## Abstract

A few years ago it was recognized that the vision of a truly low-cost, low-power radio-based cable replacement was feasible. Such a ubiquitous link would provide the basis for portable devices to communicate together in an ad hoc fashion by creating personal area networks which have similar advantages to their office environment counterpart, the local area network. Bluetooth™ is an effort by a consortium of companies to design a royalty-free technology specification enabling this vision. This article describes the radio system behind the Bluetooth concept. Designing an ad hoc radio system for worldwide usage poses several challenges. The article describes the critical system characteristics and motivates the design choices that have been made.

# The Bluetooth Radio System

### JAAP C. HAARTSEN, ERICSSON RADIO SYSTEMS B.V.

*I*n the last decades, progress in microelectronics and very large scale integration (VLSI) technology has fostered the widespread use of computing and communication devices for commercial usage. The success of consumer products like PCs, laptops, personal digital assistants (PDAs), cell phones, cordless phones, and their peripherals has been based on continuous cost and size reduction. Information transfer between these devices has been cumbersome, mainly relying on cables. Recently, a new universal radio interface has been developed enabling electronic devices to communicate wirelessly via short-range ad hoc radio connections. The Bluetooth technology — which has gained the support of leading manufacturers like Ericsson, Nokia, IBM, Toshiba, Intel, and many others — eliminates the need for wires, cables, and the corresponding connectors between cordless or mobile phones, modems, headsets, PDAs, computers, printers, projectors, and so on, and paves the way for new and completely different devices and applications. The technology enables the design of low-power, small-sized, low-cost radios that can be embedded in existing (portable) devices. Eventually, these embedded radios will lead toward ubiquitous connectivity and truly connect everything to everything. Radio technology will allow this connectivity to occur without any explicit user interaction.

This article describes the basic design and technology trade-offs which have led to the Bluetooth radio system. We describe some fundamental issues regarding ad hoc radio systems. We give an overview of the Bluetooth system itself with the emphasis on the radio architecture. It explains how the system has been optimized to support ad hoc connectivity. We also describe the Bluetooth specification effort.

## Ad Hoc Radio Connectivity

The majority of radio systems in commercial use today are based on a cellular radio architecture. A mobile network established on a wired backbone infrastructure uses one or more base stations placed at strategic positions to provide local cell coverage; users apply portable phones, or more generic mobile terminals, to access the mobile network; the terminals maintain a connection to the network via a radio link to the base stations. There is a strict separation between the base stations and the terminals. Once registered to the network, the terminals remain locked to the control channels in the network, and connections can be established and released according to the control channel protocols. Channel access, channel allocation, traffic control, and interference minimization are neatly con-

trolled by the base stations. Examples of these conventional radio systems are the public cellular phone systems like Global System for Mobile Communications (GSM), D-AMPS, and IS-95 [1–3], but also private systems like wireless local area network (WLAN) systems based on 802.11 or HIPERLAN I and HIPERLAN II [4–6], and cordless systems like Digital Enhanced Cordless Telecommunications (DECT) and Personal Handyphone System (PHS) [7, 8].

In contrast, in truly ad hoc systems, there is no difference between radio units; that is, there are no distinctive base stations or terminals. Ad hoc connectivity is based on peer communications. There is no wired infrastructure to support connectivity between portable units; there is no central controller for the units to rely on for making interconnections; nor is there support for coordination of communications. In addition, there is no intervention of operators. For the scenarios envisioned by Bluetooth, it is highly likely that a large number of ad hoc connections will coexist in the same area without any mutual coordination; that is, tens of ad hoc links must share the same medium at the same location in an uncoordinated fashion. This is different from ad hoc scenarios considered in the past, where ad hoc connectivity focused on providing a single (or very few) network(s) between the units in range [4, 5]. For the Bluetooth applications, typically many independent networks overlap in the same area. This will be indicated as a scatter ad hoc environment. Scatter ad hoc environments consist of multiple networks, each containing only a limited number of units. The difference between a conventional cellular environment, a conventional ad hoc environment, and a scatter ad hoc environment is illustrated in Fig. 1. The environmental characteristics the ad hoc radio system has to operate in have a major impact on the following fundamental issues:
• Applied radio spectrum
• Determining which units are available to connect to
• Connection establishment
• Multiple access scheme
• Channel allocation
• Medium access control
• Service prioritization (i.e., voice before data)
• (Mutual) interference
• Power consumption

Ad hoc radio system have been in use for some time, for example, walky-talky systems used by the military, police, fire departments, and rescue teams in general. However, the Bluetooth system is the first commercial ad hoc radio system envisioned to be used on a large scale and widely available to the public.

# Bluetooth Radio System Architecture

In this section the technical background of the Bluetooth radio system is presented. It describes the design trade-offs made in order to optimize the ad hoc functionality and addresses the issues listed above.

## Radio Spectrum

The choice of radio spectrum is first determined by the lack of operator interaction. The spectrum must be open to the public without the need for licenses. Second, the spectrum must be available worldwide. The first Bluetooth applications are targeted at the traveling businessperson who connects his/her portable devices wherever he/she goes. Fortunately, there is an unlicensed radio band that is globally available. This band, the Industrial, Scientific, Medical (ISM) band, is centered around 2.45 GHz and was formerly reserved for some professional user groups but has recently been opened worldwide for commercial use. In the United States, the band ranges from 2400 to 2483.5 MHz, and the FCC Part 15 regulations apply. In most parts of Europe,[1] the same band is available under the ETS-300328 regulations. In Japan, recently the band from 2400 to 2500 MHz has been allowed for commercial applications and has been harmonized with the rest of the world. Summarizing, in most countries of the world, free spectrum is available from 2400 MHz to 2483.5 MHz, and harmonization efforts are ongoing to have this radio band available truly worldwide.

The regulations in different parts of the world differ. However, their scope is to enable fair access to the radio band by an arbitrary user. The regulations generally specify the spreading of transmitted signal energy and maximum allowable transmit power. For a system to operate globally, a radio concept has to be found that satisfies all regulations simultaneously. The result will therefore be the minimum denominator of all the requirements.

## Interference Immunity

Since the radio band is free to be accessed by any radio transmitter as long as it satisfies the regulations, interference immunity is an important issue. The extent and nature of the interference in the 2.45 GHz ISM band cannot be predicted. Radio transmitters may range, for example, from 10 dBm baby monitors to 30 dBm WLAN access points. With high probability, the different systems sharing the same band will not be able to communicate. Coordination is therefore not possible. More of a problem are the high-power transmitters covered by the FCC part 18 rules which include, for example, microwave ovens and lighting devices. These devices fall outside the power and spreading regulations of part 15, but still coexist in the 2.45 GHz ISM band. In addition to interference from external sources, co-user interference must be taken into account, which results from other Bluetooth users.

Interference immunity can be obtained by interference suppression or avoidance. Suppression can be obtained by coding or direct-sequence spreading. However, the dynamic range of the interfering and intended signals in an ad hoc, uncoordinated radio environment can be huge. Taking into account the distance ratios and power differences of uncoordinated transmitters, near-far ratios in excess of 50 dB are no exception. With desired user rates on the order of 1 Mb/s and beyond, practically attained coding and processing gains are inadequate. Instead, interference avoidance is more attractive

[1] In France and Spain the exact location of the band differs, and the band is smaller.
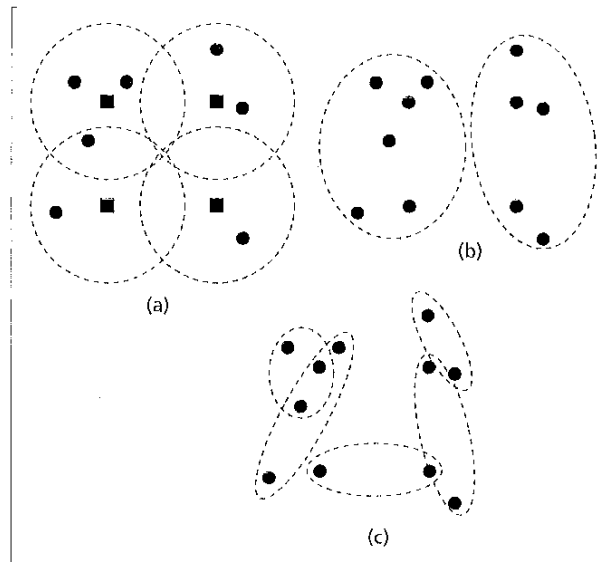


■ **Figure 1.** *Topologies for: a) cellular radio systems with squares representing stationary base stations; b) conventional ad hoc systems; and c) scatter ad hoc systems.*

since the desired signal is transmitted at points in frequency and/or time where interference is low or absent. Avoidance in time can be an alternative if the interference concerns a pulsed jammer and the desired signal can be interrupted. Avoidance in frequency is more practical. Since the 2.45 GHz band provides about 80 MHz of bandwidth and most radio systems are band-limited, with high probability a part of the radio spectrum can be found where there is no dominant interference. Filtering in the frequency domain provides the suppression of the interferers at other parts of the radio band. The filter suppression can easily arrive at 50 dB or more.

## Multiple Access Scheme

The selection of the multiple access scheme for ad hoc radio systems is driven by the lack of coordination and the regulations in the ISM band. Frequency-division multiple access (FDMA) is attractive for ad hoc systems since channel orthogonality only relies on the accuracy of the crystal oscillators in the radio units. Combined with an adaptive or dynamic channel allocation scheme, interference can be avoided. Unfortunately, pure FDMA does not fulfill the spreading requirements set in the ISM band. Time-division multiple access (TDMA) requires strict time synchronization for channel orthogonality. For multiple collocated ad hoc connections, maintaining a common timing reference becomes rather cumbersome. Code-division multiple access (CDMA) offers the best properties for ad hoc radio systems since it provides spreading and can deal with uncoordinated systems. Direct sequence (DS)-CDMA is less attractive because of the near-far problem which requires coordinated power control or excessive processing gain. In addition, as in TDMA, DS-CDMA channel orthogonality requires a common timing reference. Finally, for higher user rates, rather high chip rates are required, which is less attractive because of the wide bandwidth (interference immunity) and higher current consumption. Frequency-hopping (FH)-CDMA combines a number properties which make it the best choice for ad hoc radio systems. On average the signal can be spread over a large frequency range, but instantaneously only a small bandwidth is occupied, avoiding most of the potential interference in the ISM band. The hop carriers are orthogonal, and the interference on adjacent hops can effectively be suppressed by filter-

**■Figure 2.** *An illustration of the FH/TDD channel applied in Bluetooth.*

ing. The hop sequences will not be orthogonal (coordination of hop sequences is not allowed by the FCC rules anyway), but narrowband and co-user interference is experienced as short interruptions in the communications, which can be overcome with measures at higher-layer protocols.

Bluetooth is based on FH-CDMA. In the 2.45 GHz ISM band, a set of 79 hop carriers have been defined at a 1 MHz spacing.[2] The channel is a hopping channel with a nominal hop dwell time of 625 μs. A large number of pseudo-random hopping sequences have been defined. The particular sequence is determined by the unit that controls the FH channel, which is called the *master*. The native clock of the master unit also defines the phase in the hopping sequence. All other participants on the hopping channel are *slaves*; they use the master identity to select the same hopping sequence and add time offsets to their respective native clocks to synchronize to the frequency hopping. In the time domain, the channel is divided into slots. The minimum dwell time of 625 μs corresponds to a single slot. To simplify implementation, full-duplex communications is achieved by applying time-division duplex (TDD). This means that a unit alternately transmits and receives. Separation of transmission and reception in time effectively prevents crosstalk between the transmit and receive operations in the radio transceiver, which is essential if a one-chip implementation is desired. Since transmission and reception take place at different time slots, transmission and reception also take place at different hop carriers. Figure 2 illustrates the FH/TDD channel applied in Bluetooth. Note that multiple ad hoc links will make use of different hopping channels with different hopping sequences and may have misaligned slot timing.

## The Modulation Scheme

In the ISM band, the signal bandwidth of FH systems is limited to 1 MHz. For robustness, a binary modulation scheme was chosen. With the above-mentioned bandwidth restriction, the data rates are limited to about 1 Mb/s. For FH systems and support for bursty data traffic, a noncoherent detection scheme is most appropriate. Bluetooth uses Gaussian-shaped frequency shift keying (FSK) modulation with a nominal modulation index of $k = 0.3$. Logical ones are sent as positive frequency deviations, logical zeroes as negative frequency deviations. Demodulation can simply be accomplished by a limiting FM discriminator. This modulation scheme allows the implementation of low-cost radio units.

## Medium Access Control

Bluetooth has been optimized to allow a large number of uncoordinated communications to take place in the same area. Unlike other ad hoc solutions where all units in range

share the same channel, Bluetooth has been designed to allow a large number of independent channels, each channel serving only a limited number of participants. With the considered modulation scheme, a single FH channel in the ISM band only supports a gross bit rate of 1 Mb/s. This capacity has to be shared by all participants on the channel. Theoretically, the spectrum with 79 carriers can support 79 Mb/s. In the user scenarios targeted by Bluetooth, it is highly unlikely that all units in range need to share information among all of them. By using a large number of independent 1 Mb/s channels to which only the units are connected that really want to exchange information, the 80 MHz is exploited much more effectively. Due to nonorthogonality of the hop sequences, the theoretical capacity of 79 Mb/s cannot be reached, but is at least much larger than 1 Mb/s.

An FH Bluetooth channel is associated with a piconet. As mentioned earlier, the piconet channel is defined by the identity (providing the hop sequence) and system clock (providing the hop phase) of a master unit. All other units participating in the piconet are slaves. Each Bluetooth radio unit has a free-running system or native clock. There is not a common timing reference, but when a piconet is established, the slaves add offsets to their native clocks to synchronize to the master. These offsets are released again when the piconet is cancelled, but can be stored for later use. Different channels have different masters and therefore also different hopping sequences and phases. The number of units that can participate on a common channel is deliberately limited to eight (one master and seven slaves) in order to keep a high-capacity link between all the units. It also limits the overhead required for addressing. Bluetooth is based on peer communications. The master/slave role is only attributed to a unit for the duration of the piconet. When the piconet is cancelled, the master and slave roles are cancelled. Each unit can become a master or slave. By definition, the unit that establishes the piconet becomes the master.

In addition to defining the piconet, the master also controls the traffic on the piconet and takes care of access control. Access is completely contention free. The short dwell time of 625 μs only allows the transmission of a single packet. A contention-based access scheme would provide too much overhead and is not efficient in the short dwell time Bluetooth applies. In Bluetooth, the master implements centralized control; only communication between the master and one or more slaves is possible. The time slots are alternately used for master transmission and slave transmission. In the master transmission, the master includes a slave address of the unit for which the information is intended. In order to prevent collisions on the channel due to multiple slave transmissions, the master applies a polling technique: for each slave-to-master slot, the master decides which slave is allowed to transmit. This decision is performed on a per-slot basis: only the slave addressed in the master-to-slave slot directly preceding the slave-to-master slot is allowed to transmit in this slave-to-master slot. If the master has information to send to a specific slave, this slave is polled implicitly and can return information. If the master has no information to send, it has to poll the slave explicitly with a short poll packet. Since the master schedules the traffic in both the uplink and downlink, intelligent scheduling algorithms have to be used that take into account the slave characteristics. The master control effectively prevents collisions between the participants on the piconet channel. Independent collocated piconets may interfere when they occasionally use the same hop carrier. A type of ALOHA is applied: information is transmitted without checking for a clear carrier (no listen-before-talk). If the information is

---

[2] *Currently, for France and Spain a reduced set of 23 hop carriers has been defined at a 1 MHz carrier spacing.*

received incorrectly, it is retransmitted at the next transmission opportunity (for data only). Due to the short dwell time, collision avoidance schemes are less appropriate for FH radio. For each hop, different contenders are encountered. Backoff mechanisms are therefore less efficient.



72 bits   54 bits          0–2745 bits

Access    Packet           Payload
code      header

**■ Figure 3.** *The format of packets applied in Bluetooth.*

## Packet-Based Communications

The Bluetooth system uses packet-based transmission: the information stream is fragmented into packets. In each slot, only a single packet can be sent. All packets have the same format, starting with an access code, followed by a packet header, and ending with the user payload (Fig. 3).

The access code has pseudo-random properties and is used as a direct-sequence code in certain access operations. The access code includes the identity of the piconet master. All packets exchanged on the channel are identified by this master identity. Only if the access code matches the access code corresponding to the piconet master will the packet be accepted by the recipient. This prevents packets sent in one piconet falsely being accepted by units of another piconet that happens to land on the same hop carrier. In the receiver, the access code is matched against the anticipated code in a sliding correlator. This correlator provides the direct-sequence processing gain. The packet header contains link control information: a 3-bit slave address to separate the slaves on the piconet, a 1-bit acknowledgment/negative acknowledgmennt (ACK/NACK) for the automatic repeat request (ARQ) scheme, a 4-bit packet type code to define 16 different payload types, and an 8-bit header error check (HEC) code which is a cyclic redundancy check (CRC) code to detect errors in the header. The packet header is limited to 18 information bits in order to restrict the overhead. The header is further protected by 1/3 rate forward error correction (FEC) coding. Bluetooth defines four control packets:

- The ID or identification packet: Only consists of the access code; used for signaling
- The NULL packet: Only has an access code and a packet header; used if link control information carried by the packet header has to be conveyed
- The POLL packet: Similar to the NULL packet; used by the master to force slaves to return a response
- The FHS packet: An FH-synchronization packet; used to exchange real-time clock and identity information between the units; contains all the information to get two units hop synchronized

The remaining 12 type codes are used to define packets for synchronous and asynchronous services. These 12 types are divided into three segments. Segment 1 specifies packets that fit into a single slot, segment 2 specifies 3-slot packets, and segment 3 specifies 5-slot packets. Multislot packets are sent on a single-hop carrier. The hop carrier which is valid in the first slot is used for the remainder of the packet; therefore, there is no frequency switch in the middle of a packet. After the packet has been sent, the hop carrier as specified by the current master clock value is used (Fig. 4). Note that only an odd number of multislot packets have been defined, which guarantees that the TX/RX timing is maintained.

On the slotted channel, synchronous and asynchronous links

have been defined, as will be further explained later. The interpretation of packet type is different for synchronous and asynchronous links. Currently, asynchronous links support payloads with or without a 2/3-rate FEC coding scheme. In addition, on these links single-slot, three-slot, and five-slot packets are available. The maximum user rate that can be obtained over the asynchronous link is 723.2 kb/s. In that case, a return link of 57.6 kb/s can still be supported. Link adaptation can be applied on the asynchronous link by changing the packet length and FEC coding depending on link conditions. The payload length is variable and depends on the available user data. However, the maximum length is limited by the minimum switching time between RX and TX, which is specified at 200 μs. This switching time seems large, but allows the use of open-loop voltage controlled oscillators (VCOs) for direct modulation and provides time for packet processing between RX and TX; this is also discussed later. For synchronous links, only single-slot packets have been defined. The payload length is fixed. Payloads with 1/3-rate FEC, 2/3-rate, or no FEC are supported. The synchronous link supports a full-duplex link with a user rate of 64 kb/s in both directions.

## Physical Link Definition

The Bluetooth link supports both synchronous services such as voice traffic, and asynchronous services such as bursty data traffic. Two physical link types have been defined:

- The synchronous connection-oriented (SCO) link
- The asynchronous connectionless (ACL) link

The SCO link is a point-to-point link between the master and a single slave. The link is established by reservation of duplex slots at regular intervals. The ACL link is a point-to-multipoint link between the master and all the slaves on the piconet. The ACL link can use all of the remaining slots on the channel not used for SCO links. The traffic over the ACL link is scheduled by the master. The slotted structure of the piconet channel allows effective mixing of the synchronous and asynchronous links. An example of a channel



**■ Figure 4.** *The frequency and timing characteristics of single-slot, three-slot, and five-slot packets.*

**■ Figure 5.** *An example of mixing synchronous SCO links and asynchronous ACL links on a single piconet channel.*



**■ Figure 6.** *Frequency and timing behavior for a Bluetooth paging unit.*

with SCO and ACL links is illustrated in Fig. 5. For the SCO link and ACL link, different packet types have been defined.

## Connection Establishment

A critical design issue in ad hoc radio systems is connection establishment. How do units find each other, and how do they make connections? In Bluetooth, three elements have been defined to support connection establishment: scan, page, and inquiry. A unit in idle mode wants to sleep most of the time to save power. However, in order to allow connections to be made, the unit frequently has to listen whether other units want to connect. In truly ad hoc systems, there is no common control channel a unit can lock to in order to listen for page messages, as is common in conventional (cellular) radio systems. In Bluetooth, a unit periodically wakes up to listen for its identity. However, the explicit identity is not used, but the access code derived from this identity. When a Bluetooth unit wakes up to scan, it opens its sliding correlator which is matched to the access code derived from its own identity. The scan window is a little longer than 10 ms. Every time the unit wakes up, it scans at a different hop carrier. This is required by the regulations, which do not permit a fixed wake-up frequency, and also provides the necessary interference immunity. The Bluetooth wake-up hop sequence is only 32 hops in length and is cyclic. All 32 hops in the wake-up sequence are unique, and they span at least 64 MHz of the 80 MHz available. The sequence is pseudo-random and unique for each Bluetooth device. The sequence is derived from the unit identity. The phase in the sequence is determined by the free-running native clock in the unit. Thus, during idle mode, the native clock is used to schedule wake-up operations. It will be understood that a trade-off has to be made between idle mode power consumption and response time: increasing the sleep time will reduce power consumption, but will prolong

the time before an access can be made. The unit that wants to connect has to solve the frequency-time uncertainty: it does not know when the idle unit will wake up and on which frequency. The burden of solving this uncertainty is deliberately placed at the paging unit because this will require power consumption. Since a radio unit will be in idle mode most of the time, the paging unit should take the power burden. We first assume that the paging unit knows the identity of the unit to which it wants to connect. Then it knows the wake-up sequence and can also generate the access code which serves as the page message. The paging unit then transmits the access code repeatedly at different frequencies: every 1.25 ms; the paging unit transmits two access codes and listens twice for a response (Fig. 6).

Consecutive access codes are transmitted on different hops selected from the wake-up sequence. In a 10 ms period 16 different hop carriers are visited, which represent half of the wake-up sequence. The paging unit transmits the access code on these 16 frequencies cyclically for the duration of the sleep period of the idle unit. If the idle unit wakes up in any of these 16 frequencies, it will receive the access code and a connection setup procedure follows. However, since the paging unit does not know the phase the idle unit is using, the idle unit can equally well wake up in any of the 16 remaining frequencies in the 32-hop wake-up sequence. Therefore, if the paging unit does not receive a response from the idle unit after a time corresponding to the sleep time, it will transmit the access code repeatedly on the hop carriers in the remaining half of the sequence.[3] The maximum access delay therefore amounts to twice the sleep time. When the idle unit receives the page message, it notifies the paging unit by returning a message, which again is the access code derived from the idle unit's identity. Thereafter the paging unit transmits an FHS packet which contains all of the pager's information (e.g., identity and clock). This information is then used by both the paging unit and the idle unit to establish a piconet; that is, the paging unit becomes the master using its identity and clock to define the FH channel, and the idle unit becomes the slave.

The above-described paging process assumes that the paging unit has no knowledge at all of the clock in the idle unit. However, if the units have met before, the paging unit will have an estimate of the clock in the idle unit. When units connect, they exchange their clock information, and the time offsets between their free-running native clocks are stored. This offset is only accurate during the connection; when the connection is released, the offset information becomes less reliable due to clock drifts. The reliability of the offsets is inversely proportional to the time elapsed since the last con-

---

[3] In determining the hop carriers of the second half of the sequence, the paging unit takes into account that the clock in the idle unit also progresses. The remaining half will therefore have one carrier in common with the first half.
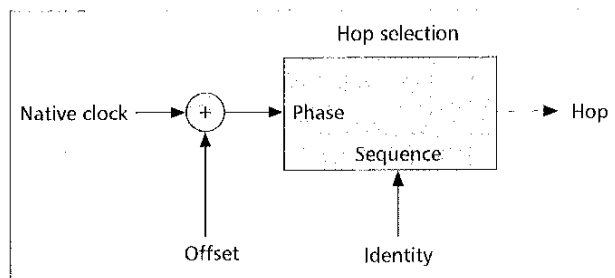
**■ Figure 7.** *The basic concept of hop selection in Bluetooth.*

nection. However, the paging unit can exploit the offset information to estimate the phase of the idle unit. Suppose the clock estimate of the idle unit in the paging unit is $k'$. If $f(m)$ is the hop in the wake-up sequence at time $m$, the paging unit will assume that the idle unit will wake up in $f(k')$. But since in 10 ms it can cover 16 different frequencies, it will also transmit the access code $a$ hop frequencies before and after $f(k')$ or $f(k' - 8)$, $f(k' - 7)$, ... $f(k')$, $f(k' + 1)$, ... $f(k' + 7)$. As a result, the phase estimate in the paging unit can be off by $-8$ or $+7$ while it still covers the wake-up frequency of the unit in idle mode. With a free-running clock accuracy of $\pm 250$ ppm, the clock estimate $k'$ is still useful at least 5 hr after the last connection. In this case, the average response time is reduced to half the sleep time.

To establish a connection, the identity of the recipient is required to determine the page message and wake-up sequence. If this information is not known, a unit that desires to make a connection may broadcast an inquiry message that induces recipients to return their address and clock information. With the inquiry procedure, the inquirer can determine which units are in range and what their characteristics are. The inquiry message is again an access code, but derived from a reserved identity (the inquiry address). Idle units also listen to the inquiry message according to a 32-hop inquiry sequence. Units that receive the inquiry message return an FHS packet which includes, among other things, their identity and clock information. For the return of the FHS packet a random backoff mechanism is used to prevent multiple recipients transmitting simultaneously.

During the page and inquiry procedures, 32 hop carriers are used. For pure hopping systems, at least 75 hop carriers must be used. However, during the page and inquiry procedures, only an access code is used for signaling. This access code is used as a direct-sequence code. The processing gain obtained from this direct-sequence code combined with the processing gain obtained from the 32-hop sequence provides sufficient processing gain to satisfy the regulations for hybrid DS/FH systems. Thus, during the page and inquiry procedures the Bluetooth system acts like a hybrid DS/FH system, whereas during the connection it acts as a pure FH system.

## Hop Selection Mechanism

Bluetooth applies a special hop selection mechanism. The hop selection mechanism can be considered a black box with an identity and clock in, and a hop carrier out (Fig. 7). The mechanism satisfies the following requirements:
- The sequence is selected by the unit identity, the phase by the unit clock.
- The sequence cycle covers about 23 hours.
- 32 consecutive hops span about 64 MHz of spectrum.

---

[4] *For 23-hop systems, a corresponding scheme is constructed with 16-hop segments and a 23-hop list.*

- On average, all frequencies are visited with equal probability.
- The number of hop sequences is very large.
- By changing the clock and/or identity, the selected hop changes instantaneously.

Note that no extra effort has been taken to make the sequences orthogonal. With only 79 hop carriers, the number of orthogonal sequences is rather limited. The first requirement supports the piconet concept where the master unit defines the hop channel by its identity and clock. The second requirement prevents repetitions in the interference pattern when several piconets are collocated. Repetitive interference is detrimental for synchronous services such as voice. The spanning requirement provides maximal interference immunity by spreading as much as possible over a short time interval. Again, this is most important for voice services. It also provides the desired features for the wake-up and inquiry sequences which are 32 hops in length. Over a larger interval, regulations require that all carriers are visited with equal probability. Since many piconets can coexist in the same area, many hop patterns must be available. This excludes the use of prestored sequences: the sequences are generated on the fly by logic circuitry. Finally, the last requirement provides flexibility to run backward and forward in the sequence by running the clock backward or forward, which is attractive in the page and inquiry procedures. In addition, it supports jumping between piconets: a unit can jump from one piconet to another by merely changing the master parameters (i.e., identity and clock). The latter requirement excludes the use of a memory in the algorithm: only combinatorial logic circuitry is used.

The selection mechanism is illustrated in Fig. 8.[4] In the first block, the identity selects a 32-hop subsequence with pseudo-random properties. The least significant part of the clock hops through this sequence according to the slot rate (1600 slots/s). The first block thus provides an index in a 32-hop segment. The segments are mapped on the 79-hop carrier list. The carrier list is constructed in such a fashion that even-numbered hops are listed in the first half of the list, odd-numbered hops in the second half of the list. An arbitrary segment of 32 consecutive list elements spans about 64 MHz. For the paging and inquiry procedures, the mapping of the 32-hop segment on the carrier list is fixed. When the clock runs, the same 32-hop sequence and 32 hop carriers will be used. However, different identities will map to different segments and different sequences, so the wake-up hop sequences of different units are well randomized. During the connection, the more significant part of the clock affects both sequence selection and segment mapping: after 32 hops (one segment) the sequence is altered, and the segment is shifted in the forward
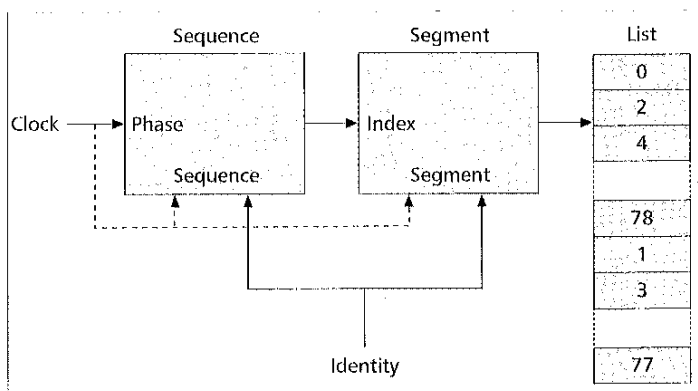


**■ Figure 8.** *The hop selection mechanism; the dashed line for the more significant clock part is used in connection mode only.*

---

**■ Figure 9.** *An example of retransmission operation in Bluetooth.*

direction by half its size (16 hops). Segments, each 32 hops in length, are concatenated, and the random selection of the index changes for each new segment; the segments slide through the carrier list, and on average all carriers are visited with equal probability. Changing the clock and/or identity will directly change the sequence and segment mapping.

## Error Correction

Bluetooth includes both FEC and packet retransmission schemes. For FEC, a 1/3-rate code and a 2/3-rate FEC code are supported. The 1/3-rate code merely uses a 3-bit repeat coding with majority decision at the recipient. With the repeat coding, extra gain is obtained due to the reduction of instantaneous bandwidth. As a result, intersymbol interference (ISI) introduced by the receive filtering is decreased. The 1/3-rate code is used for the packet header, and can additionally be applied on the payload of the synchronous packets on the SCO link. For the 2/3-rate FEC code, a shortened Hamming code is used. Error trapping can be applied for decoding. This code can be applied on both the payload of the synchronous packets on the SCO link and the payload of the asynchronous packets on the ACL link. The applied FEC codes are very simple and fast in encoding and decoding operations, which is a requirement because of the limited processing time between RX and TX. This will be further apparent in the next paragraph.

On the ACL link, an ARQ scheme can be applied. In this scheme, a packet retransmission is carried out if the reception of the packet is not acknowledged. Each payload contains a CRC to check for errors. Several ARQ schemes have been considered like stop-and-wait ARQ, go-back-N ARQ, and selective-repeat ARQ [9]. Also, hybrid schemes have been analyzed. However, to minimize complexity, overhead, and wasteful retransmissions, Bluetooth has implemented a fast-ARQ scheme where the sender is notified of the packet reception in the RX slot directly following the TX slot in which the packet was sent (Fig. 9). If the 2/3-rate FEC code is added, a type I hybrid ARQ scheme results. The ACK/NACK information is piggybacked in the packet header of the return packet. There is only the RX/TX switching time for the recipient to determine the correctness of the received packet and creating the ACK/NACK field in the header of the return packet. In addition, the ACK/NACK field in the header of the packet received indicates whether the previously sent payload was correctly received, and thus determines whether a retransmission is required or the next packet can be sent. This process is illustrated in Fig. 10. Due to the short processing time, decoding is preferably carried out on the fly while the packet is received. In addition, the simplicity of the FEC coding schemes speed up the processing. The fast-ARQ scheme is similar to the stop-and-wait ARQ scheme, but the delay has been minimized; in fact, there is no additional delay caused by the ARQ scheme.

The scheme is more efficient than go-back-N, since only failed packets are retransmitted. This is the same efficiency obtained with selective-repeat ARQ, but with reduced overhead: only a 1-bit sequencing number suffices in the fast-ARQ scheme (in order to filter out packets that are correctly received twice due to an error in the ACK/NACK field).

## Power Management

In the Bluetooth design, special attention has been paid to reduction of current consumption. In the idle mode, the unit only scans a little over 10 ms every *T* s where *T* can range from 1.28 to 3.84 s. Thus, the duty cycle is well below 1 percent. Additionally, a PARK mode has been defined where the duty cycle can be reduced even more. However, the PARK mode can only be applied after the piconet has been established. The slave can then be parked; that is, it only listens to the channel at a very low duty cycle. The slave only has to listen to the access code and the packet header (126 μs excluding guard time to account for drift) to resynchronize its clock and decide whether it can return to sleep. Since there is no uncertainty in time and frequency (the parked slave is locked to the master, similar to how cordless and cellular phones are locked to their base stations), a much lower duty cycle is achievable. Another low-power mode during connection is the SNIFF mode, in which the slave does not scan at every master-to-slave slot, but has a larger interval between scans.

In the connection state, current consumption is minimized and wasteful interference prevented by only transmitting when data is available. If no useful information needs to be exchanged, no transmission takes place. If only link control information needs to be transferred (e.g., ACK/NACK), a NULL packet without payload is sent. Since NACK is implicit, a NULL packet with NACK does not have to be sent. In longer periods of silence, the master once in a while needs to send a packet on the channel such that all slaves can resynchronize their clocks and compensate for drift. The accuracy of the clocks and the scan window length applied in the slave determines the period of this resynchronization. During continuous TX/RX operations, a unit starts to scan for the access
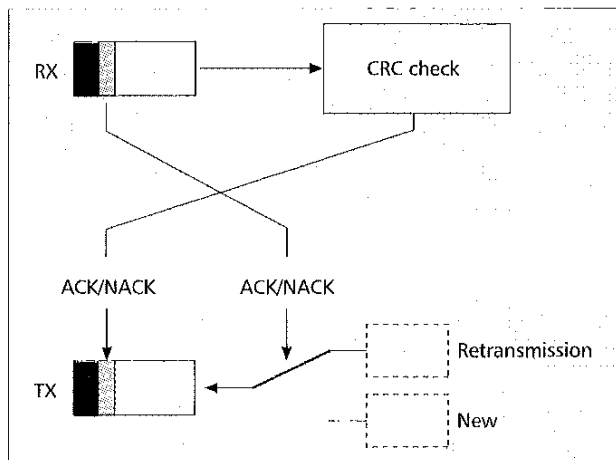


**■ Figure 10.** *ARQ mechanisms where received ACK/NAK information decides on retransmission and received payload determines transmitted ACK/NAK information.*
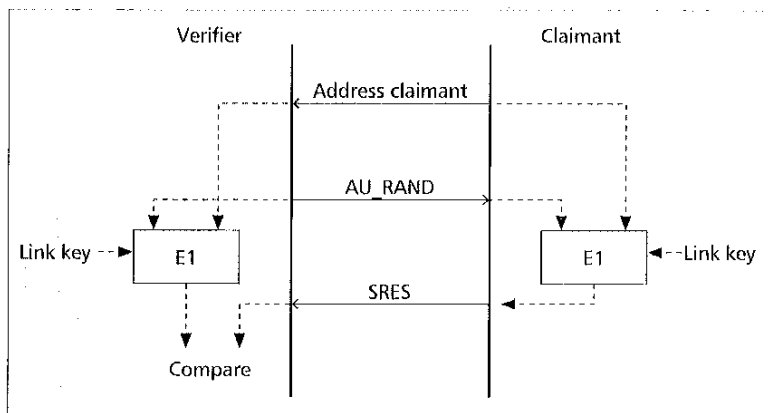
**■ Figure 11.** *The Bluetooth authentication procedure.*

code at the beginning of the RX slot. If in a certain window this access code is not found, the unit returns to sleep until the next TX slot (for the master) or RX slot (for the slave). If the access code is received (which means the received signal matches the expected access code), the header is decoded. If the 3-bit slave address does not match the recipient, further reception is stopped. The header indicates what type of packet it is and how long the packet will last; therefore, the nonaddressed recipients can determine how long they can sleep.

The nominal transmit power used by most Bluetooth applications for short-range connectivity is 0 dBm. This both restricts current consumption and keeps interference to other systems to a minimum. However, the Bluetooth radio specifications allow TX power up to 20 dBm. Above 0 dBm, closed-loop received signal strength indication (RSSI)-based power control is mandatory. This power control only compensates for propagation losses and slow fading. In the uncoordinated environment where ad hoc systems operate, interference-based power control is to say the least doubtful, especially since different types of systems with different power characteristics share the same band. Since power control cannot be coordinated among different systems, it cannot be prevented that certain systems always try to overpower their contenders, and the strongest transmitter will prevail.

## Security

Although Bluetooth is mainly intended for short-range connectivity between personal devices, some basic security elements are included to prevent unauthorized usage and eavesdropping. At connection establishment, an authentication process is carried out to verify the identities of the units involved. The authentication process uses a conventional challenge-response routine illustrated in Fig. 11. The claimant (right) transmits its claimed 48-bit address to the verifier (left). The verifier returns a challenge in the form of a 128-bit random number (AU_RAND). The AU_RAND, the claimant address, and a 128-bit common secret link key form the inputs to a computational secure hash function *E1* based on SAFER+, which produces a 32-bit signed response (SRES). The SRES produced by the claimant is sent to the verifier, which compares this result with its own SRES. Only if the two calculated SRES numbers are the same will the challenger continue with connection establishment. The authentication can be uni- or bidirectional.

In addition to the 32-bit SRES, the *E1* algorithm produces a 96-bit authenticated cipher offset (ACO). This offset is used in the encryption procedure. To prevent eavesdropping on the link, which is a danger inherent to radio communications even if the intended recipient is only at short range, the payload of each packet is encrypted. Encryption is based on stream-

ciphering; the payload bits are modulo-2 added to a binary keystream. The binary keystream is generated by a second hash function *E0* which is based on linear feedback shift registers (LFSRs). When encryption is enabled, the master sends a random number EN_RAND to the slave. Before the transmission of each packet, the LFSR is initialized by a combination of this EN_RAND, the master identity, an encryption key, and the slot number. Since the slot number changes for each new packet, the initialization is new for each packet. The encryption key is derived from the secret link key, the EN_RAND, and the ACO.

The central element in the security process is the 128-bit link key. This link key is a secret key residing in the Bluetooth hardware and is not accessible by the user. The link key is generated during an initialization phase. Two units that want to authenticate each other and establish secure links in the future have to be associated (i.e., provided with the same secret link key). An initialization phase initiated by the user is required to associate two devices. To authorize initialization, the user has to enter an identical PIN in both devices. For devices without a user interface (e.g., headsets), initialization is only possible during a short time window (e.g., after the user has pressed an initialization key). Once the initialization has been carried out, the 128-bit link keys reside in the devices and can from then on be used for automatic authentication without user interaction. In principle, the link key provides an agreement between two units. Thus, to provide security in $N$ units, $N \times (N - 1)/2$ link keys are required. Bluetooth provides methods to reduce the number of keys in certain applications. If a single unit is used by many users (e.g., a printer shared by several users), a single key is used by all users for secure communications to this single unit. In addition, methods are available to use the same encryption key for all slaves in a single piconet.

Bluetooth provides a limited number of security elements at the lowest level. More advanced security procedures (e.g., public keys, certificates) can be implemented at higher layers.

## Interpiconet Communications

The Bluetooth system has been optimized to have tens of piconets operate in the same area without noticeable performance degradation. Multiple piconets in the same area are referred to as a *scatternet*. Due to the fact that Bluetooth uses packet-based communication over slotted links, it is possible to interconnect different piconets. This means that units can participate in different piconets. However, since the radio can tune to a single hop carrier only, at any instant in time a unit can communicate in one piconet only. However, the unit can jump from one piconet to another by adjusting the piconet channel parameters (i.e., the master identity and master clock). A unit can also change role when jumping from one piconet to another. For example, a unit can be the master in one piconet at one instant in time, and be a slave in a different piconet at another instant in time. A unit can also be a slave in different piconets. However, by definition, a unit cannot be the master in different piconets, since the master parameters specify the piconet FH channel. The hop selection mechanism has been designed to allow for interpiconet communications: by changing the identity and clock input to the selection mechanism, instantaneously a new hop for the new piconet is selected. In order to make jumps between different piconets feasible, guard time has to be included in the traffic scheduling to account for the slot misalignment of different piconets. In Bluetooth, a HOLD mode has been intro-
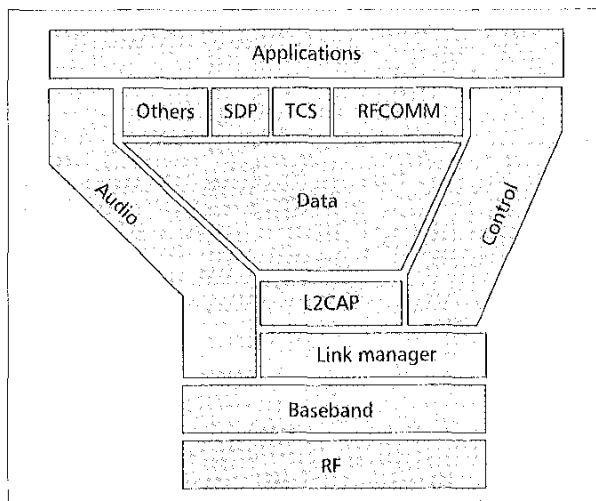
**■ Figure 12.** *The Bluetooth protocol stack.*

duced to allow a unit to temporary leave one piconet and visit another (HOLD can also be used as an additional low-power mode when no new piconet is visited during the leave). Traffic scheduling and routing in a scatternet with interpiconet communications is a challenge and still a subject for future study.

## Bluetooth Standardization

In the beginning of 1998, a Bluetooth Special Interest Group (SIG) was formed to further expand and promote the Bluetooth concept and establish an industry standard. The SIG promoters are formed by leading manufacturers of the mobile communication industry, portable computer industry, and chip integration industry: Ericsson, Nokia, IBM, Toshiba, and Intel. Version 1.0 of the specification was published in July 1999. Over 1000 companies have signed as adopters of the technology. The Bluetooth technology is royalty-free. A special certification program, including logos, is under development to guarantee Bluetooth interoperability.

The specified protocol stack of Bluetooth is shown in Fig. 12. This article has dealt mainly with the three lower layers:
• The RF layer, specifying the radio parameters
• The baseband layer, specifying the lower-level operations at the bit and packet levels (FEC operations, encryption, CRC calculations, ARQ protocol)
• The link manager (LM) layer, specifying connection establishment and release, authentication, connection and release of SCO and ACL channels, traffic scheduling, link supervision, and power management tasks

The Logical Link Control and Adaptation Protocol (L2CAP) layer has been introduced to form an interface between standard data transport protocols and the Bluetooth protocol. It handles multiplexing of higher-layer protocols, and segmentation/reassembly of large packets. The data stream crosses the LM layer, where packet scheduling on the ACL channel takes place. The audio stream is directly mapped on an SCO channel and bypasses the LM layer. The LM layer, though, is involved in the establishment of the SCO link. Between the LM layer and the application, control messages are exchanged in order to

configure the Bluetooth transceiver for the considered application. Above the L2CAP layer, RFCOMM, Telephone Control Specification (TCS), and other network protocols (e.g., TCP/IP, PPP, OBEX, Wireless Application Protocol) may reside. RFCOMM and TCS are also specified in Bluetooth and provide serial cable emulation and a cordless telephony protocol, respectively. SDP is a service discovery protocol which enables a Bluetooth unit to find the capabilities of other Bluetooth units in range. It discovers which services are available and the characteristics of these services. This can involve common services like printing, faxing, and so on, as well as more advanced services like teleconferencing, network bridging and access points, e-commerce facilities, and so on. SDP specifically addresses the Bluetooth environment; it does not specify the methods for accessing the service, for which other (non-Bluetooth) protocols can be used.

In addition to protocols which guarantee that two units speak the same language, profiles are defined. Profiles are associated with applications. The profiles specify which protocol elements are mandatory in certain applications. This concept prevents devices with little memory and processing power implementing the entire Bluetooth stack when they only require a small fraction of it. Simple devices like a headset or mouse can thus be implemented with a strongly reduced protocol stack. Profiles are dynamic in the sense that for new applications, new profiles can be added to the Bluetooth specification.

## Conclusions

In this article the Bluetooth radio system is presented. The focus is on its capabilities to provide ad hoc radio connectivity. With the restrictions set by regulations, power consumption, lack of coordination, and interference immunity, a robust radio system has evolved which provides a universal wireless interface to a large range of low-cost, portable devices. The article has also described the motivation of the various design choices.

### References

[1] M. Mouly and M.-B. Pautet, *The GSM System for Mobile Communications*, 1992.
[2] TIA/EIA/IS-136.2, "800 MHz TDMA Cellular-Radio Interface-Mobile Station-Base Station Compatibility — Traffic Channels and FSK Control Channel," Dec. 1994.
[3] TIA/EIA IS-95B, "Mobile Station-Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular Systems," 1998.
[4] IEEE 802.11, "Wireless LAN MAC and Physical Layer Specification," June 1997.
[5] ETSI RES, "High Performance Radio Local Area Network (HIPERLAN) Type 1, Functional Specifications," ETS 300 652, 1996.
[6] ETSI BRAN, "HIPERLAN Type 2, Functional Specifications," preliminary.
[7] ETSI RES, "Digital European Cordless Telecommunications (DECT), Common interface Part 1: Overview," ETS 300 175-1, 1996.
[8] "Personal Handy Phone Standard (PHS)," CRC STD-28, 1993.
[9] S. Lin and D. J. Costello, *Error Control Coding*, Prentice-Hall, 1983.

### Biography

JAAP C. HAARTSEN (jaap.haartsen@erh.ericsson.se) joined Ericsson Mobile Communications in 1991 and has since worked at sites in RTP, the United States, and Lund, Sweden in the area of wireless technology. In Sweden he worked on the foundations of the Bluetooth radio concept. Currently, he is located in Emmen, the Netherlands, where he is working with the Bluetooth system for both vcurrent and future applications. Jaap is chair of the Bluetooth air protocol group. He earned M.Sc. and Ph.D. degrees (both with honors) in electrical engineering from Delft University of Technology, the Netherlands. He holds over 25 patents.