
GridOrbit Public Display: Providing Grid Awareness in a Biology Laboratory

Juan David Hincapié Ramos

IT University of Copenhagen
Rued Langgaards Vej 7
2300, Copenhagen, Denmark
jdhr@itu.dk

Aurelien Tabard

IT University of Copenhagen
Rued Langgaards Vej 7
2300, Copenhagen, Denmark
auta@itu.dk

Jakob Bardram

IT University of Copenhagen
Rued Langgaards Vej 7
2300, Copenhagen, Denmark
bardram@itu.dk

Tomas Sokoler

IT University of Copenhagen
Rued Langgaards Vej 7
2300, Copenhagen, Denmark
sokoler@itu.dk

Abstract

We introduce GridOrbit, a public awareness display that visualizes the activity of a community grid used in a biology laboratory. This community grid executes bioinformatics algorithms and relies on users to donate CPU cycles to the grid. The goal of GridOrbit is to create a shared awareness about the research taking place in the biology laboratory. This should promote contributions to the grid, and thereby mediate the appropriation of the grid technology. GridOrbit visualizes the activity in the grid, shows information about the different active projects, and supports a messaging functionality where people comment on projects. Our work explores the usage of interactive technologies as enablers for the appropriation of an otherwise invisible infrastructure.

Keywords

Infrastructure, Infrastructure Awareness, Appropriation, Public Display, Visualization, Community Grid.

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation

General Terms

Design

Introduction

The world is supported by endless systems of infrastructure, from water supplies to communications networks. Biologists, and scientists more broadly, increas-

Copyright is held by the author/owner(s).
CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.
ACM 978-1-60558-930-5/10/04.

ingly rely on a new type of infrastructure: Computational Grids. Such an infrastructure allows scientists to execute demanding jobs on a network of distributed computers. On the down side, like most infrastructures, grids require set-up, maintenance and central control, preventing non-technical users from setting up and running such infrastructures.

Rather than relying on supercomputers, projects such as SETI@Home¹ and Folding@Home² enlist end-user computers into grids. Our work is part of a larger project creating such a community grid of end-user computers. This community grid uses peer-to-peer (P2P) technology for equally distributing tasks between its members.

The efficiency of a community grid is directly proportional to the amount of CPU power that participants are willing to donate. One central challenge is hence to motivate users to participate by installing and running the software on their local machine. And this despite the fact that only a small fraction of users actually has the need for submitting tasks to the grid, and thus benefit from the grid at a certain point in time. Because such a community grid is invisible to the end-user, such participation – or appropriation of the technology – becomes a challenge [1].

In this paper, we explore the potential of using awareness technologies to support the appropriation of a digital infrastructure. We focus on a local deployment of a community grid in a molecular biology research laboratory. Molecular biologists use bioinformatics algo-

rithms for analyzing DNA/RNA sequences of millions of bytes, which typically require substantial computational power. In this context, the community grid provides a way to improve scientific analysis without relying on external servers. The design of this awareness technology is rooted in an iterative participatory design process (see figure 1), which is based on a set of field studies of laboratory work. We describe this process in the following section.

In the second part of the article, we describe the design of GridOrbit, an application that displays an interactive visualization of the underlying activity in the grid (see Figure 2). GridOrbit runs on public displays distributed across the biology department, and shows three broad classes of information: computing power contributions, research activities, and social information.

Finally, we present our plan to evaluate GridOrbit. GridOrbit was designed with the hypothesis, currently being tested in long term deployments, that increasing local awareness of a resource sharing infrastructure will lead to broader user participation, and hence to increased contribution. More specifically, in our P2P grid deployment at the biology department, the hypothesis implies that researchers and staff will both use GridOrbit to reflect on their resource sharing habits. In doing so, we expect that they appropriate the community grid and contribute more computing power to it. Moreover, as appropriation is a collective process [2], we hope to foster discussions about the grid and the research in the lab, which could benefit not only the sharing of research resources, but the overall organization.

¹ <http://setiathome.ssl.berkeley.edu/>

² <http://folding.stanford.edu/>




Figure 1: GridPOrbit's iterative participatory design process.

Field Work

Mainwaring et al. [3] point out that the appropriation of infrastructure is a collective process where the community buys into a new way of living, and a set of conventions of practice emerges. Before building GridOrbit, we thus conducted fieldwork and participatory design activities with future users of the grid, to identify common sharing practices.

SETTING: We conducted 10 days of detailed workplace studies in a molecular biology group part of the overall project we participate in. The group comprises 3 labs and researchers of different levels of expertise.

PARTICIPANTS: Eight biology researchers participated to different participatory design sessions: two of them were post-docs, two were lab technicians, one was the professor head of the group, and the rest were PhD students. Except for the technicians and one physicist, all the participants were biologists. The technicians spend most of their time at the lab and do not have their own computer. The professor spends most of his time at the office, at meetings, or travelling, and rarely goes to the lab. The rest of the participants both work at the lab and at their office spaces.

METHOD: We conducted task-centered observations of biology work; place-centered observations of work in the laboratory and in the office; and artifact-centered observations of the use of biology tools like paper-based lab books, specimens, tubes, freezers, microscopes, and different digital tools. We also documented group meetings and casual encounters between researchers. We pursued detailed observations of 3 biologists: One simultaneously executing multiple experiments in the lab, another switching between the office

space and the lab during the day, and a third one running experiments with hazardous materials. The observations were supplemented with contextual inquiries for selected activities like running an experiment at the bench or working with software suites. One point of inquiry was the previous efforts to build and deploy resource-sharing infrastructures.

Preliminary findings

Our analysis focused on the researchers' resource sharing habits, collaborations, and mobility. We noticed:

RESOURCE SHARING: During our place-centered observations we noticed *extensive sharing of equipment and research material* (tubes, samples, electronic files) between labs. This sharing was tacitly approved as researchers from different groups would walk into any lab and grab some equipment. Organic materials, however, require explicit authorization. Particularly expensive or dangerous equipments have special rooms across the building where all biologists can use them.

PERSONAL ACTIVITIES: We observed that projects are usually small: they include one or two protocols, do not require more than 1 researcher, and can be completed in a few weeks. Many of the researchers are pursuing their PhDs, or post doctoral projects, something that impacts the projects' scale and collaboration. Also, awareness of the work of others can only happen while at the lab, not in the offices. These situations result in an *individualist work* environment, and *low awareness of the group's work*. Nonetheless, we observed efforts to foster collaboration and increase work awareness: semi-formal optional gatherings like the weekly group and the monthly department meeting.



Figure 2: Landscape metaphor

DISTRIBUTED WORK: During our task-centered observations of experiments we noticed researchers have *different workplaces for different activities*, and *switch between the lab and the office several times a day*. Experiments start at their offices by studying the literature, browsing the web, and defining the protocol. Later they migrate to the lab where the empirical work is carried out, sometimes for several experiments simultaneously. Finally, they return to their offices for analysis of results and reporting.

Design of GridOrbit

Following the user studies, we conducted participatory design workshops to co-design GridOrbit with its future users. As input to the workshops we created personas, and exposed different technologies including awareness solutions and public displays. We decided to work with public ambient displays, then created scenarios, sketched paper-based and digital user interfaces, and refined them iteratively. We ran four workshops where we evaluated and improved the current design.

Landscape Metaphor

In our current design, GridOrbit captures data from the P2P grid and visualizes grid activity by showing a landscape of windmills and light bulbs; each windmill represents a computer in the grid, and each light bulb represents a research project (all grid jobs are tagged with a project name). The level of CPU contribution by a computer defines the rotational speed of the windmill. The level of CPU power used by a project defines the glowing level of the light bulb.

The windmill & light bulb visual metaphor (see Figure 2) serves several purposes. First, we believe the concepts of a P2P grid are made easier to understand by

binding them to representations of power and consumption of a windmill farm, a concept already widely understood by the biologists. Second, the tech-savvy users are confronted with a different perspective on P2P grids, one that talks about CPU power instead of devices and protocols. Finally, we present the visualizations on public displays, and we aim for them to become common ground for understanding a shared infrastructure, and facilitating the collective process of appropriation [2]. We also considered other types of visualizations like games and trivia, but we discarded them as we hypothesize ambient information and infrastructure awareness give support for appropriation.

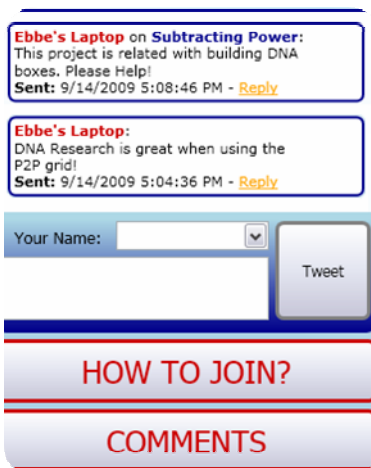
Mapping of Projects and Users

Our field studies revealed that research interests within the organization, though similar in nature, are very different from group to group and even within groups. GridOrbit creates research awareness by informing users about the different research projects in the department and their grid jobs, thereby serving two purposes. First, projects are represented as sets of tasks; this granularity makes the activities more concrete than higher-level projects. Second, publicly showing the project descriptions and activities creates opportunities for knowledge sharing.

We also considered other mappings like the groups within the department and the types of tasks, but we discarded them as too broad categories.

Notifications, Messages and TagCloud

We focused on fostering social interaction, we thus enable viewers to browse and annotate projects on the display. These annotations try to trigger social interactions by constructing topic-based conversations on the grid activity, the specific research projects, the use of

Figure 3: TweetBox component, and *How to Join* and *Comments* buttons.

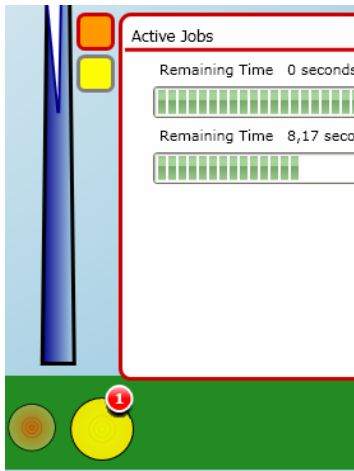


Figure 4: Project details.

public displays, and the annotations themselves. Our approach to using public displays for collaboration purposes differs from others [4, 6, 7] in taking as a basic assumption that social relations are built from what is shared, and that colleagues share both the infrastructure they use, and the research areas they work on. GridOrbit provides messaging and notifications about new messages. It also presents a TagCloud component as a gauge of users' interests.

Ambient Public Displays and Interaction Zones

Researchers' mobility around the facilities suggested us the usage of multiple ambient public displays to create awareness of the invisible P2P grid. Furthermore, Pajo et al. described how interactive displays can sustain such social interaction and enhance both creativity and productivity [4]. GridOrbit is a set of large public displays deployed in two buildings of a biology department. Using ambient public displays opens up a new design space, where ambient technologies are used to visualize infrastructure. These visualizations provide awareness of the work done in the offices. We also considered other technologies like wearables, desktop and mobile applications, but discarded them in order to keep the entry barrier low (price/training), and scalable (around 200 users). In designing the displays, we took inspiration from Prante's Hello.Wall's implicit interaction zones [5]:



Figure 5: TagCloud component.

AMBIENT ZONE: > 70cm - Presents the landscape of light bulbs and windmills (see Figure 2).

NOTIFICATION ZONE: 40-70 cm. Enables the system notifications and the TagCloud. These features aim at capturing the attention of the user (see Figure 5). The system notifications alert about new messages related

to a project in the last 24 hours. The TagCloud is made from extracting nouns from the messages.

INTERACTIVE ZONE: < 40 cm. Enables touch interaction for browsing through projects and users (see Figure 4). Figure 3 shows the TweetBox one can use to create and reply to messages, engaging in conversations. Messages are associated to users and optionally to projects, and can be deleted at any time. Selecting a light bulb will show details about the project and filter out the messages in the TweetBox (see Figure 3). Users can also interact with the TagCloud by selecting a tag, filtering out the messages in the TweetBox. Finally, the user can access information about joining the P2P grid, and leave suggestions on how to improve GridOrbit.

Implementation

GridOrbit is a .NET WPF standalone application running on two 42" touch-enabled displays enhanced with a ProximityBar. GridOrbit monitors the network's UDP traffic and captures packages transmitted by the P2P grid. Relevant information is extracted from these packages and later bound to the WPF UI components. Users interact with grid orbit with touch and an on-screen keyboard. The ProximityBar is a system of three ultrasonic (PING))) sensors controlled by an Arduino board and connected to the display via USB.

Evaluation

We carried out user testing during each design iteration. Moreover, we are currently testing GridOrbit at the molecular biology department. Our long-term deployment of GridOrbit seeks to collect data on the state of the P2P grid, the contributions to the grid, and user interactions with GridOrbit. The system logs grid state and user interaction. These quantitative studies will be

supplemented with qualitative studies targeted at analyzing the impact of GridOrbit in terms of research awareness and social interactions.

CONCLUDING REMARKS AND FUTURE WORK

We presented GridOrbit, an ambient public display for mediating the appropriation of a community grid through infrastructure awareness. GridOrbit visualizes the grid by presenting users and projects as windmills and light bulbs; it promotes the contribution of computing power by showing the current activities in the grid; it provides research awareness by displaying project data; and it supports topic-based social interactions by enabling users to send messages.

The intention is to design GridOrbit in a way that it can benefit beyond an increase in the contribution of CPU cycles to the grid. Benefits that arise from GridOrbit making hitherto invisible lab activities visible to people not directly engaged in research activities in the lab. We expect GridOrbit in the long run will contribute to an increased intra-organizational awareness. This in turn, we speculate, will bring forward numerous new opportunities to engage in social interaction across the immediate professional and organizational boundaries otherwise experienced. Whether our expectations will be rewarded relies on the extent to which people within the organization will appropriate the technology. We will evaluate these aspects over time, in a long term deployment, as participants start using GridOrbit.

REFERENCES

1. M. Chalmers and A. Galani. Seamful interweaving: heterogeneity in the theory and design of interactive systems. In DIS '04: Proceedings of the 5th conference

on Designing interactive systems, pages 243–252, New York, NY, USA, 2004. ACM.

2. P. Dourish. The appropriation of interactive technologies: Some lessons from placeless documents. *CSCW*, 12(4):465–490, 2003.

3. S. D. Mainwaring, M. F. Chang, and K. Anderson. Infrastructures and their discontents: Implications for ubicomp. pages 418–432. 2004.

4. S. Pajo, S. Zubi, A. M. Szstek, and J. Sturm. Design and evaluation of a large interactive display to support social interaction at work. In *Human-Computer Interaction – INTERACT 2009*, volume 5726 of *Lecture Notes in Computer Science*, pages 608–611. Springer Berlin / Heidelberg, 2009.

5. T. Prante, R. Stenzel, C. Röcker, N. Streitz, and C. Magerkurth. Ambient agoras: Inforiver, siam, hello.wall. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 763–764, New York, NY, USA, 2004. ACM.

6. F. Redhead and M. Brereton. Designing interaction for local communications: An urban screen study. In *Human-Computer Interaction – INTERACT 2009*, volume 5727 of *Lecture Notes in Computer Science*, pages 457–460. Springer Berlin / Heidelberg, 2009.

7. D. Snowdon and A. Grasso. Diffusing information in organizational settings: learning from experience. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 331–338, New York, NY, USA, 2002. ACM.