IBM Wireless RF LAN design and architecture

by F. J. Bauchot
F. Lanne

This paper describes the medium access control (MAC) protocol and the network integration of the IBM Wireless Radio Frequency (RF) Local Area Network (LAN) product. The MAC protocol is an adaptive, hybrid scheme relying on reservation-based and contention-based methods. It provides sustained high performances, both for light and heavy traffic conditions. Some control services of the wireless channel are also described: security, data compression, and interference management. The network integration is achieved by relying on the industry standards Open Driver Interface (ODI™) and Network Driver Interface Specification (NDIS). Wireless connectivity can either provide an extension of an existing cabled network or result in a stand-alone wireless network.

During the last few decades the complexity of data communications has significantly increased. The amount and the diversity of networking solutions available in the marketplace are far more important today than during the 1960s when mainframe-based networks were almost the only solutions offered to users. Today the customer can choose from among various subnetworking options (local area network [LAN], metropolitan area network [MAN], wide area network [WAN], etc.), network protocol stacks (Systems Network Architecture [SNA], Advanced Peer-to-Peer Networking* [APPN*], Transmission Control Protocol/Internet Protocol [TCP/IP], Sequenced Packet Exchange [SPX**], Internetwork Packet Exchange [IPX**], etc.), link controls (Ethernet, token ring, Fiber Distributed Data Interface [FDDI], asynchronous transfer mode [ATM], etc.), and so on. Moreover, all these alternatives can be mixed and matched, resulting in heterogeneous network solutions.

When a new data communications product is launched in such a complex environment, many different product design and system aspects have to be considered. This paper addresses such aspects for the IBM Wireless Radio Frequency (RF) LAN. The focus is on two major characteristics of the wireless LAN product: the Medium Access Control (MAC) protocol, which governs how multiple users can access the wireless channel, and network integration, which allows this new product to interconnect various other networks for resource-sharing purposes.

Next, the section on MAC protocol describes how a wireless LAN product can capitalize on existing MAC protocols within the constraints of a wireless channel. The section on network integration then describes how the new wireless LAN products fit into various data communication environments.

Wireless LAN MAC protocol

The evolution of LAN networking: A brief history. In the late 1960s, when teleprocessing was young, the means used to interconnect computers with remote computers and data I/O devices were primarily telephone lines, characterized by a limited bandwidth and poor reliability. They provided limited throughput and led to new data link control (DLC) protocols such as binary synchronous communi-

0018-8670/95/$3.00 © 1995 IBM

cation (BSC) and high-level data link control (HDLC) to provide upper protocol layers with reliable communication services on the low-quality channels.[1]

Later during the 1970s, new communication channels appeared, providing much better reliability. Although the bit error rate (BER) of terrestrial telephone lines was commonly in the range of 10^{-3} to 10^{-5}, it became possible to achieve BERs in the range of 10^{-7} to 10^{-9} or better.[2] Moreover, these new communication channels, which were designed to carry data rather than voice, offered much more bandwidth, but reduced the range between the communicating devices. These new communication channels, introduced as LANs, became pervasive in the 1980s. The availability of LANs has driven the paradigm shift to client/server and peer-to-peer network architectures from large mainframe-based teleprocessing.

During the early 1990s, technological innovations in electronics and packaging led to the widespread use of portable computers. Their commercial success is mainly because they can be used anywhere. They are no longer tied to a communication line, provided that the data required to run applications have been previously loaded on a fixed disk. Thus it remains necessary to provide some telecommunication line attachment (using modems or LAN adapters, for instance) to upload files from or download files to these portable computers.

The next paradigm shift in teleprocessing is to avoid use of any cable, allowing connectivity even if the computer is not cable-attached. This last shift has resulted in the wireless LAN (WLAN) products available today on the market. The end-user computer has become mobile without precluding online attachment to a server, mainframe, or peer device.

Such connectivity relies on wireless communications using electromagnetic waves. Different wireless channels are available, along with different spectrum-management techniques: infrared channels or radio channels, possibly using spread-spectrum techniques. The choice of such a channel or spectrum-management technique has been discussed in Reference 3. Whatever the choice is, the communication channel and the wireless devices are characterized by the following attributes:

- The communication range is limited (either due to regulatory constraints or the laws of physics) but sufficient to accommodate regular LAN-type connectivity.
- The available bandwidth is sufficient to sustain LAN operation but remains limited in comparison with state-of-the-art wired LANs.
- Carrier-sensing of electromagnetic waves is far less reliable than for wired channels and cannot be achieved (at an affordable product cost) while transmitting.
- Wireless devices commonly are battery-powered portable computers; it is thus desirable to allow efficient access to the channel with low power consumption.
- The channel quality (as measured by the bit-error-rate) is poor, comparable to terrestrial telephone lines.
- The communication channel is not privately owned and is thus not perceived by end users as being secure; means are needed to gain the confidence of end users for data privacy and secure network access.
- Another consequence of using a public-domain communication channel is the inability to manage and control attachment to the channel: a given user may interfere with others, so techniques are needed to avoid network collapse due to interference.
- Finally, some wireless communication techniques, such as frequency-hopping spread-spectrum,[4–6] require sophisticated synchronization procedures at the physical layer, before any piece of end-user data may be exchanged over the air.

Thus, WLAN products cannot simply rely on the latest techniques and technologies used for wired LANs. Wired communication lines have evolved in the last 30 years from high bit-error-rate, long-range, limited-bandwidth telephone lines to low bit-error-rate, limited-range, very-high-bandwidth LAN media. Wireless LAN communication lines are at some "middle point," characterized by high bit-error-rate, limited range, and medium bandwidth. As the wireless communication channel also differs from the wired ones in other aspects,[7] new schemes must efficiently accommodate these new constraints.

The next subsection describes a wireless LAN MAC protocol meeting the above constraints, which has been implemented in the recently announced IBM Wireless RF LAN product.

Description of the access method. For clarity, we assume in the following discussion that the phys-

ical layer providing services to the MAC layer is based on frequency-hopping spread-spectrum (FHSS) techniques, as justified in Reference 3. Such a physical layer allows the most flexible interference management and also puts the most constraints on the MAC layer. Under this assumption, we describe a WLAN MAC protocol that can also accommodate, with some minor modifications, other physical layers: infrared or direct sequence spread-spectrum (DSSS) RF or non-spread-spectrum RF.

The basic principles on which FHSS systems rely are simple. The total frequency band is divided into a number of frequency channels of equal width. At a given point in time, transmission can occur only within such a channel and cannot last too long on the same channel. Moreover, the energy placed on the total frequency band must be equally distributed among the channels (or a subset of the channels). These rules derive from specifications issued by regulatory bodies. The classical way to comply with them is to follow a so-called *frequency-hopping pattern*. The wireless station regularly moves from one channel to another (a move called a *frequency hop*) by following a given sequence of channels (the *pattern*) on a cyclic basis. All the wireless stations communicating among themselves must obviously follow the same frequency-hopping pattern synchronously. For instance, in the 2.4-GHz (gigahertz) industrial, scientific, and medical (ISM) band,[8] 83 channels of 1 MHz (megahertz) each are available on which transmission cannot last more than 400 milliseconds (ms). The Federal Communications Commission (FCC) requires in its document Part 15.247 that at least 75 channels among 83 be used. In this band, the frequency-hopping patterns must thus have a length falling in the range of 75 to 83. In the 2.4-GHz ISM band, the IBM Wireless RF LAN uses patterns with a length equal to 79, and transmissions last for 96 ms per channel. The IBM Wireless RF LAN also complies with regulatory constraints imposed in countries other than the United States.

When designing a WLAN MAC protocol the first rule is to try to capitalize on well-known techniques commonly used for wired LANs. Thus, a designer may try to adapt schemes such as those based on tokens (token ring or IEEE 802.5, for instance) or on collision-detection techniques (Ethernet or IEEE 802.3, for instance). In doing so, the designer quickly finds that such schemes do not match the specifics of the wireless communication channel. Indeed, when a station sends data over the air, it *broadcasts* the data so that all receivers within range of the transmitting station receive the data almost at the same time. As a result, techniques relying on passing tokens from station to station cannot be easily and efficiently accommodated. Because they are also known to be fragile, they are not good candidates for wireless LAN systems. Nor can a station detect collision while it is transmitting. It would require almost "doubling" the RF transceiver (one part dedicated for transmission, the other one for reception), thus leading to a prohibitive increase in product cost. Therefore, the designer must innovate when defining a WLAN MAC protocol, ideally by defining a *hybrid* scheme that has the benefits of known schemes, while avoiding their drawbacks. In this way an objective of the designer is to permit guaranteed bandwidth while controlling the latency.

When designing a WLAN MAC protocol, try to capitalize on well-known techniques commonly used for wired LANs.

MAC protocols can be split into roughly two groups: those that rely on some contention mechanism, and those that are contention-free. Schemes in the first group provide high efficiency, both in throughput and latency, when the channel load is low and the traffic is bursty, whereas their performances significantly degrade under high-load conditions.[9,10] In contrast, the protocols based on a contention-free scheme (such as polling or reservation[11,12]) offer very good and stable performance in high-load situations with steady-state traffic, but they do not perform as well as contention-based schemes in light-load situations. As previously mentioned, the wireless communication channel bandwidth is smaller than that commonly found in a wired LAN environment. For instance, the FCC specifies in the Part 15.247 regulation that the 2.4-GHz ISM band can be used by FHSS systems with only 1 MHz per channel. If we assume for simplicity a spectral efficiency equal to 1, it means that

Figure 1 Time frame basic structure

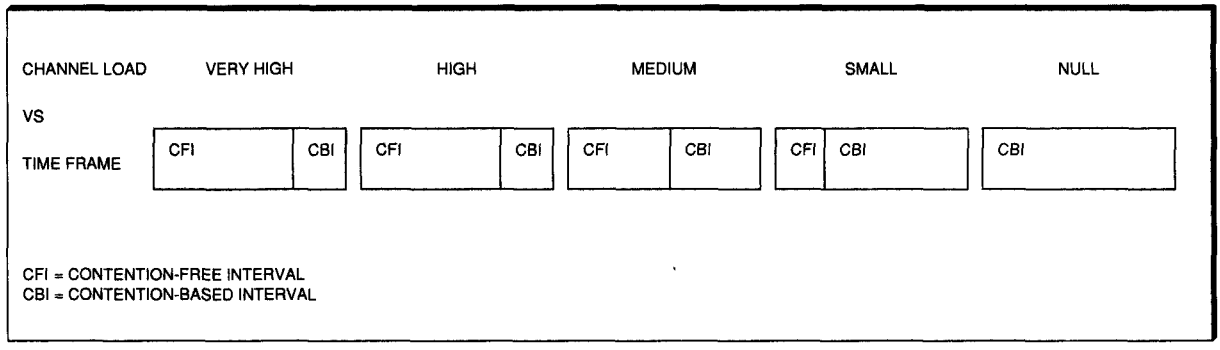

this channel cannot accommodate more than 1 Mbps (million bits per second) capacity. As a result, high-load situations will be reached in WLAN networks much easier than in wired LAN channels, such as the 10-Mbps IEEE 802.3 or the 16-Mbps IEEE 802.5.

With these considerations, it seems desirable to define a MAC protocol that offers good performance under both low-load and high-load conditions. Such a protocol must thus rely on an *adaptive* means to move smoothly from a contention-free scheme to a contention-based scheme according to the instantaneous channel load. A very first example is a reservation multiple-access protocol that uses a slotted-Aloha-type reservation channel and an adaptive retransmission probability for stable operation as originally described in Reference 12. Abramson's paper[13] collects many of the useful early papers that are background for the design and architecture of the IBM Wireless RF LAN MAC protocol, which was first described in a patent issued in 1992,[14] later followed by one issued in 1995.[15] The first publicly available description of the IBM protocol was provided by Natarajan;[16] it was later followed by an updated version[17] proposed as the MAC protocol for the IEEE 802.11 standard. This proposal was not kept by the IEEE 802.11 committee, which issued its first draft standard document in December 1994.[18]

This protocol defines a time frame (in practice, typically 100 ms in duration) consisting of two "intervals," the first one used for contention-free traffic, whereas the second one handles contention-based traffic. The capability to accommodate any type of channel load situation (and thus to remain efficient whether the load is low or high) is achieved by allowing the boundary between these two intervals to change according to the instantaneous channel load. Figure 1 illustrates the variation in interval size. Notice that the contention-free interval can be null when no traffic is handled, but the contention-based interval never collapses completely. Obviously the channel load information must be provided by a dedicated mechanism. Such a mechanism can be as simple as monitoring the current size of the queues used for transmission or reception. This mechanism, referred to as the *scheduler*, resides in a station called the *base station*. This wireless station determines the relative size of each interval for each time frame. This interval size information is broadcast by the base station to all the other stations (referred to as *remote stations*) in a protocol control packet issued at the beginning of each time frame. Such a scheme leads to a hierarchical network topology, because the base station plays a specific role. The choice of a hierarchical topology (as opposed to a distributed peer-to-peer one) is also justified by other considerations addressed in Reference 3.

With the time split as a sequence of contention-free and contention-based intervals, it is necessary to define how the channel is accessed during each of these intervals. During the contention-free interval, the channel must be accessed at any given time instant by only a single station. The classical and easy way to achieve this goal is to rely on a time division multiple access (TDMA) scheme, which splits the contention-free interval into a sequence of *time slots* and assigns each time slot to a given station. This assignment of time slots defines a *slot map* which is also broadcast by the base station to the remote stations in the protocol con-

trol header issued at the beginning of each time frame. This scheme offers several characteristics that can be turned into advantages for the system:

- For each time slot, the slot-map information specifies the transmitting station and the receiving station or set of receiving stations. Then, any station knows on a time-slot basis whether it can either transmit or receive data. This information tells each station when it can potentially enter a power-saving mode. If a given time slot is not assigned to a given station (both for the transmission and reception cases), this station can safely enter a power-saving mode as long as the time slot lasts.
- Because wireless channels have poor BER characteristics, it is necessary to limit the size of each block of data sent over the air to minimize transmission retries due to channel impairments. Two simple solutions can be chosen. The first one is to limit the size of the data frames that the sending MAC layer receives from its upper layer. This approach forces more frame exchanges between the MAC and its upper layer; it translates into higher processing requirements, leading either to an increase in product cost or to protocol performance degradation. The second solution is for the transmitting station to fragment the data frame into a series of data segments, and for the receiving station to assemble these segments back into a data frame. Each data segment (except the last one) has a fixed length. This solution avoids the drawback of the first one and can handle data frames of any size, even very large ones. In this case it is natural to allocate each time slot to a given data segment, so that a data frame, when transmitted over the air, will use a sequence of time slots. By defining the time slot as the minimal duration required to exchange a data segment (and its acknowledgment), the time slot overhead is minimized, and the channel efficiency during the contention-free interval can reach very high figures. Within each slot, the transmitting station sends a data segment to the recipient station, which is followed in the reverse direction by an acknowledgment packet issued by the receiving station upon error-free reception of the data segment. Acknowledgments are sent one-to-one with respect to data segments and can be either positive or negative to reflect the correctness of the received segment. If no acknowledgment or a negative acknowledgment is received within the time slot, the transmitting station concludes that the segment transmission has failed and that it must be retried in the next allocated time slot.

In order to build the slot-map information, the base station must be aware of the instantaneous traffic. Obviously, the base station knows if it holds any outstanding traffic ready for transmission (known as the *outbound* traffic since it goes from the base station to the remote stations), but it must also determine the current traffic originating from the remote stations (known as the *inbound* traffic). This last information is derived from *traffic reservation requests* issued by the remote stations. When a remote station wants to transmit a data frame during the contention-free interval, it must first issue a reservation request during the contention-based interval asking for the number of time slots required to send its data frame. By simultaneously monitoring its outbound transmission queue and its reservation request queue, the scheduler resident in the base station can first determine the relative size of the two intervals and then build the slot map describing how each slot of the contention-free interval is assigned to the set of remote stations. The resulting information is broadcast in the protocol control header issued by the base station when each time frame starts. The slot map consists of a sequence of quadruplets (*TYPE*, *SA*, *DA*, *NBR*), each of them corresponding to a data frame. The first field *TYPE* specifies whether the frame is inbound or outbound and whether it is transmitted in the contention-free interval or in the contention-based interval (corresponding to the periods A, B, and C as introduced later). The second and third fields *SA* and *DA* give the addresses of the source and destination stations, and the last field *NBR* specifies how many time slots are used to transmit the data frame. The slot map has a variable length that is upper-bounded by a fixed limit to match implementation constraints (memory size, for instance).

As far as the contention-based interval is concerned, the designer has the choice of selecting or not selecting a mechanism based on carrier-sensing. Our choice is to avoid carrier-sensing for the following reasons:

- First of all, sensing the carrier is not free; it asks for some circuitry in the RF transceiver.
- Second, the process of carrier-sensing on wireless channels is not fully reliable. Some back-

ground noise may be misinterpreted as the presence of valid data, and carrier-sensing suffers from the well-known problem of the hidden terminal.[19]

Considering that a slotted scheme is used in the contention-free interval, IBM's chosen access method for the contention-based interval is the *slotted Aloha scheme*,[20,21] which has several advantages: it does not require carrier-sensing, it is relatively efficient, and it is slotted. For the implementer, this last aspect is of interest, as time-slot management is then supported in the same way, whatever the interval is. The presence of "capture effect"[22] on the wireless channel allows the slotted Aloha scheme to approach the efficiency of schemes based on Carrier Sense Multiple Access (CSMA).[23] Moreover, the slotted Aloha scheme is quite simple to implement. For each time slot during the contention-based interval in which the station has data to send, it transmits (using a probability function) with a probability P_{Xmit}. It has been shown that the slotted Aloha throughput can be maximized[24] if the transmission probability P_{Xmit} is chosen equal to $1/K$, where the parameter K corresponds to the average number of contenders needing to access the channel. An algorithm running in the base station calculates an estimation of this parameter K and broadcasts it as part of the information stored in the protocol control header issued at the beginning of each time frame. The K parameter estimation derives from two measures: the probability of successful transmission in the contention-based interval (segment followed by a positive acknowledgment) and the probability of *first* successful transmission (first try segment followed by a positive acknowledgment).

An important point concerns the access of the channel during the contention-based interval. As the traffic is scheduled by the base station, it can further increase the protocol efficiency by limiting the contention-based interval to the inbound traffic, with all of its own outbound traffic not suffering potential collision over the air. This characteristic is a key advantage of this scheme because it matches typical client/server traffic patterns where the outbound data stream is predominant. This ability of the base station to transmit without contention is a major performance boost.

This slotting of the time scale is advantageous when the physical layer below the MAC layer uses an "interrupted" channel. It is, for instance, the case for FHSS systems, where the RF transceivers have to switch from one channel frequency to another on a periodic basis.[25] By choosing the frequency-hopping channel *dwell time* to be equal to a multiple of time slots, there is no loss of protocol efficiency due to the intersynchronization of both layers. Other protocols that are not slotted for a specific time must defer the beginning of a new dwell time until any outstanding transmission completes. This burdens the process of synchronization among the remote stations even more.

We illustrate these basic notions in Figure 2, which introduces the following conventions:

- The contention-free interval is further divided into two subperiods: period A, where outbound contention-free traffic occurs, and period B, where inbound contention-free traffic occurs.
- The contention-based interval is referred to as the C period.
- The protocol control header sent at the beginning of the time frame is known as the *slot header*, or SH. This header carries various control information: a *network identifier* allowing discrimination between several collocated wireless networks, a *base station address* field, the K parameter, the slot map, and other fields not addressed in this paper.

Figure 2 depicts a first possible design of the MAC protocol, close to the one originally described in Reference 17.[26] The periods A, B, and C follow each other in this order to comprise the protocol time frame whose duration is a constant (of which the dwell time for FHSS systems is an integral multiple). This scheme has very attractive performance, mainly in terms of channel efficiency under high- and low-load conditions.[27] Nevertheless, the latency does not compare well with what is commonly found in contention-based protocols, because the scheduler runs only at the beginning of any time frame. Should any outbound traffic arrive after transmission of the SH control packet, it cannot be scheduled in the current time frame and must wait for the next one to be transmitted over the air. Similarly, if inbound traffic appears at a remote station within a time frame, the remote station will be able to place a reservation request, but it cannot be honored before the end of the current time frame. Such situations are not that bad when the traffic load is high because the fresh traffic will be transmitted only when the current traffic has been transmitted. However, in low-load situations,

Figure 2 Time frame detailed structure

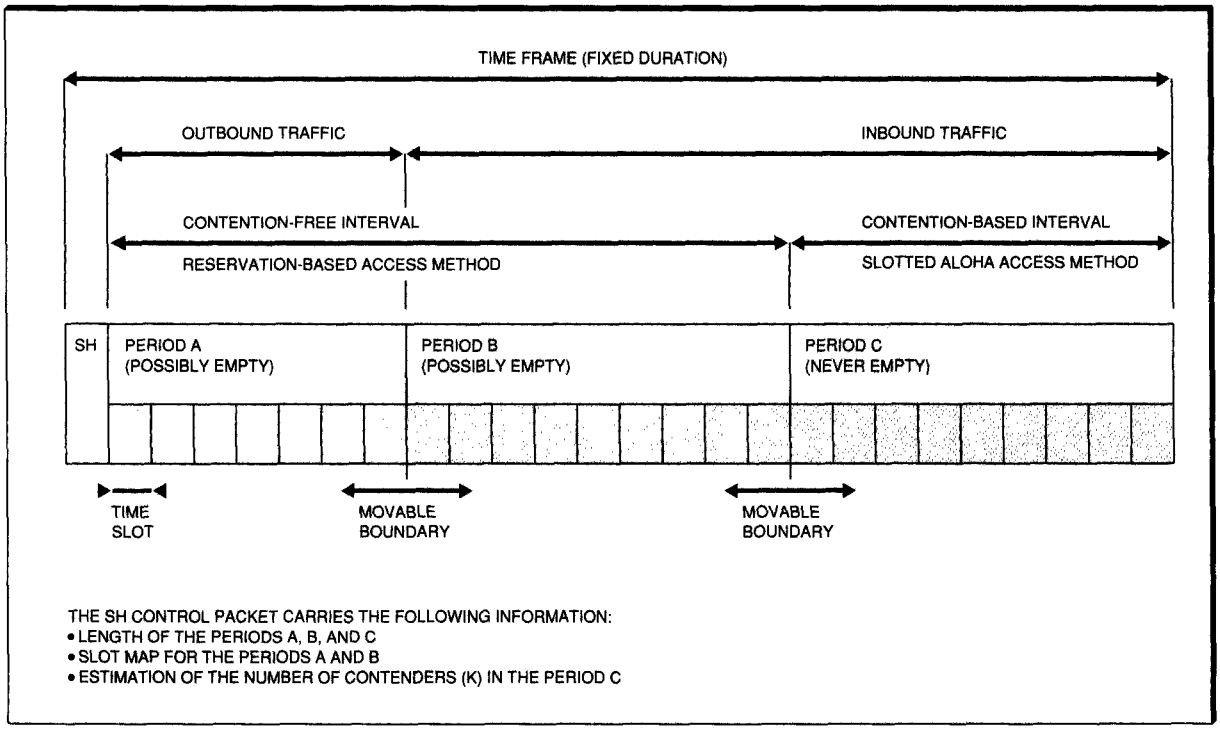

this latency limitation may limit throughput because fewer opportunities are offered to access the channel.

A simple way to fix this problem would be to decrease the duration of the time frame. Unfortunately this presents the drawback of increasing the overhead due to the scheduler processing time and the time to transmit the SH packet. The chosen solution consists of allowing the time-frame duration to vary according to the scheduled traffic. When the scheduled traffic is heavy, the time frame lasts longer (providing better channel efficiency), and when the scheduled traffic is low (or even null), the time frame shrinks to a minimal size corresponding to the smallest possible contention-based interval plus the SH header time.

Another improvement to the scheme of Figure 2 is to interleave, on a message basis, the periods A, B, and C, so that the C slots where contention occurs are not grouped at the end of the time frame, artificially decreasing the probability of collision. We illustrate this phenomenon by a practical example. Assume that the outstanding traffic, as scheduled by the base station BS, corresponds to two outbound frames A1 and A2, respectively split into L1 and L2 segments, and respectively directed to the remote stations RS1 and RS2; and assume that it corresponds to two inbound frames B3 and B4, respectively split into L3 and L4 segments, and respectively issued by the remote stations RS3 and RS4; and also corresponds to four contention-based C time slots. The Figure 2 scheme builds a slot map (i.e., assignment of time slots) corresponding to the following sequence: (A,BS,RS1,L1) (A,BS,RS2,L2) (B,RS3,BS,L3) (B,RS4,BS,L4) (C,any,BS,4). If the receiving remote stations RS1 and RS2 have to return some packets in response to the ones received, with such a scheme they will potentially collide if they choose the same time slot in the closing C period. Now if the scheduler works so that the slot map is (A,BS,RS1,L1) (B,RS3,BS,L3) (C,any,BS,2) (A,BS,RS2,L2) (B,RS4,BS,L4) (C,any,BS,2), it gives station RS1 a chance to use a slot in the first part of the C period (before transmission from RS2), while station RS2 will use the second part of the C period. In doing so, there is a limited risk of collision between RS1 and RS2.

Some performance measurements have shown that this interleaving scheme can reduce the protocol

Figure 3 Time frame structure used in the IBM Wireless RF LAN

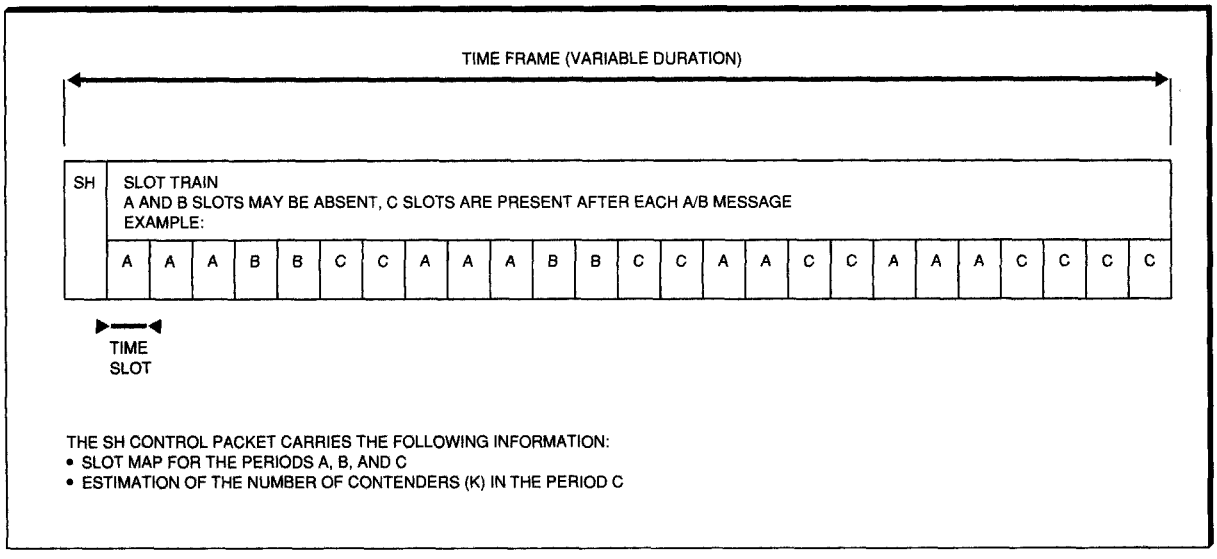

latency by more than half while the aggregate MAC protocol efficiency may gain up to 40 percent.

Figure 3 illustrates this second scheme, which is used in the IBM Wireless RF LAN. In this example, the outstanding traffic corresponds to four outbound data frames and to two inbound data frames. The resulting slot map illustrates how the outbound (period A) and inbound (period B) reservation-based time slots are interleaved with contention-based (period C) time slots.

The adaptive nature of the described scheme has shown that it is possible to remain efficient both in latency and in throughput, whatever the transmission channel load is. The capability to limit the contention for the inbound flow is another important performance improvement, mostly visible in client/server networking environments.

Description of the wireless channel control services. The previous subsection has described how stations may access the wireless channel when a data frame is ready to be transmitted. Here we describe how these data frames become eligible for transmission, and how stations provide services such as attachment to the network, data privacy, data compression, and interference management.

Network insertion. Access to wireless communication channels presents some paradoxical aspects. On the one hand, any remote station can be quite easily "connected" to the channel, simply because the channel belongs to the public domain (do not look for a plug!) and because it is sufficient to turn the power of the transceiver on to gain access. On the other hand, the transmission techniques (often imposed by regulatory bodies) are complex enough so that gaining access to the channel does not simply provide the capability to either send or receive data. Some intermediate steps are required to become *synchronized* with the wireless channel. Such steps are complex (and may take some time to complete) because the mode of operation of the physical layer used to access the wireless channel cannot be fully known by a new station wishing to enter a wireless network. We illustrate this *physical layer synchronization* procedure for frequency-hopping spread-spectrum systems.

When an FHSS remote station needs to be "inserted" into a WLAN, the information this station has is very limited. It knows at least a network identifier (discriminating a given WLAN from among several collocated ones) and also the boundaries of the frequency band used by the network. There is much more information that the station must learn "on the fly" to be synchronized:

- First, the remote station must determine the best base station to be connected to. Typically, the new remote station can be within range of sev-

eral base stations. It must use some channel quality criteria to determine the best base station, such as the strength of received signals and the instantaneous load of the base station. It is useless to remember the best base station chosen during the last network insertion, since the remote station can be mobile, and the wireless channel load can change greatly over time.

- Second, the remote station must determine the frequency-hopping pattern that the base station uses. Although the set of frequency-hopping patterns is finite and can thus be loaded in the remote stations, a given frequency-hopping pattern may evolve over time for interference management reasons, as we explain later. Thus, the remote station must learn the current frequency-hopping pattern used in the network at each network insertion. Different strategies can be used by the remote station to learn the frequency-hopping pattern. A simple solution relies on the broadcast, from the base station, of a *frequency header* (FH) control header packet at the beginning of each frequency hop, which contains the next frequencies (typically the next eight ones) along which the base station will hop. The remote station is synchronized by locking on a fixed frequency until it hears an FH control header packet. Once it has been heard, the remote station hops according to the frequencies found in each successive FH header.
- Third, the remote station must learn the time scale of the network. Even for schemes that do not rely on time slots, a station cannot transmit at any point in time. This last step gives the remote station a time base for when transmissions can be initiated.

After achieving physical-layer synchronization, the remote station, although able to receive or transmit information on the wireless network, is not yet a member of the network. At this point, it must successfully pass some security checks to be accepted as a network member. Such procedures are required for secure access to the wireless network, precluding intruders from entering it. Several steps are followed:

- First, the wireless remote station signals that it wishes to enter the network. This action is known as the *registration*[28] procedure. The wireless remote station transmits a dedicated control packet carrying the necessary information to let the receiving base station determine whether the candidate remote station can operate within the network. Up to this point, no security information has been exchanged between the candidate remote station and the base station. Registration may be denied by the base station if, for instance, some mismatch is found in the mode of operation of the candidate remote station.
- The second step deals with security to *authenticate* the candidate remote station. Authentication allows both parties (the candidate remote station and the base station) to prove that the partner is who it claims to be. It relies on the exchange of several control messages that carry identity "challenges" and "proofs." The number of exchanged messages depends on the authentication algorithm used by both parties. The scheme used by the IBM Wireless RF LAN is known as the Two-Party Protocol (2PP)[29] and requires the exchange of only three messages.
- The third step deals with *access control*. When the identity of the candidate remote station has been authenticated, some type of access checking verifies whether the candidate station is authorized to enter the network. This control can be based simply on the identity of the candidate, but it can also be enhanced by discriminating according to the day of the week, time of day, or network access point (as featured by the IBM Wireless RF LAN).
- The fourth step, completing the network insertion, is the setup of the encryption algorithm that may later be used to encipher data transmitted over the air. The purpose of this *encryption initialization* procedure is to safely agree upon an enciphering scheme and to exchange encryption keys between the base station and the remote station. Such keys are not sent "in the clear" over the air (they are themselves enciphered) and must also be authenticated to ensure that they have been sent by the proper base station.

When all the previous steps have successfully completed, the candidate remote station receives a last control packet from its owning base station specifying that it is now ready to proceed. Subsequent communications may then occur from and to the base station to carry user data.

The challenge for the designer is to perform these network insertion steps in the least amount of time, while conserving the level of security asked for by the end users. This challenge becomes even more difficult when dealing with mobile stations. In such cases, a remote station moving within the network must perform some *hand-off* mechanism (also

known as the *hand-over* mechanism) to switch from one owning base station to a new one. Because the move must be achieved even if the remote station is currently handling some traffic, it translates into more stringent performance objectives than required for an initial network insertion. The hand-off mechanism involves a resynchronization with the new base station using almost all the same steps for initial network insertion. Some of the security procedures can be reduced by capitalizing on the fact that the remote station has already been accepted as a network member. For instance, the IEEE 802.11 draft standard uses the concept of *preauthentication* to shorten the hand-off duration.[18] The hand-off mechanism induces several other networking problems such as the dynamic rerouting of information from and to the remote station. These aspects are currently being jointly studied by IBM Research and the IBM development teams.

When a remote station has successfully performed its network insertion, it is ready to send and receive information.

Data privacy. When a remote station has successfully performed its network insertion, it is ready to send and receive information over the air. Since wireless communication media involves the public domain, users are concerned with the privacy of their communications and require mechanisms to prevent eavesdropping on sensitive information on the channel. The classical answer to this concern is to encipher the data before transmission and to decipher them after reception. Today many different techniques are available to achieve data encryption; they have been intensively used in environments such as banking or hospitals.[30] In such cases, the data are enciphered typically within the applications (or in some lower layers close to the applications) where processing power is available; encryption may even be done in software. For wireless communication links, data privacy must be secured on the wireless link, so that encryption and decryption must be performed within the hardware providing wireless connectivity. We refer to this hardware as the *wireless adapter*. Such wireless adapters have limited processing power (for product-cost reasons) so that enciphering cannot be achieved by software. Some hardware-assist modules must thus be used to off-load the other hardware or software components involved in "regular" data transmission processing.

Since the enciphering-deciphering process translates into a modification of the data stream, it makes sense to couple this process with other ones handling the data stream to avoid duplicating data copies within the adapter (operations that consume both memory and processing power). Two solutions are available to the designer: perform the encryption and decryption when the data crosses either the interface with the lower layer (RF transceiver) or the interface with the higher layer (the system bus, such as ISA, MCA, or PCMCIA[31] used in various personal computer environments).

As the data frames received from higher layers are segmented by the MAC protocol before transmission, it is less costly to perform the data encryption and decryption when a data frame crosses the upper-layer interface than when the corresponding data segments cross the lower-layer interface. Indeed, the encryption-decryption cost does not vary as a linear function of the processed amount of data, but rather as the number of encryption-decryption operations that require some control overhead due to encryption key management. Another justification for performing the data enciphering in this way is the possibility of combining in the same piece of hardware the two functions corresponding to crossing the wireless adapter to host system interface (which commonly involves direct memory access [DMA] operations), and the enciphering and deciphering of data. In the next subsection we show that the support of data compression further justifies this design choice.

Data compression. Since wireless channels do not provide bandwidth comparable to wired LANs, some means must be used to recover, even partially, from this channel capacity limitation. Data compression is a good answer, with proven efficiency in various data communication systems such as terrestrial line modems.[32] Some algorithms such as Lempel-Ziv[33] can provide compression ratios higher than 1:10, typically falling in the range of 1:1.5 to 1:3. Such techniques can artificially dou-

ble or even triple the channel bandwidth, as seen by the higher layers of the communication protocol stack. Data compression can be easily and efficiently implemented at a limited cost by using dedicated hardware chips. Any software implementation of the compression function would not only consume great processing power, but would last long enough to lose a significant amount of the benefit brought by compression itself (reduction of the transmission duration).

Interference characterization is performed by cyclically looking at the set of used frequencies.

The interesting point is how compression is combined with other treatments of the data stream. Up to now, we have seen that data frames, when received from the upper layers across the host-system bus interface, are first encrypted, then stored in memory, then sent over the air as a sequence of data segments. Compression algorithms compact data because they are in general not purely random but rather a sequence of bit patterns occurring several times. The more redundant the data stream is, the more efficient the compression. An easy way to exploit the data redundancy is to feed the compressor larger frames. For instance, if the data segments are 256 bytes long, the expected compression ratio for a segment will be lower than for a full data frame whose length may be larger than 1K bytes. Thus, the compression-decompression process should be performed on the frames before segmentation. The order in which compression and encryption must be done is even more obvious because the effect of encryption is to "whiten" (or decorrelate) the data stream, thereby almost eliminating any redundancy in the original data stream. If encryption is performed first, the compression will perform very badly, and in some cases the data may even expand!

The right sequence of operations followed by a piece of data (source bus-crossing, compression, encryption, segmentation, transmission, reception, assembly, decryption, decompression, target bus-crossing), can now dictate the appropriate wireless adapter architecture. Figure 4 represents the hardware design of the IBM Wireless RF LAN adapter.

Interference management. Because the wireless channel operates in the public domain, those who comply with the local regulatory constraints are allowed to use this channel for their own purposes. It means that a given user cannot control the access to the wireless media, even within the user's own premises. A wireless device must be able to work (without significant performance loss) even in the presence of other wireless devices that appear as *interferers*. In fact, the regulatory constraints are set up to allow the best *coexistence* of multiple users relying only on distributed control techniques. The spread-spectrum techniques imposed by the FCC for the ISM band in Part 15.247 of its regulations are a good example of such distributed techniques. FHSS and DSSS have been designed to allow several wireless devices to share a common frequency band in an efficient way. Nevertheless, these techniques have their own limits so that the aggregate capacity decreases when the number of interferers increases. To partially avoid this limitation, it is desirable to adapt the mode of operation of the device so that it accounts for the nature of the interferers it is aware of. Assuming again that the wireless device uses the FHSS technique, reducing interference involves three successive steps: (1) monitoring the RF environment, (2) characterizing the interferer, if any, and (3) applying the right corrective action to limit the effect of the interferer.

RF environment monitoring can be simply achieved by maintaining in the base station, for each channel in the frequency-hopping pattern, a parameter corresponding to the transmission retry ratio. Many other sophisticated measurements could be logged with adequate electronics, but since the only objective of a communication device is to send data without retry to the recipient device, the retry ratio is the best indicator characterizing a given channel frequency as good or not.

Interference characterization is performed by cyclically looking at the set of used frequencies. Different cases may appear: for example, for all the frequencies, the retry ratios show good results; or a single frequency gives a bad retry ratio; or several frequencies give a bad retry ratio; or finally,

Figure 4 Wireless adapter hardware architecture

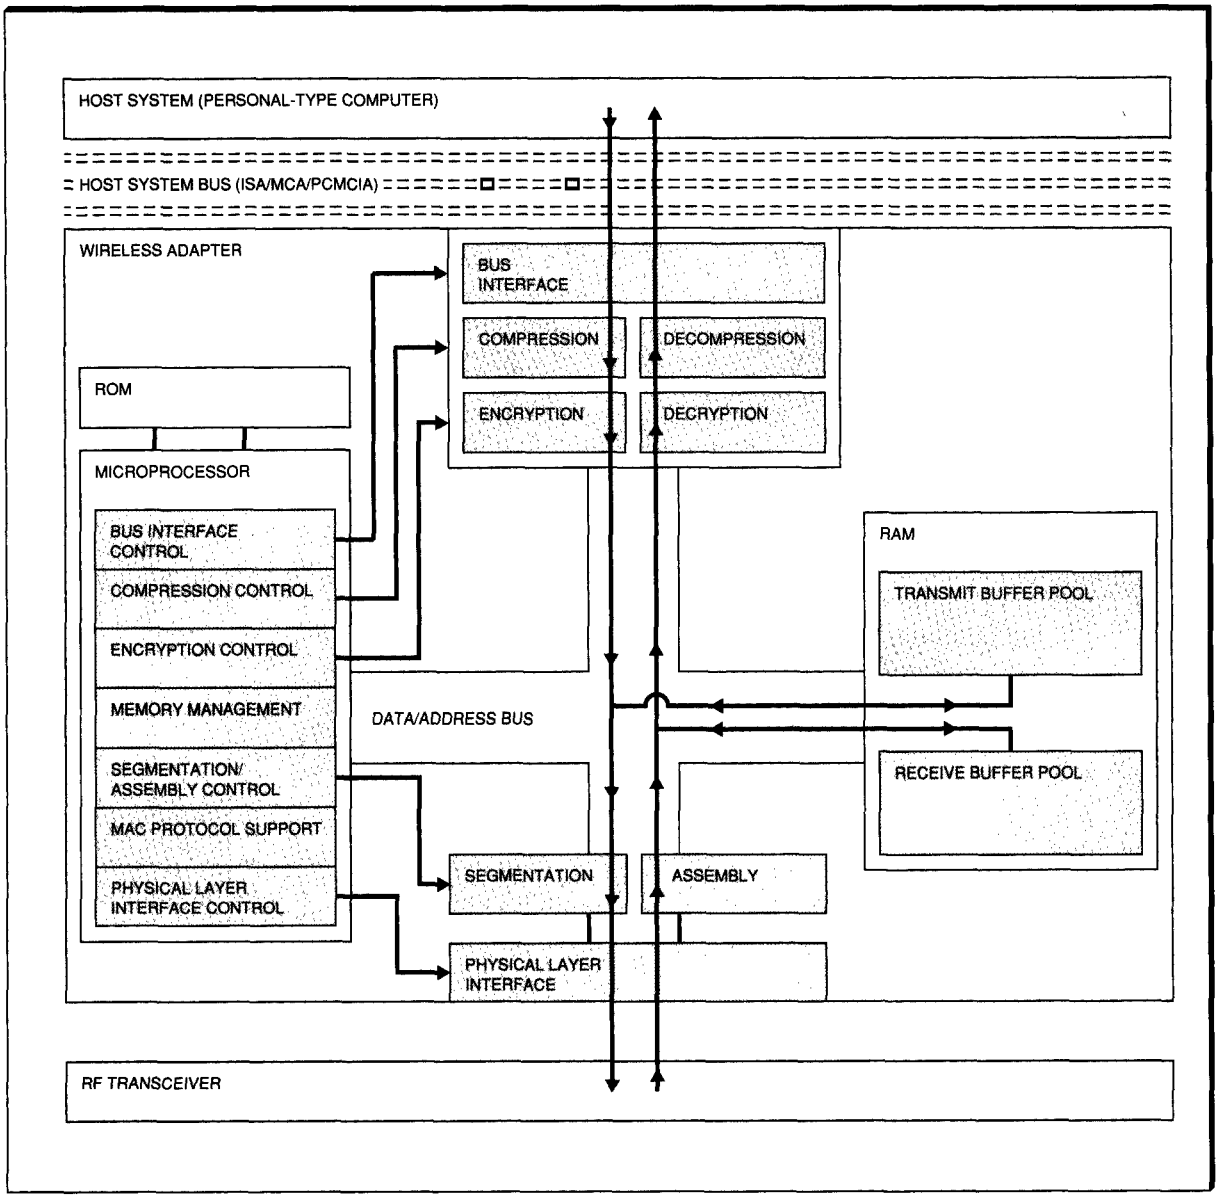

all the frequencies show a retry ratio that is gradually degrading along the time scale. Each of the previous cases corresponds to a different type of interferer and can be reduced by a specific means:

- When all the frequencies show good retry ratios, the current frequency-hopping pattern must be kept unchanged. The system does not suffer from interference resulting in transmission impairment.
- When a single frequency experiences interference, corrective action consists in replacing this noisy frequency by a *spare* frequency. Such spare frequencies are available because the regulatory bodies require use of only a subset of the total number of available channels. For instance,

Figure 5 Interference management strategy examples

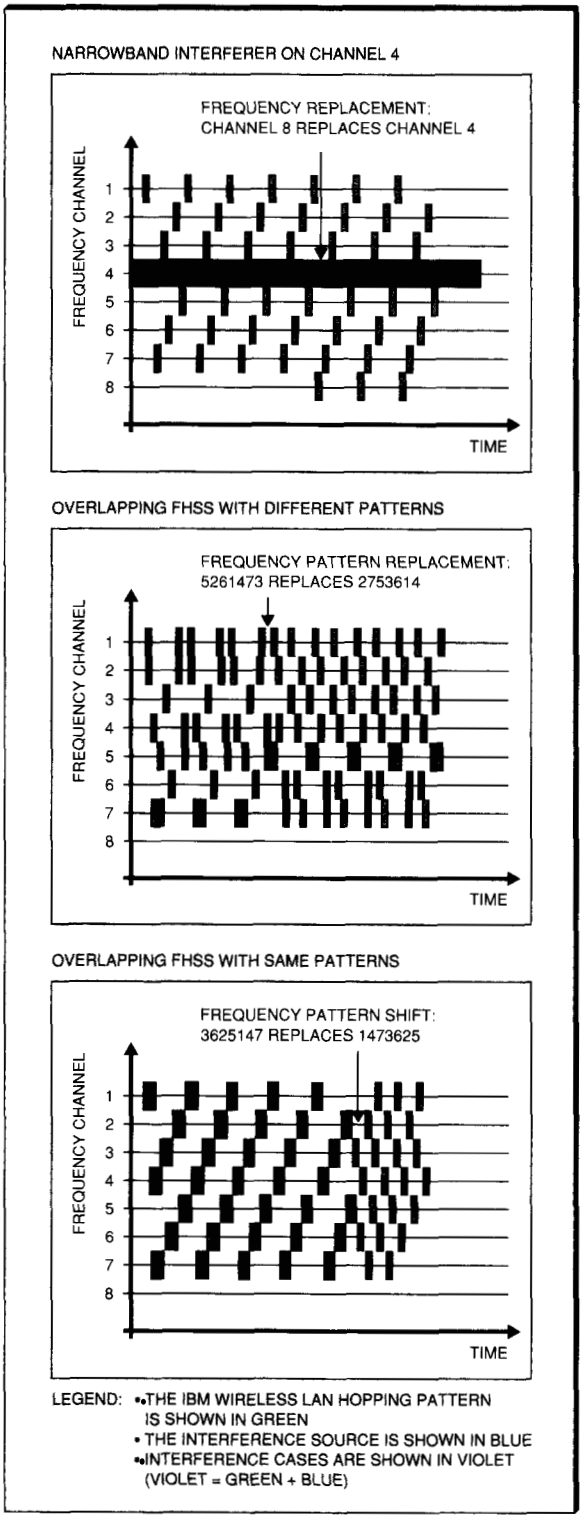

the FCC Part 15.247 specifies that at least 75 channels out of 83 must be used for the 2.4-GHz ISM band, allowing up to eight spare channels. Single frequency interference cases may appear if a narrowband primary user of the ISM band transmits within range of the wireless device.

- When several frequencies experience interference, corrective action consists in changing the complete frequency-hopping pattern. Such interference cases may appear if two FHSS systems are partially overlapping and do not use orthogonal hopping patterns.
- When all the frequencies show increasing retry ratios, it means that some collocated devices are currently using the same frequency-hopping pattern as the station experiencing interference. Indeed, if the two identical patterns are slowly shifting at different speeds (due to different oscillator characteristics), it translates into a phase offset that will gradually decrease to zero. When the phase offset reaches some threshold, all the frequencies of the pattern become affected, and all the retry ratios show homogeneously degraded results. The corrective action is simply to perform a phase shift on the frequency-hopping pattern so that the phase offset becomes large enough to avoid frequency overlapping for each channel.

Figure 5 illustrates the three interference situations as well as the corresponding corrective actions. In this example, we have assumed that the frequency band has eight different channels and that the frequency-hopping pattern uses channels 1–7, with channel 8 as a spare. The frequency-hopping pattern used by the IBM Wireless RF LAN is shown in green, whereas the interference source is shown in blue. Each time the IBM Wireless RF LAN and the interference source occupy the same frequency channel, they interfere, and it is represented by violet.

In the top example, a narrowband interferer continuously uses channel 4. The corrective action is the replacement of the noisy channel 4 by the spare channel 8, so that the frequency-hopping pattern 1873625 replaces the frequency-hopping pattern 1473625. In the middle example, two channels, 3 and 6, are sensed as noisy. The corrective action is the replacement of the whole frequency-hopping pattern 5261473 by the new one 2753614. In the bottom example, the same frequency-hopping pattern is used by two overlapping systems. The cor-

rective action is to shift the current frequency-hopping pattern so that 3625147 replaces 1473625.

For the three interference situations, the corrective action leads to a modification of the current frequency-hopping pattern (single frequency replacement, or frequency-hopping pattern replacement, or frequency-hopping pattern phase shift). The base station that performs the interference management must inform the remote stations of the frequency-hopping pattern modification. It is automatically done by updating, after the change, the next frequencies to come that are broadcast as part of the FH control header packet.

Network integration

Having shown how stations can communicate among themselves when equipped with the IBM Wireless RF LAN adapter, we now discuss how the wireless path can be tied to existing communication protocol stacks to provide an end-to-end networking solution. Two aspects are important to the customer in weighing the merits of moving to the wireless technology: compatibility with existing networking software, and compatibility with the existing network infrastructure.

Compatibility with networking software. An end user would not be reluctant to change the network adapter card on his or her workstation as long as network resources can still be accessed and shared through the same user interface and with an equivalent or better level of performance. In other words, the user wants the adapter card to be compatible with the network operating system running on his or her workstation.

Protocol stacks (NetBIOS*, TCP/IP, SNA) are generally software programs running on the processor of the workstation. In contrast, lower-level protocols (e.g, the wireless LAN MAC protocol) are generally implemented either in a network adapter card housed in an expansion slot of the workstation, or in the device driver of the network adapter card, or in both of them. The protocol stack uses the services provided by the wireless adapter and its device driver to have a message transferred. Figure 6 is a sketch of this scheme. It is at this level that compatibility with a network operating system needs to be achieved.

Notice that the protocol stack does not interface directly with the network adapter but to a network

Figure 6 Network resource access from a workstation

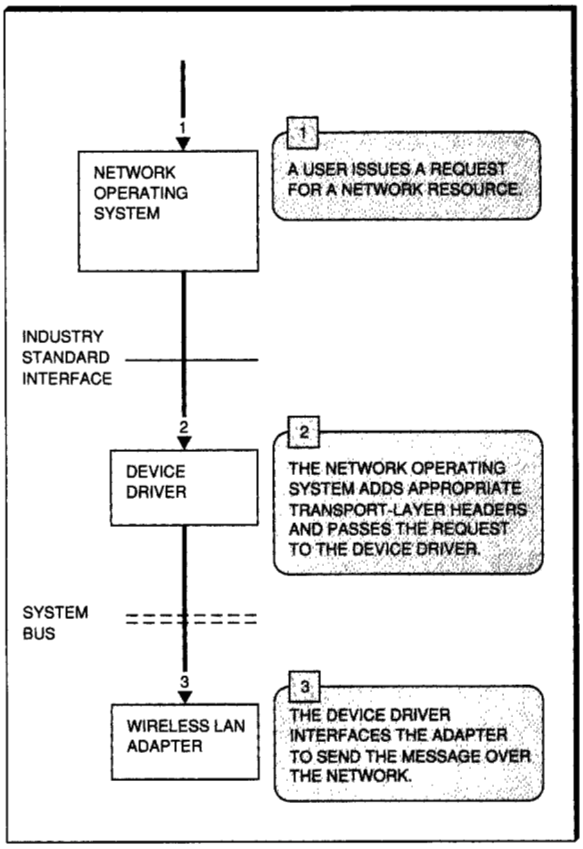

adapter driver, a program that is responsible for handling the network adapter card. To achieve compatibility with most of the network operating systems in the marketplace, the IBM Wireless RF LAN provides network adapter drivers that follow the two major industry standards defining the interface between the protocol stack and the network adapter driver.

In the LAN environment, two main *de facto* standards for the network adapter card interface are today available: Network Driver Interface Specification (NDIS)[34] and Open Driver Interface (ODI**).[35] NDIS was jointly developed by the 3Com and Microsoft Corporations. NDIS separates protocol handling from hardware manipulation by defining functions that protocol stacks and network adapter drivers should provide to one another. In addition, NDIS describes how protocol stacks and

network adapter drivers should receive their configuration data and how they should be associated (for example, TCP/IP by means of a token-ring card). ODI was developed by Novell, Inc. as the network adapter card interface used in its NetWare** networking offering. ODI offers the same level of functionality as the NDIS interface.

Besides the compatibility aspect, a key advantage of following these standards is that hardware and software are clearly isolated. Choosing the network adapter card and choosing the network operating system are two separate decisions that use independent criteria.

Another important advantage of those standards is that they usually integrate multiplexing, allowing several protocol stacks to use a single network adapter card at the same time, or a single protocol stack to use several network adapter cards at the same time. The multiplexing feature has particular value from a user's perspective:

- If several protocol stacks can use a single network adapter card at the same time, it means that a user at a workstation can do several things at the same time, for example, run a 3270 emulation to access OfficeVision* mail on an IBM host (using SNA) and download some spreadsheets from an IBM LAN Server* (using NetBIOS) over a single network adapter card.
- If a protocol stack can use several network adapter cards at the same time, it allows some network interconnection functions. For example, router or gateway software will be able to obtain a frame from one of its network adapter cards and transfer it to another one.

By providing NDIS- and ODI-compliant network adapter cards, the IBM Wireless RF LAN can be used in the majority of today's available network operating systems such as: Artisoft Lantastic**, IBM LAN Server, Microsoft LAN Manager**, Microsoft Windows for Workgroups**, Novell NetWare, and Novell NetWare Lite**.

Compatibility with the infrastructure. Two configurations are possible for a wireless LAN, the single cell configuration and the extension of an existing wired network. Those two configurations correspond to two different customer environments.

Consider a customer who does not use any LAN today and decides to invest in a wireless LAN because of its various advantages over a wired LAN for linking just a few workstations together. This customer will build a single cell configuration. The single cell configuration is a pure wireless network, where cabling is not used at all. Each workstation is equipped with an adapter for the IBM Wireless RF LAN and has the required component of a network operating system installed on it. One of the wireless stations operates as the base station, whereas the others are remote stations.

In contrast, a customer who already has a wired network may decide to extend an existing network with wireless cells, because cabling is not practical in some locations (in a hospital, for example) or because the required extension is only temporary (a department store at Christmastime, for example).

The wireless extension of a wired network has some implications. Not only must the remote stations communicate with one another but also with workstations and servers on the wired LAN. In the department store example, wireless point-of-sale terminals deployed temporarily will need to access the price server located within the wired network of the store.

To permit interconnection between wireless cells and a wired network, the IBM Wireless RF LAN allows the customer to choose the solution best-suited to that customer's environment. Bridging, routing, or some gateway approaches are all available. We give three examples of possible configurations: bridging from the wireless cell to a token-ring network, IP or IPX routing from the wireless cell to various LAN or WAN types, and connecting the wireless cell to an SNA backbone through a gateway.

Like a token-ring network or an Ethernet, a wireless cell operates as a *LAN segment*. Bridging is the simplest way to interconnect LAN segments. The IBM Wireless RF LAN provides a source-routing bridge program to interconnect a wireless cell and a token-ring network. In such a configuration, the base station is equipped with two network adapters and drivers for both the token-ring network and wireless sides. The source-routing bridge program of the IBM Wireless RF LAN is an Operating System/2* (OS/2*) application installed in the base station working with both drivers to bridge the traffic. Such a configuration is very powerful, since a bridge can deal with virtually all protocols. An ex-

Figure 7 IBM LAN Server NetBIOS traffic bridged in the base station

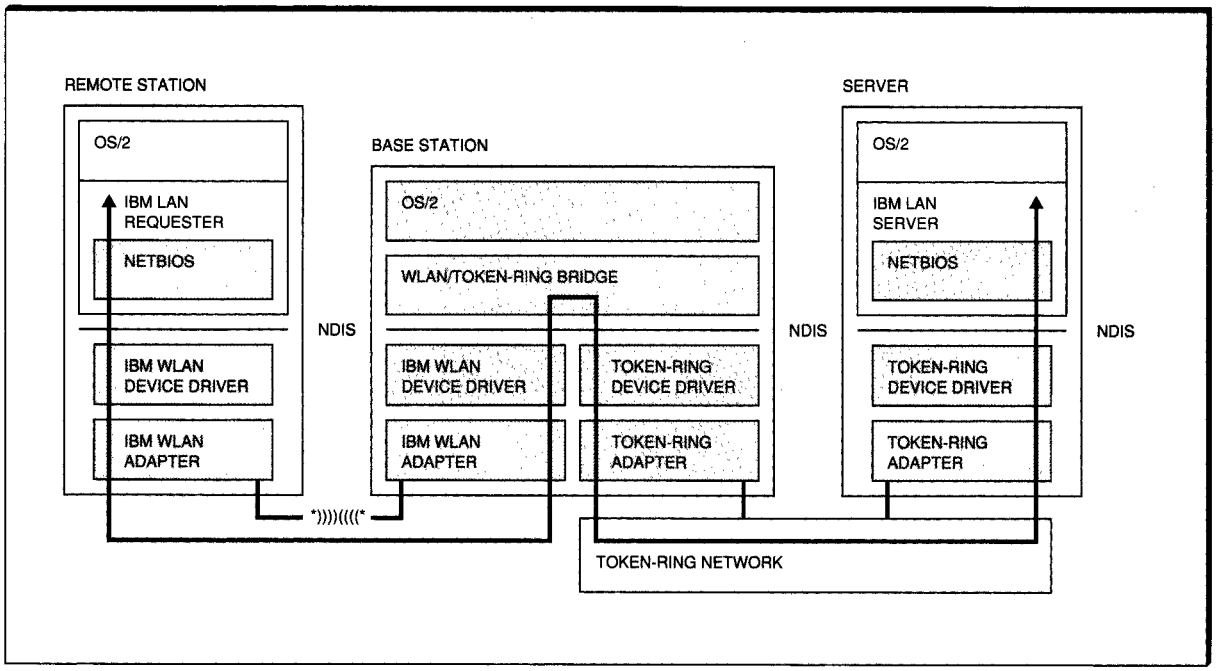

ample of such a configuration is shown in Figure 7, which shows the IBM LAN Server as the network operating system and NetBIOS as the protocol used between the wireless client and the wired server.

Another classical way to interconnect LAN segments is by routing, which is also available for the IBM Wireless RF LAN. The Wireless RF LAN provides network adapter drivers for the base station that can operate in both the NetWare and OS/2 operating systems. Therefore, the IBM Wireless RF LAN can take advantage of router programs running in those environments to make the base station a router between the wireless cell and various types of LANs (Ethernet, token ring, PCNet*) or WANs. For example, IP routing is available with the IBM program product TCP/IP for OS/2,[36] and IPX routing is available with the Novell IPX router for NetWare. Such a configuration appears in Figure 8 where a Network File System** (NFS**) server application relies on the User Datagram Protocol/Internet Protocol (UDP/IP) stack.

For customers using an SNA backbone between their various sites, the IBM Wireless RF LAN provides a solution allowing the base station to connect the wireless cell to that backbone. Together with the gateway function of the IBM Communications Manager/2*, the IBM Wireless RF LAN allows remote stations to communicate with an IBM host directly through their base station. Figure 9 depicts a configuration in which the base station appears to its downstream remote stations as an SNA type 4 node communications controller and to the host as an SNA type 2.0 node that supports one or more logical units (LUs).[37]

Compliance with industry standards for a network card interface gives the wireless LAN user the freedom to choose the networking products that best fit the user's needs. As a result the IBM Wireless RF LAN product can be "married" with virtually any networking program product, offering a very attractive alternative to wired networking solutions.

Conclusion

The physical nature of the wireless channel and the constraints imposed by regulatory bodies have been key factors influencing the design and definition of a new, hybrid MAC protocol used in the

Figure 8 TCP/IP traffic routed in the base station

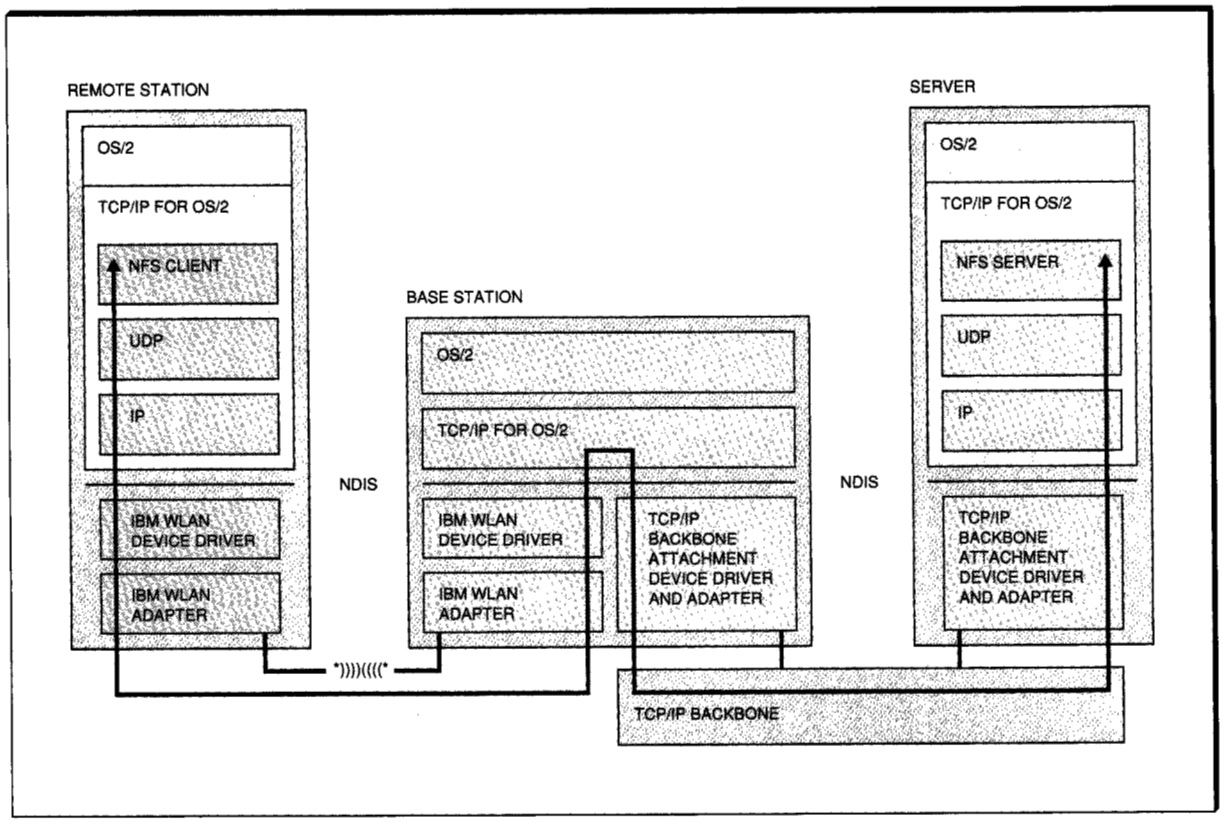

IBM Wireless RF LAN. The adaptive nature of this MAC protocol allows sustained high performance for various traffic conditions, under the constraints of a frequency-hopping spread-spectrum system operating in the 2.4-GHz ISM band. Short latency and high throughput are no longer antagonistic attributes and can be provided by the same MAC protocol. The specifics and the user perception of the wireless channel have also led to providing MAC level services such as security, data compression, and interference management.

The integration of a wireless LAN product within existing data communication networks can be achieved by relying on standard network adapter card interfaces, such as NDIS and ODI. Compliance with these interfaces allows the IBM Wireless RF LAN to be easily integrated into regular LAN environments. The end user can thus fully obtain the advantages of wireless communications (ease of installation, mobility, cost) and still rely on his or her favorite applications and networking software. The wireless connectivity can provide either an extension to an existing cabled infrastructure network, thanks to bridging, routing, or gateway functions, or a stand-alone networking solution where cables are completely eliminated.

Acknowledgments

The authors thank F. LeFevre and L. Revardel, both from CER IBM La Gaude, for their contributions to this paper. The authors also acknowledge the IBM Research Division work that preceded the final protocol. Because many people from different IBM Research laboratories contributed to this effort, we prefer to thank all of the contributors in the IBM laboratories in Hawthorne and Yorktown Heights, New York, and in Zurich as a whole, rather than to try to list all of the individuals to

Figure 9 Base station used as an SNA gateway

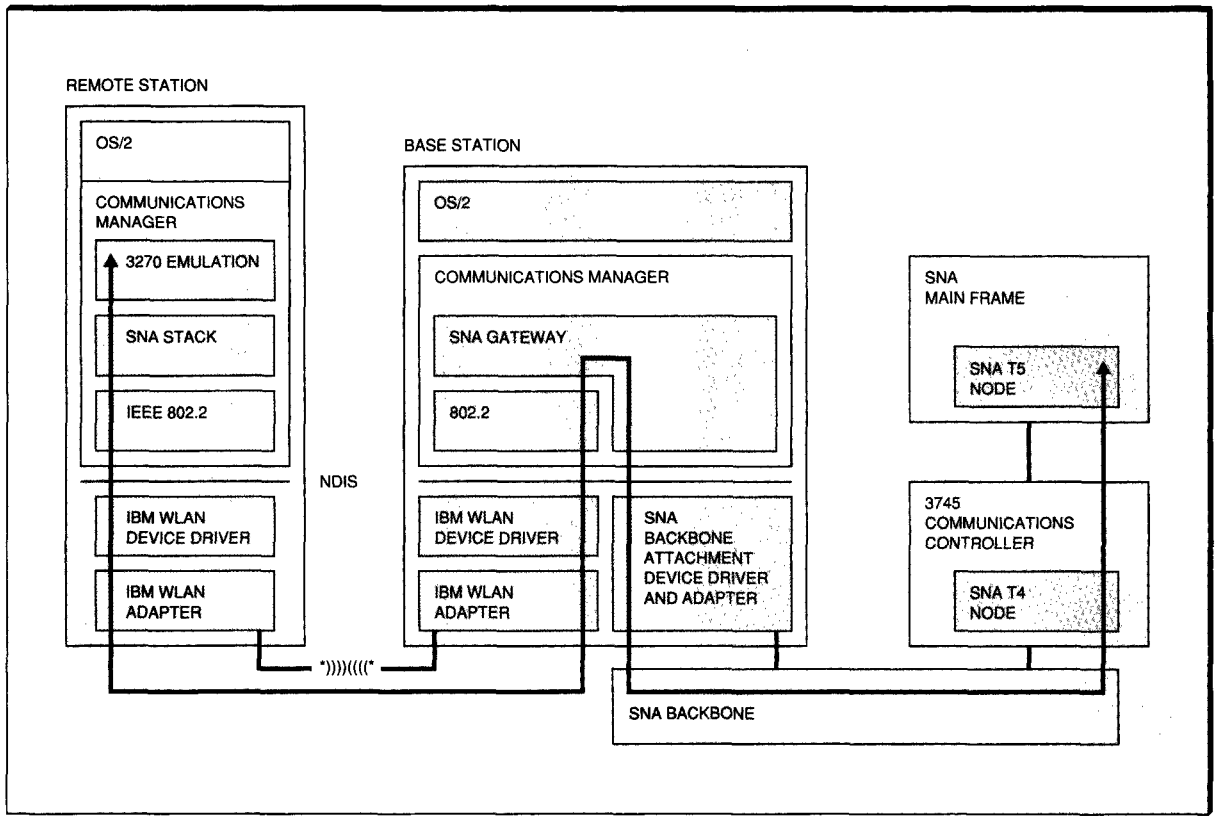

whom we are grateful. We also thank G. D. Schultz from the IBM Research Triangle Park site for his editorial suggestions and the anonymous reviewers who have significantly corrected and clarified our initial effort.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Novell, Inc., Artisoft Corporation, Microsoft Corporation, or Sun Microsystems, Inc.

Cited references and notes

1. C. Macchi et al., *Téléinformatique*, ISBN 2-04-016907-5, DUNOD Informatique, Paris (1987).
2. *Data Communications, Networks, and Systems*, T. C. Bartee, Editor, ISBN 0-672-22235-3, Howard W. Sams & Co., Indianapolis (1985).
3. D. F. Bantz and F. J. Bauchot, "Wireless LAN Design Alternatives," *IEEE Network Magazine* **8**, No. 2 (March/April 1994).
4. C. E. Cook et al., *Spread-Spectrum Communications*, ISBN 0-87942-170-3, IEEE Press, Piscataway, NJ (1983).
5. R. Skang and J. F. Hjelmstad, *Spread Spectrum in Communication*, IEEE Telecommunications Series 12, ISBN 0-86341-034-0 (1985).
6. G. R. Cooper and C. D. McGillem, *Modern Communications and Spread Spectrum*, McGraw-Hill Book Co., New York (1986).
7. W. C. Y. Lee, *Mobile Communications Design Fundamentals*, ISBN 0-672-22305-8, Howard W. Sams & Co. (1986).
8. *FCC Part 15.247*, Document 47 CFR 15.247, U.S. Government Printing Office, Washington, DC.
9. F. A. Tobagi, "Multiaccess Link Control," *Computer Network Architectures and Protocols*, P. E. Green, Jr., Editor, Plenum Press, New York (1982).
10. R. Rom and M. Sidi, *Multiple Access Protocols—Performance and Analysis*, Springer-Verlag, New York (1990).
11. C. Heide, *The CODIAC Protocol, Centralized or Distributed Integrated Access Control (CODIAC), A Wireless MAC Protocol*, IEEE Document P802.11/93-54.
12. N. M. Mitrou, Th. D. Orinos, and E. N. Protonotarius, "A Reservation Multiple Access Protocol for Microcellular Mobile Communication Systems," *IEEE Transactions on Vehicular Technology* **39**, No. 4, 340–351 (November 1990).
13. *Multiple Access Communications: Foundations for Emerging Technologies*, N. Abramson, Editor, IEEE Press, Piscataway, NJ (1993).

14. D. F. Bantz, R. T. Cato, C.-C. Huang, *Broadcast-Initiated Bipartite Frame Multi-Access Protocol*, U.S. Patent 5,123,029 (June 16, 1992).
15. H. Ahmadi, D. F. Bantz, F. J. Bauchot, A. Krishna, R. O. LaMaire, and K. S. Natarajan, *Adaptive Medium Access Control Scheme for Wireless LAN*, U.S. Patent 5,384,777 (January 24, 1995).
16. K. S. Natarajan, "A Hybrid Medium Access Protocol for Wireless LANs," *Proceedings of the 1992 IEEE International Conference on Selected Topics in Wireless Communications*, Vancouver, B.C., Canada (June 25-26, 1992), pp. 134–137.
17. F. J. Bauchot and K. S. Natarajan, *Wireless LAN MAC Protocol: 2nd Update*, IEEE Document P802.11/93-62 (1993).
18. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Document P802.11/D1 (December 1, 1994).
19. F. A. Tobagi and L. Kleinrock, "Packet Switching in Radio Channels: Part II—The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution," *IEEE Transactions on Communications* **COM-23**, No. 12, 1417–1433 (December 1975).
20. L. G. Roberts, "Aloha Packet System with and without Slots and Capture Broadcast Channels," *Computer Communications Review* **5**, 28–42 (April 1975).
21. L. Kleinrock and S. S. Lam, "Packet-Switching in a Multiaccess Broadcast Channel: Performance Evaluation," *IEEE Transactions on Communications* **COM-23**, No. 4, 410–423 (April 1975).
22. The capture effect denotes a situation where simultaneous data packet transmissions on the channel do not translate into packet loss at the receiver owing to different channel attenuations from the sources.
23. R. O. LaMaire, A. Krishna, and H. Ahmadi, "Analysis of a Wireless MAC Protocol with Client-Server Traffic and Capture," *IEEE Journal on Selected Areas in Communications* **12**, No. 8, 1299–1313 (October 1994).
24. R. O. LaMaire, *Performance of a Reservation Multiple-Access Protocol*, IEEE Document P802.11/92-108.
25. Most continuous channels are in fact periodically interrupted, due to the need to broadcast synchronizing information.
26. The slot header SH was originally designed as three different headers (AH, BH, and CH) sent at the beginning of each period.
27. R. O. LaMaire, A. Krishna, and H. Ahmadi, "Analysis of a Wireless MAC Protocol with Client-Server Traffic," *Proceedings of Infocom'93* (March/April 1993), pp. 429–438.
28. In IEEE 802.11 terminology, it is also known as the "Association."[18]
29. R. Molva, G. Tsudik, E. Van Herrenweghen, and S. Zatti, "KryptoKnight Authentication and Key Distribution Systems," *Proceedings of ESORICS'92*, Toulouse, France (October 1992).
30. D. W. Davies and W. L. Price, *Security for Computer Networks*, ISBN 0-147-90063-X, John Wiley & Sons, New York (1984).
31. ISA, MCA, and PCMCIA stand for Industry Standard Architecture, Micro Channel® Architecture, and Personal Computer Memory Card Industry Association, respectively. They are the *de facto* bus architecture standards used in the various IBM-type personal computer environments.
32. S. E. Turner, "Small Is Beautiful: How V.42bis Cuts Cost," *Data Communications International* (December 1990), pp. 83–86.
33. T. C. Bell, J. G. Cleary, and I. H. Witten, *Text Compression*, Prentice-Hall Advanced Reference Series, Prentice-Hall, Inc., Englewood Cliffs, NJ (1990).
34. *3Com/Microsoft LAN Manager Network Driver Interface Specification*, Version 2.0.1 FINAL (October 8, 1990).
35. *NetWare from IBM: Network Protocols and Standards*, GG24-3890-00, International Technical Support Center, IBM Corporation, Austin, TX (August 1992); available through IBM branch offices.
36. *TCP/IP Tutorial and Technical Overview*, GG24-3376-03, IBM Corporation; available through IBM branch offices.
37. *LAN Concept and Products*, GG24-3178-03, IBM Corporation; available through IBM branch offices.

Accepted for publication March 28, 1995.

Frédéric J. Bauchot *C.E.R. IBM France, Le Plan du Bois, 06610 La Gaude, France (electronic mail: fbauchot@vnet.ibm.com).* Dr. Bauchot received his Engineering Diploma from the Ecole Nationale Supérieure des Télécommunications at Paris in 1981. He then worked two years at the Centre National de Recherche Scientifique and received a Docteur Ingénieur degree in digital telecommunications and signal processing. He joined IBM La Gaude in 1984 and has held various development and architecture positions in signal converter products (modems), SNA concentrator products, and communications controller products. Since 1991, Dr. Bauchot has been the lead architect of the Mobile Networking Products Unit in La Gaude. He holds 12 patents in the field of wireless communications and is a senior member of the IEEE.

Fabien Lanne *C.E.R. IBM France, Le Plan du Bois, 06610 La Gaude, France (electronic mail: flanne@vnet.ibm.com).* Mr. Lanne received his degree "Diplome d'ingénieur en génie électrique" in 1985 from the Institut National des Sciences Appliquées, Rennes, France and his degree "Diplome d'ingénieur en informatique avancée" in 1986 from the Ecole Supérieure d'Electricité, Rennes, France. He joined IBM La Gaude in 1987 and held a development position in communications controller products before joining the Mobile Networking Products Unit as an architect.

Reprint Order No. G321-5574.