficiently short-ranged, as is the case in a variety of relevant scientific applications, it's possible to truncate the potential so that particle interactions beyond a certain distance r_{cut} can be ignored. Each integration step can then be computed with $O(N)$ operations. Similar methods also exist for long-range potentials. The price to pay for such a drastic speedup is that a *neighborlist*—a bookkeeping of which particles are close to one another—is needed at each time step. Implementing the neighborlist in the GPU presents a significant challenge, but we've developed efficient solutions,² as has J.A. van Meel and colleagues.³

Our early GPU implementation,² dubbed HOOMD for highly optimized objected-oriented MD and freely available under an open source license, showed spectacular speedups compared with traditional distributed memory clusters. It became clear that making GPU MD available for general users was well worth the effort. This goal reached an important milestone on 13 August 2008, when we released HOOMD v0.7.0. Besides several additional features, HOOMD's most important facility is a scripting system that makes configuring and running MD simulations available to general users. Figure 1 shows a HOOMD script that runs a standard Lennard-Jones MD simulation. Other examples are described in the tutorial provided

continued on p. 10

Life project released a study that found that nearly 70 percent of Americans who use the Internet use some form of cloud computing—mostly for email service and photo sharing.

The report was released in Washington, D.C., during a Google policy forum. Panelist Michael R. Nelson, a visiting professor at the Center for Communication, Culture, and Technology at Georgetown University and former technology policy advisor to President Clinton, commented that cloud computing is “as important as the Web was 15 years ago,” while Ari Schwartz, vice president of the Center for Democracy and Technology, added that the government will have to address users' expectations of privacy as cloud computing continues to grow (*The Los Ange-*

les Times, 12 September 2008; <http://latimesblogs.latimes.com/technology/2008/09/most-web-users.html>).

Wing acknowledged that researchers might have difficulty with the idea of data security when they don't know where their data are located. “But there's definitely a trust relationship between the user and the provider of the cloud service,” she says. “As researchers, we're still working out those kinds of issues.”

To guarantee that data are correct, cloud providers would have to be in control of all the sources of data that users are sifting through, says Roy Campbell, coprincipal investigator of the CCT, and it's not clear how that would be done. Solving this problem will “breed a whole new sort of activity about data provenance, and make data prove-

nance more important.” He predicts that statistics will be applied to show where the data are coming from and ensure data validity.

Cloud Semantics

Much of cloud computing these days involves superficial Web searching, Campbell says, and scientists are interested in going deeper—understanding natural language, for instance, or using Web content to aid machine learning. “That’s clearly the next step in Web searches—organizing things and looking at this data with a deeper understanding.”

At Carnegie Mellon, a CluE-funded project grew out of efforts to build Web databases for intelligent language tutoring. Computer scientist Jamie Callan explains that a computer can help people learn languages by offering them reading material that they’re interested in, tailored to suit their language level. Because computers aren’t good at generating meaningful language texts on the fly, his team harvests documents from the Web and filters them for grammar and content. Out of every 1,000 documents they collect, they might discard all but one or two. When they first began the project, their software sifted through Web pages on a small number of student computers, a process he calls “a real pain in the neck. We ran computers for months collecting documents and filtering them down, and we never had a large enough data set.”


Then they took the project to Yahoo’s M45 academic cluster. They collected 200 million documents in less than two months, and reduced them to a database of roughly 6 million that could be used in the classroom.

“Once we had access to this big data

set, we realized we could do all sorts of things that we couldn’t do before,” he says. “We became interested in whether we could create a large Web data set that could be broadly shared with the research community. We’ll use the Google–IBM cluster to do that.” Other researchers will be able to use the database to do text mining, as Callan’s group did, while his team will use it to explore ways of improving search efficiency. The goal this time is to gather a billion documents and index them by subject to make searching faster. Then users could perform a targeted search of only a small subset of the database, but get results that were just as accurate as if they had searched the whole thing.

With Callan’s indexing, clouds could compute results faster and use less electricity. That’s a big deal, since a typical server farm might consume up to 100 watts of electricity per square foot, and the newest generation of densely-packed server farms—such as the four that Microsoft is currently constructing in Iowa, Illinois, Texas, and Washington—could consume thousands of watts per square foot. According to *BusinessWeek*, reducing power consumption in a typical data center by 25 percent could result in US\$4.5 million a year in savings (“It’s Too Darn Hot,” 20 March 2008; www.businessweek.com/magazine/content/08_13/b4077060400752.htm).

The ability to search cloud data will spur some changes that are already under way in computer science, Campbell says. There’s a “human element” to computation, as exemplified by Google. When users query Google, they get a list of results ranked in part by user interest. “You’re factoring in a human element, and assuming that the human element is doing the right thing,” he says. “For scientists,



QUANTUM COMPUTING WITHOUT MAGIC

Devices

Zdzislaw Meglicki

“*Quantum Computing without Magic* gives a refreshing and down to earth approach to quantum information as well as quantum devices which will have a fundamental impact on the technology of the 21st century.”

— Raymond Laflamme, Canada Research Chair in Quantum Information, and Director, Institute for Quantum Computing, University of Waterloo

Scientific and Engineering Computation series
448 pp., \$35 paper



USING OPENMP

Portable Shared Memory Parallel Programming

Barbara Chapman, Gabriele Jost, and Ruud van der Pas

foreword by David J. Kuck

“This book makes OpenMP knowledge available to the average programmer in an understandable, easy-to-apply fashion, while still providing information for those who wish to dive more deeply into the subject.” — Larry Meadows, CEO, the OpenMP ARB

Scientific Computation and Engineering series
384 pp., 35 illus., \$35 paper

continued from p. 8

```
from hoond_script import *

# create 100 random particles of name A
init.create_random(N=100, phi_p=0.01, name='A')

# specify Lennard-Jones interactions between
particle pairs
lj = pair.lj(r_cut=3.0)
lj.pair_coeff.set('A', 'A', epsilon=1.0,
sigma=1.0, alpha=1.0)

# integrate at constant temperature
integrate.nvt(dt=0.005, T=1.2, tau=0.5)

# run 10,000 time steps
run(10e3)
```

Figure 1. A HOOND script to run a molecular dynamics simulation of a Lennard-Jones liquid. The second and third commands specify the cut-off's value that determines the neighborlist and sets the coefficients for the Lennard-Jones potentials. The second to last command specifies the integration step as well as the necessary parameters for setting up the thermostat. The last command runs the simulation for 10^3 steps. Computer-savvy readers might recognize that the HOOND scripts are written in Python.

at the HOOND Web page (www.ameslab.gov/hoond). Typical simulations in HOOND with a single GTX 280 GPU perform at a level equivalent to the computer power of a 42-core computer cluster.

GPUs are powerful yet small, easy to maintain, consume less energy, and are much cheaper than CPU clusters. In addition, only a simple graphic card upgrade can convert nearly any computer desktop to the equivalent of a 42-core cluster. Over the past five years, the computational horse-

power and memory bandwidth of GPUs has grown at a much faster rate than CPUs, so we expect GPUs to be even more efficient than their CPU counterparts in the future. For all these reasons, it's inevitable that GPU computing will soon become the commonplace solution for small and medium MD simulations and other scientific applications.

Acknowledgments

We acknowledge Chris Lorenz for his contributions to this article's initial stages. We also thank C. Philips and R. Sknepnek for their valuable contributions to HOOND and Nvidia for providing us with numerous GPUs through the Professor Partnership Program. HOOND is funded by the US National Science Foundation grant DMR-0748475 (previously DMR-0426597) and by the US Department of Energy through the Ames Lab under contract DE-AC02-07CH11358.

References

1. D. Frenkel and B. Smit, *Understanding Molecular Simulations*, Academic Press, 2002.
2. J.A. Anderson, C.D. Lorenz, and A. Travasset, "General Purpose Molecular Dynamics Simulations Fully Implemented on Graphics Processing Units," *J. Computational Physics*, vol. 227, no. 10, 2008, pp. 5342–5359.
3. J.A. van Meel et al., "Harvesting Graphics Power for MD Simulations," *Molecular Simulation*, vol. 34, no. 3, 2008, pp. 259–266.

Joshua A. Anderson is a graduate student in the department of physics and astronomy at Iowa State University and also works for the Ames Lab. Anderson has a BS double-major in computer science and physics from the Michigan Technological University. He is a member of the American Physical Society. Contact him at joaander@ameslab.gov.

Alex Travasset is an associate professor in the Department of Physics and Astronomy at Iowa State and an associate scientist at the Ames Lab. Travasset has a PhD in physics from the Universitat de Barcelona. Contact him at trvsst@ameslab.gov.

Reach Higher

Advancing in the IEEE Computer Society can elevate your standing in the profession.

GIVE YOUR CAREER A BOOST • UPGRADE YOUR MEMBERSHIP

www.computer.org/join/grades.htm

it's difficult to think about how to accommodate these issues in experiments, but it's moving the whole of this area toward the social sciences—psychology, as well as other disciplines."

Wing agrees that cloud computing will enable researchers to look at data in new ways. "You can explore new kinds of questions in a cloud," she says, "that you wouldn't dare to on a supercomputer."



Pam Frost Gorder is a freelance science writer based in Columbus, Ohio.