A. Project Summary

When humans interact with each other and their environment, there are breakdowns and disconnects that naturally result due to the mismatch in their understanding of the current situation. Often, all the information needed for understanding is available, but for a number of reasons it is too difficult to deal with it. For example, in disaster situations, information arrives too fast for emergency crews to understand it all. This problem is often exacerbated by computers because they only have a minimal understanding of human dialogue and interaction with the environment. Cell phones, for example, interrupt us without considering the fact that we may be in an important meeting.

The breakdowns and disconnects that often occur in human activity, both naturally-occurring and computer-mediated, can be smoothed out by making computers more aware of the physical and social contexts they are a part of. This project seeks to investigate user-centered methods and technologies to assist designers and application developers in the design, prototyping, and evaluation of representative real-world context-aware applications, focusing on issues that have previously not been adequately addressed, including usability, privacy, and scalability. This work focuses on (1) exploring new user-centered methods and tools for designing, prototyping, and evaluating context-aware applications, (2) understanding the nature of privacy concerns with respect to context-aware systems and developing mechanisms for addressing these concerns, and (3) providing an infrastructure to help programmers acquire, model, store, query, and protect context data across a large number of users, sensors, and applications. Considerable progress is required in all three of these areas before any one of them can be solved satisfactorily.

These methods and technologies will be evaluated by developing three socially relevant applications, including (1) emergency response support for coordinating, communicating, and allocating resources, (2) an augmented wheelchair providing word prediction for individuals with significant motor impairments, and (3) collaborative design spaces supporting groups engaged in long-term design activities.

We have already taken several important steps towards realizing this vision. Previously, we built a context toolkit, which allows application developers to create context-aware applications without having to understand the details of the sensors used to capture the context. This experience led us to develop a richer, context-aware infrastructure and services model that will scale to the types of applications envisioned above. This new architecture, called the Context Fabric, has been designed and we are beginning to implement a prototype to evaluate the approach. We have also had a considerable amount of experience in building context-aware applications, as well as technology that allows disabled users to communicate using computers. Finally, we have had a considerable amount of experience in developing design tools, both single user and collaborative tools.

Our research plan is to first survey and observe the target users for the target applications to see what cues and forms of context can help these users achieve their goals. We will continue to implement and deploy the first version of our Context Fabric infrastructure for testing. Next, we will start to design, prototype, and evaluate the applications described above, while continuing to work on design tools that support the easy creation of this style of application. We will use the results of these evaluations to refine the design and implementation of our infrastructure and then move to a wider-scale deployment. We will test this wider-scale deployment by developing the emergency response application, which pushes hard on the scalability dimension of the problem. Finally, we will distribute and test our final version of the infrastructure, tools, and applications that we develop to interested parties at other institutions.

We will evaluate this research by performing laboratory and field experiments, and by tracking outside use and acceptance of the work. These evaluation techniques will try to show that our approach is appropriate for developing successful context-aware applications that solve real human problems.

Exploring context-aware computing and developing new methods, tools, and infrastructure for their design and evaluation will lead to five major research contributions. First, the exploration of context-aware computing by researchers, students, and professional designers using these new tools will result in a body of knowledge about context-aware applications and interaction techniques that work well for a variety of application domains. Second, by using these methods and tools, students will become educated on how to build and apply context-aware techniques to address the problems of tomorrow's user interfaces. Third, the work will contribute new techniques for evaluating context-aware applications. Fourth, the architecture we create will allow developers to build context-aware applications that are scalable and robust, rather than simple laboratory proofs of concept. Finally, the infrastructure and prototyping tools we develop will allow professionals to more easily create interfaces that are usable by groups that are currently ignored, such as the disabled, as well as by mainstream users working in non-traditional settings.