Convergence of distributed computing infrastructures with focus on Desktop Grids

Peter Kacsuk and Robert Lovas
Computer and Automation Research Institute,
Hungarian Academy of Sciences (MTA SZTAKI)
Budapest, Hungary
rlovas@sztaki.hu

DEGISCO is supported by the FP7 Capacities Programme under contract nr RI-261561

The Grid Ecosystem

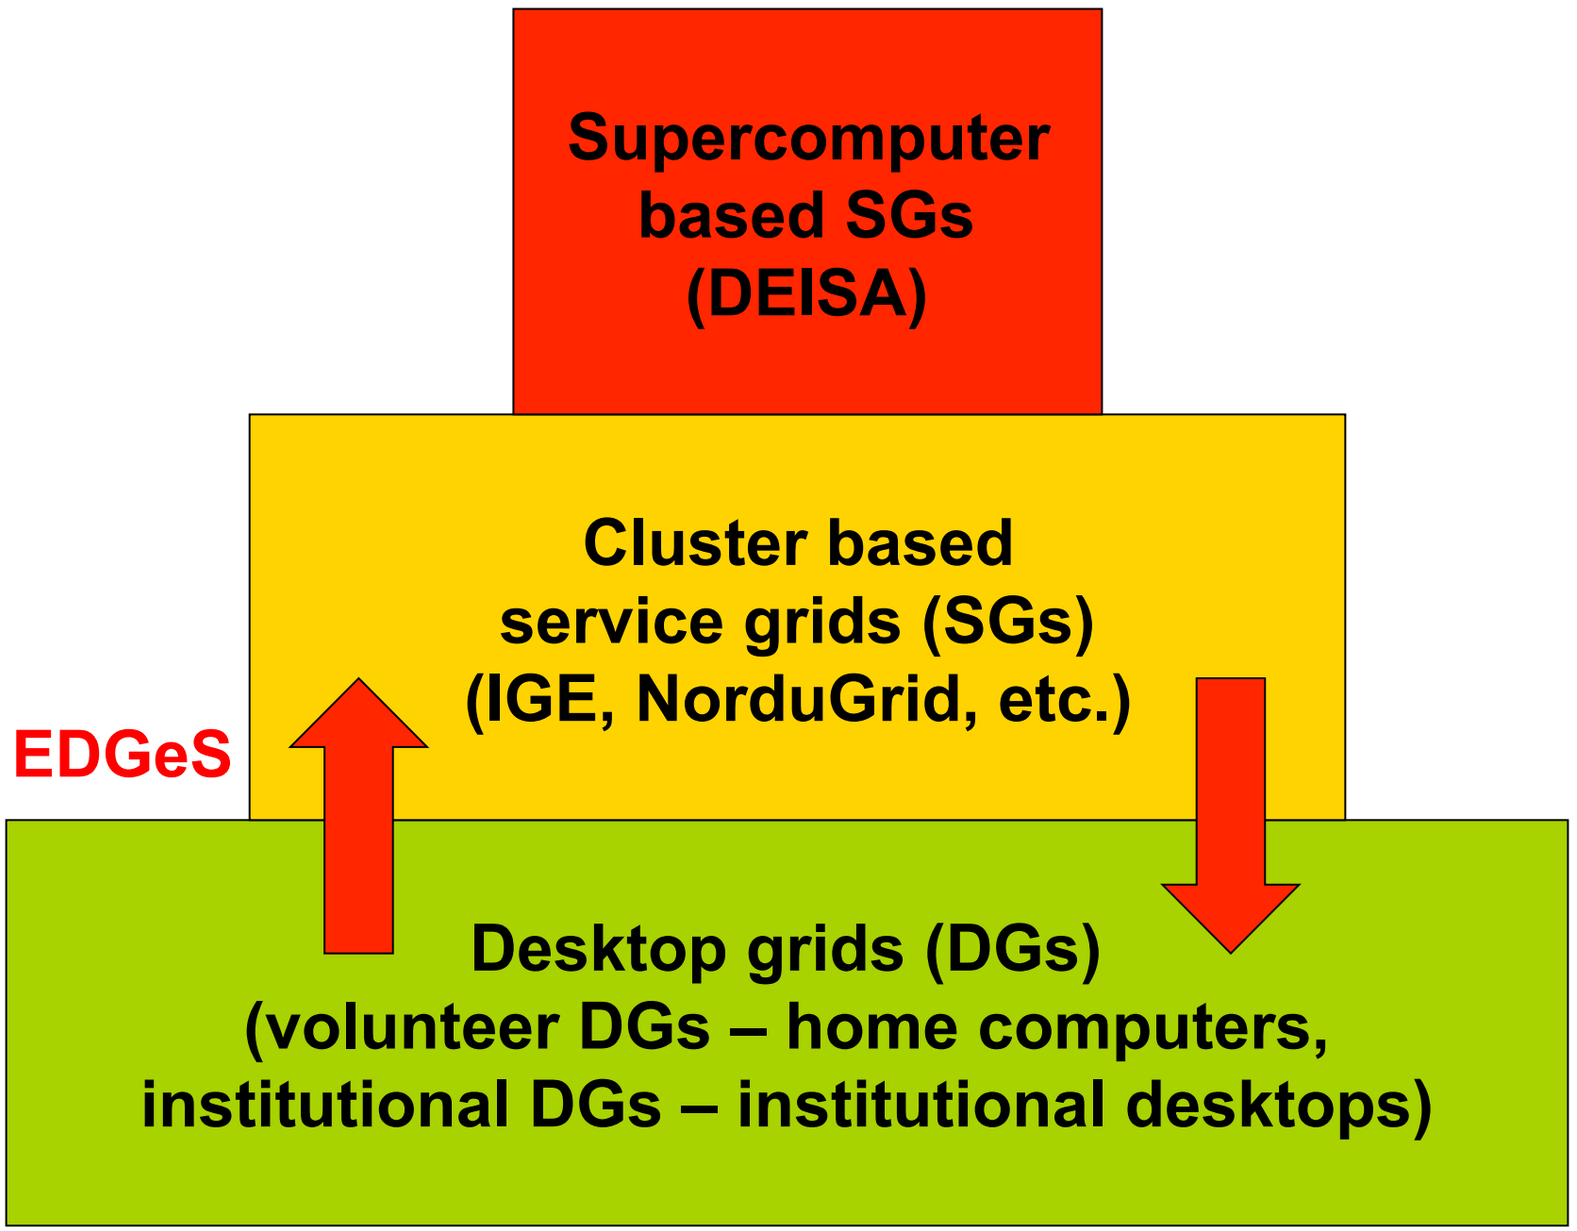

- **Very expensive,**
- small number of sites and very large number of cores
- **MPI appls**
- **Moderately expensive,**
- **moderate number of sites and CPUs**
- **any appls**
- **Inexpensive,**
- **very large number of CPUs (~10K – 1M)**
- **Bag of task appls**

EU FP7 projects on desktop grids: EDGeS –> EDGI and DEGISCO

EDGeS

- **DG<->SG integration:**
  - **gLite → BOINC, XtremWeb**
  - **BOINC, XtremWeb → gLite**
- **Compute intensive applications**

2008 - 2010

further developed by

supported by

EDGI

- **ARC, Unicore, Clouds**
- **QoS with Clouds**
- **Data intensive apps**
- **SG->DG direction support**

DEGISCO

- **Disseminate EDGeS results world-wide**
- **Green IT aspects**

2010 - 2012

Approaches to distributed HPC/HTC computing

- **Cluster computing:**
  - **dedicated** computers operated in a **single location (organisation)** under a common queuing system
- **Grid computing:**
  - **separate organizations** agree to share their **dedicated** computing resources (supercomputers, clusters, servers, PCs)
- **Desktop grid (DG) computing:**
  - Local DG: desktop PCs **within an organization** (such as a department or university) are used as **non-dedicated** computing resource (i.e. spare capacities and CPU cycles)
  - Global (Volunteer) DG: similar to desktop grid computing but computing resources are volunteered **by the public**.
- **Cloud computing:**
  - **third-party** (company) provides access to computers on a pay-as-you-go basis.

Master/slave parallelism and parameter studies

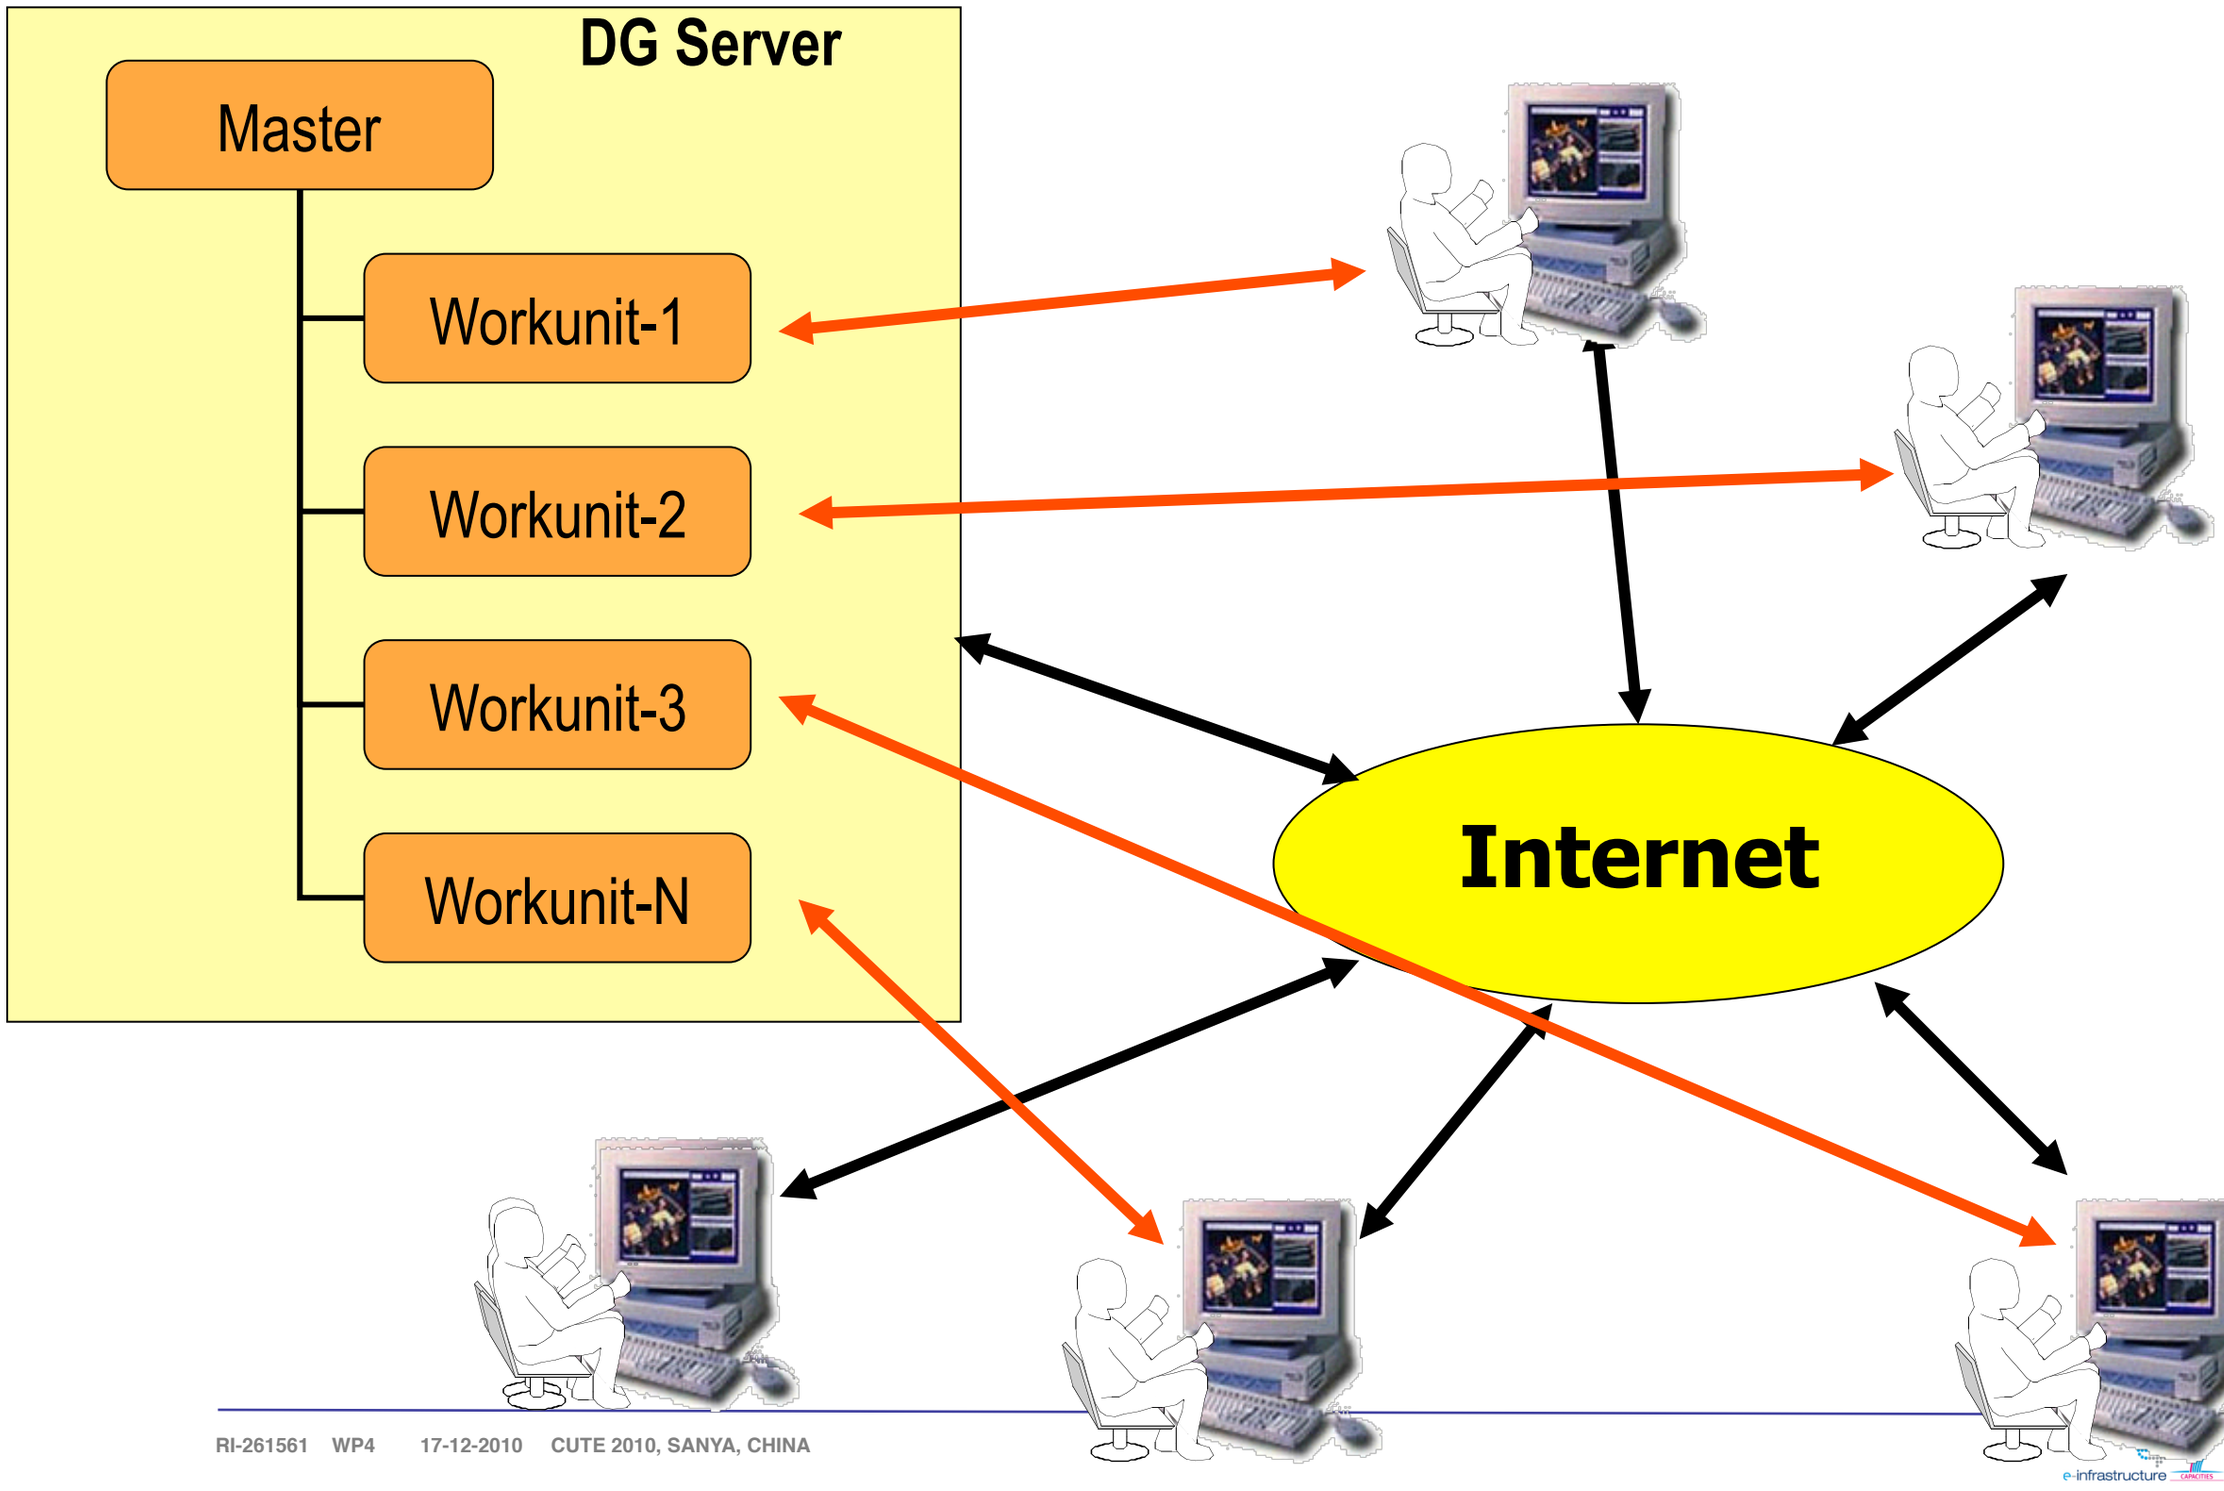

Desktop Grid model

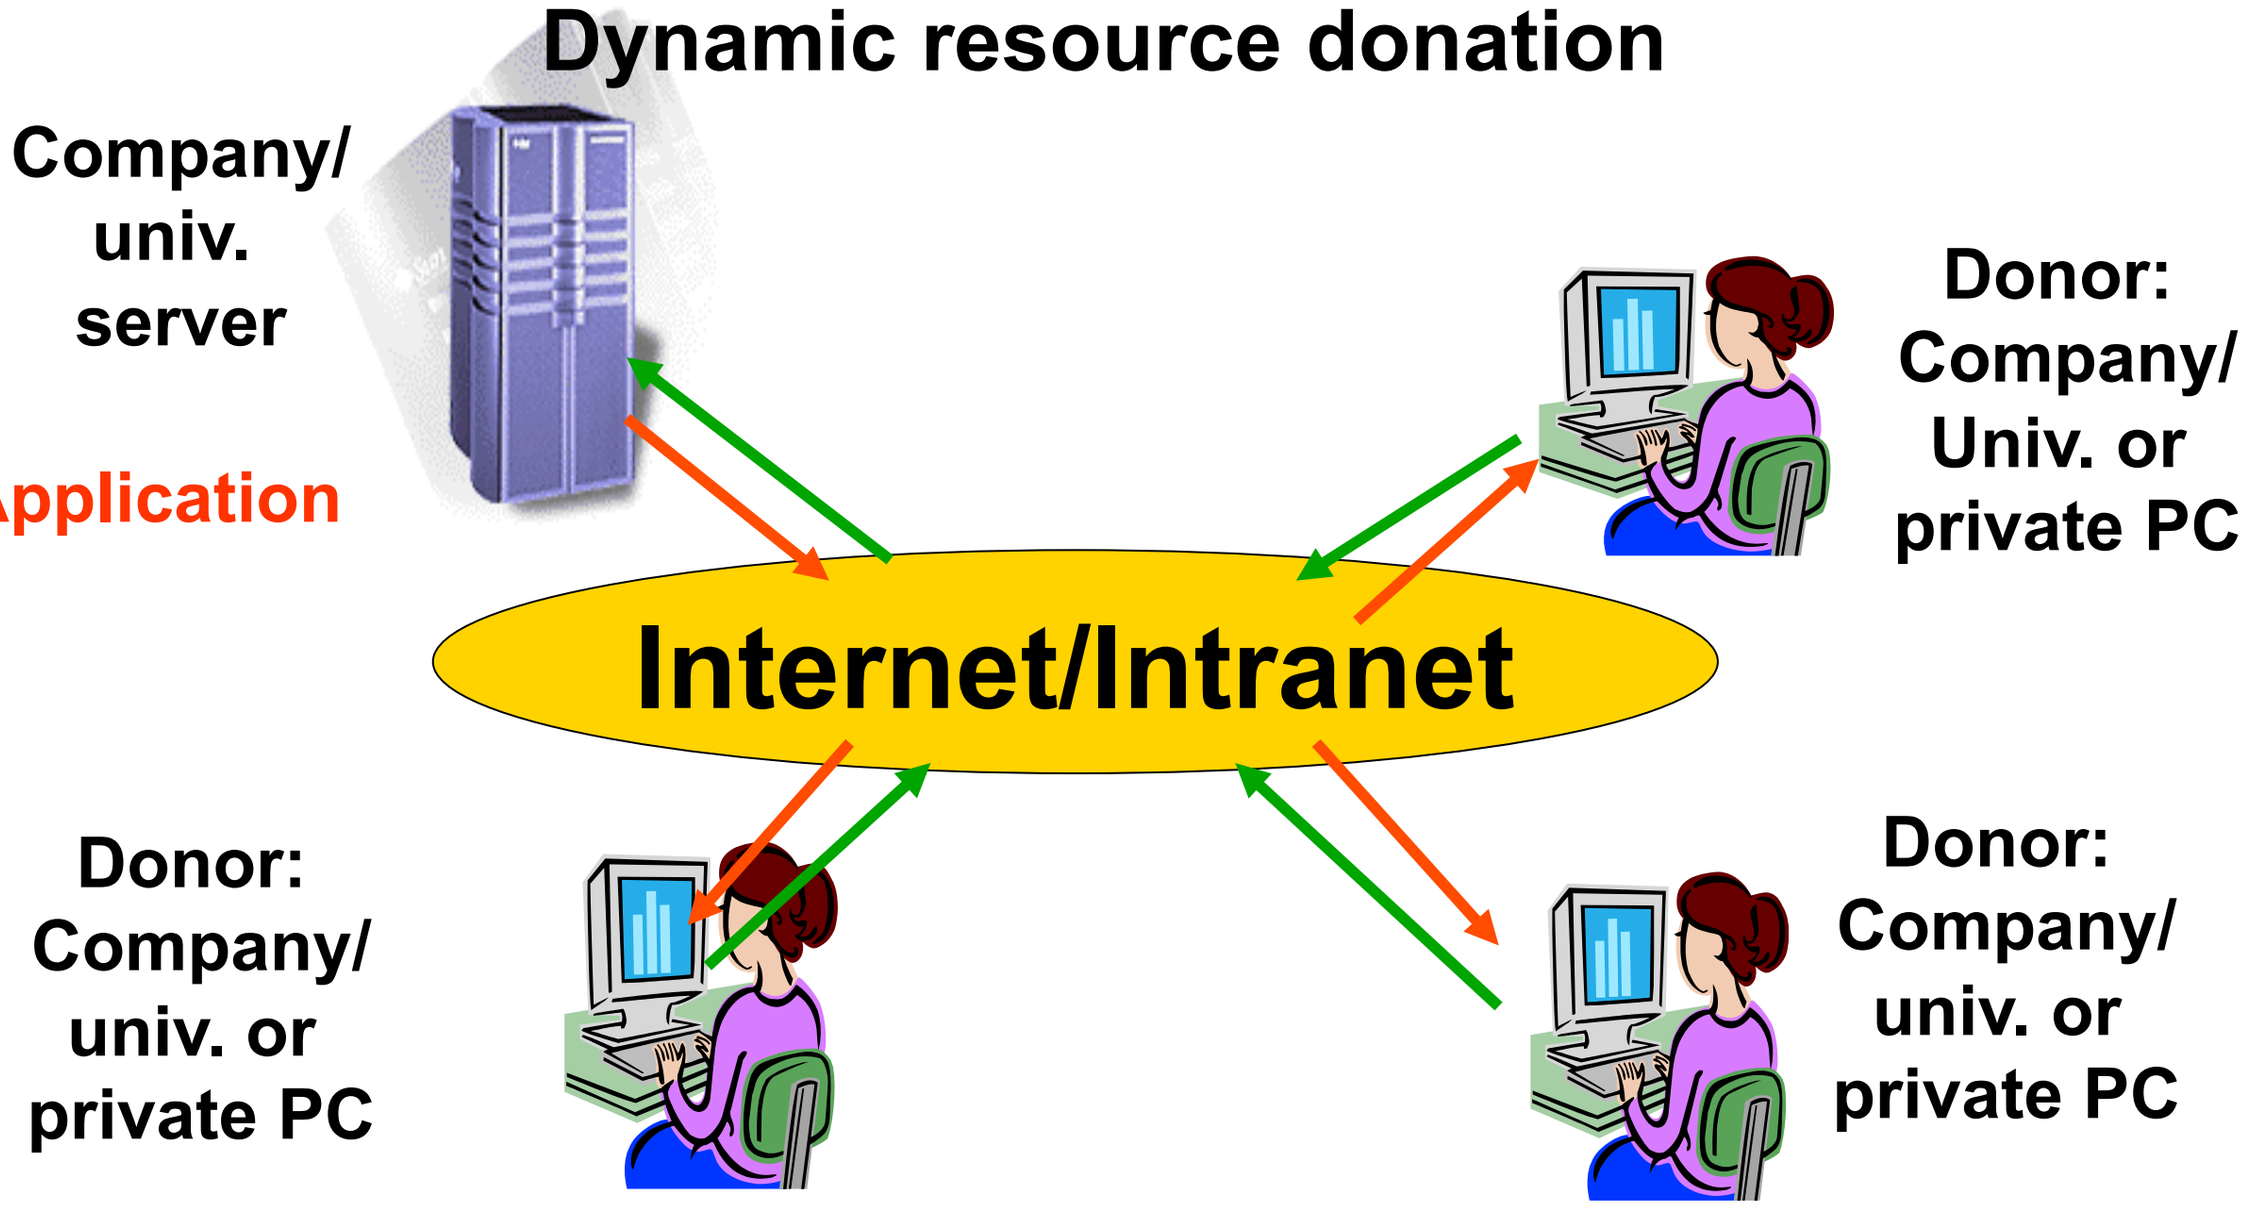

Characteristics of the desktop Grid model

- Anybody can donate resources
- Heterogeneous resources, that dynamically join and leave
- **One or a small number of projects** can use the resources
- Asymmetric relationship between (D) donors and (U) users:

  U << D
- **Advantage**:
  - Donating a PC is extremely easy
  - Setting up and maintaining a DG server is much easier than installing the server sw of utility grids
- **Disadvantage**:
  - Dynamic job submission is not supported
  - Supported application is static (typically very long running application)

Breakthrough in 1999

- **SETI@home** from University of California-Berkeley analyzes data from the Arecibo radio telescope, looking for artificial signals from space.

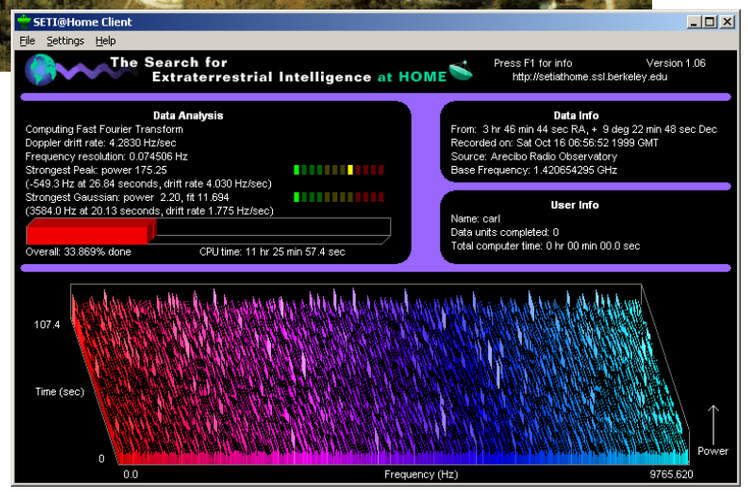

- Search for Extra-Terrestrial Intelligence in radio signals
- 3.8M volunteers from 226 countries (2004)
- 1200 CPU years/day
- 38 TFlops sustained performance (Japanese Earth Simulator 32 TFlops)

Potential of volunteer computing

- Number of privately owned PCs:
  - currently 1.5 billion
  - grow to 2 billion by 2015
- This resource pool is self-financing, self-updating and self-maintaining
  - people buy new PCs, upgrade system software, maintain their computers, and pay their electric bills
- Participant PCs are state-of-the-art general purpose computers:
  - 100 million GPUs (tens of thousands GPUs are already being used for volunteer computing)

BOINC middleware

- In 2002, the BOINC project was established (with funding from the National Science Foundation)
  - To develop general-purpose middleware for volunteer computing, making it easier and cheaper for scientists to use.
- Client-Server approach
  - Most important technical challenges and functionalities will be discussed later.

- Supported class of applications: bag-of-task, parameter sweep/ paramater study, monte carlo simulation, master-worker.
  - At client level: GPUs with e.g. CUDA, multicore processors with SMP programming model.

- Further Desktop Grid middleware solutions:
  - Condor
  - XtremWeb-HEP from INRIA/IN2P3
  - OurGrid (Peer-to-peer)

Some volunteer grids

- World Community Grid - IBM (http://www.worldcommunitygrid.org/) *300.000 PCs*
- Leiden Classical Grid - Education on Grid (http://boinc.gorlaeus.net/) *17.000 PCs*
- SZTAKI Desktop Grid – Hungarian apps (http://www.desktopgrid.hu/) *70.000 PCs*
- AlmereGrid (http://almeregrid.nl) 3.000 PCs
- PS3GRID - Based on Playstations (http://ps3grid.net/)

Performance

- ~1 million computers are ***actively*** participating in volunteer computing → supplying ~10PetaFLOPS of computing power
- Fastest supercomputers supply few PetaFLOPS
- Largest grids (e.g. European Grid Infrastructure) have the hundreds of thousands of hosts.
- In terms of throughput, volunteer computing is competitive.
- Near-term potential of volunteer computing goes well beyond Exa-scale.

Cost effectiveness

- **For scientists, desktop grid computing is cheaper than other paradigms.**
- A medium-scale volunteer project (10,000 computers, 100 TeraFLOPS) has got the overall cost of $20.000 - $200.000 USD by using:
  - a single server computer
  - and one or two staff.
- An equivalent CPU cluster costs at least an order of magnitude more.
- Cloud computing is even more expensive.
  - However, studies suggest that cloud computing is cost-effective for hosting volunteer computing project servers.

Types of Desktop Grids

- **Global Desktop Grid**
  - Aim is to collect resources world-wide for grand-challenge scientific problems
  - Could be volunteer (e.g. SETI@home) and dedicated (e.g. CancerGrid)
- Examples:
  - SETI@home
  - SZTAKI Desktop Grid (SZDG) global version at:
    http://szdg.lpds.sztaki.hu/szdg/
- **Local Desktop Grid**
  - Aim is to enable the quick and easy creation of grid for any community (company, univ. city, etc.) to solve their own applications
  - Could be dedicated (e.g. Wmin DG) and volunteer
- Example:
  - SZTAKI Desktop Grid (SZDG) local version

Local Dedicated DG based on SZDG local version: University of Westminster (London, UK) as a best practice

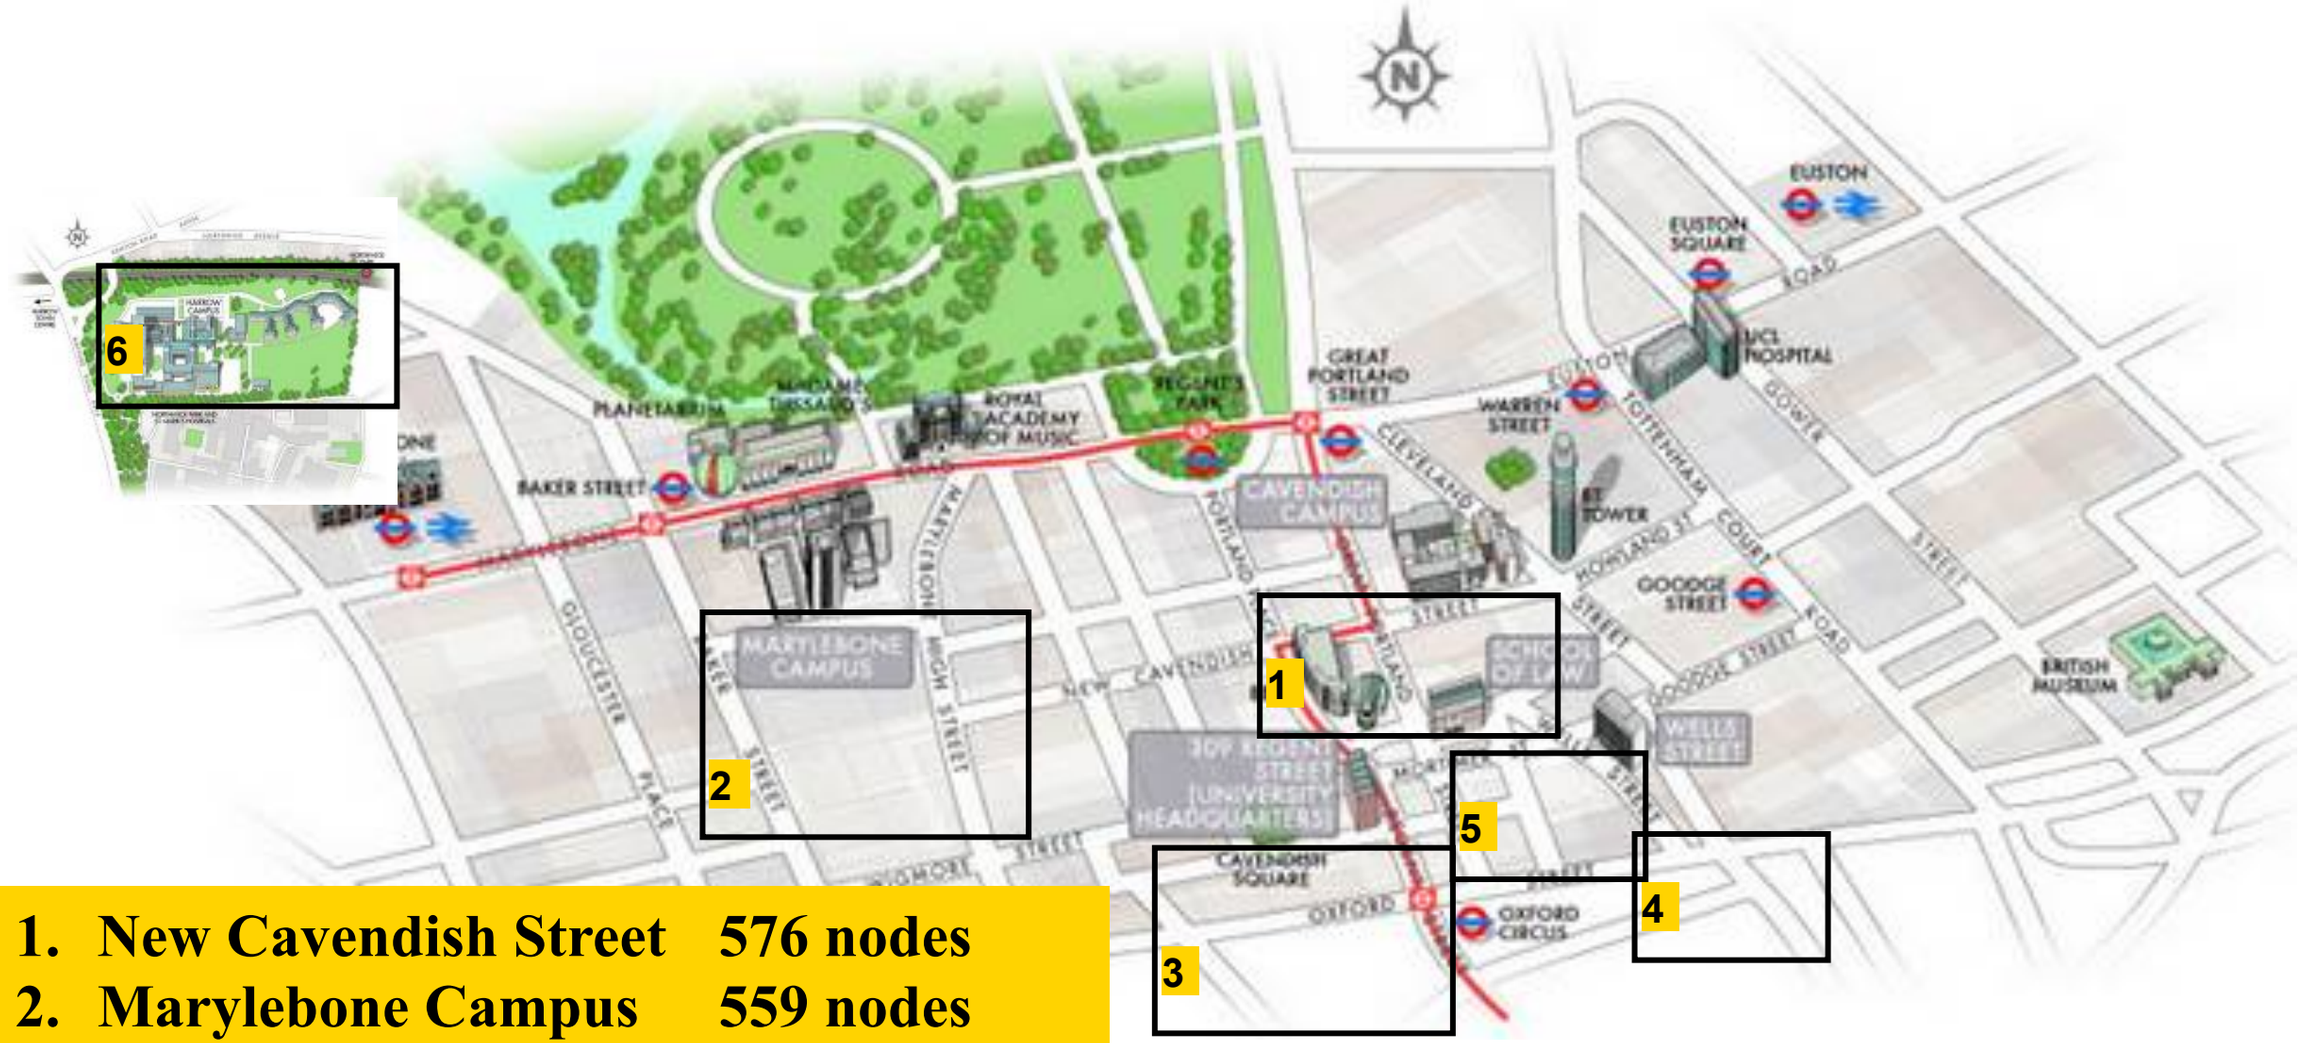

1. New Cavendish Street 576 nodes
2. Marylebone Campus 559 nodes
3. Regent Street 395 nodes
4. Wells Street 31 nodes
5. Little Tichfield Street 66 nodes
6. Harrow Campus 254 nodes

Total: 1881 nodes

Lifecycle of a node:

1. PCs basically used by students/staff
2. If unused, switch to Desktop Grid mode
3. No more work from DG server -> shutdown (green solution)

Scientific adoption

- Volunteer computing has not yet been widely adopted. Less than 100 research groups are currently using volunteer computing.

- Cluster and grid computing are much more widely used by scientists, and local desktop (enterprise) grids are also gaining more attention.

- Some projects founded by the European Commission in FP7 have started to bring closer to each other these communities (EDGeS, EDGI & DEGISCO).

Energy efficiency

- The FLOP/Watt ratio of a PC is lower than that of a supercomputer, and it is tempting to conclude that volunteer computing is less energy-efficient than supercomputing.

- However, in cold climates, for example, energy used by a PC may replace energy used by a space heater, to which the PC is thermodynamically equivalent.

- Other **Green factors** are to be studied by the **DEGISCO** project.

Human factors

- In volunteer computing these factors are particularly crucial and complex. Several surveys have been done: Why do people volunteer?
  - **Support scientific goals:** such as curing diseases, finding extraterrestrial life, or predicting climate change.
  - **Community:** some volunteers enjoy participating in the online communities and social networks.
  - **Competition:** some volunteers are interested in the performance of computer systems, and they use volunteer computing to quantify and publicize the performance of their computers.
- There are attempts to commercialize volunteer computing by paying participants, directly or via a lottery, and reselling the computing power.
- Local (enterprise) Desktop Grids have commercial success (Platform, Parabon) at several companies (e.g. Novartis).

Human factors (cont)

- To attract and retain volunteers, a project must perform a variety of human functions:
  - **Web content** describing its research goals, methods, and credentials.
  - Periodic **updates** on its scientific progress.
  - It must manage the **moderation** of its web site's message boards to ensure that they remain positive and useful.
  - **Publicize** itself by media.

Volunteers must trust projects, but projects cannot trust volunteers.

Security Issues

- ***Security 1:*** What if hackers break into a project server and use it to distribute malware to the attached computers?
  - BOINC prevents this by requiring that executables be **digitally signed** using a secure, **off-line** signing computer.

- ***Security 2***: What if hackers create a fraudulent project that poses as academic research while in fact stealing volunteers' private data?
  - This is partly addressed by account-based **sandboxing**: applications run under an unprivileged user account and typically have no access to files other than their own input and outputs.
  - In the future, stronger sandboxing may be possible using **virtualization** technology.

Future… and current projects

- Volunteer computing has demonstrated its potential for high throughput scientific computing. However, only a small fraction of this potential has been realized. Moving forward will require progress in three areas:

1. Increased participation: The volunteer population has remained around 500,000 for several years. Can it be grown by an order of magnitude or two?

- Effective use of social networks:

  → **WEB2GRID project, IDGF**

- Computer manufacturers or software vendors bundle BOINC with other products → **IDGF**

Future… and current projects (cont)

2. ***Increased scientific adoption***: The set of volunteer projects is small and fairly stagnant.
   - It would help if **more universities and institutions created umbrella projects**, or if there were more support for higher-level computing models, such as workflow management systems
   - promotion of volunteer computing by **scientific funding agencies**
   - increased acceptance of volunteer computing by the **HPC and computer science communities**

→ **EDGeS, EDGI, DEGISCO projects**

3. ***Tracking technology***: Today, the bulk of the world's computing power is in desktop and laptop PCs, but in a decade or two it may shift to **energy-efficient** mobile devices. Such devices, while docked, could be used for volunteer computing.

→ **EDGI, DEGISCO projects**

End of 1st part

Summary: Volunteer Grid computing can help “democratize” the access to High Performance/Throughput Computing

Acknowledgement

The first part of the presentation is particularly based on David P. Anderson (University of Berkeley):

Volunteer Computing – The ultimate cloud (Crossroads, Spring 2010)

Thank you for your attention …

For more information please visit the Websites:

http://degisco.eu

http://edgi-project.eu/

http://desktopgridfederation.org

DEGISCO is supported by a Grant from the European Commission's FP7 IST Capacities programme under grant agreement RI-261561.