GridOrbit – An Infrastructure Awareness System for Increasing Contribution in Volunteer Computing

Juan David Hincapié-Ramos, Aurélien Tabard, Jakob E. Bardram
IT University of Copenhagen
DK-2300 Copenhagen, Denmark
{jdhr,auta,bardram}@itu.dk

ABSTRACT
The success of a volunteer computing infrastructure depends on the contributions of its users. An example of such an infrastructure is the Mini-Grid, a local peer-to-peer system used for computational analysis of DNA. The speed of analysis increases as more users join the Mini-Grid. However, the invisible nature of such an infrastructure hinders adoption, as it is difficult for users to participate in an infrastructure they are not aware of. This paper introduces GridOrbit, a system designed to increase user awareness, fostering contributions to this infrastructure. We designed GridOrbit using a participatory design process with biologists, and subsequently deployed it for use in a biology laboratory. Our results indicate that the number of contributors to the Mini-Grid increased with the use of awareness technologies. In addition, our analysis presents their motives and behaviors. Finally, a characterization of user interaction with GridOrbit emerged, which enabled us to understand how awareness systems can be better designed. We see GridOrbit as an example of a broader class of technologies designed to create 'Infrastructure Awareness' as a means to increase the contributions to technological infrastructures.

Author Keywords
Volunteer Computing, Infrastructures, Infrastructure Awareness, Public Displays, Ambient Displays.

ACM Classification Keywords
H.5.3 Collaborative computing

General Terms
Experimentation

INTRODUCTION
Volunteer computing is a powerful way to conduct computation-intensive data processing by harnessing computing resources from large numbers of geographically distributed individuals. The most prominent examples of volunteer computing initiatives are SETI@Home and Folding@Home[1], which seek to gain computational power by enlisting end-user computers like PCs and game consoles.

Volunteer computing infrastructures however, like other infrastructures, are often invisible to their end-users [16, 22], and this invisibility poses a fundamental obstacle to their adoption [5]. A core question is how end-users can become aware of such an invisible infrastructure, and start participating and contributing to it. Existing volunteer computing projects rely on individual and social motivations, leading to efforts in building communities, setting up competitions, and rewarding users who participate [18]. Thus, there is a substantial overhead associated with recruiting users who will donate CPU cycles 'for free', and this in turn becomes a core challenge for the scientists using the infrastructure.

In this paper, we explore the use of awareness technologies to recruit contributors to a volunteer computing infrastructure in a molecular biology research laboratory. Molecular biologists use the infrastructure to execute bioinformatics algorithms for analyzing DNA/RNA sequences of millions of bytes. This infrastructure uses peer-to-peer (P2P) technology for distributing tasks to computers within the organization. This implies that the infrastructure requires many users to participate. Thus, one central challenge is to motivate users to contribute despite the fact that only a minor part of them have the actual need of submitting tasks.

To facilitate recruitment, we aimed at increasing the visibility of the P2P grid. We engaged in a user-centered design process with biologists, which resulted in the design of a system named GridOrbit. GridOrbit displays an interactive visualization of the underlying activity in the infrastructure on public screens (see figure 3), and provides users with feedback about their contributions through notifications on their personal computers (see figure 4).

GridOrbit was designed with the hypothesis that increasing awareness of a local resource sharing infrastructure will lead to broader participation and increased contribution. We tested this hypothesis in a one-month deployment of the system. Results show that public display visualizations and personal notifications can be used for creating an awareness of an otherwise invisible infrastructure. We further show that awareness of the activity of an infrastructure through both public and personal displays supports the recruitment of new contributors. Based on this case, we expand the notion of *Infrastructure Awareness* and discuss how it can improve the adoption and thereby the value of voluntary infrastructures.

[1]http://setiathome.ssl.berkeley.edu/ http://folding.stanford.edu/

RELATED WORK

Research on the participative aspects of volunteer computing focuses on how to attract and maintain contributors. Anderson, who leads the SETI@home and BOINC projects [1], proposes the use of visualizations through a screen-saver application to give feedback about participation. He also proposes credits as a way to provide immediate gratification and foster competition through public rankings. Based on the credit system, different participant-oriented websites emerged, with features like user profiles, team formation, and public message boards. These websites serve the double purpose of attracting and maintaining contributors, and also providing "cheap customer support". Likewise, Mowbray reflects on the participative aspects of OurGrid [17], and argues that successful volunteer computing infrastructures need to build a network of human social interactions to function well. BitTorrent's faster downloads [8] or OurGrid's Network of Favors [7] also aim at maintaining existing contributors by ensuring fairness among contributors.

A different strategy to engage volunteers is to offer an "understandable and fun" game experience. With FoldIt, Cooper et al. [10] propose a multiplayer puzzle-like game that leverages human spatial reasoning ability to engage non-experts in making protein structure predictions.

Krebs studied the motivations of volunteers in a large-scale volunteer computing infrastructure [14]. In her study, users ranked pre-defined motivations, and considered intrinsic motives like solidarity, enjoyment, and supporting good causes as the most important ones to start contributing. Krebs proposes a series of improvements for volunteer computing including: (1) overview and statistics, (2) better communications about what is going on behind the scenes, (3) increased awareness, and (4) on-line communities. Similarly, Nov et al. [18] investigated the motivational and non-motivational factors impacting participant contribution. Their findings also suggest that personal motives like enjoyment, reputation and value of the project influence enrollment. However, Nov et al.'s show that these personal motives do not improve long term contribution whereas social factors do. For instance, while contribution levels decrease over time, if contributors are part of a team, their contribution levels tend to remain stable. From this study, the authors define a set of implications for the design of volunteer computing infrastructures including: (1) presenting up-to-date individual contribution levels, (2) increasing the number of channels for communicating rankings, (3) displaying contributions from other team members, and (4) providing secondary activity channels to sustain long term engagements.

Beyond volunteer computing infrastructures, the Ubicomp and CHI communities have also discussed the adoption of technological infrastructures. Here the main goal is to help users understand the infrastructures they interact with. For instance, Lim et al. studied the automatic generation of explanations for improving the intelligibility of technological infrastructures, and thus support their adoption [15]. In order to elicit the behavior of an intelligent home infrastructure, Vermeulen at al. also proposed a system to visualize the inner-workings of the infrastructure and override them if necessary [23]. Preliminary tests suggest that users value the ability to see what is happening inside the system and that this ability helps them to correct their mental models.

In a similar fashion, Chetty et al. designed the HomeWatcher to visualize the invisible usage of home bandwidth, and studied its consequences in traditional households [6]. Their findings suggest that public visualizations of shared infrastructure usage helps users to understand the technology (in terms of current capacity and availability), and, at the social level, reveal the household politics and mediate power relations. By using the HomeWatcher, users experience a better bandwidth use in the household.

Within Ubicomp, recent research have considered how a person's behavior might be changed through ambient displays. For example, Rogers et al. [21] used ambient displays to inform in playful ways about stair and elevator usage, and to induce people to use the stairs. Their findings suggest that people were not aware of their behavioral change while the logged data showed a significant difference. They conclude that ambient displays could induce positive attitudes and behavioral change in situations with multiple choices, and that this influence occurs without conscious thought.

In this paper we focus on the recruitment for a *local* volunteer computing infrastructure by raising awareness of the infrastructure within a specific organization. We argue that awareness technologies are an efficient way to provide personal and group feedback, and that this feedback fosters contribution. Our hypothesis is based on Nov et al.'s recommendations to create local communication channels [18], Krebs' increased awareness recommendation [14], and Rogers et al.'s behavioral effects of awareness displays [21].

THE MINI-GRID CASE

The Mini-Grid is a local volunteer computing infrastructure, harnessing the CPU power of contributors' desktop and laptop computers [2]. The main motivation of the Mini-Grid is to allow biologists to run computationally intensive DNA/-RNA analyses without computing clusters. Running this kind of analyses is an important part of their work, as they simulate experiments and analyze experimental results. Our biologists share their time between planing and analysis work, and lab work. Faster computation allows for more interactive analysis work, more extensive exploration, and more laboratory experiments.

Like other volunteer computing infrastructures, the Mini-Grid's main challenge lies in the recruitment of contributors. At the architectural level, the Mini-Grid is designed as a symmetric system, based on a P2P network. This means that computers can act both as contributors of CPU power and submitters of tasks. Symmetric infrastructures like the Mini-Grid and OurGrid [7] are different from traditional volunteer computing infrastructures. The ability for *any contributor to also submit tasks to the Mini-Grid* changes the assumptions of traditional volunteer computing. For example, there is no shared common goal (like the projects in BOINC), so users can join only to submit tasks and get a free-ride.

At a practical level, the usage of the Mini-Grid by molecular biologists is influenced by two factors: *(1) invisibility*, the Mini-Grid is a software component and biologists have no way to know whether other researchers are contributing to it, or if they need help. This challenge is related to Poole et al.'s argument that invisibility hinders trust and adoption, keeping users from forming correct mental models of infrastructures [19]. *(2) alternative solutions*, the Mini-Grid is a new technology and biologists already developed ways of executing computationally demanding tasks. Biologists run their tasks on high-end computers over long periods of time, or use cloud services such as the ones from NCBI[2].

Using the Mini-Grid

The Mini-Grid works as a plug-in to the CLC Workbench[3] (CLC-WB), a bioinformatics software suite. To run an analysis or simulation on the Mini-Grid, the user selects the algorithm s/he would normally use and, if the algorithm is grid-enabled[4], ticks an extra box for running on Mini-Grid. To participate contributors simply run the CLC Workbench.

Once submitted, the computational task is broken down in smaller instances distributed to executing computers over the network. Upon completion, the task's results are returned to the source. The speed at which a batch of tasks is completed thus depends on the current capacity of the Mini-Grid, i.e. the number of computers currently contributing CPU power. As more computers join the Mini-Grid, the computation time of a batch of tasks decreases.

DESIGN METHODS AND RESEARCH

To better understand how to support the adoption of the Mini-Grid, we engaged in a participatory design process with biologists from the molecular biology department of a large university [13]. The design process comprised a set of five detailed workplace studies (each lasting 2 days, top line in fig. 1) and a series of participatory design workshops (bottom line in fig. 1) over a period of one year .

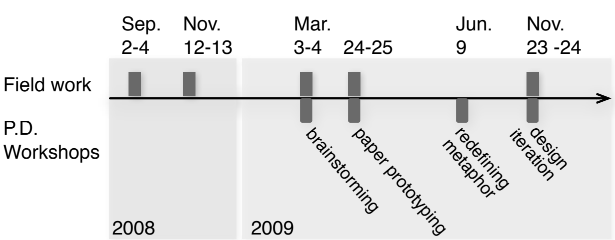

Figure 1. Field work and participatory-design workshops.

We conducted task-centered observations of biology work; place-centered observations of work in the laboratory and in the office; and artifact-centered observations focusing on the use of digital and physical research resources [4]. At this stage we worked with eight biology researchers: three PhD students, two post-docs, two lab technicians and the professor heading the group. We further documented group meetings and casual encounters between researchers. We observed three biologists in detail throughout the day; one executing simultaneous lab experiments, another moving between office space and laboratory, 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 to predict and analyze experimental results.

After the initial fieldwork, we initiated a series of participatory design workshops aimed at co-designing interactive technologies to support adoption of the Mini-Grid. During the workshops we created personas, and discussed different technologies including awareness systems and public displays.

Observations and Implications for Design

Based on our fieldwork and design workshops we summarized the implications for design in three categories: the role of awareness in sharing, the use of public and personal displays, and the Mini-Grid-related awareness information.

Sharing and Awareness

During our place-centered observations we paid particular attention to the sharing of equipment and research materials (tubes, samples, protocols). The resources within the department are a shared infrastructure which researchers rely on to pursue their research, and sharing is common among researchers from different groups. Nevertheless, to benefit from it, they have to maintain an awareness of the resources, their availability, owner, location and conditions of use. For cheap equipment, awareness is achieved through informal systems such as verbal notification or post-it notes. Equipment that is more scarce due to expense or ownership requires more complex awareness systems. Expensive equipment relies on a public booking system to enable distributed use. Genetic material is owned by individual researchers, and therefore its use by others must be negotiated.

The Mini-Grid can be viewed as a similar shared infrastructure, where computing power is a personal resource to be shared among researchers, like tubes or materials. But unlike such physical resources, computing power is not tangible, and hence cannot embody the same sharing practices. In order to benefit from the existing sharing practices, biologists should have an awareness of the Mini-Grid similar to that of the other equipment they share. Awareness technologies are good canditates to provide users with this knowledge.

Public and Personal Displays

Task-centered observations revealed that biologists are mobile when carrying out experiments. Biologists start an experiment in the office by studying the relevant literature, running simulations, and defining protocols. They then move to the laboratory to carry out the empirical work, sometimes running several experiments simultaneously, while using remote computers somewhere else in the building to process experimental data. When done, they return to the office for analyzing results and reporting.

[2]http://www.ncbi.com

[3]http://www.clcbio.com/

[4]Only the PPfold algorithm was grid-enabled during the study.

Laboratories embed different codes of conduct from and office space. The lab is regulated by safety norms requiring users to wear lab coats and gloves while working in a sterile environment. For example, researchers are not supposed to bring in their laptops or to use smart-phones while running experiments. Furthermore, objects in the room cannot be removed before being sterilized and checked for radioactivity.

These observations highlight that technologies for providing awareness to biologists need to take into consideration the specific work context. While in the office, personal displays can provide biologists with awareness of the Mini-Grid. When biologists move out of the office, public displays located in the open spaces like the lab, the cafeteria or the corridors seem more appropriate.

Awareness Cues

A central part of the participatory design process was devoted to identifying the important information cues that users could be interested in with relation to the Mini-Grid. These cues could then be displayed on the personal and public displays. We identified the following cues:

Grid Capacity – The current capacity of the Mini-Grid in terms of connected computers or available computing power. This cue would provide users with an awareness of the availability of the Mini-Grid and the need for user participation.

Grid Activity – The current activities taking place in the Mini-Grid in terms of submitted and executing tasks. This cue would help users reflect upon usage of the Mini-Grid at a given moment in time, and in the recent past.

People and Computers – The people contributing to the Mini-Grid. This cue would allow users to relate the activity on the Mini-Grid to real persons they know, thereby providing an awareness of colleagues' activity.

THE GRIDORBIT AWARENESS SYSTEM

The participatory design process led to the design of GridOrbit, a technology for providing biologists with an awareness of the Mini-Grid. Derived from our fieldwork, GridOrbit is made up by two elements for public and office spaces. The first is a public display for casual and mobile users. The second is a notification system for personal computers for the office workspace.

The GridOrbit public display

The GridOrbit public display is designed to provide the aforementioned awareness cues, and to attract passersby to the Mini-Grid to start contributing. The GridOrbit public display visualizes the locally running Mini-Grid as a cloud of computers (see figure 2). Computers currently connected are shown in color, while off-line computers are gray and progressively fade away after a week of inactivity. The tasks submitted and executed are respectively represented by green and red concentric rings, forming an aura. Their thickness depends on the number of submitted and executed tasks, with the red rings always in the center. The aim of this cloud metaphor is to convey awareness about the capacity (in terms of number of computers) and the current activity level (in terms of tasks being executed at the moment) of the Mini-Grid. It also illustrates the collective nature of the Mini-Grid as made up of these pooled contributions. Moreover, highlighting the activity in terms of auras makes it easier to understand the concepts of submitting and executing tasks.

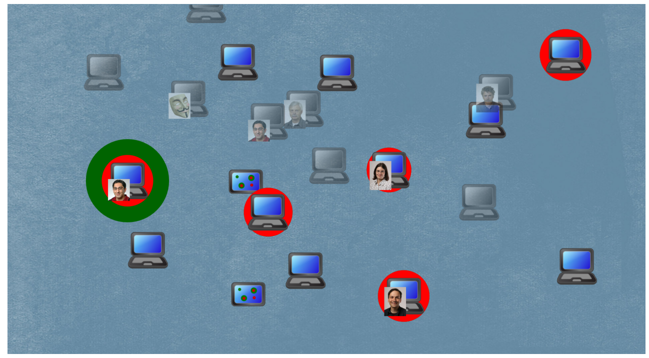

Figure 2. GridOrbit visualizes the Mini-Grid as a cloud of computers.

When users install the Mini-Grid plug-in on their computer, they can assign a custom name to the computer, which is then displayed in GridOrbit. Users can also use the GridOrbit to attach a picture of themselves to their computer. This functionality conveys awareness about people and computers, by making the individual participants visible.

As a way to foster social interaction, GridOrbit supports posting messages to the display. Messages are associated with computers. Messages are known as Tweets and are accumulated in the TweetBox. Messages are broadcasted to all running instances of GridOrbit and can hence be read on every public display. GridOrbit analyzes the content of the messages, extracts recurrent words, and presents them through a TagCloud component, reflecting the interests of the users. By selecting a tag, users filter out the messages in the TweetBox.

For interacting with the GridOrbit public displays, we defined three interaction zones that modulate the level of detail in the information visualized; as passersby get closer to the screen, more detailed information is available. Our interaction zones are inspired by the Hello.Wall implicit interaction zones [20]. The user is in the *Ambient zone* when more than 70cm away from the display. In this zone, GridOrbit shows a cloud of computers and their auras as shown in figure 2. When the user is between 40 and 70cm away from the display (illustrated as the yellow area in figure 3B), the user is in the *Notification zone*. In this zone, GridOrbit presents the TagCloud made from the messages. When the user stands in front of the display (shown as the red area in figure 3B), s/he is in the *Interactive zone*. In this zone, touch interaction is enabled and used for looking up detailed information about each computer, looking at the TagCloud, and reading or writing messages using an on-screen keyboard. Users can also obtain information about joining the Mini-Grid, and leave suggestions on how to improve GridOrbit.

The GridOrbit notification system

To complement GridOrbit public displays, we designed a notification system to provide awareness to contributors working in their offices. The GridOrbit notification system pro-

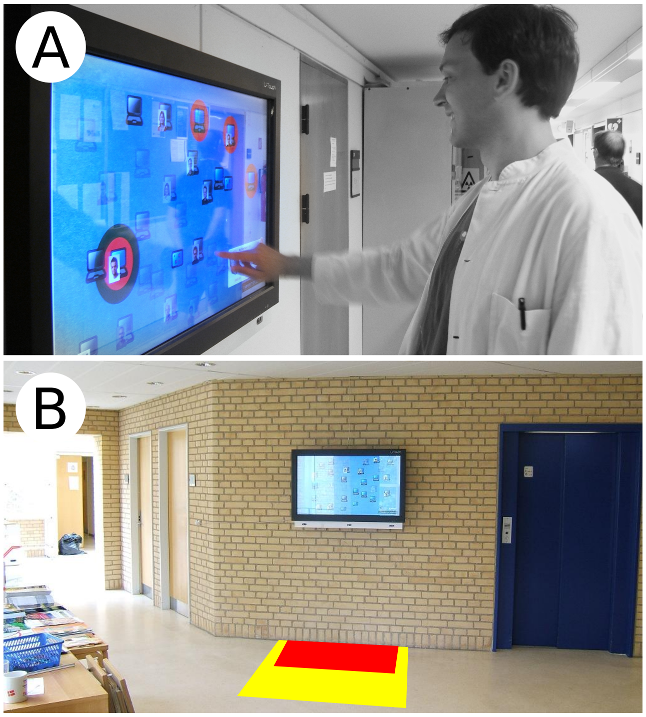

Figure 3. GridOrbit public displays are deployed at two different places on the campus; (A) across the corridor from the cafeteria, and (B) next to the elevator in the local mail area. The interaction zone is marked in red. The notification zone is marked in yellow.

vides more personalized feedback about the contributors' activity on the Mini-Grid using standard desktop notifications. Its goal is twofold. Since the Mini-Grid does not run as an always-on daemon but as a plug-in to a software suite, we wanted to remind users to contribute while in front of their computers. We also wanted to investigate if different notification strategies would lead to different level of contribution.

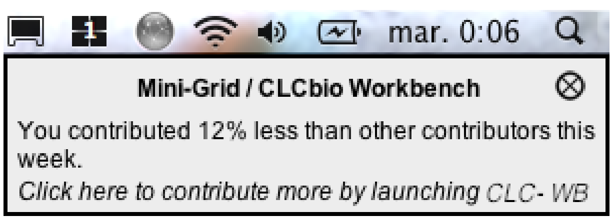

Figure 4. A message displayed by the GridOrbit notification system.

We tested two different types of feedback strategies: (i) personal and (ii) norm-activation [12]. The personal strategy promotes contribution to the Mini-Grid by notifying volunteers of their actual contributions, relying on the volunteers' reflection and call to action to contribute more. The personal strategy displayed messages such as: "*You contributed for less than XX hour(s) to the Mini-Grid this week*". The norm-activation strategy promotes contribution to the Mini-Grid by comparing the volunteer's contribution to the rest of the group, relying on the users willingness to stand out or at least not be left out of the group. The norm-activation strategy displayed messages such us: "*You contributed ZZ% less than other contributors this week*", as shown in figure 4. Messages popped-up every 20 to 120 minutes to test the impact of notification frequency on participants' contribution. The frequency changed on a daily basis and participants could not modify it, but could de-activate the notifications.

We used negative phrasing to encourage more contribution and because we wanted to have similar sentence construction in both strategies. For the personal message strategy, we measured contribution by monitoring the CLC Workbench. We used a Wizard-of-Oz technique to define the values used in the norm activation messages[5], presenting every contribution as less than the mean, even if contribution was superior to the average. In order to still be realistic, particularly toward the people who where dedicating a computer to the Mini-Grid most of the time, the values varied between 5 and 25%. These values were designed to be believable for participants with both high and low levels of contribution.

Implementation

The GridOrbit public displays are 40" touch-enabled screens enhanced with a custom designed ProximityBar for detecting the distance from the display to users standing in front of it. This ProximityBar is composed of three PING))) ultrasonic sensors controlled by an Arduino board and connected to the display via USB. The public displays run a .NET WPF standalone application which monitors the UDP control messages and captures all task related packages sent by the Mini-Grid. Relevant information is extracted from these packages, mapped into the data model, which again is bound to the UI components.

The GridOrbit notification system is implemented in Java and runs as a daemon on Linux, MacOS X and Windows. It uses a modified version of JToaster[6] for generating the notifications. The GridOrbit notification system measures contribution by monitoring when the CLC Workbench was active on a computer, hence we also used it to measure volunteers' contribution to the Mini-Grid throughout the deployment.

DEPLOYMENT

We deployed the Mini-Grid and GridOrbit in the molecular biology department of a large Danish University over a period of one month to investigate two aspects of the awareness technology: (i) What is the impact of awareness technology on the number of volunteers participating in the Mini-Grid and the number of hours they contributed? (ii) How do users relate to, and interact with, such awareness technology?

In order to measure the impact of awareness technologies on the contribution to the Mini-Grid we followed a Quasi-Experimental design [9]. We divided the deployment into three phases as illustrated in Figure 5. The first phase (10 days) served as a baseline with the Mini-Grid deployed without awareness technology; during the second phase (7 days), the GridOrbit public displays were turned on; and during the third phase (13 days), the GridOrbit personal notification sub-system was enabled on users' computers.

[5]Many variables can be used to measure contribution: CPU use, time the computer was available, contribution time effectively used, number of tasks computed, etc. These different metrics require access to low level information we did not have access to, and generate very different contribution results.

[6]http://jtoaster.sourceforge.net/

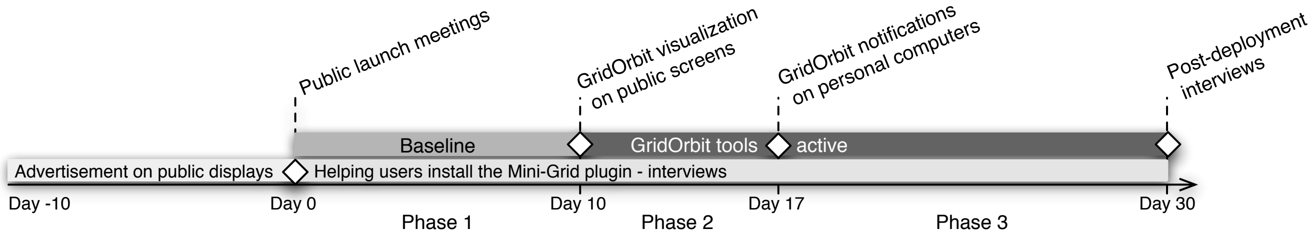

Figure 5. Timeline of our one-month deployment

The deployment started with two public meetings explaining biologists how to use the Mini-Grid. These meetings took place in the two buildings of the department. A GridOrbit public display was deployed in both buildings, which are in 15 minute walking distance of each other. Overall, 28 people attended these meetings and we gathered a list of 10 attendees interested in contributing to the Mini-Grid. After the workshop, we met with 8 of the potential participants in order to help them install the Mini-Grid plug-in and the GridOrbit notification system. Participants who joined later installed the Mini-Grid plug-in without our help, but did not install the GridOrbit notification system.

To decrease novelty effects, we turned on the public displays 10 days before the initial meetings. The displays showed an announcement of these upcoming meetings. During phase 1, the public displays only showed some background information on the Mini-Grid and information on how to join it. During phase 2 and 3, the public displays visualized the Mini-Grid and its activity. During phase 3, the personal notification system displayed messages about users' contribution to the Mini-Grid. Only the 8 initial participants had the personal notifications as it required a computer-specific set-up, and because we wanted to compare contribution with and without GridOrbit from day-0 to day-31.

Across all phases, we captured both quantitative and qualitative data. In order to study contribution to the Mini-Grid, we captured the number of computers connected, and the number of task submissions and executions per computer. To study usage patterns, we captured data on the visits and interactions with the GridOrbit public displays and the users' responses to the notification messages. Finally, we carried out semi-structured interviews with a selected group of participants when the deployment went from one phase to another (illustrated as diamonds in figure 5).

GridOrbit registers a visit when a user stands by the screen for more than 2 seconds, within either the notification or interaction zones. Users observing GridOrbit from a distance greater than the notification zone could not be tracked by the ProximityBar. An interaction is logged when the user touches the screen. GridOrbit does not identify users.

RESULTS AND ANALYSIS

At the end of the 30 day deployment, 35 different computers had contributed to the Mini-Grid at least once. The GridOrbit public displays had received 5022 visits, out of which 592 involved interacting with the screen. Researchers submitted 103 task batches to a total of 7264 individual tasks. Out of the 35 contributing computers, 8 submitted at least 1 task batch to the Mini-Grid. Finally, users posted 14 messages through the messaging functionality. Below we present results related to both the impact of GridOrbit on adoption, as well as insight into how users were using it.

Adoption Impact

We label the number of computers that installed the Mini-Grid plug-in and contributed at least once as *potential capacity*. Similarly, we label the average number of computers connected to the Mini-Grid at any given time *actual capacity*. Figure 6 shows the potential and actual capacity of the Mini-Grid. The figure shows that both potential and actual capacity increased after introducing the awareness technologies. During the first (baseline) phase, the potential capacity reached 19 computers. After introducing the awareness technologies, 16 new computers joined the Mini-Grid, increasing the potential capacity to 35. Moreover, the actual capacity was 5.33 ($\sigma = 2.66$) computers in the baseline phase. After introducing the awareness technologies, the actual capacity increased to 8.02 ($\sigma = 3.56$) computers: a 2.69 change ($p < 0.001$). This data represents a 75% increase in potential capacity, and a 51% increase in actual capacity.

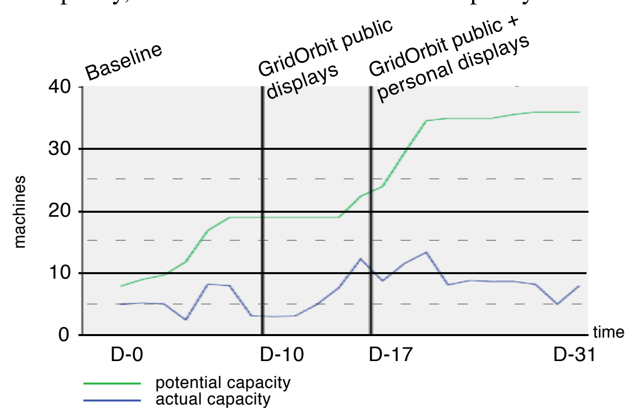

Figure 6. Potential and actual capacity as seen during the deployment.

Further analysis focusing on users logged on the personal computer shows no significant difference in actual capacity between phase 2 and phase 3. While we can observe an increase from a mean contribution time of 10h57 to 14h22 per day, the standard deviation is very high in both cases (~8h20) because of volunteers' contribution patterns: either a few hours a day or the entire 24 hours. We obtained similar results when analyzing the impact of the GridOrbit notification system's strategies on contributions. Analysis shows no significant difference in contributions resulting from the strategies.

Based on interviews and logs from personal computers with the notification system installed, we are able to analyze the contribution patterns on a per-computer basis. This analysis revealed 5 distinct ways of contributing to the Mini-Grid:

Bootstrapping Computers – Biologists involved in the project dedicated two off-the-shelf desktop computers running 24/7, to guarantee a minimal capacity.

Dedicated Secondary Computers – 3 biologists with desktop computers rarely used in their office decided to dedicate them to the Mini-Grid. Figure 7 shows an always-on mode of contribution. Interestingly, Participant 2 switched from having his computer on only during working hours to an always-on mode when GridOrbit started to display messages about the computer's contribution.

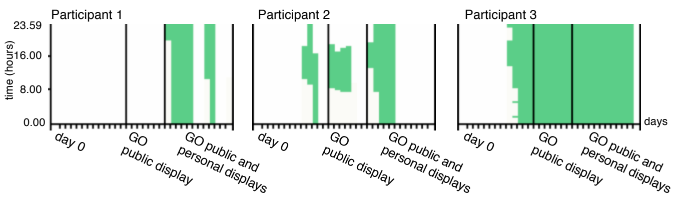

Figure 7. Contribution pattern of dedicated secondary computers, *(white is inactive; green is sharing resources.)*

Work Computers – 3 biologists installed the Mini-Grid plug-in on their office desktop computers, which they used for everyday work. Figure 8 shows a hybrid pattern; always-on when they were working in the lab and hence not using the office desktop computer (green), mixed with periods off-the-grid when working intensively in their offices (red).

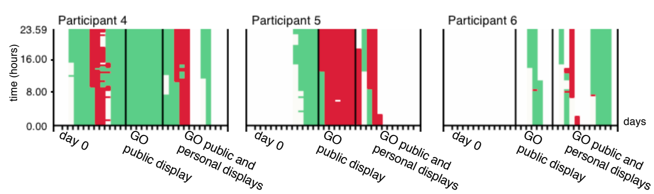

Figure 8. Contribution pattern of work computers *(white is inactive; green is sharing resources; red is active not sharing resources.)*

Intermittent use – 2 biologists installed the Mini-Grid plug-in on their laptop, i.e., main computer. Figure 9 shows a very intermittent contribution to the grid.

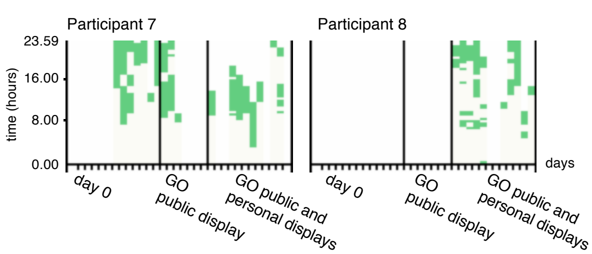

Figure 9. Contribution pattern of laptops *(white: inactive; green: sharing resources.)*

Shared Public Computers – Finally, biologists deployed the Mini-Grid on public computers, in the department library and on the desktop computers available in the different labs.

During pre- and post-deployment interviews, we identified a diverse set of motivations for contributing to the Mini-Grid. For the few researchers using the Mini-Grid-enabled algorithms, getting as much computer power as possible was the primary drive. Combined with a curiosity about the power of the Mini-Grid and testing the limits of system, this led to configure several public computers. Another set of participants did not use Mini-Grid-enabled algorithms, but was interested in the potential of the Mini-Grid and thus wanted to participate and experiment with it. Finally, some participants only joined to contribute to the Mini-Grid, not benefit from it in terms of faster calculations. Being recognized as a contributor and being involved in an experimental technological project was the main drive. This recognition was reinforced when participants started to associate their own portrait to the computers displayed in GridOrbit. This feature helped more anonymous people like lab technicians to be visible as part of the research activities and projects of the lab.

Using GridOrbit

We looked into how GridOrbit was used in everyday work, particularly in terms of visits to the public displays. We examined the visitor data in relation to the interaction zones, and type of interaction. We removed outliers: visits of less than 1.5 seconds (twice the time to walk past the screen) and more than 70 seconds. In relation to the interaction zones, 39.57% of the visitors walked into the interaction zone, and 60.43% stayed in the outer notification zone. In relation to the type of interactions, only 9.87% of the users interacted with GridOrbit via touch, with an average user session lasting 17.5 seconds ($\sigma = 1.5$). Crossing these two datasets shows that 29.7% of the users who walk into the interaction zone (very close to the display) do not interact.

This data reveals three different visitor behaviors, which we label (i) *aware*, (ii) *curious*, and (iii) *explorer*. There seems to be a 60–30–10 distribution between these three behaviors in front of the public display.

We analyzed the effects of the dynamic elements of GridOrbit's visualization on the visitors' behavior. The visualization of GridOrbit changes in response to task executions, representing activity; the more tasks being executed on the Mini-Grid, the bigger the green and red auras around the involved computers are. The visualization also changes in response to potential capacity in terms of the total number of computers contributing to the Mini-Grid; the more computers contributing to the Mini-Grid, the more computers are shown in GridOrbit. The potential drawback of this visualization is that the display may be cluttered when capacity increases.

Table 1 shows the relationship (r) between the number of visitors per hour, and the visualization of activity and capacity of the Mini-Grid. The top row in the table shows that, overall, the number of visits to GridOrbit correlates positively to the activity level (auras), but negatively to the capacity (computers). The following rows show that both correlations hold true for the three visitor behaviors identified previously. The correlations are significant at the levels shown in the table, meaning both activity and capacity effectively account for variations in the number of visitors.

Visitors	Activity	Capacity
Total	$r = 0.212$ (p<0.002)	$r =$ -0.116 (p<0.1)
Aware	$r = 0.186$ (p<0.01)	$r =$ -0.135 (p<0.05)
Curious	$r = 0.202$ (p<0.005)	$r =$ -0.181 (p<0.01)
Explorer	$r = 0.133$ (p<0.06)	$r =$ -0.114 (p<0.1)

Table 1. Correlational data between visitors of GridOrbit and changes in visualizations of activity level and capacity.

We can conclude from this analysis that activity visualization seems to have attracted more visitors, whereas visualizing a higher capacity seems to have decreased the number of visitors. Thus, adopting the Mini-Grid can lead to awareness stagnation; the more users adopt the Mini-Grid, the more cluttered GridOrbit will look, which decreases the number of visitors. The efficiency of awareness technologies designed to support the recruiting of contributors may hence decrease as they start to look cluttered. Therefore, the design of infrastructure awareness visualizations should take into account issues of scaling.

DISCUSSION

To explore the participative aspect of local volunteer computing infrastructures we designed GridOrbit as a means to increase users' awareness of the Mini-Grid's activity. Our results suggest that awareness technologies increases the use of local volunteer computing infrastructures. In this section we reflect on three issues: first, the relation of our results to other studies of volunteer computing infrastructures. Second, the impact of infrastructure design on the recruitment of contributors. And third, how our experience can be expanded to support other cases of infrastructure adoption.

Recruiting and Awareness Technologies

Volunteer computing infrastructures like BOINC and OurGrid use rewards and community websites to attract and maintain contributors. We depart from these approaches by proposing awareness technologies as a way to generate genuine interest and provide users with feedback on their contributions. Our results show that, as suggested by Nov et al.'s [18] study, providing feedback on contribution and increasing the number of feedback channels are effective ways to recruit volunteers. These results validate Rogers et al's [21] *ambient influence* in a situation where a socially desired behavior is induced through ambient technologies.

While other works highlight the importance of supporting communication to sustain contribution, passersby and Mini-Grid contributors rarely used the messaging functionality on the GridOrbit public display (14 messages total). This behavior could be explained by the relatively small number of potential contributors (approximately 300 people), and the social dynamics that such a small co-located group affords. This suggests that, during the initial phases of deployment, local infrastructures like the Mini-Grid can rely on the existing social dynamics and communication channels to recruit contributors in small and co-located groups of users. In this perspective, one of GridOrbit's effects was to sustain word-of-mouth as an informal communication channel. Awareness technologies should focus on fewer and simpler features, providing awareness through ambient information visualizations.

Our conclusions on recruitment are more relevant for local groups, rather than large and distributed groups in which creating and supporting social dynamics can prove valuable as they might not already exist. Furthermore, our month-long study was too short to provide insights on the effects of long term contributions. It might be, as documented by others [18], that contribution decreases over time, and that secondary engagement channels can modulate the decrease. Another area to explore is the exposure of historic contribution/usage data, and estimated future demand for capacity.

Infrastructure Design and Participation

The Mini-Grid's design as a P2P infrastructure influenced adoption and contribution. By design, each Mini-Grid-enabled computer can both execute and submit tasks to the Mini-Grid. This feature prompted some researchers to join the Mini-Grid not only to help but also to explore how they could benefit from it. Some researchers contributed to determine whether they could improve their own algorithms with the use of the Mini-Grid.

More broadly, our study adds to the discussion [11] on how the architecture of infrastructures shapes the way people will participate in them. It is important to focus not only on how to design tools or strategies to recruit contributors, but also on how the infrastructure itself could be improved to better sustain the contribution of its participants. For instance, during our participatory workshops we identified how the Mini-Grid could be improved to be more engaging and more efficient. One improvement we discussed was to broadcast the results of a task execution. This would both increase interest and willingness to participate, and would also be a way to get results quickly for someone interested in the results of a complex task already executed.

Infrastructure Awareness

Current approaches for recruiting contributors, or solving the problems arising from the infrastructures' invisibility, aim at visualizing the infrastructure itself and providing detailed feedback about users' contributions. These approaches assume that users put all of their attention on the feedback system. Providing detailed information in the center of users' attention is ideal for learning, troubleshooting, triggering reflection, and accomplishing difficult tasks. However, it also requires effort and attention from contributors.

Our work with the Mini-Grid, GridOrbit's public displays and personal notifications explore a calmer approach (in Weiser's sense of calm technology). In this approach, which we call Infrastructure Awareness, systems provide feedback while staying on the periphery of users' attention.

Infrastructure Awareness is a feedback mechanism on properties of technological infrastructures provided in the periphery of user's attention. Feedback is conveyed through representations of the state or changes of such properties. Infrastructure Awareness Systems aim to convey this information on the periphery of users' attention by means of ambient technologies. The infrastructure awareness system can either be embedded in the infrastructure itself, or be an entirely separate system. The latter case is relevant for infrastructures that cannot themselves provide the awareness, or legacy systems that did not consider doing so.

We theoretically ground Infrastructure Awareness on Benford and Fahlén's spatial model of awareness [3]. This model uses the notions of 'nimbus' and 'focus': nimbus is what an object projects about itself, i.e. the things that can be known about an object; focus is what an user is interested in. Benford and Fahlén argue that awareness takes place when the user's focus meets the object's nimbus i.e. when the user is interested and has access to what the object projects about itself.

We modeled the problem of the invisibility of infrastructures in terms of this awareness model, as presented in figure 10A. In our model there are two obstacles to awareness: first, the user's focus (U in figure 10) does not meet with the infrastructure's nimbus (I in figure 10); and second, the infrastructure's invisibility (shown in dashed lines in figure 10) represents the users' incapacity to see the it. An Infrastructure Awareness System (IA) seeks to overcome these difficulties by extracting data from the infrastructure's nimbus and translating it into information that the user is interested in and can acquire by means of technologies like ambient displays. This is shown in figure 10B.

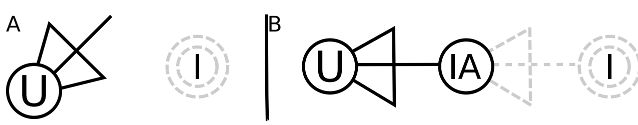

Figure 10. Infrastructure Awareness' awareness model.
(The circle represents the object's nimbus. The cone represents the object's focus and its orientation.)

In creating Infrastructure Awareness Systems, designers should not only elicit users' interests (their focus), but also which properties of the infrastructure can be acquired. Infrastructure Awareness can also be used as tool to frame participatory design activities as we describe in [13], and applied for the design of GridOrbit. Furthermore, this deployment showed how relevant Infrastructure Awareness can be for everyday and mid-term engagement with infrastructures.

Infrastructure awareness systems should support the different visitor behaviors (aware, curious, and explorer) that we observed in our deployment. While further deployments should look at the proportions with different visitors and in different situations, we speculate similar proportions can be found. To support these visitors, infrastructure awareness systems can provide information at different granularities. Peripheral feedback, achieved with representations at high levels of abstraction, can support adoption and participation. More attention-demanding feedback, achieved with representations at lower levels of abstraction (higher fidelity), can support in-depth interaction. Attention-demanding feedback could be used to support learning, troubleshooting, triggering reflection, or accomplishing difficult tasks.

CONCLUSION

We presented GridOrbit, an infrastructure awareness system designed to foster participation to the Mini-Grid, a local volunteer computing infrastructure. During our deployment, we observed that the GridOrbit public displays and personal desktop notifications led to an increase in both the total number of computers that joined the grid, and the average number of computers available at any point in time (respectively potential and actual capacity).

We identified four main motivations for biologists to join the Mini-Grid: getting analysis results faster, exploring the potential benefit of the infrastructure, being recognized as a contributor, and helping their colleagues. We also identified five contribution patterns: bootstrapping computers and secondary work computers with a 24/7 contribution; primary work computers and shared public computers with an interrupted 24/7 contribution pattern when the computers were used for other tasks; and laptops with a very intermittent contribution pattern.

Analyzing how passersby engaged with the GridOrbit public displays, we distinguished three types of visitors: aware, curious, and explorer. By looking further into visitors' interaction with the screens, we identified that visualizing activity rather than the potential capacity increased users' engagement with GridOrbit .

Reflecting on the design and use of GridOrbit, we discussed how the Mini-Grid and its awareness technologies depart from other volunteer infrastructures and their recruitment models. Because the Mini-Grid is P2P (hence symmetric) and local, it led us to foster participation through ambient technologies rather than competitions or games. This suggests that HCI researchers should not only focus on designing systems for engaging contributors, but should also consider the architecture of infrastructures as a means to shape how infrastructures are used. Finally, to generalize our work, we proposed the notion of infrastructure awareness, and defined an infrastructure awareness model extending Benford and Fahlén's spatial model of awareness.

ACKNOWLEDGEMENT

We thank N. Venkataraman, E. Andersen, Z. Sukosd and all the participants from Aarhus University. This research has been funded by the Danish Agency for Science, Technology, and Innovation under the project "PC Mini-Grids for Prediction of Viral RNA Structure and Evolution", #09-061856.

REFERENCES

1. D. P. Anderson. BOINC: A system for public-resource computing and storage. In *GRID '04: Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*, pages 4–10, Washington, DC, USA, 2004. IEEE Computer Society.

2. J. E. Bardram and N. Venkataraman. The mini-grid framework: Application programming support for ad-hoc, peer-to-peer volunteer grids. In *Proceeding of the GPC 2010 Conference.*, May 2010.

3. S. Benford and L. Fahlén. A spatial model of interaction in large virtual environments. In *ECSCW'93: Proceedings of the third European Conference on Computer-Supported Cooperative Work*, pages 109–124, Norwell, MA, USA, 1993. Kluwer Academic Publishers.

4. J. Blomberg, M. Burrell, and G. Guest. The human-computer interaction handbook. chapter An ethnographic approach to design, pages 964–986. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 2003.

5. M. Chalmers and A. Galani. Seamful interweaving: heterogeneity in the theory and design of interactive systems. In *DIS '04: Proceedings of the 2004 conference on Designing interactive systems*, pages 243–252, New York, NY, USA, 2004. ACM Press.

6. M. Chetty, R. Banks, R. Harper, T. Regan, A. Sellen, C. Gkantsidis, T. Karagiannis, and P. Key. Who's hogging the bandwidth?: the consequences of revealing the invisible in the home. In *CHI '10: Proceedings of the 28th international conference on Human factors in computing systems*, pages 659–668, New York, NY, USA, 2010. ACM.

7. W. Cirne, F. Brasileiro, N. Andrade, L. B. Costa, A. Andrade, R. Novaes, and M. Mowbray. Labs of the world, unite!!! *Journal of Grid Computing*, 4(3), 2006.

8. B. Cohen. Incentives build robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer systems*, volume 6, pages 68–72, 2003.

9. T. D. Cook and D. T. Campbell. *Quasi-Experimentation: Design and Analysis Issues for Field Settings*. Houghton Mifflin Company, 1979.

10. S. Cooper, A. Treuille, J. Barbero, A. Leaver-Fay, K. Tuite, F. Khatib, A. C. Snyder, M. Beenen, D. Salesin, D. Baker, and Z. Popović. The challenge of designing scientific discovery games. In *FDG '10: Proceedings of the Fifth International Conference on the Foundations of Digital Games*, pages 40–47, New York, NY, USA, 2010. ACM.

11. W. K. Edwards, M. W. Newman, and E. S. Poole. The infrastructure problem in HCI. In *CHI '10: Proceedings of the 28th international conference on Human factors in computing systems*, pages 423–432, New York, NY, USA, 2010. ACM.

12. J. Froehlich, L. Findlater, and J. Landay. The design of eco-feedback technology. In *CHI '10: Proceedings of the 28th international conference on Human factors in computing systems*, pages 1999–2008, New York, NY, USA, 2010. ACM.

13. J. D. Hincapié-Ramos, A. Tabard, and J. Bardram. Designing for the invisible: user-centered design of infrastructure awareness systems. In *DIS'10: Proceedings of the 8th ACM Conference on Designing Interactive Systems*, pages 302–305, New York, NY, USA, 2010. ACM.

14. V. Krebs. Motivations of cybervolunteers in an applied distributed computing environment: MalariaControl.net as an example. *First Monday*, 15(2-1), 2010.

15. B. Y. Lim, A. K. Dey, and D. Avrahami. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 2119–2128, New York, NY, USA, 2009. ACM.

16. S. D. Mainwaring, M. F. Chang, and K. Anderson. Infrastructures and their discontents: Implications for ubicomp. In *6th International Conference in Ubiquitous Computing. Ubicomp 2004*, pages 418–432. September 2004.

17. M. Mowbray. How web community organisation can help grid computing. *Int. J. Web Based Communities*, 3(1):44–54, 2007.

18. O. Nov, D. Anderson, and O. Arazy. Volunteer computing: a model of the factors determining contribution to community-based scientific research. In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 741–750, New York, NY, USA, 2010. ACM.

19. E. S. Poole, C. A. Le Dantec, J. R. Eagan, and W. K. Edwards. Reflecting on the invisible: understanding end-user perceptions of ubiquitous computing. In *Ubicomp '08: Proceedings of the 10th international conference on Ubiquitous computing*, pages 192–201, New York, NY, USA, 2008. ACM.

20. 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.

21. Y. Rogers, W. R. Hazlewood, P. Marshall, N. Dalton, and S. Hertrich. Ambient influence: can twinkly lights lure and abstract representations trigger behavioral change? In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, Ubicomp '10, pages 261–270, New York, NY, USA, 2010. ACM.

22. S. L. Star. The ethnography of infrastructure. *American Behavioral Scientist*, 43(3):377–391, November 1999.

23. J. Vermeulen, J. Slenders, K. Luyten, and K. Coninx. I bet you look good on the wall: Making the invisible computer visible. In *AmI '09: Proceedings of the European Conference on Ambient Intelligence*, pages 196–205, Berlin, Heidelberg, 2009. Springer-Verlag.