Design and Implementation of a WLAN/CDMA2000 Interworking Architecture

Milind M. Buddhikot, Girish Chandranmenon, Seungjae Han, Yui-Wah Lee, Scott Miller, and Luca Salgarelli, Bell Laboratories, Lucent Technologies

ABSTRACT

The combination of 3G and WLAN wireless technologies offers the possibility of achieving anywhere, anytime Internet access, bringing benefits to both end users and service providers. In this article we discuss interworking architectures for providing integrated service capability across widely deployed 3G CDMA2000-based and 802.11-based networks. Specifically, we present two design choices for integration: *tightly coupled* and *loosely coupled*, and recommend the latter as a preferred option. We describe in detail the implementation of a loosely coupled integrated network, which provides two kinds of roaming services, *SimpleIP service* and *Mobile-IP service*. We present in detail two new components used to build these services: a network element called a *WLAN integration gateway* deployed in WLAN networks and a client software on the mobile device. For a mobile device with interfaces to both technologies, our system supports seamless handoff in the presence of overlapping radio coverage.

INTRODUCTION

The Internet has emerged as an all-pervasive technology that continues to experience tremendous growth and popularity. Our ever increasing dependence on this technology has sparked interest in making it truly ubiquitous — available *anywhere*, *anytime*. Recent advances in wireless technologies will help realize this vision. To this end, there are two technologies gaining momentum. *Wireless local area networks* (WLANs) based on the IEEE 802.11 standards [1], also known as WiFi, are popular in enterprise networks, homes, and public hotspots such as airports and hotels. WLAN enables wireless networks that support data rates of 1–54 Mb/s over small areas of a few thousand square meters. *Wireless wide area networks*, based on third-generation (3G) standards such as code-division multiple access 2000 (CDMA2000) [2], on the other hand, support peak rates from 144 kb/s to 2.4 Mb/s and offer connectivity over a wide area of the order of several square kilometers.

Given the complementary characteristics of WLAN (faster short-distance access) and CDMA2000 (slower long-range access), it is compelling to combine them to provide ubiquitous wireless access. Such integration can bring significant benefits to service providers and end users. It will allow CDMA2000 service providers to economically offload data traffic from wide-area wireless spectrum to WLANs in indoor locations, hotspots, and other areas with high user density. Providing WLAN hotspot access as a value-added service can increase their customer base. For WLAN service providers, integration will bring them a larger user base from partner CDMA2000 networks, without having to win them through per-customer service contracts. Also, the customers will benefit from enhanced performance through the greater coverage, higher data rate, and lower overall cost of such a combined service.

Figure 1 illustrates a conceptual view of the integrated public wireless networks that will offer such a service. End user devices such as laptops, palmtops, and phones that can access networks based on both technologies are already becoming available. A user of this integrated network would prefer to have exactly *one service subscription* with *one service provider*, typically called its *home network provider*. When a user subscribes to such a service, credentials in the form of authentication information (e.g., shared secret keys), profile information (e.g., class of service, minimum bandwidth), and accounting information will be stored in a network-based authentication, authorization, and accounting (AAA) server called the *home AAA*. This single account will enable the user to access data and voice ser-

0163-6804/03/$17.00 © 2003 IEEE

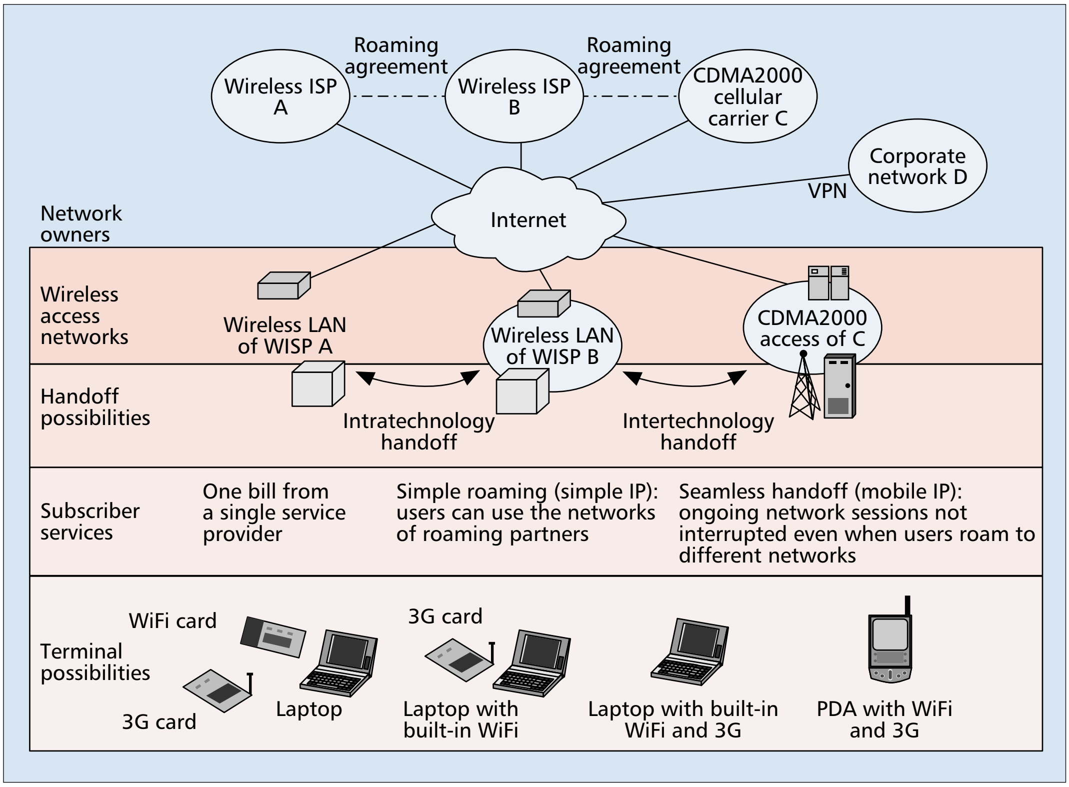

■ **Figure 1.** *CDMA2000/WLAN integration: the big picture.*

The PDSN in the architecture aggregates data traffic from multiple RNCs and interfaces the RAN to a packet switched network. The PDSN terminates a PPP connection and maintains session state for each mobile node in its serving area.

vices anywhere,any time, receive exactly *one billing statement*, *roam* freely among all networks with which the user's provider has agreements, and get similar *quality of service*. Using a single account on different networks in this way requires that network providers be able to authenticate each other's users and obtain their service profile parameters. This is enabled by roaming agreements established among service providers using AAA protocols, such as RADIUS [3] or DIAMETER [4], and AAA broker networks.

The emerging integrated public wireless networks will offer two roaming services: simple IP service and Mobile IP service. Simple IP service offers integrated billing and subscriber profiles, but does not guarantee session continuity across network boundaries. Mobile IP service additionally enables seamless handoffs between networks to preserve ongoing sessions.

The goal of this article is to first describe design options available to build an integrated network, and then present details of design and implementation of the abovementioned two integrated services using the preferred approach. Specifically, we discuss two possible approaches, tightly coupled and loosely coupled, and advocate the latter as the preferred approach. We describe our implementation of the two services in the loosely coupled architecture, which includes two new components: a network element called an Integration of Two Access Technologies (IOTA) gateway and a service access software on the mobile device. We primarily focus on integration of WLAN and 3GPP2 CDMA2000 networks. However, the loose integration architecture and design choices for the integration gateway and client software described here apply to integration of WLAN and 3GPP UMTS networks as well.

CDMA2000 AND WLAN BACKGROUND

In this section we will provide a brief background on the architecture of CDMA2000 and 802.11 WLAN networks.

OVERVIEW OF CDMA2000 NETWORK

Figure 2a illustrates the basic architecture of 3G-1x and 1xEV-DO CDMA2000 networks. The radio access network (RAN) in CDMA2000 networks consists of multiple base stations (BSs) each connected to a radio network controller (RNC) by T1/T3 links. The RNC manages several concurrent Radio Link Protocol (RLP) layer 2 sessions with mobile nodes (MNs) and performs per-link bandwidth management functions. The 144 kb/s per carrier throughput in 3G-1x is shared among multiple active MNs, though at any given instant, a single MN may be allocated full data rate. When an MN moves from one RNC to the other, the on-going RLP session is torn down and a new session is established with the visited RNC.

The packet data serving node (PDSN) in the architecture aggregates data traffic from multiple RNCs and interfaces the RAN to a packet-switched network. The PDSN terminates a Point-to-Point Protocol (PPP) connection (Fig. 2a) and maintains session state for each MN in its serving area. PPP header and payload compression can be negotiated between the PDSN

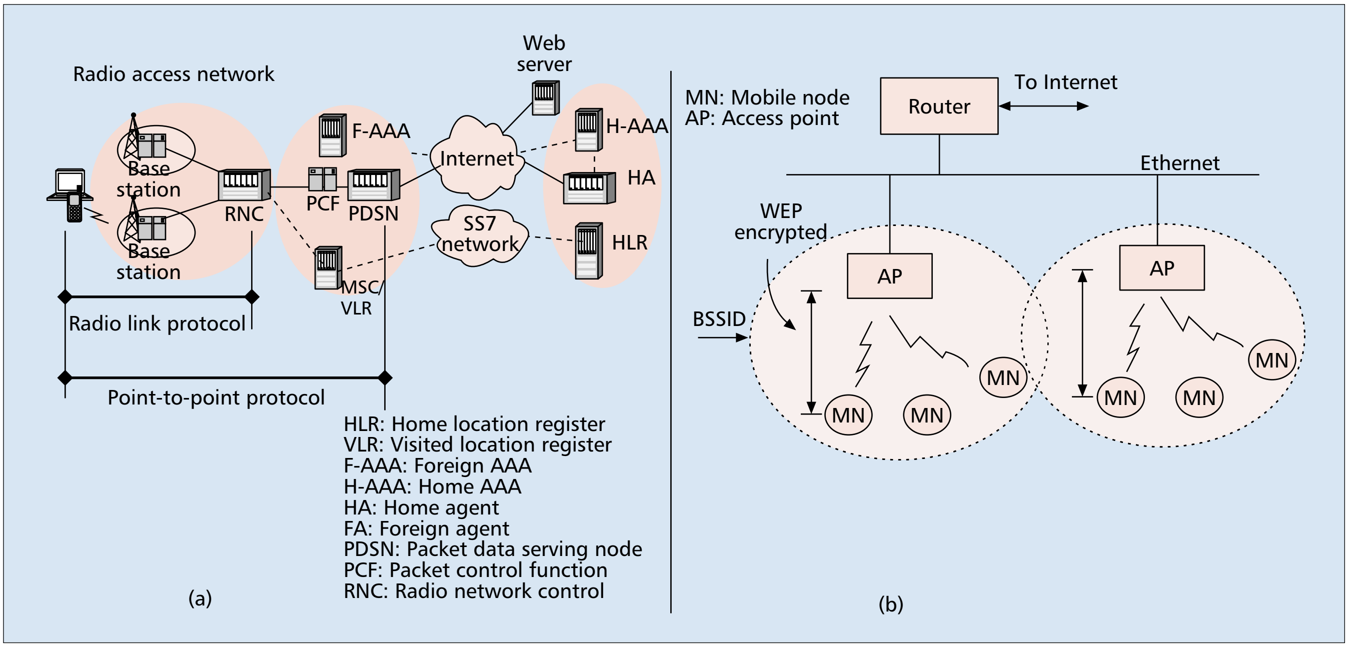

■ **Figure 2.** *Basic architecture of a) a CDMA2000 network and b) an 802.11 network.*

and the MN. The hierarchical architecture and the radio access protocols of CDMA2000 enables mobility within the serving area of the PDSN, by keeping PPP connections alive.

The PDSN is required to support two modes of IP operation: simple IP and mobile IP. In simple IP mode, if the MN moves from one PDSN to another, the PPP connection must be reestablished, and a new IP address is acquired. This requires the user to reestablish all their data sessions. In mobile IP mode, the PDSN implements the foreign agent (FA) functionality defined in Mobile IP [5], allowing cross-PDSN mobility.

From a data networking point of view, the 1xEV-DO architecture is similar to that of 3G-1x, in that it uses PPP between the MN and the PDSN, and provides mobility within the serving area of the PDSN. However, 1xEV-DO offers data-only service with up to 2.4 Mb/s downstream bandwidth, and relies solely on the AAA server for authentication.

Overview of WLAN 802.11

The WLANs standardized in the IEEE 802.11 standards [1] support two modes of operation: infrastructure mode and ad hoc mode. Of these two modes, infrastructure mode network architecture, illustrated in Fig. 2b, resembles the wide-area cellular networks and is of most interest to wireless Internet service providers (WISPs). It consists of an access point (AP) that performs three functions:

- It implements one or more of the 802.11 radio interface protocols, frequency hopping spread-spectrum (FHSS), direct sequence spread-spectrum (DSSS) or orthogonal frequency-division multiplex (OFDM).
- It implements the carrier sense multiple access with collision avoidance (CSMA/CA) medium access control (MAC) protocol and performs packet store-and-forward to coordinate communications of MNs in the cell characterized by a basic service set ID (BSSID).
- It interfaces the cell to a packet-switched network such as Ethernet, and therefore implements layer 2 packet forwarding functions such as bridging.

To access the 802.11 LAN, the MN first authenticates to the AP and then associates with it to obtain an association identifier. The packet transmissions between the AP and the MN can be optionally protected using a symmetric key-based RC4-based encryption called Wired Equivalency Privacy (WEP). The basic authentication and encryption mechanisms in the 802.11 standard have been shown to be inadequate to meet the design goals of confidentiality, integrity, and access control. The newer 802.11i standard employs 802.1x Port Based Access Control protocol [6, 7] for MN authentication, Temporal Key Integrity Protocol (TKIP) for dynamic rekeying of encryption keys, and optional Wireless Robust Authentication Protocol (WRAP) that employs AES encryption to eliminate these limitations.

Architectural Choices

The WLAN and 3G integration architecture is characterized by the amount of interdependence it introduces between the two component networks. Two candidate integration architectures, *tightly coupled* and *loosely coupled interworking* have been considered in the research and standards communities [8, 9].

Tightly-coupled Interworking

The rationale behind the tightly coupled approach is to make the WLAN network appear to the 3G core network as another 3G access network. The WLAN network would emulate functions that are natively available in 3G radio access networks. In this architecture, utilized by WISP no. 1 in Fig. 3, the WLAN gateway network element introduced to achieve integration appears to the upstream 3G core as a packet control function (PCF) in the case of a CDMA2000 core network. The WLAN gateway hides the details of the WLAN network to the 3G core, and implements all the 3G protocols (mobility management, authentication, etc.) required in a 3G radio access network. MNs in this approach are required to implement the corresponding 3G protocol stack on top of their standard WLAN network cards, and switch from one physical layer to the next as needed. All the traffic generated by clients in the WLAN network is injected using 3G protocols into the 3G core network. These networks would share the same authentication, signaling, transport, and billing infrastructures, independent from the protocols used at the physical layer on the radio interface.

However, this approach presents several disadvantages. Since the 3G core network directly exposes its interfaces to the WLAN network, the same operator will typically be required to own both the WLAN and the 3G network. In fact, in this case, independently operated WLAN islands could not be integrated with 3G networks without explicit physical connectivity to the 3G core network. Today's 3G networks are being deployed using carefully engineered network planning tools, and the capacity and configuration of each network element is calculated using mechanisms that are very specific to the technology utilized over the air interface. By injecting the WLAN traffic directly into the 3G core, the setup of the entire network, as well as the configuration and design of network elements such as PDSNs, have to be modified to sustain the increased load.

The configuration of the client devices also presents several issues with this approach. First, as described earlier, the WLAN cards would need to implement the 3G protocol stack. It would also mandate the use of 3G-specific authentication mechanisms for authentication on WLANs, forcing WLAN providers to interconnect to the 3G carriers' SS7 network to perform authentication procedures. This would also imply the use of WLAN cards with built-in 3G credentials.

For the reasons described above, the complexity and high cost of reconfiguration of the 3G core networks and WLAN gateways would force operators that chose the tightly coupled approach to become uncompetitive to WLAN-only ISPs.

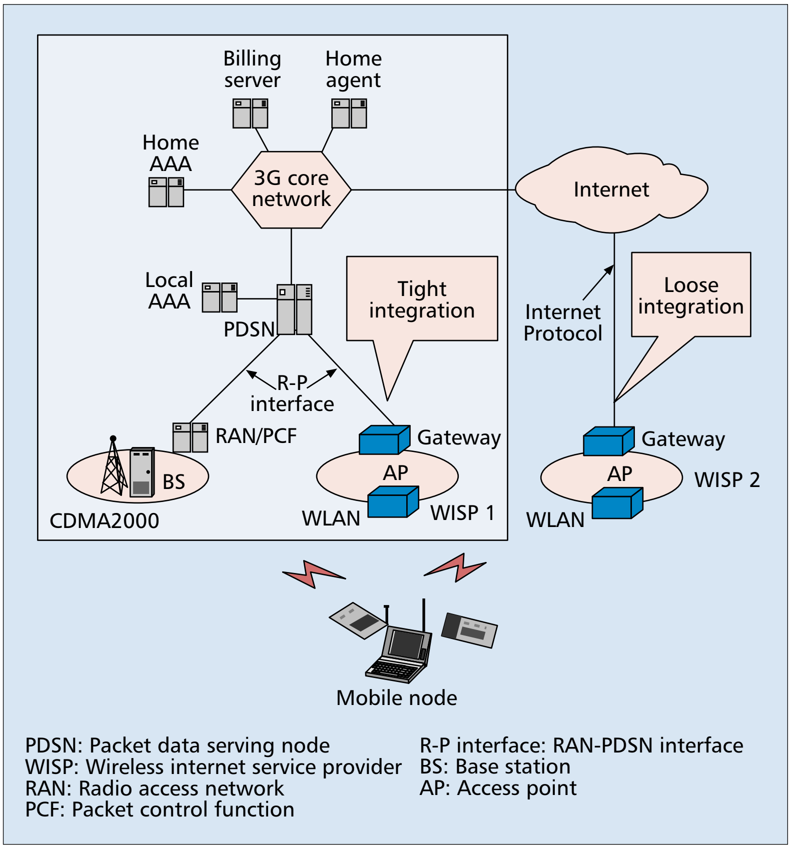

Figure 3. *3G and WLAN integration: tightly coupled vs. loosely coupled architectures.*

Loosely Coupled Interworking

Like the previous architecture, the loosely coupled approach calls for the introduction of a new element in the WLAN network, the WLAN gateway. However, in this design (WISP no. 2 in Fig. 3), the gateway connects to the Internet and does not have any direct link to 3G network elements such as PDSNs or 3G core network switches. The user population that accesses services of the WLAN gateway may include users that have locally signed on, as well as mobile users visiting from other networks. We call this approach loosely coupled interworking because it completely separates the data paths in WLAN and 3G networks. The high-speed WLAN data traffic is never injected into the 3G core network, but the end user still experiences seamless access.

In this approach, different mechanisms and protocols can handle authentication, billing, and mobility management in the 3G and WLAN portions of the network. However, for seamless operation to be possible, they have to interoperate. In the case of interoperation with CDMA2000, this requires that the WLAN gateway support Mobile IP functionalities to handle mobility across networks, as well as AAA services to interwork with the 3G's home network AAA servers. This will enable the 3G provider to collect the WLAN accounting records and generate a unified billing statement indicating usage and various price schemes for both (3G and WLAN) networks. At the same time, the use of compatible AAA services on the two networks would allow the WLAN gateway to dynamically obtain per-user service policies from their home AAA servers, and to enforce and adapt such policies to the WLAN network.

There are several advantages to the loosely coupled integration approach. First, it allows independent deployment and traffic engineering of WLAN and 3G networks. 3G carriers can benefit from other providers' WLAN deployments without extensive capital investments. At the same time, they can continue to deploy 3G

It should be clear that the loosely-coupled approach offers several architectural advantages over the tightly-coupled approach, with virtually no drawbacks. Therefore, it has emerged as a preferred architecture for the integration of WLAN with 3G networks.

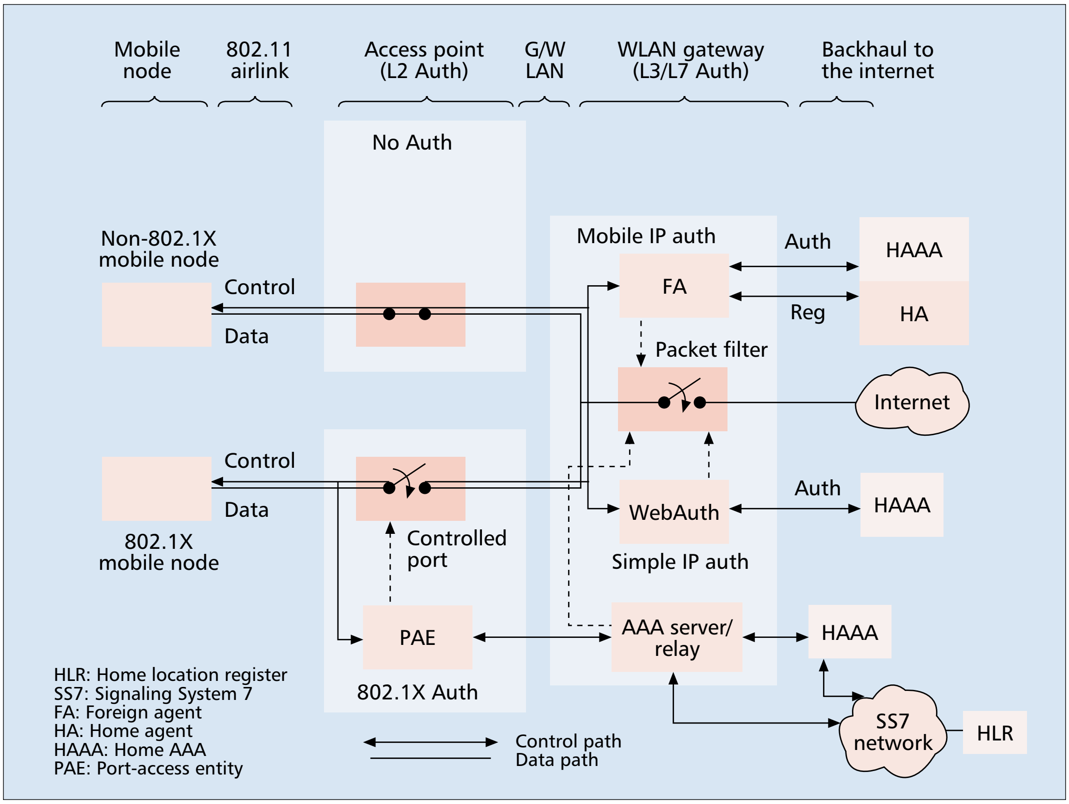

Figure 4. *The authentication model for a WLAN gateway.*

networks using well established engineering techniques and tools. Furthermore, while roaming agreements with many partners can result in widespread coverage, including key hotspot areas, subscribers benefit from having just one service provider for all network access. They no longer need to establish separate accounts with providers in different regions, or covering different access technologies. Finally, unlike the tightly coupled approach, this architecture allows a WISP to provide its own public WLAN hotspot, interoperate through roaming agreements with public WLAN and 3G service providers, or manage a privately installed enterprise WLAN.

It should be clear that the loosely coupled approach offers several architectural advantages over the tightly coupled approach, with virtually no drawbacks. Therefore, it has emerged as a preferred architecture for the integration of WLAN with 3G networks, and we will use it as a reference throughout the rest of the article.

Authentication and Privacy

A WLAN gateway should provide Internet access to only legitimate users, and therefore must support user authentication at one or more protocol layers. In the WLAN link layer, three authentication and/or access control methods are possible:

- *Static filtering based on MAC address*: In this method, WLAN APs drop traffic of all hosts except those of certain preconfigured network devices. Typically filtering rules are specified using the layer 2 address (a.k.a. MAC or hardware address) of the network device.
- *WEP of the 802.11b standard [1]*: In this method, WLAN APs verify that the end host knows a shared secret in the form of a 40- or 104-bit WEP key, which is used for all network devices accessing the same AP.
- *The 802.11i standard [7]*: 802.11i is a newer standard for access control that allows dynamic per-user per-session authentication and encryption keys and stronger packet encryption.

The first two methods are not suitable for use in a public environment, whereas the third method is not backward compatible with legacy APs and MNs that do not have 802.11i support.

In a public environment with dynamic user population, exhaustive and static configuration of all APs with a list of MAC addresses is infeasible.

The main problem with WEP is that the same key is shared by all users using the same AP. In public environments, it is very difficult to securely distribute and revoke this key for a dynamic user population. Furthermore, since the same key is also used for encryption, all authenticated users can snoop on each other's traffic. Apart from this problem, there are well-known attacks on the flawed WEP encryption algorithm [10].

802.11i is considered a significant improvement for the public environment. It employs the IEEE 802.1x port access control standard that specifies the use of Extensible Authentication Protocol (EAP over LAN (EAPOL) [6] between the MN and AP to perform per-session user authentication. EAPOL encapsulates EAP packets, which an AP can transfer to a service provider home AAA using the RADIUS AAA protocol. This allows the use of any well-known EAP schemes such as EAP-TLS [11], EAP-SIM [12], EAP-AKA [13], or EAP-SKE [14] to authenticate a MN. Additionally, individual per-user session keys, used for encryption and integrity protection, are derived and distributed during the authentication exchange with the home AAA server. This eliminates the need for any preconfiguration of keys and MAC addresses in WLAN APs, and only requires a security association between the user and its home service provider. The 802.11i standard also specifies TKIP that defines a key derivation procedure to derive encryption, authentication, and integrity protection keys and a WEP-compatible encryption enhancement to fix known flaws in WEP. The TKIP specification improves WEP authentication and encryption to acceptable levels and provides graceful migration of existing infrastructure and client devices. The 802.11i standard also describes an optional Wireless Robust Authentication Protocol (WRAP) that uses strong 128-bit AES encryption, which is attractive in the long term.

Figure 4 illustrates the authentication model a WLAN integration gateway should implement to support authentication at layers 2, 3, and 7. The model relies on the dynamic packet filters that use information that includes MAC address, IP source and destination address, and TCP/UDP source and destination ports. The dynamic filters are updated based on the status of user authentication.

The authentication path and the corresponding dynamic packet filters used depend on the service mode. For mobile IP mode, the authentication is done as part of the Mobile IP registration, in which the MN registers through the FA to the home agent (HA). During the registration, the MN presents to the FA evidence that it knows the MN AAA key, which is a shared secret between the MN and the home AAA (H-AAA). Until the registration succeeds, the FA inserts packet filters that block all other MN traffic.

For simple IP mode, the MN's authentication procedure is triggered by the first Web access of the user. The HTTP access will be intercepted by the packet filter, and it will be redirected to a Web authenticator in the gateway. The authenticator presents to the user a secured login page over an HTTPS connection instead of the original Web page the user requested. The user enters her username and password to login. The authenticator authenticates the user by consulting the H-AAA.

In our model, a non-802.11i MN can connect through the AP without any layer 2 authentication. However, it cannot connect to the Internet unless it has successfully authenticated with the gateway.

On the other hand, an 802.11i-capable MN needs to authenticate with both the AP and the gateway for access to the Internet. Note that in simple IP mode, layer 3 authentication is redundant and can be eliminated to allow single logon for a user if the integration gateway monitors layer 2 authentication traffic and unblocks MN traffic on successful authentication.

Note that certain EAP schemes such as EAP-SIM and EAP-AKA rely on SIM/USIM cards on the MN and corresponding credentials stored at a home location register (HLR) on a Signaling System No. 7 (SS7) network. In this case, either the gateway or the H-AAA must interface to an SS7 network to communicate with the HLR.

One important final remark about the WLAN integration gateway is that it may support IPSec or SSL virtual private networks (VPNs) and provide fast encryption to support privacy at layer 3 and above. This may be important if the 802.11 WLAN is operated without layer 2 encryption and authentication, and only layer 3 or higher authentication is employed. Note that Web authentication takes place over an HTTPS connection, and therefore the username and password information cannot be snooped by a malicious user in the WLAN network. Similarly, MIP registration has built-in replay protection. However, in both cases, lack of layer 2 encryption allows for MAC and IP address spoofing between successful authentications.

In our model, a non-802.11i mobile node can connect through the access point without any layer-2 authentication. However, it cannot connect to the Internet unless it has successfully authenticated with the gateway.

Two Integrated Services

In this section we describe the basic simple IP service and a more advanced mobile IP service.

Simple IP Service

In case of 802.11i-enabled deployments, the client first performs layer 2 authentication, and then requests an IP address from the local WLAN integration gateway. If the authentication is successful, the gateway provides a private or public IP address through DHCP and configures appropriate features such as QoS guarantees and Network Address Translation (NAT). Optionally, for non-802.11i deployments, Web-based authentication can be used.

Note that in the event of mobility, the user acquires a new IP address, and ongoing sessions are lost. Therefore, simple IP service is most appropriate for environments with limited mobility where layer 2 mobility mechanisms satisfy mobility needs. One key advantage of this service is that it does not need specialized client software for service access.

Mobile IP Service

The goal of mobile IP service is to preserve user sessions when a user roams among heterogeneous networks of different providers with different access technologies. We employ two basic ideas to achieve this goal:

- Use of Mobile IP in the WLAN gateway
- Intelligent interface selection at the client in the presence of overlapped coverage between CDMA2000 and WLAN networks

We elaborate on these ideas in detail below.

IP does not provide native support for mobility; when a host moves and attaches to a different physical network, its IP address must change, forcing all transport-level sessions with any Internet hosts to break. Mobile IP [5], standardized in the Internet Engineering Task Force (IETF), addresses this problem; it allows an Internet host to keep a fixed address called a *home address* (HoA). To deliver packets to the current point of attachment, Mobile IP employs

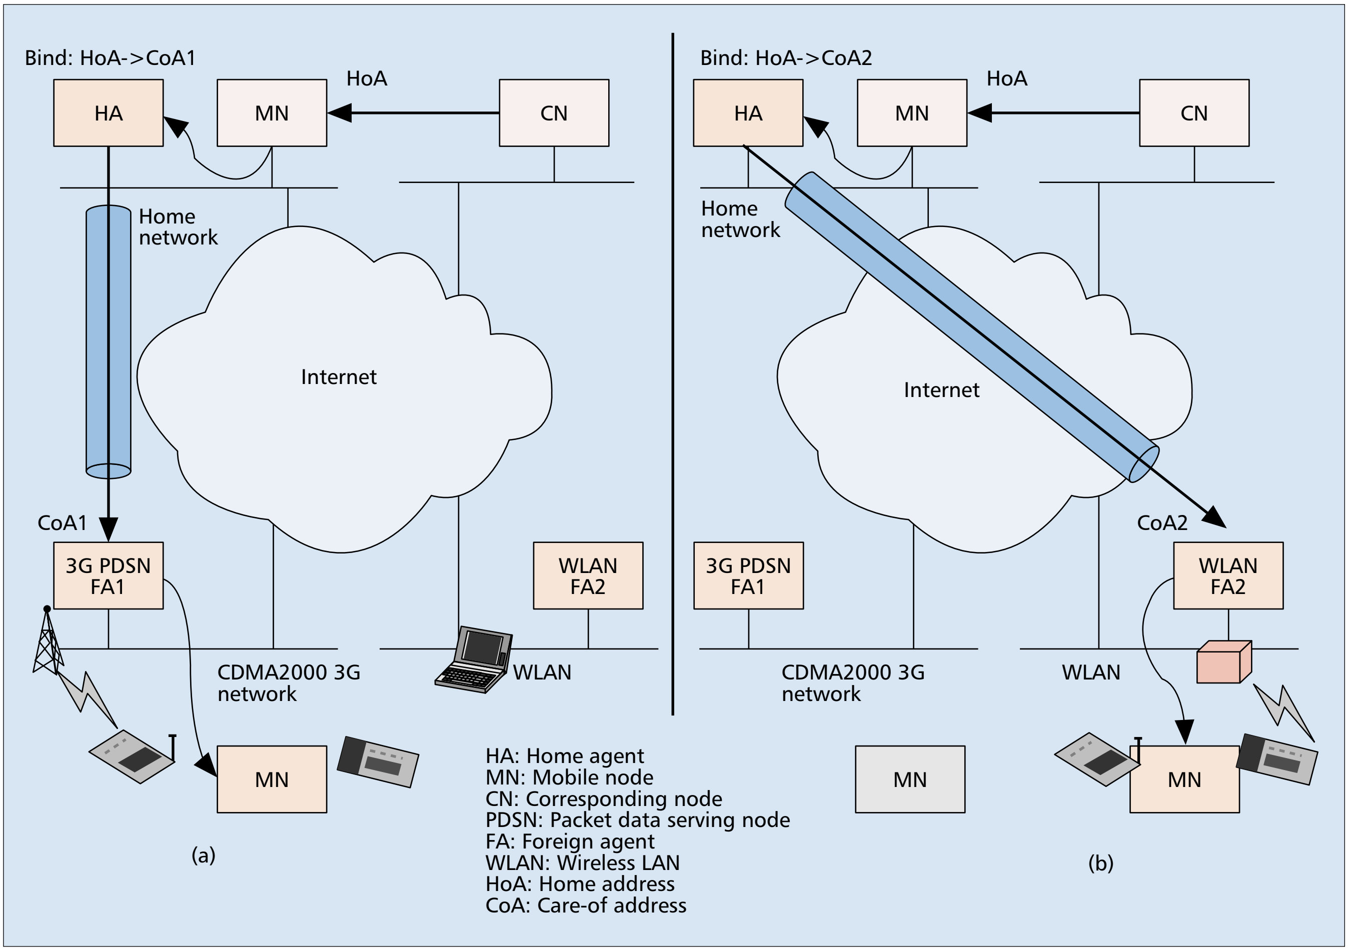

■ **Figure 5.** *Using Mobile IP for intertechnology handoff: a) MN at a 3G network; b) MN at a WLAN.*

two network elements: a home agent (HA) in the home network and a foreign agent (FA) in the visited network. When in the foreign network, an MN discovers a local FA and registers the address of the FA as a care-of address (CoA) with its HA, creating the binding Bind: HoA → CoA (Fig. 5a). The HA intercepts a packet from any correspondent node (CN) to the MN, encapsulates it in another IP packet, and tunnels it to the FA. The FA decapsulates and delivers the original packet to the MN. Since the MN maintains its HoA, all its transport protocol sessions are preserved.

The CDMA2000 standard incorporates Mobile IP to achieve inter-PDSN handoff. The PDSN node in the CDMA2000 network implements the FA.

Therefore, a natural way to implement intertechnology handoff between CDMA2000 and WLAN networks is to implement Mobile IP functionality in the WLAN gateway and MN. The scenario is depicted in Fig. 5b.

The MN performs session handoffs in two cases: when it loses signal on the wireless link currently in use or when it finds a better wireless link that can provide better performance. For example, it will switch from 3G to WLAN when it acquires WLAN, and switch from WLAN to 3G when it loses WLAN signal. To avoid service disruption and packet loss during service handoff, the MN can exploit any overlapped 3G and WLAN coverage. It can keep both network interfaces active. While using the current network link (e.g., 3G), it can use the noncurrent network link (e.g., WLAN) to prepare a handoff in the background. Figure 6a depicts this scenario. Specifically, it shows the WLAN signal observed by the client over time. At t_1, when the signal strength exceeds the threshold, H, the client will attempt to use the WLAN airlink. Similarly at time t_2, when the signal strength drops below the threshold L, the client will revert to the 3G airlink. Two thresholds, H and L, are used to avoid unnecessary handoffs that can result in poor connection. Switching to a different airlink involves several steps: discovery of a local FA, Mobile IP registration with the FA over the new airlink, creation of new tunnels at the HA, and setting up a packet filter in the gateway. If the client completes these steps before losing the signal on the current interface, the delays with these steps (indicated by Δ_1 and Δ_2 in Fig. 6) can be masked, and handoff will appear instantaneous to the client application. For incoming traffic, packets that are in the network will continue to arrive at the old interface, to which the MN still listens. For outgoing traffic, the TCP/IP stack on the MN will start using the new airlink immediately after the atomic operation of airlink switching. As a result, packet loss due to handoff is minimized.

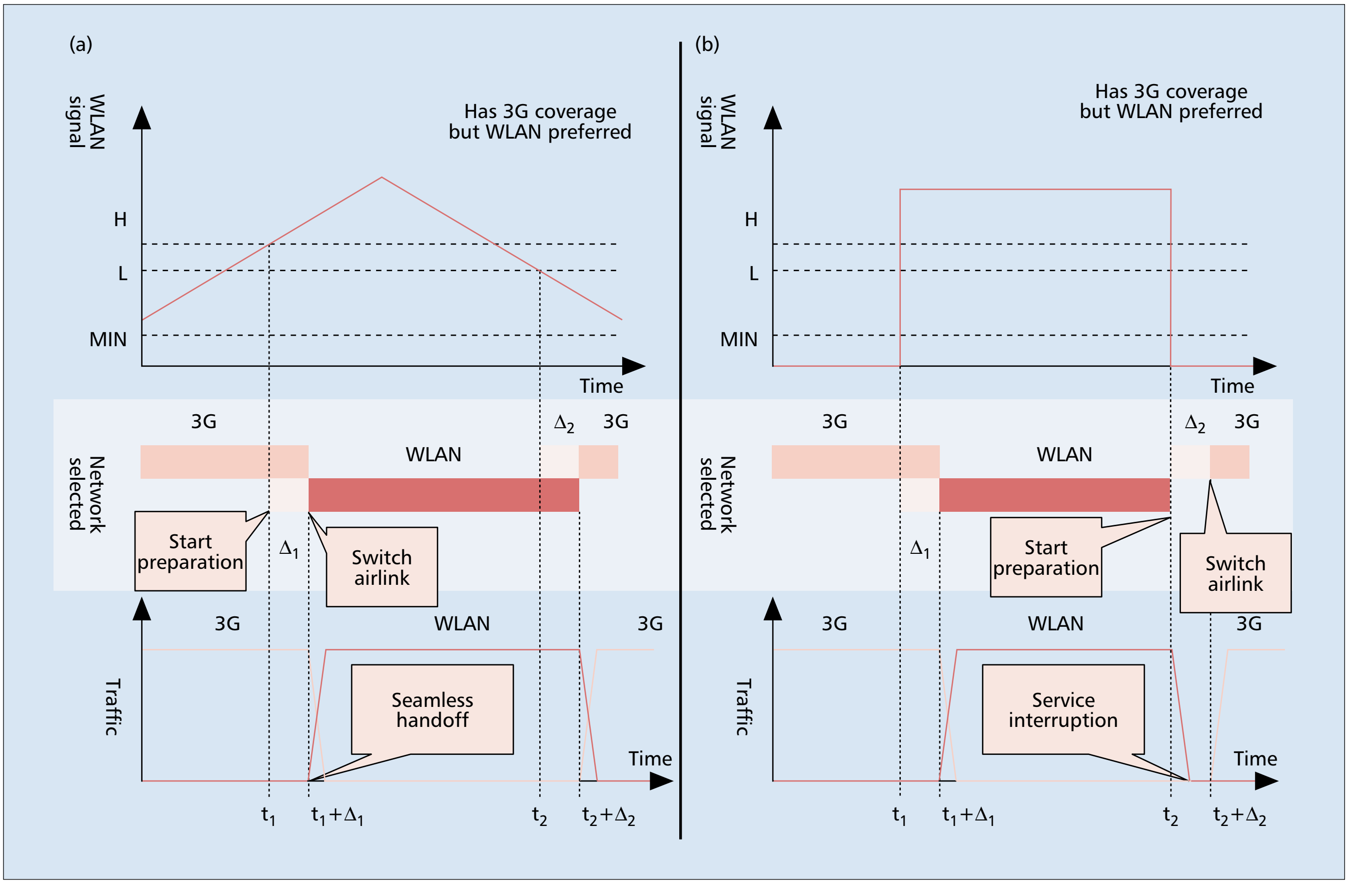

Figure 6. *Scenarios for a) overlapped and b) nonoverlapped radio coverage.*

Of course, in the absence of overlapped coverage, there will be service interruption and packet loss. This is illustrated in Fig. 6b, where the disruption occurs between t_2 and $t_2 + \Delta_2$. This incurred delay is determined by the performance of the gateway, HA, and HoA, as well as the network latency among them. Typically, Δ_2 is on the order tens to hundreds of milliseconds [9].

Note that the use of Mobile IP can worsen the performance of Web sessions in the presence of a Web cache outside the WLAN gateway. Figure 7a illustrates the case where requests from the client are transparently directed to a Web cache.

For a cache miss, the cache forwards the requests to the Web server and obtains a response. For a cache hit, the cache would already have the response in its own local disk. In either case, the cache would forward the response back to the users. In mobile IP service, requests coming from users would appear to have come from their HAs. Therefore, the cache would forward the response back to their home networks, where the HA would tunnel the response back to the gateway. As a result, while the cache was intended to reduce the traffic on the backhaul link, in this setup it would not eliminate any traffic even for cache hits. In fact, the presence of the cache would double the traffic volume on the backhaul for cache misses.

Figure 7b illustrates the scenario where the Web cache is an integral part of the WLAN gateway. Since the gateway is aware of the MN's presence and its use of mobile IP service, it instructs the cache to forward the Web response directly to the client.

The IOTA Implementation

Based on the loosely coupled architecture described earlier, we built a prototype system, IOTA, with two primary components: the integration gateway and the multi-interface mobility client. Another example implementation designed specifically for enterprise can be found in [15].

IOTA Gateway

The IOTA gateway integrates a number of subsystems, as shown in Fig. 8d: a RADIUS AAA proxy, a Mobile IP FA (MIP FA), a DHCP server, a dynamic firewall, a QoS module, an integrated Web cache, and an accounting module. These IOTA subsystems rely on an ondisk database to store persistent information about each client's session and accounting records. Depending on the hardware configuration, the gateway box can have a built-in 802.11b access point, or connect to external 802.11b AP(s) through a gateway LAN.

The IOTA gateway uses the in-kernel Linux `iptables` service to perform dynamic packet filtering, packet mangling, and NAT functions. User-space IOTA subsystems use the IOTA

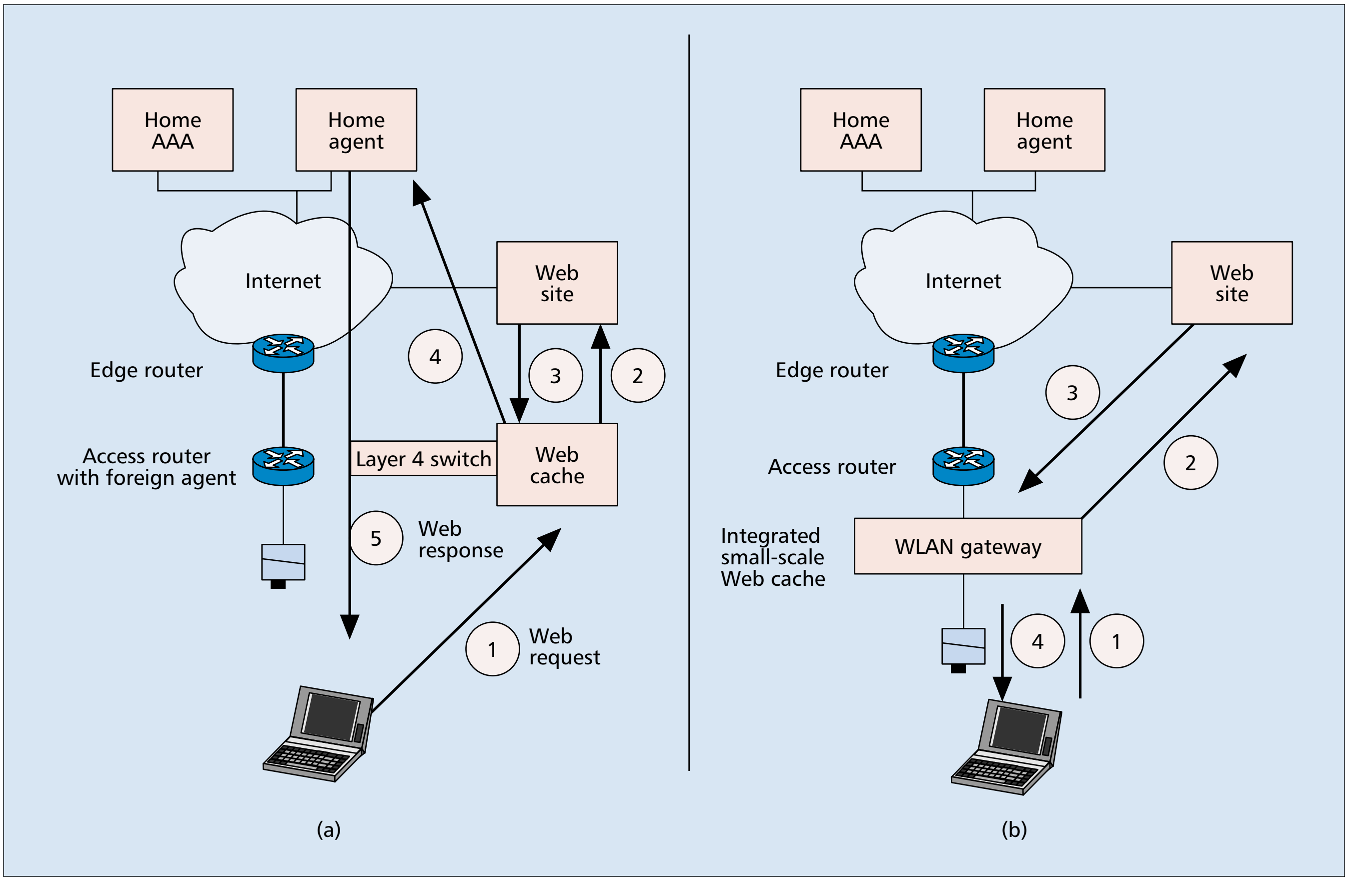

Figure 7. *The WLAN gateway integrates a Web cache, minimizing Mobile IP overheads: a) inefficient routing without integrated Web cache at the FA; b) direct routing with the WLAN gateway: integrated Web cache and FA.*

Packet Filter (IPF) library to interact with `iptables`. Dynamic packet filtering is primarily used to achieve controlled access to the Internet for wireless clients, but it also implements certain firewall functions to prevent attacks from malicious users. Dynamic packet mangling redirects unauthenticated simple IP users' Web request to the local Web authenticator, but it also redirects some other traffic such as DNS lookup traffic. The NAT function allows assignment of private IP addresses for wireless clients within the WLAN. These private IP addresses are not routable in the public Internet but alleviate the demand on publicly routable IP addresses. When packets of these clients travel out to the Internet, the NAT function will translate these private IP addresses to public IP addresses.

The IOTA gateway software runs on off-the-shelf computers running the Linux operating system (Fig. 8b).

Multi-Interface Mobility Client

Supporting the mobile IP service across WLAN and CDMA2000 networks requires a client software that can perform Mobile IP signaling with the FA and HA. Such a client must also intelligently select and sign the user onto the best access network depending on the network conditions. This latter feature is particularly useful in mobile IP mode.

We implement the multi-interface client software for Linux and Windows 2000/XP. There are three components for the software: a graphical user interface (GUI) and a mobility client in the user space, and a client driver in the kernel space. Our current implementation supports 802.11b, CDMA2000, and Ethernet interfaces. The GUI and software architecture for the client are shown in Figures 8a and 8c, respectively.

The mobility client detects the presence of new networks and initiates link-level associations: PPP connections for CDMA2000 networks and LAN associations for WLAN APs. It periodically scans all interfaces and measures observed signal strength. It uses an intelligent switching algorithm that accounts for signal strength and priority of different airlinks to avoid spurious network switching, often termed *bouncing*. It handles the intertechnology handoff between different networks with minimal packet loss using MIP. The client also provides WLAN specific functionality, such as supporting preferred network lists, WEP key configuration, and selection of WLAN APs based on signal strength. Our layering of IPsec over Mobile IP enables users who are signed onto an enterprise VPN to maintain their sessions while moving through the integrated network. The client GUI allows the user to monitor the status of the physical interfaces and configure the MIP profiles and network interfaces.

The multi-interface mobility client driver is implemented in the kernel below the network

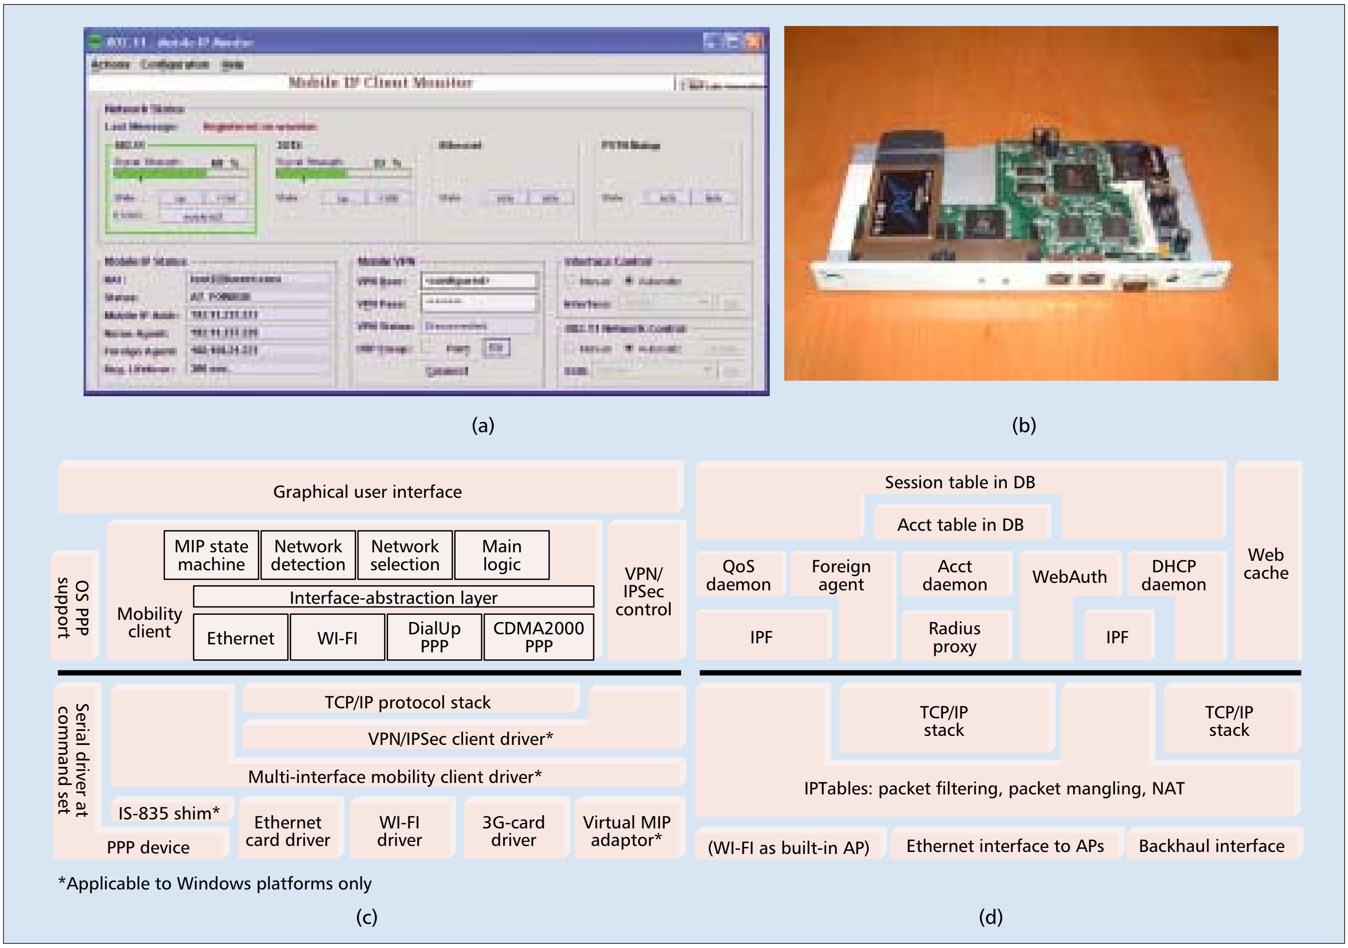

Figure 8. *IOTA implementation: a) IOTA client GUI; b) IOTA gateway; c) IOTA client architecture; d) IOTA gateway architecture.*

protocol stack and offers the abstraction of a single virtual nonmobile interface to the OS protocol stack.

Using these three components together, the client software provides an illusion to networking applications on the MN that the node is always on the same network, even though the node may actually be moving across network boundaries of different access technologies.

Conclusions

Integrated WLAN/CDMA2000 services will benefit both service providers and users. A loosely coupled network architecture that allows independent deployment and growth of each network will emerge as the preferred way to implement such services. Using Mobile IP and AAA protocols, a service provider can support the two access technologies with a single home infrastructure for authentication and mobility management, and allow interoperator roaming.

A typical implementation for loosely coupled architecture requires a WLAN integration gateway and mobility client software. The gateway supports simple IP and mobile IP services, and implements an authentication model that allows various layer 2 and layer 3 authentication schemes, preventing unauthorized users from accessing the public WLAN network. In the mobile IP mode of operation, the mobility client achieves seamless intertechnology handoffs without requiring user intervention.

We believe that the technologies described in this article may foster rapid deployment of integrated services and growth of ubiquitous high-speed wireless data.

Acknowledgments

We would like to thank Salim Virani, Dharani Vilwanathan, and Jian Cai for their contributions to the IOTA project. We would also like to thank Peretz Feder, David Benenati, and Reuven Shapira from Lucent Technologies for frequent insightful discussions on this topic and their help in laboratory setup for our prototype testing.

References

[1] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," ANSI/IEEE Std 802.11: 1999 (E) Part 11, ISO/IEC 880211, 1999.

[2] "TIA/EIA/IS-835B cdma2000 Wireless IP Network Standard," 3GPP2, 2000.

[3] C. Rigney *et al.*, "Remote Authentication Dial In User Service (RADIUS)," IETF RFC 2865, June 2000.

[4] P. Calhoun *et al.*, "Diameter Base Protocol," IETF Internet draft, Apr. 2002. draft-ietf-aaa-diameter-10.txt, work in progress.

[5] C. Perkins, Ed., "IP Mobility Support for IPv4," IETF RFC 3344, Aug. 2002.

[6] "Local and Metropolitan Area Networks: Standard for Port Based Network Access Control," Tech. rep., IEEE P802.1x, Jan. 2001.

We believe that the technologies described in this article may foster rapid deployment of integrated services and growth of ubiquitous high speed wireless data.

[7] "Part 11: Wireless MAC and Physical Layer Specifications: Specification of Enhanced Security," Tech. rep., IEEE P802.11i, Nov. 2002.

[8] A. Salkintzis, C. Fors, and R. Pazhyannur, "WLAN-GPRS Integration for Next-generation Mobile Data Networks," *IEEE Wireless Commun.*, Oct. 2002.

[9] M. Buddhikot *et al.*, "Integration of 802.11 and Third Generation Wireless Data Networks," *IEEE INFOCOM 2003*, Apr. 2003.

[10] N. Borisov, I. Goldberg, and D. Wagner, "Intercepting Mobile Communications: The Insecurity of 802.11," *7th Annual Int'l. Conf. Mobile Comp. and Net.*, July. 2001.

[11] B. Aboba and D. Simon, "PPP EAP TLS Authentication Protocol," IETF RFC 2716, Oct. 1999.

[12] H. Haverinen, "EAP SIM Authentication," IETF Internet draft, draft-haverinen-pppext-eap-sim-03.txt, Feb. 2002, work in progress.

[13] J. Arkko and H. Haverinen, "EAP AKA Authentication," IETF Internet draft, draft-arkko-pppext-eap-aka-03.txt, Feb. 2002, work in progress.

[14] L. Salgarelli *et al.*, "EAP SKE Authentication and Key Exchange Protocol," IETF Internet draft, Apr. 2002. draft-salgarellipppext-eap-ske-01.txt, work in progress.

[15] H. Luo *et al.*, "Integrating Wireless LAN and Cellular Data for Enterprise," *IEEE Internet Comp.*, Mar.–Apr. 2003, pp. 25–33.

BIOGRAPHIES

MILIND M. BUDDHIKOT (mbuddhikot@lucent.com) is a member of technical staff in the Center for Networking Research at Bell Laboratories, Lucent Technologies (Bell Labs-Lucent), Holmdel, New Jersey. He holds a D.Sc. in computer science (July 1998) from Washington University, St. Louis, Missouri, and an M.Tech. in communication engineering (December 1988) from the Indian Institute of Technology (IIT), Bombay. His current research interests are in the areas of systems and protocols for public wireless networks, authentication and dynamic key exchange, and multimedia messaging and caching. He has authored over 20 research papers and eight patent submissions on the design of multimedia systems and protocols, layer-4 packet classification, MPLS path routing, and authentication and dynamic key exchange. He served as a co-guest-editor of *IEEE Network*'s March 2001 Special Issue on Fast IP Packet Forwarding and Classification for Next Generation Internet Services. He has served in the capacity of tutorial chair for IEEE LCN '94, and '95, as a publicity chair for NOSSDAV '97, and as a program committee member for MMCN 2001, 2003, IEEE ICNP2002 and 2003, and IEEE LCN '93–2000 conferences.

GIRISH CHANDRANMENON received his B. Tech. degree in 1991 from IIT, Madras, India, his M. S. in 1994 from the University of New Hampshire, Durham, and his D.Sc. in 1999 from Washington University. He joined Bell Labs in 1999. He is currently a member of technical staff in the High Speed Mobile Data Research Department at Bell Labs. His research interests are design and implementation of efficient network protocols; in particular, his research currently focuses on wireless and mobile networking protocols for data and voice.

SEUNGJAE HAN is a member of technical staff of the Wireless Research Laboratory, Bell Labs. He received B.S. and M.S. degrees in computer engineering from Seoul National University, Korea, in 1989 and 1991, respectively, and a Ph.D. degree in computer science and engineering from the University of Michigan, Ann Arbor, in 1998. His research interests include QoS networks, wireless networks, and fault-tolerant systems.

YUI-WAH (CLEMENT) LEE received his Ph.D. degree in computer science and engineering from the Chinese University of Hong Kong in 2000. Before that, he received his M.Phil. and B.Sc. degrees in electrical and electronic engineering from the University of Hong Kong in 1993 and 1988, respectively. He is currently a member of technical staff in the Networking Research Laboratory at Bell Labs. He was a visiting scholar in Carnegie Mellon University in 1996–1997. His research interests are system design and implementation in general, and mobile computing and networking in particular.

SCOTT C. MILLER is director of the High Speed Mobile Data Research Department at Bell Labs, Holmdel, New Jersey. He has B.S. and M.S. degrees in electrical engineering from Cooper Union, New York, New York. His current research involves the integration of 802.11 and 3G wireless data service and related mobile networking issues concerning seamless mobility, authentication, security, roaming, and accounting. Prior to his work on 802.11/3G integration, he led several systems research efforts in wireless applications, implementing novel systems for wireless messaging, speech-driven directory services, wireless instant messaging, carrier-based content billing, and multimedia content adaptation

LUCA SALGARELLI received his Laurea (Dr.Eng. degree) in electronic engineering from Milan Polytechnic University in 1995, and his M.Phil. in computer science from CEFRIEL, Milan, in the same year. He was a researcher in the Networking Department of CEFRIEL from 1995, where he worked in the areas of broadband IP networks and QoS provisioning. Since 1998 he has been with Lucent Technologies, first in the Wireless R&D Department in the United Kingdom, and then in the Networking Laboratory of Bell Labs Research, Holmdel, New Jersey. His activities cover design, development, and evaluation of systems and protocols for data networks, in particular when they involve mobility, QoS provisioning, and security. He is currently on a leave of absence from Bell Labs, teaching both graduate and undergraduate courses at the University of Brescia, Italy.