

PAWNs: Satisfying the Need for Ubiquitous Secure Connectivity and Location Services

Paramvir Bahl¹, Anand Balachandran², Allen Miu³, Wilf Russell¹, Geoffrey M. Voelker², Yi-Min Wang¹

¹ Microsoft Research

² University of California, San Diego

³ MIT

Abstract

The dawning of the 21st century has seen unprecedented growth in the number of wireless users, applications, and network access technologies. This trend is enabling the vision of pervasive, ubiquitous computing where users have network access anytime, anywhere, and applications are location-sensitive and context-aware. To realize this vision, we need to extend network connectivity beyond private networks, such as corporate and university networks, into public spaces like airports, malls, hotels, parks, arenas, etc. – those places where individuals spend a considerable amount of their time outside of private networks.

In this article, we argue that wireless LAN technologies are the ideal mechanism for extending network connectivity to these public places, and enabling location and context-aware applications in them. However, implementing and deploying public-area wireless networks (PAWNs) present a number of practical challenges, including network security, privacy, authentication, mobility management, and provisioning of key services. We discuss these challenges as a general problem for PAWNs, and then describe a PAWN that we have designed, implemented, and deployed called CHOICE that addresses them. We describe the architecture and components of CHOICE, the service models it supports, and the location services and context-aware applications that we have implemented and deployed in it.

1 Introduction

The dawning of the 21st century has seen unprecedented growth in the number of wireless users, applications, and network access technologies. This trend is enabling the vision of pervasive, ubiquitous computing where users have network access anytime, anywhere, and applications are location-sensitive and context-aware. To realize this vision, we need to extend network connectivity beyond private networks, such as corporate and university networks, into public spaces like airports, malls, hotels, parks, arenas, etc. -- those places where individuals spend a considerable amount of their time outside of private networks.

In this article, we argue that the substantial performance benefits of wireless LANs make them ideal for public-area wireless networks (PAWNs). Although next-generation cellular networks will undoubtedly play a role in providing wide-area long-range service, advances in indoor, short-range wireless communication technology and the proliferation of lightweight, hand-held devices with built-in high-speed radio access have made wireless LAN

deployments increasingly common. Based upon the IEEE 802.11 standard [1], wireless LANs are emerging as the ideal solution for providing high-speed connectivity in private networks and, to a limited extent so far, public places. As a result, network connectivity at 11 Mb/sec is becoming commonplace, and this data rate is expected to grow 10-fold in the next three years [2].

However, wireless LANs alone are not sufficient for implementing public-area wireless networks, and there are a number of challenges in making them a viable platform for PAWNs. First, network access in public places must be able to support a wide range of service models, from free access to paid connectivity to differentiated quality of service. As a result, PAWNs must support mechanisms that enable providers to implement a wide range of policies for providing user access to the wireless network. For policies that restrict user access (i.e., non-free access), PAWNs must authenticate users before providing them service, and they must secure the network against unauthorized and malicious access. For policies that require payment, PAWNs must provide mechanisms for implementing accounting and billing services. Second, to support location-sensitive and context-aware applications, PAWNs must also provide mechanisms for both determining and disseminating location and other contextual information about users to their applications. Finally, PAWNs must keep all personal information, such as communication traffic and contextual information like user location, private and secure.

Over the past year, we have designed, implemented, and deployed a public-area wireless network called CHOICE that addresses these challenges in one integrated system [3]. CHOICE is a service platform on which any number of network providers can offer network services, enabling users to choose the kind of service that best fits their needs. It supports two kinds of service models, one that enables free access to local intranet services like a local Web server, and a second, enhanced service model that supports billed access with full Internet connectivity with various quality of service options. It uses a network admission server in conjunction with a global authentication database to authenticate users and grant them access to the wired network. It uses a traffic control gateway to perform per-packet verification to enforce authorized access, quality of service policies, and

accounting constraints. Finally, it supports various mechanisms for determining user location and propagating that information to location-sensitive and context-aware applications. Altogether, CHOICE demonstrates that wireless LANs are a compelling technology for providing high-performance wireless network access in public places.

The rest of this article is organized as follows. In Section 2, we discuss the deployment considerations for implementing a public-area wireless network, particularly one that supports context-aware applications. In Section 3, we describe the architecture, service models, and location services and applications of CHOICE, a PAWN we have implemented and deployed. In Section 4 we motivate why a PAWN such as CHOICE is practical and a commercially viable alternative to other network architectures such as cellular networks. Finally, we summarize in Section 5.

2 Public Area Wireless Networks

In this section, we begin by discussing the deployment issues that make PAWNs different from network deployments in home and enterprise settings. We then describe the types of access services that PAWNs can be expected to provide. Finally, we discuss the issues in providing location services in PAWNs.

2.1 Deployment Issues

PAWNs present many challenges to the network designer. First and foremost among them is the lack of an implicit trust relationship between the users and the public network. In contrast, home and enterprise networks have pre-arranged trust relationships with their users, and so access to the network can be conveniently granted through a pre-configured common key. Public networks need to provide access to unknown users who may not have visited the network before. This necessitates the use of a formal authentication mechanism that enables users to identify themselves to the network. Authentication also makes the users accountable for the services that they use in the network, thereby providing a convenient means for billing. Also, all users may not prefer the same mode of authentication. Hence the network must have a provision to support multiple authentication options. Finally, the authentication process must be end-to-end secure. In other words, no one except the user and the authenticating entity must be privy to personal information such as username, password, credit card numbers, etc.

Furthermore, public networks are exposed to malicious users and are thus vulnerable to many kinds of attacks. Therefore, in order to secure the host organization, there is a need to perform access control in addition to user authentication to prevent unauthorized users from accessing the network. The access control mechanism should guard against the most common modes of attack, like dictionary

attacks on passwords, replay attacks, and IP and MAC address spoofing.

A third aspect that differentiates PAWNs from private networks is that of service differentiation. In other words, since access is provided in an “individual-centric” manner, it is possible to offer enhanced services to users who are willing to pay extra to get better network service. For example, service providers may offer higher bandwidth connections and privileged access to a local music/entertainment repository for the savvy user willing to pay more for such service. Also, a new set of transactional and collaborative applications that exploit the knowledge of user locations can be deployed in the public setting according to the users’ needs and preferences. In other words, users can “shop” for multiple levels of network services. We describe the aspects of service differentiation and location-based services in the next two subsections.

2.2 Access Services within PAWNs

A PAWN is built under the premise that so long as a user’s identity can be established by the host organization and he or she agrees to the payment options for the services received, the network should grant access to the user. However, treating every individual independently offers new potential for the host organization to offer enhanced services to users who desire them. We describe these services in more detail below.

2.2.1 Bandwidth Allocation

Bandwidth in the wireless domain is a scarce, shared resource. This scarcity is accentuated by the fact that the demand for bandwidth is hard to predict in a public setting, which is characterized by a very transient user population. In order to satisfy the bandwidth demands of all users, the host organization will have to implement a quality of service (QoS) policy to manage and allocate the bandwidth in a scalable manner. Alternatively, bandwidth allocation could be handled through service policies that may have been pre-negotiated between the host organization and other corporations, effectively dividing users into various service classes.

2.2.2 Security Provisioning

In Section 2.1, we stated that one of the issues in deploying PAWNs is securing the host organization against malicious attackers. Additionally, a PAWN may also provide security services to the user for the purposes of data integrity. Authorized users should be able to choose varying levels of security for their data. Again, as in the case of bandwidth, this could also be configured via a policy where users belonging to a particular organization would be assured a certain level of encryption of their data for a pre-negotiated cost.

2.2.3 Billing and Accounting

The goal of the billing and accounting service is to (a) enable the host organization to create flexible charging plans and bill users accurately for the amount of bandwidth that they use, and, (b) estimate aggregate system demand from members of various organizations that have negotiated service packages for their members, thereby segregating users into different service classes based on the extent of use of the network resources. If the demand from users of a certain organization was particularly high, it could change the users to a higher service level and negotiate a new agreement with the organization.

2.2.4 Mobility Management

One of the inherent characteristics of a PAWN is that the users are mobile. A user may set her mobile device to hibernate at home, and then move to a public area serviced by a PAWN. After such a migration, the user's device may have to be reconfigured for network access in the PAWN. This is due, for example, to the changes in security and authentication models in PAWNs versus private networks. Although the user may use Mobile IP [4] to handle network-layer mobility, the host still needs to be dynamically configured to operate properly when switching among public and private networks.

2.3 Context and Location Services Within PAWNs

A PAWN is typically characterized by users who are mobile but who often crowd around areas of interest (e.g. restaurants, mall rest areas, airport terminals, etc.). This gives the potential to the host organization to offer users services that exploit knowledge of their geographic location. Such location-aware services enable users to interact with their immediate environment and the system to mirror the surrounding environment intelligently to user devices.

2.3.1 Issues and Differences

Context and location-aware applications have been well researched but generally in enterprise settings [5] [6]. Typically, these systems have relied on technologies such as badges and emitters (based on IR technology) that are attached to users, equipment, and building walls that enable the system to create a "location map" of the environment using a network of IR-sensors.

PAWNs open up a new environment for supporting location-aware applications for the following reasons. First, a PAWN is characterized by users who are previously unknown to the system. Therefore, we cannot often expect that these users would be equipped with special devices for determining location information. Second, PAWNs generally span large areas where range-limited sensor technologies like IR scale poorly. Thus, providing extensive

coverage would involve high installation and maintenance costs.

Ideally, a PAWN would implement and deploy location-aware services that require limited user intervention so that location services can be provided transparently. One way to do this is to complement the already useful data-networking capability provided by RF wireless LANs with a location capability so that the users require no extra hardware to use the service. Further, the granularity and accuracy of location information needed in a PAWN is very different from that in enterprise settings, where location information is typically used for indoor surveillance applications and online collaboration between users. Since PAWNs are characterized by frequently roaming users, the location information has to be updated at a corresponding higher rate. Also, some of these applications (see below) may involve financial transactions. Therefore security of location information and privacy of users is an essential consideration.

2.3.2 Determining Location

As already mentioned, central to each location-aware application is the ability to determine user locations with a useful degree of accuracy. There are multiple ways one could use an RF data network to determine user location. Below we describe three algorithms that can compute location information with varying levels of accuracy.

(a) Association with the Access Point

The naïve approach is to determine the access point (AP) in the wireless network that the user is associated with. Upon entering the network, the user can detect the identity of this AP (IP or MAC address) from the MAC-level beacons that are periodically sent by all APs. In this approach, the wireless device of the user is programmed to associate with the AP from which the strongest signal is heard. However, knowledge of the AP-level association merely indicates that the user is within communication

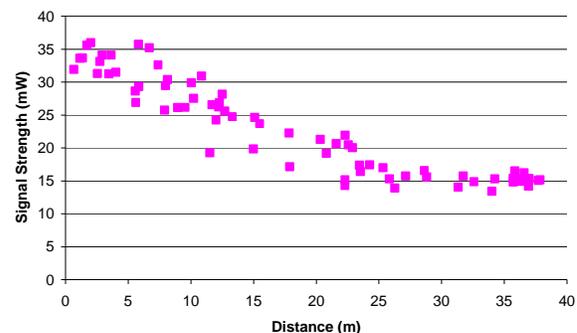


Figure 1: Empirically recorded profile of signal strength as a function of distance between transmitter and receiver.

range of the AP, and at best gives the radius of the largest circle around the AP within which the user could be located.

(b) Using Signal Strength of AP beacons

In order to improve the location estimate obtained from AP association, the identity of the AP can be used in conjunction with the strength of the beacon signal from the AP. The beacon signal strength can be used to estimate the user's radial distance from the AP using a simple radio propagation model that characterizes how the signal strength varies with distance in a radio channel. Indoor signal propagation characteristics can be affected by the presence of walls and obstructions; consequently the radio propagation model has to take these into account before estimating the user's location. The location estimate could be obtained using either theoretically computed or empirically determined models [7]. As an example, Figure 1 shows the variation of signal strength as a function of the separation between transmitter and receiver in an Aironet 802.11 wireless LAN [8].

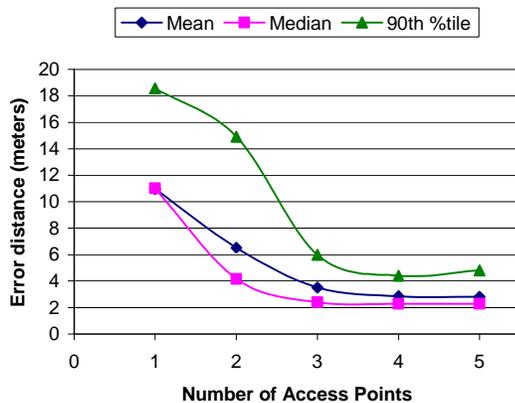


Figure 2: The effect of the number of access points on the error in the location estimate.

(c) Using Signal Strength From Multiple APs

The estimate from the above method can be further improved by using the signal strength values from a network of APs. The idea is that, as a user moves about the network, there is a clear trend in the observed signal strength of the beacons at the receiver and this trend is independently observed for every AP within range. As the user moves, the location estimation algorithm computes a tuple $(ss1, ss2, ss3)$, of the signal strength received from multiple APs. An algorithm then matches this set with the closest set in a pre-computed signal-strength database that lists the measured signal strengths at various fixed locations in the network. The coordinates corresponding to the closest matching set is guessed to be the user's location. Figure 2 shows that the error distance between the actual



Figure 3: A user of the CHOICE network at a mall in Bellevue, Washington.

and estimated locations improves considerably as more APs are considered in the estimation process, but that the effect tapers off beyond 3 APs. There are many enhancements to the above technique that account for dynamic user mobility and dynamic changes in the RF environment, which considerably reduce the error in the location estimate. For these enhancements, we refer readers to [9].

3 The CHOICE Network

Over the past year, we have designed, implemented, and deployed a public-area wireless network called CHOICE at a public mall. Figure 3 shows a user using the CHOICE network in the Crossroads Shopping Center in Bellevue, Washington.

We are aware of some recent work in the areas of user registration [4], authentication and authorization [10], and security [11] [12] in foreign networks. However, to the best of our knowledge, the CHOICE architecture is the first working system that integrates the task of authentication, access verification, policy-based service provisioning, and location services in a lightweight, hardware- and protocol-agnostic manner. We describe the system components, access services and location-aware services of the CHOICE network below.

3.1 System Architecture and Components

Figure 4 illustrates an organization of the CHOICE network, a wireless network that we have proposed for large public areas such as airports, university campuses, shopping malls, and convention centers. This network consists of a global authenticator, an admission server, one or more traffic control gateways, a client module and a policy manager. We briefly describe the functions of each component below. For a more detailed system-level description of this public-area network architecture and all of its components and interactions, we refer readers to [13].

3.1.1 Global Authenticator

One way of authenticating unknown users in a public network is to use a trusted database that is globally available. The global authenticator maintains a database of all valid users who subscribe to its service, and it is used to establish the identity of users in an end-to-end secure manner. Users can choose one of many available authenticators like E-cash systems, credit card organizations (e.g., MasterCard), digital certificate agencies, or databases maintained by particular businesses and clubs (e.g., Gold Club Frequent Fliers).

3.1.2 Network Admission Server

The network admission server (NAS) restricts access to the public network until the user is authenticated and allows authorized access through the traffic control gateway (described below) upon successful completion of authentication. When a user first enters a PAWN, the DHCP server running on the NAS provides an IP address and the local Web server redirects the connection to the global authentication service of her choice. The user's client detects the existence of the CHOICE network service through service broadcasts (see Section 3.2.4). If not already present, the client module needed to detect this service can be downloaded from the CHOICE Web server, which is always identified by the URL <http://choice/>. At this point, the NAS performs IP-address filtering on every incoming packet; any packet with a destination address other than the DHCP server, the Web server, or the authenticator is dropped.

Upon authentication, the NAS provides the user and the traffic control gateway with a $(key, token)$ pair and a key_id , which are valid for a finite amount of time and renewable afterward. The key_id is an index into an array of valid $(key, token)$ pairs that have been handed out to users. The key is used for encryption/decryption and the $token$ is the value that is tagged to every packet before encrypting it with the key . Thus, the encrypted tag provides a cryptographic binding between the user and the packet so that the network can identify the source of the packet and determine the packet's access rights and privileges. After obtaining the $(key, token)$ pair from the NAS, the user has authorized access to the network resources; the system then redirects all user communication through the traffic control gateway.

3.1.3 Traffic Control Gateway and Client Module

The traffic control gateway (TCG) handles verification and enforces policies on a per-packet basis for users authorized by the NAS. In addition to checking whether each packet is encrypted with the correct key and tagged with the corresponding $token$, the TCG also interacts with the policy manager (see below) to implement policies that may be negotiated between users and the host organization.

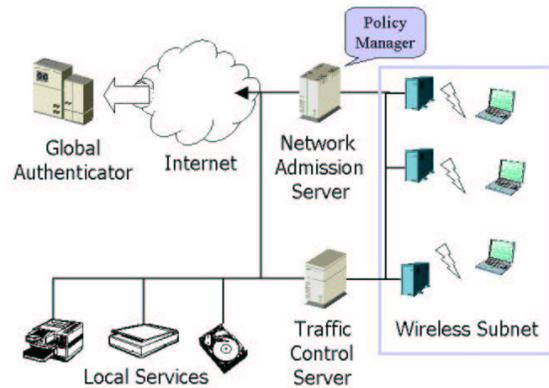


Figure 4: The architecture of the CHOICE public-area wireless network.

The TCG may either be located one hop into the access network (at the border router) or can be built into the wireless access points.

The client module is a software component resident on user devices that tags all outgoing packets with the $(key, token)$ pair obtained from the NAS. The client module can be downloaded from the host organization's Web server. This feature in our architecture enables users to freely access the Internet from any PAWN without requiring any additional support on their devices or any modifications to the protocol stack.

3.1.4 Policy Manager

The goal of the policy manager is to enable service differentiation. The policy manager allows the host organization to set policies that may be pre-negotiated with other corporations. For example, a corporation may negotiate a service package with a PAWN deployed in a local airport such that, whenever any of its employees access the airport wireless LAN, the airport would provide a certain level of bandwidth at a fixed cost. The policy manager maintains entries in its policy table to enable differentiated bandwidth allocation for each organization that has negotiated such a service. Alternatively, the policy manager can also maintain general policies that segregate users into different classes based on access cost. When users visit the PAWN, they may be able to choose the best available access package after assessing them against the criteria of security, bandwidth, and cost.

3.2 Access Services within CHOICE

We now describe how the CHOICE architecture enables the implementation of differentiated services that were introduced in Section 2.2. These services are fulfilled at a cost to the user and are based on a differentiated

charging model, which would vary according to the level of service the user requires.

3.2.1 Differentiated Bandwidth Allocation

We envision that each user who subscribes to the differentiated bandwidth service in CHOICE indicates her range of bandwidth expectation (b_{\min} , b_{\max}), to the network. The resource allocation algorithm tries to provide every user with her b_{\min} requirement and shares the remaining available resource equally among all users. In order for the network to honor the data rate requirements of users, the host organization has to meet two requirements: (a) ensure that users are provided with their negotiated data rate during their sessions, and (b) ensure that no user consumes more than their allocated share of bandwidth. To ensure the first requirement, the host organization performs admission control on new connections and ideally implements a fair scheduling algorithm on admitted connections [14]. The latter requirement is achieved at the TCG, which performs per-packet verification and thus keeps track of the bandwidth consumed by each user over time. Users can also change their bandwidth requirements during a session by renegotiating the service with the host organization.

3.2.2 Security Provisioning

In CHOICE, we support multiple levels of security embodied in basic, medium and enhanced modes of security. The basic mode provides minimum encryption of the security token that is tagged to every outgoing packet. Medium encryption encrypts packet headers as well, and full encryption encrypts the entire packet. The client module running on the user's wireless device can dynamically change the encryption algorithm used to encrypt the packet. Our architecture does not preclude the use of alternative higher-layer security mechanisms like IPSec [12].

3.2.3 Billing and Accounting

While the TCG implements per-packet verification for purposes of security, it automatically incorporates per-packet accounting for each user. Accounting for the amount of bandwidth used by each user is achieved as a result of per-packet processing at the TCG. The accounting information gathered at the per-packet level can then be handed to third party accounting and charging systems that would then be responsible for billing [15]. We note here that CHOICE does not advocate any particular pricing scheme; it only provides the mechanisms and the flexibility to the host organization for implementing different policies.

3.2.4 Mobility Management Service

We have mentioned that users need to be able to seamlessly manage and configure their devices when they enter and leave a PAWN (section 2.2.4). In CHOICE, we

implement this using a network discovery service, where the network broadcasts beacons that contain a unique network identifier, and the IP addresses of the NAS and TCG. Thus when the user first enters the PAWN, the client module uses the information contained in the broadcast beacon to establish the initial connection to the Web server and prompts the user to begin authentication [16]. After the authentication succeeds, the client module receives and stores the (key , $token$) pair, enables packet tagging, and sets the default gateway to the advertised TCG. When the user returns to the home network, the client module no longer receives any beacons and, after a timeout, disables packet tagging and restores the host's default network settings to gain access in the home network. Note that the client module still has the un-expired (key , $token$) pair stored in a table indexed by the advertised network identifier. Should the user decide to return to the same PAWN again, the client module can simply re-enable packet tagging and provide seamless network access without the need for another authentication.

3.3 Location-Sensitive and Context-Aware Applications in CHOICE

In this section we describe some of the location applications that have been built to use the CHOICE network. The applications use a combination of the techniques described in Section 2.3 for determining the user location. Our purpose in building, deploying and describing these applications is to show how businesses can use PAWNs to offer additional services over and beyond basic Internet access and how the CHOICE network enables such services.

3.3.1 WISH (Where IS Harry)

WISH is an application deployed on the CHOICE network that enables CHOICE users to look for other people who are in their vicinity and who have allowed their name and location to be made publicly available to the system. The URL <http://wish/> always resolves to a Web page on the CHOICE Web server that includes the names of WISH users, their interests, tag lines, etc., and a map pointing out the location of each user. The idea is to encourage social interaction between people who may not know each other but who may share several common interests. Subscription to this service is completely voluntary.

The WISH system, shown in Figure 5, consists of a client software module that sits on a library customized for wireless devices (WiLIB) and a stateless server module. The control of location information dissemination is left solely with the user. The WISH client software, running on the user's handheld device, periodically extracts from its radio-frequency (RF) wireless network card the identity of the AP that the device is associated with and the strength of

the signals received from the AP (Sections 2.3.2 (a) and (b)). It then sends this information along with the user's name and activity status to a WISH server. The WISH server maintains an RF signal propagation model and a table that maps APs to a physical location. Using the information provided by the client, the WISH system is able to determine the user's real-time location to within a few meters of where he or she is. A confidence percentage is associated with each estimate. Note that for ease of deployment, we do not use the technique of extracting

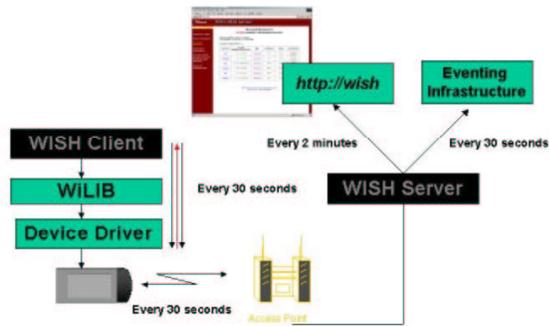


Figure 5: The WISH service architecture

signal strength from multiple APs in order to determine the user location (Section 2.3.2 (c)).

3.3.2 Location-based Buddy List

Buddy lists are fairly common these days. For example, both AOL and MSN messenger provide a buddy list service in their instant messaging software [17] [18]. We have taken this concept a step further by including the notion of location into buddy lists.

Location-based buddy lists are best explained with the help of an example. Say two friends living in different parts of the country arrive at a common airport the same time. They are scheduled to take their connecting flights in a few hours. They are traveling on different airlines and arrive at different gates so they are not initially in contact with each other. Normally, unless they had chatted earlier, they would have passed each other and not met. However, with CHOICE and the location-based buddy list software they get an alert that notifies them that their buddy is in the vicinity with directions on how to find him. Several similar examples can be described, but it suffices to say that location is a fairly useful addition to an already popular buddy list service.

Figure 6 shows the architecture for our location-based buddy list service. When a user first connects to the PAWN via CHOICE, his pre-configured buddy list is extracted and

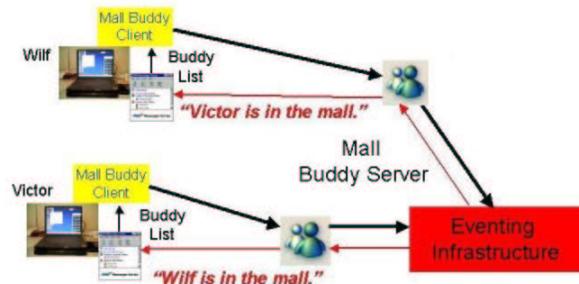


Figure 6: location-based buddy list architecture

sent to a backend eventing server [19]. The eventing server already knows who the user is (via CHOICE authentication) so it stores this information in a local database. Additionally the WISH client running on the user's machine periodically updates the eventing server with the user's location information (see Section 3.3.1). The process is repeated when a new user (the buddy) connects to the same PAWN. The eventing server sees a match and dispatches an instant message alert to both users notifying them that they are near each other. It also sends them a map showing them each other's position.

3.3.3 OnSale Mall Buddy Server

A third application that we have implemented on the CHOICE network is a personalized sale announcement system based on location. This is a subject-based publish/subscribe eventing system based on user profiles and product categories. Figure 7 shows the service architecture for this system. Just as in the case of the buddy list, after a user connects to a PAWN via CHOICE, his profile is extracted and sent to a backend eventing server. In this case, a user profile includes his shopping interests, such as sports-goods, food, clothing and apparel, electronics, etc. The back-end server stores this information for as long as the user stays connected¹. When a local business owner decides to put an item on sale, he or she goes to a local administrator's Web page (a service provided by the building owner) and adds the sale information to this web page. The sale information includes the name of the store, the name of the item, the original price and a sale price. The vendor then selects the category under which this item belongs. When the vendor inputs the information, the Web server synchronously invokes the variable-update API of the eventing server, which causes the firing of a CHANGE event. On finding a category match between the "onsale" item and user profiles, the eventing server generates an instant messaging alert and sends this information to all interested users. For example,

¹ Connection status is maintained via WISH location updates, which act as a "keep-alive" signal for the eventing server.

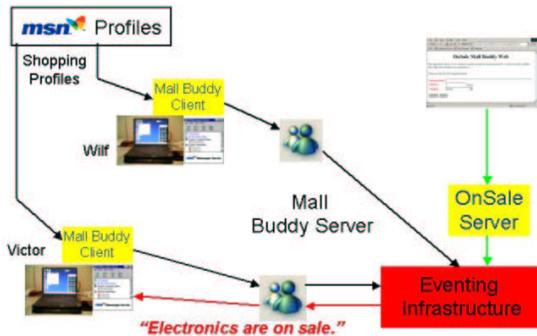


Figure 7: OnSale mall buddy service architecture

when an electronic item is put on sale, users interested in electronics get an alert with the name of the item, its sale price, and the store that has put it on sale. A map containing directions from where the user is to the store is provided with this notification.

Having discussed the Internet access and location services aspect of CHOICE, we now turn our attention to how this technology promotes different business models and encourages the ubiquitous deployment of PAWNs.

4 Looking into the Crystal Ball of PAWN Deployment

If PAWNs are to become ubiquitous, there has to be a business model that generates revenue and thus encourages businesses to install them on premises. As mentioned in the introduction, we firmly believe that the growing need of users to stay connected will drive PAWN deployment. Initially PAWNs will be deployed in large hotels, conference centers, sports arenas, and airline lounges. Once users become accustomed to their availability, though, PAWNs will spread to other public places as well. Competitive pressures between building owners trying to attract businesses and customers to their premises will drive PAWN deployment.

So how will this deployment occur? One model that is currently popular is that of small wireless service provider (WSP) companies making deals with local businesses and deploying the PAWN infrastructure in these locations. A potential problem with this model is that the smaller WSPs generally do not have deep pockets and hence the number of places they are able to deploy PAWNs is limited to the point of not being attractive to large sets of users. Due to the lack of a large customer base the subscriber fee is often high and cost-prohibitive for the average user. Additionally, with smaller WSPs there is a risk of Internet access being offered by a confusing and unpredictable

patchwork of providers. Roaming agreements between these WSPs can alleviate some issues but these have to be worked out and progress is slow since no one knows which of the smaller WSPs will survive over the longer term.

A second model for deployment is similar to the one in the cellular world, in which a few very large cash-rich WSPs deploy thousands of PAWNs world-wide. The advantage for the user is that signing up with one large WSP provides assurance that he or she will have access to enough PAWNs to make their service useful. As of this writing, the adoption of this deployment model is unclear as it is not known whether or not the large WSPs will be willing to spend the money on a wireless service that competes with cellular data service. Having already spent billions of dollars in buying spectrum and installing the cellular infrastructure the larger WSPs are currently focused on recouping their expenses from subscribers. Consequently, they may not be overly eager to spend more money in deploying PAWNs.

A third model, the one that we prefer and which we think has a good chance of success, is one where local businesses install their own PAWN and make them available to all customers who visit their premises. This model distributes the infrastructure cost while allowing users to connect to the Internet at a large number of places. User authentication can be carried out by a globally available third-party for a small fee to the local business. However, for this model to succeed, technology is needed that allows businesses to be creative in how users are authenticated, how they can generate revenues by offering personalized services, and how they can protect themselves against malicious users.

We believe that a building owner who has deployed a PAWN would like to have a finer-grain control on how his network is used. Towards this end, we consider two service models. The basic service model includes free access to the local Intranet services and resources like the organization's Web portal page with links to resident businesses and services like an indoor navigation system that directs the user through the building. Such basic service does not require users to authenticate themselves but does require the user to have a valid IP address. The second enhanced service model is used to generate revenues by charging the user according to a differentiated charging plan based on the level of service the user opts for. This model includes services like Internet access, location-based buddy lists, notification of on-going sales according to pre-defined preferences, etc., all of which require the user to authenticate herself to the network.

The CHOICE network technology that we have built, deployed, tested and described in this paper supports both of these service models. Furthermore, CHOICE assures businesses a nominal level of security. We believe that customers are concerned about the quality and security of

their connections, while businesses are concerned about free-riders who are pilfering bandwidth from their real customers. Businesses would rather not offer any wireless connectivity than offer a connectivity that is low quality, insecure, and unreliable because then they risk losing a lot more angry, frustrated, and dissatisfied customers who experience poor connections while in their store. CHOICE is designed to avoid this kind of situation by giving business owners a degree of control over how they allocate and manage their network resources and maximize their customer's satisfaction. It is a self-contained software solution that is easy to install and maintain, another important consideration for business owners.

In summary, we claim that PAWNs with CHOICE makes a compelling story from a business perspective since everyone involved with it benefits. In particular, end-users benefit because it gives them a viable choice in how they access the Internet while giving them options for changing the amount of bandwidth and services they get from the network. The hardware vendors benefit, as they are able to sell more wireless hardware to businesses. The service providers benefit as their resources are bought and used. The building owners benefit as they use PAWNs to stay competitive while attracting more visitors or customers to their buildings. With CHOICE they can offer additional personalized services either free or for a nominal price to their customers. Finally, the software vendors benefit as they can sell new types of functionality over this network. For all these reasons we feel that the business case for the deploying a PAWN with CHOICE is a strong one and it is in the interest of the public building administrators to deploy this technology in their buildings.

5 Summary

To further realize the vision of pervasive, ubiquitous computing, we must extend high-speed network connectivity beyond private networks into public places. In this article, we motivate the use of wireless LANs as a basis for extending network connectivity to public places. However, using wireless LANs to implement public-area wireless networks (PAWNs) raises a number of interesting challenges that have to be overcome to make such networks practical. We discussed these issues, and described a PAWN that we have designed, implemented, and deployed called CHOICE that addresses these challenges in one integrated system. Specifically we discussed:

- Service models: The network can provide various types of service to users. CHOICE supports a wide range of service models, from free access to local services to paid connectivity, possibly with quality of service guarantees.
- Authentication: The network must authenticate users to authorize access under a given service model. CHOICE supports multiple authentication modes like E-cash

systems, credit cards, digital certificates, etc. Once authenticated, CHOICE returns a security token to the user's mobile device for identification.

- Access enforcement: The network must ensure that users can only access network connectivity and services for which they are authorized. CHOICE uses per-packet filtering at the access points to identify packets from authorized users and block unauthorized access.
- Policy enforcement: Depending upon the service model, the network might have to provide guarantees, such as minimal bandwidth. CHOICE uses the network admission server to centrally control bandwidth allocation, and leverages the per-packet filtering mechanism to enforce service policies.
- Billing and accounting: To support service models with charges, the network must account and bill for access and service. Once again, CHOICE leverages the per-packet filtering mechanism to maintain fine-grained accounting of network utilization by users.
- Security and privacy: Since it is a shared, public medium, the network must support data integrity and security. CHOICE supports three levels of security negotiated during authentication: encryption of the security token alone, the packet headers, or the entire packet.
- Location services: To support location-sensitive applications, the network must be able to determine user locations and disseminate location information to applications on the mobile device. CHOICE uses RF signal intensity maps of overlapping base stations, combined with predictive heuristics, to determine user location.

CHOICE has been available to users for over a year now and it has successfully demonstrated that wireless LANs are a viable and compelling technology for providing high-performance wireless Internet access and personalized location services in public places.

References

- [1] IEEE 802.11b/D3.0, "Wireless LAN Medium Access Control (MAC) and Physical (PHY) Layer Specification: High Speed Physical Layer Extensions in the 2.4 GHz Band", 1999.
- [2] R. V. Nee, G. Awater, M. Morikura, H. Takanashi, M. Webster, and K. W. Halford, "New high-rate wireless LAN standards," IEEE Communications Magazine, Vol. 37, no. 12, pp. 82-88, Dec. 1999.
- [3] The CHOICE Network Project, Sept. 1999, <http://www.mschoice.com/>
- [4] IETF Working Group on IP Routing for Wireless/Mobile Hosts,

- <http://www.ietf.org/html.charters/mobileip-charter.html>
- [5] A. Harter and A. Hopper, "A new location technique for the Active Office," *IEEE Personal Communications* 4, 5, Oct. 1997.
 - [6] R. Want, A. Hopper, V. Falcao, and J. Gibbons, "The Active Badge Location System," *ACM Transactions on Information Systems*, Vol. 40, No. 1, pp. 91-102, January 1992.
 - [7] S. Y. Seidel, and T. S. Rapport, "914 MHz path loss prediction Model for Indoor Wireless Communications in Multi-floored buildings," *IEEE Trans. on Antennas & Propagation*, February 1992.
 - [8] P. Bahl and V. N. Padmanabhan, "RADAR: An In-Building RF-based User Location and Tracking System", In *Proc. IEEE Infocom 2000*, April 2000.
 - [9] P. Bahl, V. N. Padmanabhan, and A. Balachandran, "A Software System for Locating Mobile Users: Design, Evaluation, and Lessons," *MSR-TR-2000-12*, Feb 2000.
 - [10] D. Estrin, J. C. Mogul, and G. Tsudik. "VISA Protocols for Controlling Inter-Organization Datagram Flow," *IEEE Journal on Selected Areas in Communication* 7(4), 486-498, May 1989.
 - [11] IEEE Draft P802.1x/D1, "Port Based Network Access Control," September 1999.
 - [12] R. Atkinson, "Security architecture for the Internet Protocol", *IETF RFC 2401*, Nov. 1998.
 - [13] P. Bahl, A. Balachandran, and S. Venkatachary, "The CHOICE Network – Broadband Wireless Internet Access in Public Places", *MSR-TR-2000-21*, Feb. 2000. (Also in *Proc. IEEE ICC 2001*)
 - [14] N. H. Vaidya, P. Bahl, and S. Gupta, "Distributed Fair Scheduling in a Wireless LAN", In *Proc. ACM MobiCom 2000*, July 2000.
 - [15] B. Aboba, J. Arkko, and D. Harrington, "Introduction to Accounting Management, *IETF RFC 2975*, Oct 2000.
 - [16] A. Miu and P. Bahl, "Dynamic Host Configuration for Managing Mobility between Public and Private Networks, In *Proc. Usenix USITS 2001*, March 2001.
 - [17] MSN Messenger Service, <http://messenger.msn.com>
 - [18] AOL Instant Messenger, <http://www.aol.com/aim/>
 - [19] Yi-Min Wang, P. Bahl, and W. Russell, "The SIMBA User Alert Service Architecture for Dependable Alert Delivery," *International Conference on Dependable Systems and Networks*, July 2000.