Pervasive Help @ Home: Connecting People Who Connect Devices

Mark W. Newman and Mark S. Ackerman

School of Information and Electrical Engineering and Computer Sciences Department
University of Michigan
Ann Arbor, MI 48109 USA
{mwnewman,ackerm}@umich.edu

Abstract. We propose a system to allow users of domestic pervasive computing technology to receive indirect help regarding the composition and control of their systems of devices from other users with similar configurations. Our work extends previous work on end-user composition and collaborative assistance and is based around the notion of a shared repository of compositions (called the Composition Community) from which end-users can receive information about typical and best-practice service compositions that are contributed by other community members. A Composition Monitor automatically determines the composition of each user's network, allowing contributions to the repository to be made without explicit effort. A Composition Visualizer allows users to view and compare different compositions in order to choose the most appropriate one for their needs. The goal of our work is to bring together the knowledge and collective problem solving skills of diverse users who are currently isolated from each other in their separate homes.

Keywords: pervasive computing, home computing, end-user composition, collaborative help

1 Introduction

Each home is a technological island. New devices and services enter each home and populate niches that is particular to the specific not only to their specific functions but to the other technologies already present, with which they must coexist, interoperate, or compete. The home is also an island in the sense that each home's inhabitants are isolated from other home's inhabitants, even as they each struggle to comprehend, configure, and control the various technologies with which they each live.

The goal of our project is to allow users to receive help relevant to the control and configuration of pervasive computing technology in the home. A centerpiece of our approach is to allow users to benefit from the work of other users with similar environment compositions. If we can find a way to harness the experiences of the vast range of users of ubiquitous environments, the ability to find a solution to any problem that arises becomes increasingly tractable. Using this approach, a user wishing to change the functionality of their environment—whether to add new capabilities or to recover from a breakdown by restoring it to a previously attained level of functioning—can take advantage of information about the configurations of large numbers of others users' environments that have been more or less successful at providing the desired functionality. This amounts to a new form of collective help that can be of great value, perhaps even essential, to the ability of users with a wide range of skill, motivation, and patience to manage and control their individual ubiquitous computing environments.

The purpose of this position paper is to argue the need for collaborative help facilities in pervasive environments, briefly survey the related work, and outline requirements for a system that would provide help to users who will increasingly need to configure, manage, and control pervasive computing systems in their homes.

2 Related Work

The primary goal of this research is to support user control over pervasive computing environments by providing help with composition and maintenance. There are two key areas of related work that must be considered: attempts to support user control and composition in ubiquitous computing environments, and techniques to provide help to users of interactive systems.

2.1 End-User Composition in Ubiquitous Computing Environments

A central focus of our work is supporting end-user composition. By "composition" we mean the work that must be done in order to configure the constituent components of a ubiquitous computing environment to ensure that commands will function properly, information will be delivered properly, and automated behaviors will execute properly.

Supporting the ability of users to govern the automated behavior of ubiquitous computing environments has received some attention in prior literature. The Jigsaw Editor [13], CAMP [16], and iCAP [5] all allow end-users to create rules that associate the behaviors of different components within a ubicomp environment. In each of these systems, a major focus was on exploring the value of novel interaction techniques (respectively: assembling jigsaw puzzle pieces, composing magnetic poetry, and informally sketching *if-then* action rules) and to learn more about what users would like to be able to do with an environment supporting context-aware and/or sensor-driven interaction. The projects did not address the maintenance of the system or its behavior over time, facilities to allow users to receive help with the composition task, or the ability of users to understand current and possible future states of an environment when the environment configuration(s) depicted were not recently created by themselves.

Newman's previous work on end-user composition was concerned with allowing users to access command and composition capabilities from a single point of control. Additionally, this work explored the creation, modification, and invocation of reusable templates to allow end-users to create (and potentially share) custom configurations of devices and services for particular interactions [9-11]. While this work provides a foundation for the current project, it did not address the need of users to receive help when creating or modifying compositions, nor did it provide facilities for visualizing and understanding complex pervasive environments.

2.2 Help and User Assistance

Help has been an HCI research problem for over 20 years. During the 1980s, help was largely seen as answering immediate questions about applications and occasionally operating systems. Help was largely documentation-centric, and initial research focused on how to design and deliver help documentation so that it could be used more effectively (e.g., [7, 12]). In the 1990s, with the advent of networks and CSCW, emphasis began to shift away from formal documentation, prepared in advance, to ad-hoc queries to other people and research began to explore the possibilities of using computer-mediated communication (CMC) systems such as email and bulletin boards for help (e.g., [6, 14]). Ackerman's Answer Garden series ([1, 3]) was one such attempt to combine hypermedia, information, and people into a help network.

Recently, help as a research area has become somewhat moribund. The standard solution for users seems to be to google for answers to problems or to ask on a myriad assortment of webboards. For many problems, particularly time-sensitive and complex problems, these solutions are not optimal.

While some research continues (e.g., investigations of Web-based online Q&A communities [19] and wikis [15], tools to support authoring and distillation of informal communication [8], and studies of collaborative help system [2, 17], there have been as yet no attempts to consider help as a context-aware collaborative problem in pervasive environments.

A goal of our project is to “restart” user assistance within HCI and within the pervasive computing research community. Little, if anything, has been done recently to consider how help might be brought into the next generation of computational environments. Without this functionality, users will be left to their own, facing problems that will become increasingly difficult and complex. It is, in our opinion, unlikely that users will be able to adopt open ubiquitous computing environments without some assistance.

3 Collaborative Help in the Pervasive Home

Our proposed system (shown in Figure 1) consists of three interrelated elements:

1. Facilities (e.g., user interfaces) to allow users to understand compositions, especially the differences between otherwise similar compositions and the areas where deviations from previous states have occurred.

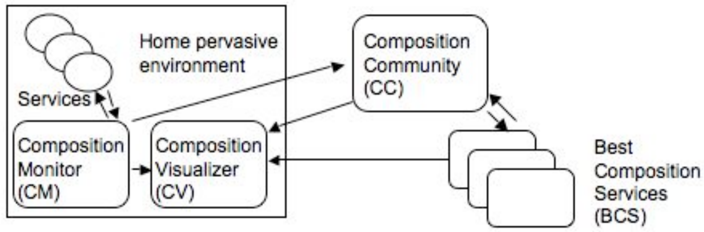

Figure 1: An overview of a proposed system for delivering collaborative help to users of pervasive computing environments. We plan to develop facilities for visualizing (CV), provisioning both community- and expert-generated help (CC and BCS), and automatically detecting and sharing compositions (CM).

2. Shared repositories that allow users to benefit from the collective work of other users who have successfully configured their environments, along with techniques for identifying compositions in the repository that are similar to the composition for which help is being requested. shows the technical components that will implement (1) and (2).
3. Mechanisms for automatically collecting information about users' compositions and depositing the information in a shared repository.

3.1 Visualizing Compositions

For users to identify opportunities for improvement or repair and make informed decisions about how to operationalize changes to their environments, they must be able to rapidly view the configuration and status of their environments. In addition, they must be provided with information about how their configurations differ from other possible configurations that may be "better" in some way (e.g., featuring additional functionality, operating more smoothly). Thus, the first component of our system is the Composition Visualizer (CV), which allows end-users to reason about their configurations and compare it against others'.

In order to illustrate the usefulness of CV, we present a pair of scenarios that show how such a tool could be used for diagnosing an emergent problem and for modifying an environment to support new functionality.

Scenario 1: A Broken Thermostat

On Thanksgiving morning, Joe realizes his living room's context-dependent thermostat no longer seems to work. He would really like to get the system going again before his Thanksgiving guests arrive to a 40 degree[1] living room. Joe doesn't really know how to debug his house—he just wants it to work. Using CV, Joe immediately sees that his "People tracker" service (a service that detects which rooms are occupied by people) has failed to communicate with the thermostat in the past 48 hours—a deviation from the previous several months when 100s of messages per day

[1] The reader may assume Celsius or Fahrenheit—the story works equally well either way.

passed between them. By exploring with CV, Joe is able to determine that there is a discrepancy between the thermostat's recently upgraded firmware and the version of the People Tracker service he is running. Using this information, Joe is able to get both services in sync and repair the problem (see Figure 2).

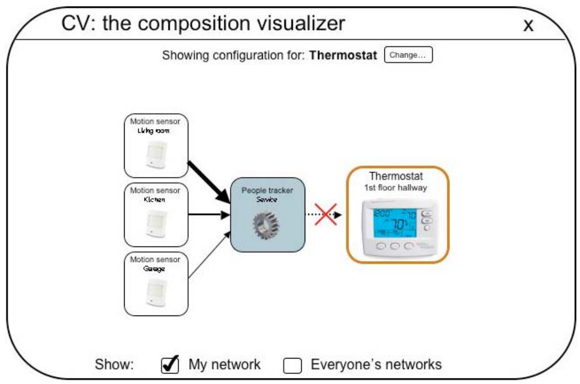

(a)

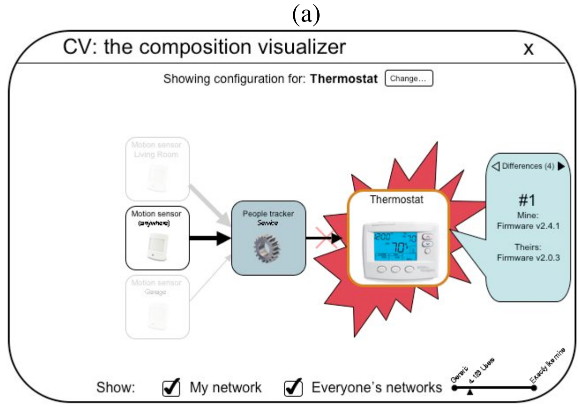

(b)

Figure 2: In this illustration of a possible sequence of interactions with a composition visualizer (CV), a user ("Joe") views a composition involving his thermostat, and notices that a communication failure has been detected between his "People tracker" service and the thermostat (a). He then elects to see relevant composition information from "Everyone's networks" and explores the discrepancy between the state of his thermostat and the most commonly found state for similar thermostats (b). By exploring further, he is able to see that the problem lies in the fact the he is still running an older version of the person tracker service. Based on this information, he decides to upgrade the person tracker service to the latest version.

Scenario 2: Augmenting a Context-Aware Alarm Clock
In a conversation at the grocery store, Margaret learned from her neighbor Tom that he had figured out how to automatically reset his children's alarm clocks on snow days. She'd like to know how to add that functionality to her own home. Using CV, she is able to explore a variety of compositions that involve her son Max's alarm clock. She is intrigued by several suggestions for associating traffic conditions and airline flight status information with wake-up time, but passes them by and eventually finds a composition (perhaps Tom's) that uses information from the local Snow Day Alert web service to disable the alarm clock when the schools are closed. She happily adds the new composition to her own environment.

Goals for a Composition Visualizer
As can be seen from these scenarios, an effective visualization tool that allows users to understand task-relevant relationships and details can enable users to create and modify compositions that would otherwise be quite difficult to work with. Moreover, obtaining help from document-centric sources or web-based discussion boards would be unlikely to provide an actionable solution for either of these scenarios.

The key issues in designing and evaluating CV will be determining what information about compositions (their own and others') is most important to present to users, developing effective ways of representing the most relevant information, and providing users with the necessary controls to efficiently interact with and refine the visualization to obtain the information they need. As indicated by the scenarios described above, we believe that anomalies (e.g., divergence from previous known "good" states or divergence from commonly-used patterns) are a particularly significant type of information to display. For example, it is valuable for a user to know when the current state of a composition has diverged from a previously known "good" state. Another key challenge will be to find effective ways to encapsulate or "black box" portions of the network in particular situations so that the visualizations do not become overwhelmingly complex while still providing users with the information they need in order to troubleshoot and/or augment their environments.

3.2 Providing for User Assistance in Composition

Given that the goal of CV is to allow users to be made aware of anomalous configurations found within their environments and to explore new functional capabilities that from which they could benefit, we turn now to the question of where such information could come from. The Composition Community (CC) provides storage and retrieval mechanisms for meta-data about services, and more importantly, about compositions. It also allows power users or others to annotate services and compositions.

A service in each home (the "Composition Monitor" described below) will upload the current configuration of that home's environment in a privacy-preserving manner to the CC. Once information about a number of compositions is available, these data can be used to determine whether a user's configuration or desired configuration is feasible—i.e., whether it is used by others and if others have reported breakdowns or problems.

Even with such a facility, a typical user may be unable or uninterested in finding new services and compositions if doing so requires going through long lists of potential options. They will be interested, however, in finding new functionality that will best serve them. While we plan to exploit social filtering and recommendation mechanisms to allow popular compositions and services to bubble to the top, we also see great value in allowing “power users” with greater skill, patience, and motivation to hand-craft compositions that others can use. Thus we propose the Best Composition Service (BCS) to allow potential power users to create "best practices" or "cool functionality" compositions by stringing together services and then publishing these for others to incorporate into their own systems. Any user can then go to a BCS, obtain a set of compositions that provide a desired functionality, and then using CC, check the composition against her current environment. Again, when everyone's configurations are available, some users can add additional functionality with relatively little cost.

3.3 Monitoring and Sharing Compositions

In order to automatically monitor and construct composition models of users’ environments, we depend on emerging frameworks for device and service interoperation that are based on networked communication (e.g., UPnP [18] or DPWS [4]). Because of the nature of communication among networked services, we are able to construct a model of the interrelationships among services without requiring modification of the services themselves. Messages on the network are used to perform communication among services such as invoking operations, subscribing to event notifications, and sending event notifications. Thus, we plan to implement a composition monitor that captures all network traffic and analyzes it to build a graph of service communication. By analyzing the temporal relationships among communications, we can subdivide the graph into sets of services whose communication patterns are functionally related.

Once sets of functional compositions have been determined, they are uploaded to a shared composition repository. The repository will be located as part of the Composition Community server, and all Composition Visualizers will be able to access such data.

4 Conclusion

As pervasive technology makes its way into the home, each home’s inhabitants will be forced to contend with configuration and maintenance challenges of ever-increasing complexity. Helping users overcome these challenges is crucial to the ultimate success of pervasive computing in the home.

In this position paper, we have laid out our proposal for a system to deliver help to otherwise isolated users as they attempt to configure and control their technological environments. Our system builds upon prior research in end-user composition and collaborative help by providing mechanisms for visualizing compositions (including one's own and those of others), as well as accessing and contributing to shared repositories of composition information. We believe such a system is novel and will represent a large step forward in helping users solve difficult composition problems in pervasive computing environments.

References

1. Ackerman, M.S. Augmenting Organizational Memory: A Field Study of Answer Garden. ACM Transactions on Information Systems, 16(3). pp. 203-224, (1998).
2. Ackerman, M.S. and Haverson, C. Organizational Memory: Processes, Boundary Objects, and Trajectories. Computer Supported Cooperative Work; The Journal of Collaborative Computing, 13(2). pp. 155-189, (2004).
3. Ackerman, M.S. and McDonald, D.W. Answer Garden 2: Merging Orgnaziational Memory with Collective Help. *In Proc. ACM Conference on Computer-Supported Cooperative Work (CSCW '96)*. pp. 97-105. Boston, MA, USA (1996).
4. Chan, S., et al. Devices Profile for Web Services Specification. http://specs.xmlsoap.org/ws/2006/02/devprof/devicesprofile.pdf.
5. Dey, A.K., Sohn, T., Streng, S. and Kodama, J. iCAP: Interactive Prototyping of Context-Aware Applications. *In Proc. the 4th International Conference on Pervasive Computing (Pervasive '06)*. Dublin, Ireland (2006).
6. Finholt, T.A. Accessing Organizational Expertise through Computer-Mediated Communication (Ph.D. Dissertation). Carnegie-Mellon University (1990).
7. Marchioni, G. and Schneiderman, B. Finding Facts vs. Browsing Knowledge in Hypertext Systems. International Journal of Man-Machine Studies, 21(1). pp. 70-79, (1988).
8. Nam, K.K. and Ackerman, M.S. Arkose: Reusing Informal Information in Online Communities. *In Proc. ACM Group Conference (Group 2007)*. pp. 137-146. Sanibel Island, FL, USA (2007).
9. Newman, M.W. Now We're Cooking: Recipes for End-User Service Composition in the Digital Home. IT @ Home Workshop at CHI '06, Montreal, Quebec, Canada (2006).
10. Newman, M.W., Elliott, A. and Smith, T.F. Providing an Integrated User Experience of Networked Media, Devices, and Services Through End-User Composition. *In Proc. the International Conference on Pervasive Computing (Pervasive 2008). To Appear*. Sydney, Australia (2008).
11. Newman, M.W., et al. Designing for Serendipity: Supporting End-user Configuration of Ubiquitous Computing Environments. *In Proc. the International Conference on Designing interactive Systems (DIS 2002)*. pp. 147-156. London, UK (2002).
12. Remde, J.R., Gomez, L.M. and Landauer, T.K. SuperBook: an automatic tool for information exploration--hypertext? : *In Proc. the ACM conference on Hypertext (Hypertext '87)*. pp. 175-188. Chapel Hill, North Carolina, United States (1987).
13. Rodden, T., et al. Between the dazzle of a new building and its eventual corpse: assembling the ubiquitous home. *In Proc. the International Conference on Designing Interactive Systems (DIS 2004)*. pp. 71-80. Cambridge, MA, USA (2004).
14. Sproull, L. and Kiesler, S. Connections: New Ways of Working in the Networked Organization. MIT Press, Cambridge, MA, USA (1991).
15. Stevens, G. and Weidenhofer, T. CHIC - A Pluggable Solution for Community Help in Context. *In Proc. NordCHI '06* (2006).

16.Truong, K.N., Huang, E.M. and Abowd, G.D. CAMP: A Magnetic Poetry Interface for End-User Programming of Capture Applications for the Home. *In Proc. the 6th International Conference on Ubiquitous Computing (Ubicomp 2004)*. pp. 143-160. Nottingham, UK (2004).
17.Twidale, M. Where am I and Who am I? Issues in Collaborative Technical Help. *In Proc. ACM Conference on Computer-Supportive Cooperative Work (CSCW '04)*. pp. 378-387. Chicago, IL, USA (2004).
18.UPnP Forum. UPnP Device Architecture. UPnP Forum (2000).
19.Zhang, J. and Ackerman, M.S. Expertise Networks in Online Communities: Structure and Algorithms. *In Proc. the 16th International Conference on the World Wide Web (WWW '07)*. Banff, Calgary, Canada (2007).