**Personalized Search Here and Now in Large-Scale Pervasive Computing Networks**

**Christine Julien**
The University of Texas at Austin

**Background and Experience of the Participant.** Dr. Christine Julien is an Associate Professor of Electrical and Computer Engineering at the University of Texas at Austin, where she is a member of the Center for Advanced Research in Software Engineering and the founder and director of the Mobile and Pervasive Computing Group. Her research interests involve software engineering concerns for mobile and pervasive computing, and she has led the development of significant middleware systems, communication protocols, and sensing abstractions to support developers in these challenging environments. She has served on program committees of various pervasive computing conferences and published in pervasive computing related conferences and workshops.

**Vision.** *As we aim to support Internet-style capabilities in pervasive computing networks, understanding, managing, and using the unpredictability associated with query processing in these environments becomes of paramount importance.* In pervasive computing networks, especially as we scale to the high density environments we envision, people and machines embedded in the environment consume vast quantities of data. The available information has high degrees of spatiotemporal relevance, and the information, people, devices, and the networks that connect them are all highly dynamic. At the same time the events that the data represents are often unpredictable and transient. However, users demand the ability to find spatiotemporal relevant information in the here and now. This requires addressing several key research issues, including:

- *Eliciting requirements*: we must identify and understand user requirements of personalized search in large-scale pervasive computing networks;
- *Supporting efficient in-network search processing*: we must design protocols that can limit search scope to handle the large-scale networks, handle the inherent dynamics during query processing and to use knowledge to bootstrap search processing;
- *Redefining and supporting relevance determination*: we must enable searches to identify aspects of results that contribute to increased relevance and to use that information to more efficiently, effectively, and productively resolve searches; and
- *Enabling long-lived searches*: we must make it possible for users to monitor conditions in the large-scale pervasive computing network, even as conditions change over time.

*Supporting personalized spatiotemporal search in dynamic, large-scale pervasive computing environments requires abstractions and protocols that consider uncertainty and dynamics.*

**Evidence.** In our previous work, we have developed a formal model of dynamic query processing and an associated self-assessing query protocol that can process a search query inside the pervasive computing network without relying on the Internet for infrastructure support. This self-assessing protocol not only returns a query result but also communicates a label signifying the "quality" of the result based on monitoring the dynamics that occur during query processing. This work sets the stage for applying in-network processing to personalized search of large-scale pervasive computing networks.

**Extended Vision.** Pervasive computing networks made possible by wireless and sensor technologies enable people, devices, and machines to intimately cooperate on tasks. Fluidity and high degrees of dynamics define these large-scale networks; volatile information sources generate high volumes of short-lived data. As these trends grow, we

must understand how traditional models of interaction, such as search, will evolve given the rapid convergence of the individual and the digital world.

Searching a network is a basic activity in our lives, yet searching a pervasive computing network introduces new requirements. *Spatiotemporal locality*, knowing what is happening *here* and *now*, is paramount. New challenges arise in efficiently and effectively resolving what information is *relevant* in a network as active, variable, and tangible as a pervasive computing system, where the volume of information is too great to be shipped to a central resource, the lifespan of the information is short, and the ratio of data used to data available is minute. Moreover, the benefit of infrastructure and sophisticated advance indexing is typically not available. Existing mobile applications enable the perception of local search, but they search relatively static information for which it is reasonable to rely on the highly indexed Internet. Applications that generate dynamic data do so at a low rate and only generate data likely to be consumed. Pervasive computing applications generate high volumes of data with short lifespans, much of which may not be used, motivating decentralized and scalable spatiotemporal search in pervasive computing networks.

Consider a hospital. While researching a patient's symptoms, a physician may use the Internet to navigate electronic texts, consult colleagues, review the patient's medical history, etc. Most of the information is relatively static. When a patient unexpectedly goes into cardiac arrest, the situation becomes a rapidly evolving one, and the physician's needs change dramatically. He may need to know which nurses have experience with the patient, if and where necessary supplies are on hand, what the patient's vitals were immediately before the arrest and how they are changing, etc. This example demonstrates the need to extract information from the immediate surroundings and also how understanding what information is relevant is a critical aspect of a pervasive computing application.

While one might argue that such functionality can be provided by cataloging information and using traditional Internet search, the nature of pervasive computing networks fundamentally changes the game. Every *thing* in the environment may have several associated sensors (e.g., everything with an RFID tag for identity information, each instrument with an accelerometer for detecting movement, IV sensors that monitor flow rate, etc.). A likely use case of these sensors may be periodically ``checking in on'' each sensor (e.g., once an hour) and cataloging that information centrally. However, in supporting dynamic pervasive computing network applications like the one described above, there is the potential for so much data to be generated and there is no way to know which pieces of that data are going to be useful to a search (much less where the searcher will be located when he searches for it). Uploading all of this data to the Internet is cost prohibitive (in terms of bandwidth), and cataloging it for efficient search is infeasible; this is already a recognized problem in military sensor systems

Given the escalating ubiquity of sensing and mobile devices, pervasive computing networks will play a central role in how we obtain our daily information in the very near future. This vision aims to develop a personalized search engine for the here and now that will enable us to investigate the new requirements of large-scale pervasive computing environments with the specific goal of capturing how search expectations are affected when users and their applications interact directly with the physical environment. The technical advances embodied in the examination of search in pervasive computing networks will have far-reaching impact on the way we live and work in a digitally rich spatial neighborhood populated by people, machines, and actuators.