Neighborhood-Consistent Transaction Management for Pervasive Computing Environments*

Filip Perich, Anupam Joshi, Yelena Yesha, and Timothy Finin

Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
1000 Hilltop Circle,
Baltimore, MD 21250, USA
{fperic1,joshi,yeyesha,finin}@csee.umbc.edu

Abstract. This paper examines the problem of transaction management in pervasive computing environments and presents a new approach to address them. We represent each entity as a mobile or static semi-autonomous device. The purpose of each device is to satisfy user queries based on its local data repository and interactions with other devices currently in its vicinity. Pervasive environments, unlike traditional mobile computing paradigm, do not differentiate between clients and servers that are located in a fixed, wired infrastructure. Consequently, we model all devices as peers. These environments also relax other assumptions made by mobile computing paradigm, such as the possibility of reconnection with a given device, support from wired infrastructure, or the presence of a global schema. These fundamental characteristics of pervasive computing environments limit the use of techniques developed for transactions in a "mobile" computing environments. We define an alternative optimistic transaction model whose main emphasis is to provide a high rate of successful transaction terminations and to maintain a neighborhood-based consistency. The model accomplishes this via the help of active witnesses and by employing an epidemic voting protocol. The advantage of our model is that it enables two or more peers to engage in a reliable and consistent transaction while in a pervasive environment without assuming that they can talk to each other via infrastructure such as base stations. The advantage of using active witnesses and an epidemic voting protocol is that transaction termination does not depend on any single point of a failure. Additionally, the use of an epidemic voting protocol does not require all involved entities to be simultaneously connected at any time and, therefore, further overcomes the dynamic nature of the environments. We present the implementation of the model and results from simulations.

1 Introduction

Maintaining data consistency between devices in distributed mobile environments has always been, and continues to be, a challenge. These environments represent networks composed of stationary and mobile nodes that share a subset of a global data repository. The devices use their network connectivity to exchange data with other nodes in the

* This work was supported in part by NSF awards IIS 9875433 and 0209001, and DARPA contract F30602-00-2-0591.

V. Mařík et al. (Eds.): DEXA 2003, LNCS 2736, pp. 276–286, 2003.

network. In order to operate correctly, devices involved in a transaction must ensure that their data repositories remain in a consistent state. While stationary nodes often embody powerful computers located in a fixed, wired infrastructure, mobile nodes represent devices with low-bandwidth communication, limited battery life and with limitation to other resources. Consequently, transacting devices must accommodate mobility and, in turn, possible failures due to a network disconnection. The challenge of providing data consistency is especially substantial for pervasive computing environments.

Pervasive computing environments extend the traditional concept of mobile networks [10,18]. Mobile devices in pervasive computing environments consist of hand helds, wearables, computers in vehicles, computers embedded in the physical infrastructure, and (nano) sensors. A device satisfies user queries by relying on its local data repository and data available in other devices in its vicinity. Additionally, every device is equipped with short range ad-hoc networking technologies such as Bluetooth [3]. The ad-hoc networking technology allows mobile devices to spontaneously interact with other devices, both fixed and mobile, in their vicinity. For example, two cars passing each other on the street can establish a network connection and exchange data while within range of each other. At the same time, pervasive computing environments do not guarantee any infrastructure support, a crucial requirement for traditional mobile systems [6,7,24]. Hence unlike traditional mobile computing, the pervasive environment does not differentiate between mobile clients and servers located in a fixed, wired infrastructure. Instead we model all devices as peers and any two devices may engage in a transaction, a case not covered by traditional mobile computing paradigm. In the mobile computing paradigm, only one transacting device, the mobile client, is allowed to move during a transaction. This allowed previous solutions to rely on the help of the infrastructure by using mobile support stations as proxies; however, there is no default infrastructure support in pervasive computing environments. Additionally, pervasive environments relax other assumptions made in mobile computing paradigm. As all devices may be mobile, the vicinity of each device is likely to change in both spatial and temporal dimensions. This not only limits data and data source availability but the serendipitous nature of the environment also limits the possibility of reconnection between transacting devices. In pervasive computing environments, there is no guarantee that all devices wishing to transact may be concurrently available and that two disconnected devices will *meet* again. For example, when two people serendipitously meet at an airport and agree to exchange a song for a micro-payment, their electronic wallets must be updated correctly even when one person leaves the airport before the transaction completes [1]. Consequently, transacting peers must either trust each other or rely on a third party. In a traditional mobile paradigm, the third party is a server located in a fixed, wired infrastructure. In that case, transacting peers must send all relevant data to the server before they disconnect. This may not always be possible in pervasive environments. Instead, an alternative approach is to use other peers in the environment as third parties. This raises the issue of trust since there is no guarantee that these peers will behave correctly. We address the issue via the use of a random witness selection policy which reduces the probability of obtaining malicious witnesses. In summary, the change in perception of mobile devices, together with other characteristics of pervasive computing environments, limits the use of traditional mobile transactions.