Dr. Calton Pu

CCB, room 269
Georgia Tech, Atlanta
GA 30332-0280
USA

Web: *www.cc.gatech.edu/~calton/*
E-mail: *calton.pu@cc.gatech.edu*
Tel: +1-404-385-1106
Fax: +1-404-385-2295

Biography

Calton Pu was born in Taiwan, but grew up in Brazil. He received his PhD from University of Washington in 1986 and served on the faculty of Columbia University and Oregon Graduate Institute. Currently, he is holding the position of Professor and John P. Imlay, Jr. Chair in Software at the College of Computing, Georgia Institute of Technology. He is leading the Infosphere project, part of the DARPA Information Technology Expedition program. Infosphere is a joint project between Georgia Tech and OGI, plus a number of other collaborators. Infosphere builds on his previous and ongoing research interests. First, he has been working on next-generation operating system kernels to achieve high performance, adaptiveness, security, and modularity, using program specialization, software feedback, and domain-specific languages. This area has included projects such as Synthetix, Immunix, Microlanguages, and Microfeedback, applied to distributed multimedia and system survivability. Second, he has been working on new data and transaction management by extending database technology. This area has included projects such as Epsilon Serializability, Reflective Transaction Framework, and Continual Queries over the Internet. His collaborations include applications of these techniques in scientific research on macromolecular structure data, weather data, and environmental data, as well as in industrial settings. He has published more than 30 journal papers and book chapters, 100 conference and refereed workshop papers, and served on more than 40 program committees, including the co-PC chair of SRDS'95, co-general chair of ICDE'97, and co-PC chair of ICDE'99. He has served as an associate editor of IEEE TKDE, DAPD, and IJODL. His research is currently funded by NSF, DARPA, Intel, and other sources.

Research Projects

Infosphere Project

URL: *www.cc.gatech.edu/projects/Infosphere*

OBJECTIVE

The Infosphere project will develop concepts, techniques, and tools for the next generation systems software in pervasive computing environments. Systems software for such environments must support end-to-end quality of service (QoS) for users and developers, for example, in terms of performance, availability, maintainability, and survivability. The core abstraction is called Infopipe, which supports information flow through a variety of environments with these QoS properties. These pervasive computing environments include high speed networks such as the Next Generation Internet on one end of spectrum, and limited bandwidth wireless connections on the other end. These environments are also characterized by continuous changes that demand stable and responsive systems software adaptation. Infopipe software will support the generation and composition of code to support information flow through these changing environments. Infosphere focuses on the challenges of bringing fresh information from a variety of sensors to new applications such as personalized fresh information delivery for mission-critical operations in urban terrain.

APPROACH

Infosphere contains two major areas of research. The first area, closer to the users, is the set of next generation information-driven applications making full use of fresh information in a pervasive computing environment. Examples of these new applications include real-time sensor-based remote control, networked embedded systems, real-time decision support, ubiquitous personal guidance, micro-region weather forecasting and virtual presence. These services will support real world applications such as precision farming through rough weather and rescue missions in remote areas or hostile territory. To fulfill this vision, proper systems software must support and control the information flow from sensors to applications. The design and construction of this systems software, called Infopipe, is the subject of our second and main area of research.

The design of Infopipes leverages on our experience from the Quasar and Continual Queries (CQ) projects at OGI, as well as the Event Channel and Aware Home projects at Georgia Tech. Quasar project has built and released systems software tools supporting broad-sense quality of service (QoS) using adaptive resource management techniques based on Microfeedback. Continual Queries project has built and released software tools supporting Internet scale information update monitoring and filtering. Event Channel software supports publish/subscribe communications at the middleware level. Aware Home is a real-time sensor laboratory. Infopipes will support QoS properties declared by users. The system will manage resources according to these QoS declarations and competition for these resources. The information-driven applications will use Infopipes as underlying communications infrastructure, and CQ concepts and software as filtering stages aware of application semantics.

In addition to Quasar, CQ, and Event Channel software, Infopipes also will use specialization and domain specific language (DSL) techniques, software, and tools developed in the Synthetix project. An Infopipe will be described by DSL microprograms that declare its input and output type (syntax and semantics), as well as QoS attributes, constraints, and trade-offs. Using these DSL microprograms, the automated generation and composition of Infopipe software will leverage on specialization tools to simplify and speed up the resulting software. Infopipe composition includes the computation of the QoS properties supported by the composite Infopipe, calculated from the QoS properties of the component Infopipes.

The next generation information-driven applications such as precision farming and rescue missions in rough weather depend on timely delivery of fresh information from sensors to information utilities such as micro-region weather forecast models, and finally to end users. Stale sensor information would be of little use in weather prediction. Applications such as Aware Home (a real-time sensor laboratory), real-time decision support, and virtual presence will help us evaluate Infopipe software and tools for Infopipe software generation. The evaluation will include the software generation process, the effectiveness of generated Infopipe software, and the satisfaction of QoS properties by Infopipes such as efficiency, availability, and survivability.

Recent Publications

Dylan McNamee, Jonathan Walpole, Calton Pu, Crispin Cowan, Charles Krasic, Ashvin Goel, Perry Wagle, Charles Consel, Gilles Muller and Renauld Marlet, *Specialization tools and techniques for systematic optimization of system software*, ACM TOCS, Vol. 19, No. 2, May 2001, Pages 217-251