e science TALK
e-ScienceBriefings
Talking about e-science

Desktop grids: Connecting everyone to science

Today's personal computers are powerful but, most of the time, a large proportion of their computational power is left unused. A desktop grid takes this unused capacity, no matter what its location, and puts it to work solving scientific problems. With over 1 billion desktop computers in use, desktop grids can offer a low cost, readily available computing resource for scientists while allowing citizens across the world to contribute to scientific research. Together with grids and supercomputers, desktop grids can be a useful complement to the e-Infrastructure landscape.

What is a desktop grid?

Desktop grids fall into two categories - local and volunteer. While local desktop grids are comprised mainly of a set of computers at one location, a business or institute for example, the resources in a volunteer desktop grid are provided by citizens all over the world.

Today researchers are using desktop grids to simulate protein folding *(Folding@home)*, find ways to provide clean water *(IBM's World Community Grid)*, and model climate change *(Climateprediction.net)*. Scientific problems that can be split up into small tasks and sent to different computers for computation are perfect for solving with desktop grids. These projects farm out tasks to computers located across the globe which send the results back to scientists once they are complete.

Desktop grids can open up computing to researchers that otherwise would be unable to access such large amounts of computing power. For example the Citizen Cyberscience Centre has been set up to help scientists in developing countries access the power of internet-based volunteer networks. Initiatives such as Africa@Home and Asia@Home have also encouraged more researchers in these regions to use desktop grids.

David Anderson, BOINC director – *"BOINC is being used by over 50 volunteer computing projects, doing research in everything from quantum mechanics to cosmology. About 430,000 PCs from all over the world, many of them equipped with GPUs, participate in these projects. Together they supply about 6 PetaFLOPS of computing power."*

A large majority of volunteer computing projects are based on open source software called BOINC. BOINC allows scientists to plug their own projects into the software, so volunteers can easily download and run applications on their computer. The BOINC client, used by the volunteers, can be configured to run only when the PC is not in use, often as a screensaver, or to run at the lowest priority while the PC is in use. Other desktop grid middlewares include XtremWeb, developed by INRIA/CNRS, which is mainly used to manage computations on desktop computers within an organisation.

Bringing physics home

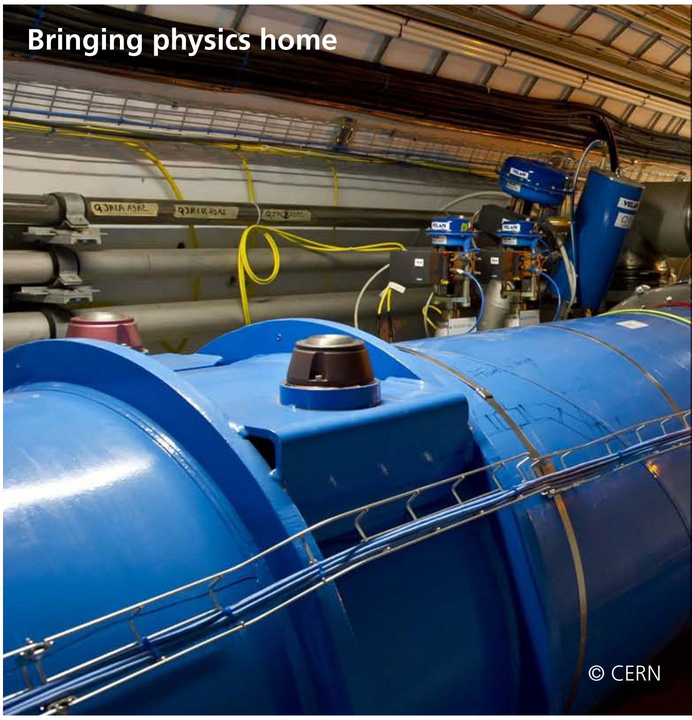

© CERN

LHC@home 2.0 aims to bring the world's largest particle accelerator into your home. The platform – an extension of the already successful LHC@home – allows volunteers to connect to CERN-based research projects simply by donating their extra computing power. The project Test4Theory, for example, simulates high-energy particle collisions which scientists can compare to real-life collisions, such as those occurring in the Large Hadron Collider (LHC).

"My dream is to be able to establish a 'virtual LHC', which would require being able to generate 40 million events per second, as much as the real LHC, running at full steam," says Peter Skands, the lead scientist behind Test4Theory. "We estimate that it would take somewhere between 10 000 and 100 000 connected computers to achieve this, a combined amount of computing power that we have only faintly begun to imagine, since we started working with LHC@home 2.0. With the enthusiasm we have seen in the public so far, there definitely appears to be awesome possibilities for what we can do with this platform."

Nicole Vasapoli, BOINC volunteer - *"I'm a meteorologist so was originally interested in donating my computing time for climate research as it was related to my work. But, as time went on, I began to take part in many other exciting life science projects, to assist scientists in solving pressing problems. I've found that BOINC is free, easy to install and it's an entertaining way to be part of scientific progress."*

Netting malaria

Anyone going on holiday to a malaria-affected country will often head to their doctor for a course of malaria tablets. But for those who live in countries at risk, taking preventative medicines is impractical.

So what do citizens of these areas do? They use mosquito nets or insect repellent, treat their houses with insecticide, or get rapid treatment in the event of becoming ill.

While none of these methods are perfect, each can make a big impact given the widespread nature of the disease. Malaria is preventable and treatable but three quarters of a million people die from it every year, putting healthcare services in affected countries under enormous strain.


Thomas Omondi / UK Department for International Development

Healthcare providers and governments need to find ways to determine the most effective combination of treatments for their area. They can use mathematical models to simulate the effectiveness of different combinations of malaria control, and work out what the best solution is for a given situation.

In 2003 researchers at the Swiss Tropical and Public Health Institute started running malaria models to answer these questions. Starting with just 50 of their own PCs, they soon opened up the project via BOINC and malariacontrol.net was born. Today 50,000 people contribute computing time to the project, through over 70,000 PCs. In total, malariacontrol.net and its volunteers have notched up over 10,000 CPU years helping healthcare professionals fight malaria.

Enhancing other e-infrastructures

Desktop grids are just one of a number of ways in which researchers can access computing capacity. They can provide a useful complement to the other facilities in the e-infrastructure landscape. While supercomputers are able to solve a wide variety of complex computational problems they are expensive and are limited to a relatively small number of researchers. Cluster-based grids can provide a cheaper solution for more researchers, but for a more limited set of applications. Assuming the computers making up a desktop grid are already paid for, they can open up computational research to more scientists at an even lower cost.

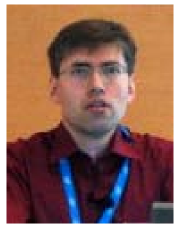

Mikhail Posypkin, Institute for Systems Analysis of Russian Academy of Sciences *"Desktop grids offer a cost-effective alternative to supercomputers or service grids. Unlike supercomputers or service grids large desktop grids are almost free: all you need is a meaningful distributed application. Presently desktop grids can contribute their resources to existing grid infrastructures thus producing a really powerful combined distributed computing infrastructure."*

The EDGeS (Enabling Desktop Grids for e-Science) project, which has now finished, worked to connect desktop grids to the wider EGEE (Enabling Grids for E-sciencE) European infrastructure. Using the gLite middleware the project defined common policies to integrate desktop grids into existing EGEE service grids. Today EDGI (the European Desktop Grid Initiative) is continuing this work by connecting desktop grids into the European Grid Infrastructure (EGI). As well as focusing on gLite, EDGI aims to build bridges to the UNICORE and ARC middlewares.

Desktop grids have high capacity but not a guaranteed quality of service as the available computing power depends on which computers are not being used. EDGI hopes to include a cloud in the infrastructure, which can be used as and when necessary. This will allow the service to be used by applications which need to be completed within a specific deadline.

By integrating desktop grids with other e-infrastructures, researchers can run applications across different types of computing resources, matching parts of the computational problem to the most suitable execution environments. For example, some parts of an application can be run on a desktop grid, and others on a high-end supercomputer. Collaborations between a number of infrastructure providers, such as those established through the International Desktop Grid Federation (IDGF) are already in place to enable these applications. On the international scale the DEGISCO project aims to export and share desktop grid knowledge outside of the EU, while policy bodies like e-IRG are preparing the setup of the legal and political frameworks.

Vicky Huang, ASGC – *"Asia is a geographically large region, with diverse and scattered resources (technologies and facilities) coupled with the general problem of insufficient investment from government in academic hardware supply. As such, the concept of desktop grids and volunteer computing is very suitable and useful to popularise in the Asia-Pacific region through initiatives such as DEGISCO."*

Desktop computing challenges

Desktop grids can provide a variety of different benefits, however their use raises a number of challenges. A Desktop Grids for eScience Road Map produced by the DEGISCO project in July 2011 took a closer look at some of the following issues:

- **Supporting a desktop grid:** Aside from having to develop applications that can run across a number of heterogeneous systems, the distributed nature of a desktop grid poses unique problems. As volunteers provide the resources, it is difficult to test and fix applications.
- **Making it green:** Desktop grids are often touted as a 'green solution' as they use computing resources already in existence. However, in reality, determining whether a desktop grid is green, or not, is complex. How volunteers choose to donate their computing time plays a big part in this – adding on a CPU load to a machine running at a low capacity doesn't cost much energy, but using a computer that would otherwise be switched off does. Even the country a machine is running in can make a real difference. Connecting a computer in a hot country such as Dubai to a desktop grid is likely to use more energy, as the machine needs to be kept cool.

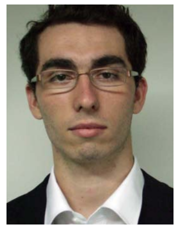

Morgan Duarte, BOINC volunteer - *"I'm sharing my computing resources with BOINC to help solve tomorrow's challenges and be more involved in scientific progress and our future. I believe it is an efficient way to use our continuously increasing computers' power without affecting my own personal use. For me, volunteering my computer for science is very rewarding."*

- **Local policies:** Desktop grids are subject to the local ICT policies at the institute or organisation that is hosting the donated computer. For example, if a company chooses to switch off computers at night, this can affect the availability of the desktop grid.
- **Evolving hardware:** Today increasing numbers of people are accessing the internet through new technologies such as mobile phones instead of PCs. In the future this evolving situation could have consequences for the desktop grid concept as it currently stands.

Leslie Versweyveld, AlmereGrid & IDGF *"The special feature desktop grids have to offer is that they are already part of the e-Infrastructures landscape. We are actually sitting on a huge source of computational power that is largely left unused in numbers of universities, research institutes, companies, home offices and households. It is already there, we only need to tap into it, technically gain access to it and transform this enormous resource of computational power to fuel e-science research in all possible areas. Basically, it is a mere question of ecological recycling."*

When citizens become scientists

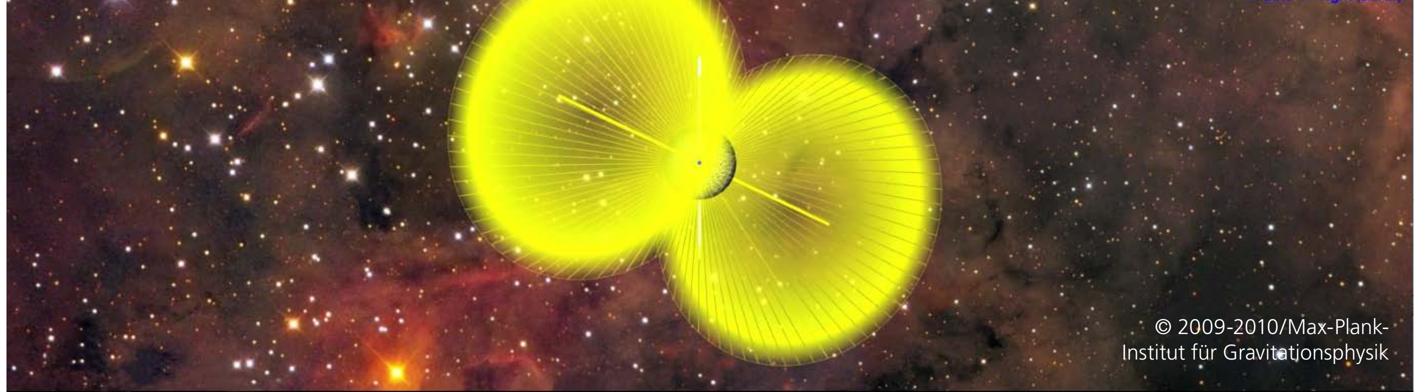

© 2009-2010/Max-Planck-Institut für Gravitationsphysik

When Einstein@Home discovered a new pulsar its discovery wasn't credited to astronomers, but to its volunteers - Daniel Gebhardt, from Mainz, Germany, and husband-and-wife team Chris and Helen Colvin of Ames, Iowa.

Pulsars are highly magnetised, rotating neutron stars that emit a beam of electromagnetic radiation. Einstein@Home, a BOINC project, was originally set up to search for gravitational waves in data from the US LIGO Observatory.

However in March 2009, the project also began to use its volunteers' computers to search for signals from radio pulsars in observations from the Arecibo Observatory in Puerto Rico.

The new pulsar, discovered in October 2010 and named PSR J2007+2722, was the first deep-space discovery by Einstein@Home. Since then a further seven pulsars have been discovered by Einstein@Home volunteers, showing how donating your computer can make a real difference.

Managing a desktop grid

Unlike supercomputers or cluster-based grids, desktop grids have an extra component that needs to be managed – their volunteers. Using volunteers to donate computing time forms the basis of all volunteer desktop grids, and can create positive links between citizens and science.

Francois Grey, Citizen Cyberscience Centre
"In my view, the most revolutionary aspect of volunteer computing is the public participation. Far from being passive, many participants turn volunteer computing into a serious hobby. Some contribute to debugging the software, others help newcomers in the forums, still others set up teams and events to encourage more participation. I predict that ultimately, this will lead to public involvement in setting the agenda for the research that is carried out using public resources. Just as has already happened for journalism on the web, the distinction between amateur and professional will start to blur."

The first step - recruiting volunteers - needn't be a difficult one. When the project LHC@home began, its creators thought it would attract no interest. However one thousand people downloaded the application in the first 24 hours with no publicity effort at all. Often volunteers are interested in the area of science they are contributing towards such as searching for new drugs or ways to generate clean water. AlmereGrid has taken recruitment one step further, by setting up a 'city grid' intended to reach out to volunteers that may not be traditionally interested in donating computing time. AlmereGrid has partnered with local and national companies across Almere in the Netherlands to disseminate information on volunteer computing and get more people interested in the topic.

While projects do not need to pay volunteers to use computing resources, they do need to keep volunteers informed. To ensure volunteers' interest is sustained over a project's lifetime they should be provided with feedback and information on how the project's research is progressing.

Glossary

CPU: Central Processing Unit; a microprocessor (a processor on an integrated circuit) inside a computer that can execute computer programs.
GPU: Graphics Processing Unit; a device that renders graphics for a computer. GPUs have a highly parallel structure that makes them more effective than general-purpose CPUs for some complex processing tasks.
Quality of service: the ability to guarantee a certain level of performance.

Scan this QR code into your smart phone for more on this e-ScienceBriefing

Fundraising through science

The Charity Engine has ambitions to be a worldwide computer. Launching in summer 2011, Charity Engine will provide volunteers' computing time to a collection of hand-chosen projects and raise money for charities at the same time. By joining Charity Engine, its volunteers will also have the chance to win a cash prize of up to a million dollars, every few weeks.

Charity Engine raises funds for its associated charities, as well for its prize draws, by selling volunteers' computing time in bulk to science and industry. Its volunteers are not asked to support any particular science project they simply agree to let Charity Engine send ethical work to their PCs.


www.sxc.hu

"Our volunteers are joining to make computer-generated charity donations and prize draw entries, they might not actually care about the science," says Mark McAndrew, founder of Charity Engine. "But that's fine, because all that idle, wasted computing power will make Charity Engine the ultimate supercomputer - and we love the science."

For more information:

AlmereGrid: *www.almeregrid.nl*
BOINC: *http://boinc.berkeley.edu*
Charity Engine: *www.charity-engine.org*
Citizen Cyberscience Centre: *www.citizencyberscience.net*
DEGISCO: *www.degisco.eu*
Desktop Grids for eScience - A Road map: *http://bit.ly/DEGISCOroadmap*
EDGI: *www.edgi-project.eu*
Einstein@Home: *http://einstein.phys.uwm.edu*
e-IRG: *www.e-irg.eu*
IDGF: *www.desktopgridfederation.org*
LHC@home: *http://lhcathome.web.cern.ch*
Malariacontrol.net: *www.malariacontrol.net*
XtremWeb: *www.xtremweb.net*
EGI: *www.egi.eu*
iSGTW: *www.isgtw.org*
e-ScienceTalk: *www.e-sciencetalk.org*

e-ScienceTalk is co-funded by the EC under FP7
INFSO-RI-260733