Peer-to-Peer Policy Management System for Wearable Mobile Devices

Michael Massimi
Department of Computer Science
The College of New Jersey
Ewing, NJ 08628
massimi2@tcnj.edu

Ursula Wolz
Department of Computer Science
The College of New Jersey
Ewing, NJ 08628
wolz@tcnj.edu

Abstract

Wearable computing devices are obviously made all the more powerful due to their ability to be innately mobile. With this in mind, spontaneous peer-to-peer networks may emerge amongst several users wearing individualized devices. Although security of wireless networks in general has been explored, we propose a system wherein users are protected from one another and from abuse of communally owned resources through the use of a peer-to-peer policy management system called PPPM.

1. Introduction

Wearable devices are intrinsically mobile, with a clear need for peer-to-peer models of communication. In particular, it is not a far leap to imagine spontaneous, ad hoc networks being created between users who move in and out of local networks. We propose a system that mediates resource allocation to particular users within these networks. Such ad hoc networks, formed on-the-fly between a group of users and a passive commodities server, break the traditional mold of having an omnipotent system administrator who reigns over the goings-on of the network clients. Such systems have been referred to as WPANs (Wireless Personal Area Networks), and "allow a number of independent data devices to communicate" [1]. A crucial question for WPANs is how to determine the trustworthiness of the devices that enter the network to maintain security while at the same time preventing abuses of local resources in a fair manner.

We describe an architecture called Peer-to-Peer Policy Manager (PPPM) that supports the creation of a dynamic and equitable security policy management system that is based on peer consensus of resource management. Assuming that all parties involved are of equal stature, but not necessarily of equal trustworthiness, it is proper to have them police themselves akin to a "neighborhood watch." A key element of this approach is that it can be rapidly applied if the resources in question are either ambiguously owned or no single entity can be easily appointed as a system administrator.

2. Background and Related Work

Although there has been research into the security aspects of wireless computing in general, relatively little work has been done in investigating the dynamic networks that may accompany mobile devices with regard to protecting users from one another. There tends to be a general consensus that "personal area networks will proliferate early in the millennium" and standards are already being devised for it [1]. Just as there is a need for standards, there is also a need for security, as is evidenced by the wireless encryption package that accompanies the 802.11 wireless routers [see 2]. PPPM attempts to protect users from one another in a spontaneous network through reliance upon a sense of generalized trust.

The WebSplitter system [3] developed at IBM permits multiple users on different devices to view portions of the same web page at the same time. The authors note that "[a] partial view of an XML page is first constructed based on a user's access privileges to the content in the page." In our PPPM model, the access privileges in question are defined through the policing of one's peers. The WebSplitter system is also concerned with resource management; for example, particular devices lacking multimodal features such as audio or video capabilities can take advantage of resources nearby, like speakers residing in the room. The problem is how to automatically mediate resources between a group of users whose trustworthiness and priority are not well-defined. For example, what if more than one user wanted access to those speakers at the same time? How does the ad hoc network determine who gets the speakers, and for how long? We begin to address these questions through the architecture presented here.

3. System Design

The PPPM system is comprised of three different entities. First, there are individual users' **devices**. For our current implementation goals, these are handheld PDAs or cellular phones that are network-ready. We anticipate that the same design principles could be applied to more exotic wearable devices. The operating systems' kernel

1530-0811/03 $ 17.00 © 2003 IEEE

contains a small section that provides operability for the policy management system, an extension of the work seen in [4] regarding distributed firewalling systems. Although putting the policy client into the kernel is useful insofar that it prevents casual mistreatment, there still exists a possibility for fraud (see "Limitations and Weaknesses"). Each device has software that allows it to submit requests for status changes of its policy or that of others. Additionally, it must be able to accept changes to its policies. For example, if a group of devices notice a resource abuser in their midst, they can "vote" to have access reduced or denied. When the group mediates a change to a particular device, the software automatically obtains the new policy, overwrites the old one, and applies the changes on the target device. It is assumed that the human user can communicate policy demands and participate in the mediation process as needed.

Second, there exists any number of **shared resources**. These may include printers, audio speakers (as in the WebSplitter example above), scarce storage space, or local information sources. We assume that many of the users are in contest for these resources at any given time, just as processes in an operating system vie for control of resources. Resources are ambiguously or communally owned and no individual user can claim the device forever (just as a process cannot indefinitely claim a CD-ROM drive as its own in an operating system).

Finally, there is a **policy manager**. This is software that can determine policy based on input received from the devices in the WPAN. The architecture of this system is a simple rule-based one, where rules are created based on group mediation principles. Ultimately, any device should be able to establish itself as policy manager (a group leader if you will) through various equitable models (e.g. an election). In our example, suppose that on a college campus a WPAN is established between several colleagues who mediate between themselves how resources should be shared. Each time they enter the network, the policies are pushed to each device over the network from the centralized policy manager upon, and the policy manager is responsible for storing the policy between sessions. The question of session-to-session policy is addressed by the independent policy manager, but becomes more complex if WPAN are truly ad hoc. Individual devices would need to keep information on the trustworthiness of known peers.

4. Implementation

Implementation of the PPPM system consists of developing prototypes of the device, the policy manager, and the communications software. A major consideration is the rule base and formulas for determining what policy changes to enact. We are currently implementing a protocol for systematically testing the policy manager on "canned" devices that exhibit particular traits and can be manipulated in an experiment setting.

5. Limitations and Weaknesses

Although this system provides the benefits of a democratic computing environment, there exist potential points of exploitation.

The policy client responsible for enforcement resides within the kernel of the operating system. Placing it here is necessary because we gain the ability to access resources that are addressable only by the device. For example, we may want to control sound volume during a presentation so an ongoing talk cannot be disrupted. A weakness of placing the policy client here is that highly sophisticated users can still compromise the kernel.

A number of techniques for exploiting the system are possible. Generally, vote manipulation poses the largest threat. This can be mitigated through encryption, user validation, and other general security systems. Simply because there exist measures to dampen the effects does not mean, however, that the system will be tamper-proof.

Some users may attempt to lure others into joining WPANs that are not reputable. We envision that people will have the same approach towards joining a WPAN as one does when running a program on a personal computer. The user should trust the policy manager and all other users to act responsibly.

6. Future Work

Usability studies must be conducted on the responses of the real users in participating in this particular type of system. Further tweaking of rule formulae may be necessary as well.

7. References

[1] R.C. Braley, I.C. Gifford, R.F. Heile, "Wireless Personal Area Networks: An Overview of the IEEE P802.15 Working Group", Mobile Computing and Communications Review, Volume 4, Number 1, ACM Press, New York, January 2000, pp. 26-33.

[2] J. Allen, J. Wilson, "Securing a Wireless Network", User Services Conference, ACM Press, New York, 2002, pp. 213-215.

[3] R. Han, V. Perret, M. Naghshineh, "WebSplitter: A Unified XML Framework for Multi-Device Collaborative Web Browsing", CSCW Conference, ACM Press, New York, December 2000, pp. 221-230.

[4] S. Ioannidis, A.D. Keromytis, S.M. Bellovin, J.M. Smith, "Implementing a Distributed Firewall", Conference on Computer and Communication Security, ACM Press, New York, 2000, pp. 190-199.