Title: Network Systems to Support Spatio-Temporal Context-Awareness
Author: Gustavo de Veciana, Dept of ECE. The University of Texas at Austin

Participant's background. Gustavo de Veciana received his Ph.D. in electrical engineering from the U.C. Berkeley in 1993 and is currently a Professor at the Department of Electrical and Computer Engineering at U.T. Austin. His research focuses on the analysis and design of wireless and wireline telecommunication networks; architectures and protocols to support sensing and pervasive computing; applied probability and queueing theory. He has been the recipient of numerous awards including: and NS CAREER Award 1996, co-recipient of the IEEE William McCalla Best ICCAD Paper Award for 2000, co-recipient of the Best Paper in ACM Transactions on Design Automation of Electronic Systems, Jan 2002-2004, co-recipient of the Best Paper in the International Teletraffic Congress (ITC-22) 2010, and of the Best Paper in ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems 2010. Dr. de Veciana has served as editor for the IEEE/ACM Transactions on Networking. He served as the Director and Associate Director of the Wireless Networking and Communications Group (WNCG) at the University of Texas at Austin, from 2003-2007. In 2009 he was designated IEEE Fellow for his contributions to the analysis and design of communication networks.

Proposed Vision. A significant, and largely unaddressed, networking challenge is the widespread use of networked systems to harness potentially massive flows of spatio-temporal events and information to deliver timely and locally-relevant context-awareness to devices/applications and their users. Events and information generated by humans, their devices, agents, sensors, and machines, can significantly enhance an applications' knowledge of resource availability, a proximal agent, user's intents, and state of the environment. Usage scenarios for this type of dynamic spatio-temporal information include *resource discovery* and *allocation*, e.g., locating an available parking spot, allocating frequency bands for cognitive radio network, and sharing compute resources. More complex usages might involve the evaluation of predicates over the flow of events, e.g., to determine the existence of a certain configuration of events or resources at a given location. Evaluation of such predicates can, for example, be used to create "virtual environments" wherein devices' behaviors within a given space-time scope are predicated on the existence of certain conditions, e.g., different staging areas in an emergency relief effort, a playscape that monitors children's movements, or a quiet zone where phones are automatically silenced. One can further envisage complex computations and services that can update spatio-temporal state to capture higher-level characteristics of the environment, or underlying spatio-temporal dynamics as might be used to learn or infer critical aspects of the underlying systems. We see two key research directions as critical towards developing easily deployable scalable infrastructure to support pervasive computing at these large scales:

1. Study of shared distributed/peer-to-peer infrastructures for enabling real-time context awareness. To achieve economies of scale future pervasive computing systems might (should) be built on *shared* network wireless/wired infrastructure(s) supporting the capture, sharing, and computation of spatially and temporally relevant contextual information. In the case where spatio-temporal events/information that are short-lived and relevant only to devices/users in close proximity it is unlikely centralized solutions will scale due to poor amortization of communications overheads. The key then will be to study distributed wired peer-to-peer overlays and/or wireless ad hoc networks to make context available close to where it is likely to be consumed.

2. Theoretical limits for achievable spatio-temporal context-awareness. As part of this vision we believe there is a potential to develop theoretical results akin to those that have driven the development of information theory. Indeed Shannon's capacity result established the limits on how much information could be transferred over a communication link, with more recent results extending this to different types of networks. We believe we should strive to establish *contextual-capacity limits* for pervasive computing systems. In other words given an infrastructure subject to a spatio-temporal loads, e.g., what fraction of the relevant context will a device/user be able to recover given a particular communications and storage infrastructure.