

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 July 2002 (25.07.2002)

PCT

(10) International Publication Number
WO 02/058334 A2

- (51) International Patent Classification⁷: **H04L 12/28**
- (21) International Application Number: PCT/US02/01559
- (22) International Filing Date: 18 January 2002 (18.01.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/262,558 18 January 2001 (18.01.2001) US
60/288,294 2 May 2001 (02.05.2001) US
60/288,296 2 May 2001 (02.05.2001) US
60/333,844 28 November 2001 (28.11.2001) US
- (71) Applicant (for all designated States except US): **STRIX SYSTEMS, INC.** [US/US]; Suite 150, 310 North Westlake Boulevard, Westlake Village, CA 91362 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **BEASLEY, James** [US/US]; 634 Verdemont Circle, Simi Valley, CA 93065 (US). **DOMBROWSKI, Dennis** [US/US]; 12132 London Grove Court, Moorpark, CA 93021 (US). **KUIKEN, Matthew** [US/US]; 13903 Washington Avenue, Hawthorne, CA 90250 (US). **MERGENTHAL, Wade** [US/US]; 5255 Goldenridge Court, Camarillo, CA 93012 (US). **STEPHENS, Spencer** [US/US]; 4538 Ledge Avenue, Toluca Lake, CA 91602 (US).
- (74) Agents: **DALEY-WATSON, Christopher, J.** et al.; Perkins Coie LLP, P.O. Box 1247, Seattle, WA 98111-1247 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— without international search report and to be republished upon receipt of that report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*



WO 02/058334 A2

(54) Title: WIRELESS BASE STATION TO BASE STATION SYNCHRONIZATION IN A COMMUNICATION SYSTEM, SUCH AS A SYSTEM EMPLOYING A SHORT-RANGE FREQUENCY HOPPING SCHEME

(57) Abstract: In a network, a method is disclosed for wirelessly exchanging communications with at least one mobile unit. The network includes first and second base stations units coupled to the network, and may include a system controller. The method includes: receiving a communication signal from the second base station unit, wherein the first and second base station units are configured to employ a wireless communications protocol, and wherein the wireless communications protocol does not provide for handoff of communications links between base station units; at the first base station unit, determining if the second base station unit has been synchronized based on the communication signal; at the first base station unit, if the second base station unit is synchronized, then synchronizing an internal clock based on the synchronized second base station, wherein the synchronizing is performed without assistance from a system controller.

WIRELESS BASE STATION TO BASE STATION SYNCHRONIZATION IN A
COMMUNICATION SYSTEM, SUCH AS A SYSTEM EMPLOYING A SHORT-
RANGE FREQUENCY HOPPING SCHEME

CROSS-REFERENCE TO RELATED APPLICATION(S)

This application claims the benefit of U.S. Provisional Patent Application Nos. 60/262,558 filed January 18, 2001; 60/288,296 filed May 2, 2001; 60/288,294 filed May 2, 2001; and 60/333,844 filed November 28, 2001.

BACKGROUND

The disclosed embodiments relate to wireless systems and networks.

Cellular mobile telephony provides voice data links between users of mobile devices and fixed devices on a network. It gives users mobility without regard to how they are actually connected to the network. This is done by providing access points or base station units that can hand off the connections of mobile devices without interruption of the service. 2G (second generation) digital mobile phone service and AMPS (analog mobile phone service) are examples of such telephone networks.

Many limitations exist with such telephony networks. For example, while cellular mobile telephony networks generally provide service over large areas, such networks have limited capacity for the number of user devices that can be connected within the area covered by each BSU.

The bandwidth provided to users of this service is generally considered low bandwidth or "narrow-band access." Thus, large data applications, such as transferring of large data files, cannot be effectively performed using such networks.

Providers of this service must acquire rights to use regulated radio spectra, often at high investment costs. Thus, employing such networks can be quite expensive for many applications or users.

This type of service is often unavailable or unreliable in certain areas, such as office buildings, convention centers, and train stations, due to physical properties of these areas and/or due to the density of users seeking access to the service.

The Personal Communications Network ("PCN") provides similar features to users of mobile devices, including voice and data links to a network, while providing mobility. PCN has a user model similar to that for cellular mobile telephony, so user behavior for one can be carried over to the other. PCN does not have some of the same limitations as cellular telephony. It offers wider bandwidth, or "broadband access," and can provide greater availability with higher reliability in those particular areas that cellular telephony cannot. The RF spectra used by PCN is unlicensed, and no special access rights are required.

PCN uses wireless networking technology, including IEEE 802.11 and 802.11b, which use direct-sequence spread spectrum, and Bluetooth, which uses frequency-hopping spread spectrum. Importantly, however, the Bluetooth wireless standard does not support movement from one area of coverage or "cell" to another.

Currently, the most well-known attempt to provide features that PCN offers in a mobile environment is referred to as 3G- (third generation) wireless. 3G-wireless, however, is constrained by factors that do not constrain PCN implementations. These include heavy investment for the acquisition of spectrum usage rights.

The above wireless technologies require knowledge of which base stations in a network have coverage overlap with other, neighboring base stations within the network. Such overlap is a prerequisite for base station to base station handoff occurring without links getting dropped. Each base station is typically given the knowledge of base stations to which they can hand off. One example of this is the GSM system, in which each base station is under the control of a base station controller. See, "Mobile Communications," Jochen Schiller, pp. 84-112 (Addison-Wesley, 2000). In the GSM system, the base station controllers and the devices that coordinate them (mobile services switching center) are what decide when and where to hand off. A network can

have a potentially large number of base stations, so it can require significant time and effort to inform each base station about which other base stations can receive handoffs from it.

Overall, there is a need for a wireless networking system that solves the above problems while supplying the above benefits, as well as providing other advantages.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a schematic drawing showing a handoff of a mobile user device from one base station unit to a neighboring base station unit under one embodiment of the invention.

Figure 2A is a block diagram illustrating a single subnet Internet Protocol ("IP") architecture where all base station units and mobile units are assigned addresses within this single subnet.

Figure 2B is an enlarged block diagram of one of the base station units shown in Figure 2A.

Figure 3 is a block diagram of a multiple subnet architecture where base station units and mobile units are assigned addresses that fall within multiple IP subnets, and mobile units belong to a separate subnet dedicated to mobile units and their point of presence.

Figure 4 is a block diagram illustrating an intrasubnet handoff from one base station unit to another under the network of Figure 3.

Figure 5 is a block diagram illustrating an intersubnet handoff of a mobile unit from a base station unit in one subnet to a base station unit in another subnet under the network of Figure 3.

Figure 6 is a block diagram showing logical and actual routes of packets from a host or server to a mobile unit under the network of Figure 3.

Figure 7A is a timing diagram showing interference in unaligned prior art methods.

Figure 7B is a timing diagram showing an aligned method under an embodiment of the invention.

Figure 8 is a packet timing diagram employing an embodiment of the invention for assigning transmit and receive slots within each slot pair.

Figure 9A is a flow diagram illustrating functionality performed by a seed base station under a base station synchronization process.

Figure 9B is a flow diagram illustrating functionality performed by base station units following functionality performed by the seed base station.

Figure 9C is a flow diagram illustrating details regarding individual base station synchronization.

Figure 10 is a schematic diagram illustrating synchronization from a seed base station to neighboring base stations under a first synchronization wave.

Figure 11 is a schematic diagram illustrating a first step under a second synchronization wave.

Figure 12 is a schematic diagram illustrating a second step under the second synchronization wave.

Figure 13 is a schematic diagram illustrating a third step under the second synchronization wave.

Figure 14 is a schematic diagram illustrating a third synchronization wave.

Figure 15 is a schematic diagram illustrating synchronization when a base station is activated or re-activated after initial synchronization has been established.

Figure 16 is a block diagram illustrating software components employed by two neighboring BSUs.

Figure 17A is a data diagram illustrating a neighbor list record.

Figure 17B is a data diagram illustrating a synchronization server record.

Figure 18 is a block diagram of an alternative embodiment illustrating synchronization performed between base stations over wired media.

In the drawings, the same reference numbers identify identical or substantially similar elements or acts. To easily identify the discussion of any particular element or act, the most significant digit or digits in a reference number refer to the Figure number in which that element is first introduced (e.g., element 204 is first introduced and discussed with respect to Figure 2).

Note: the headings provided herein are for convenience and do not necessarily affect the scope or interpretation of the invention.

DETAILED DESCRIPTION

Embodiments of the invention, described below, permit network access points or base station units ("BSU") within a cellular network to synchronize communications. Via higher than normal operational RF power, and having RF controllers convert transmit slots into receive slots, wireless system base stations' transmission slot edges are time-aligned, thereby reducing data throughput loss due to channel collisions. Synchronization initialization propagates outward, autonomously, from an origin base station to neighboring base stations to form base station synchronization relationships. Thereafter, synchronization is performed periodically among the base stations according to the formed base station synchronization relationships.

Adjacent base stations are referred to here as "neighbors," "neighboring base stations," or "neighbor BSUs." For the sake of clarity, the term "neighbor BSU" is defined here as: a relationship between one BSU and nearby BSUs; when a BSU is near enough to a second BSU such that a mobile unit ("MU") linked to the second BSU can be handed off to the first BSU, then the first BSU is a neighbor BSU with respect to the second BSU. An "owner" refers to the base station that currently links an MU to a network; thus, a "neighbor" also refers to a different base station monitoring its own link quality with that MU on behalf of its owner.

The inventors found that prior systems provide for base station unit synchronization. However, such prior systems generally fall into one of two groups. The first group uses complex communication protocols, such as codivision multiple access ("CDMA") whereby pilot signals or other communication overheard in such protocols facilitate synchronization. Examples of such systems in this first group may be found in U.S. Patent Nos. 6,219,347, 6,212,398, 6,119,016, 6,018,667 and 5,875,402.

The second group of systems employ the use of mobile units to synchronize base stations, such as: establishing a priority or rank between

base stations units or provide a sync signal to mobile units (e.g., U.S. Patent No. 5,613,211); mobile units becoming synchronized with a network and providing such synchronization information to base stations (e.g., U.S. Patent Nos. 5,898,929 and 5,818,814); analyzing traffic with mobile units or others from adjacent base stations (e.g., U.S. Patent No. 6,112,100); sending special requests to a mobile unit for synchronization (e.g., U.S. Patent No. 6,185,429); and, adding functionality to mobile unit to permit it to initiate synchronization (e.g., U.S. Patent No. 5,625,653). Certain other prior art systems apparently employ specialized wired networks, such as circulating a data structure between base stations to permit synchronization (e.g., U.S. Patent No. 5,068,916), and employing a wired computer LAN (e.g., U.S. Patent No. 5,408,506, which provides a interesting summary of single- and multiple-master LAN systems, beginning in column 2). No prior system is believed to permit base station units to automatically and autonomously synchronize with each other in a relatively simple communication system employing a seed base station unit that propagates synchronization outwardly therefrom.

Described below first is a suitable system in which aspects of the invention may operate. Second, slot patterns and allocation among radios in a base station with slot patterns employed by a mobile unit are discussed. Third, a method of modifying slot usage is provided. Fourth, a method of increasing throughput by aligning slot edges and details on one embodiment for synchronizing base stations are provided. Fifth, software and functionality details are provided. Finally, alternative embodiments are presented.

Aspects of the invention are described in detail below with respect to the Bluetooth specification. Bluetooth refers to a specification designed to standardize wireless transmission between a wide variety of devices, such as personal computers ("PCs"), cordless telephones, headsets, printers, personal digital assistants ("PDAs"), etc. Bluetooth acts as a "virtual cable," whereby a computer can automatically use a mobile telecommunications device (such as a cell phone) to send data wirelessly, such as exchange e-mail, transmit data to a fax machine, etc. Bluetooth operates in the unlicensed 2.4 GHz spectrum using frequency-hopping spread spectrum, where data packets are spread across the

Bluetooth spectrum at a nominal rate of 1,600 hops per second to lessen interference and fading. According to the Bluetooth specification, Bluetooth devices operate their antennas at one of three different maximum antenna power levels, *i.e.*, 1 mW, 2.5 mW and 100 mW. The nominal link range is 10 meters, and the gross data rate is 10 Mbps, although increases may be possible. Bluetooth can support both synchronous connection oriented ("SCO") links for voice and asynchronous connectionless links ("ACL") for packet data. Details on Bluetooth may be found at <http://www.bluetooth.com>, <http://www.palowireless.com>, "Bluetooth Revealed," Miller & Bisdikian, Prentice Hall PTR, 2001, and "Bluetooth Demystified," Muller, McGraw-Hill, 2001, and in particular, "Specification of the Bluetooth System," version 1.1: <http://www.bluetooth.com/developer/specification/specification.asp>.

The following description provides specific details for a thorough understanding of, and enabling description for, embodiments of the invention. However, one skilled in the art will understand that the invention may be practiced without many of these details. In other instances, well-known structures and functions have not been shown or described in detail to avoid unnecessarily obscuring the description of the embodiments of the invention.

Suitable System

As introduced above, the Bluetooth specification defines a protocol for device-to-device and device-to-network communication in a small area. The size of the area is determined by the maximum range over which these devices can communicate and is a function of their radio performance. Communication is not possible if an MU is located outside the maximum range. When such devices communicate, they first establish a link and then maintain that link for the duration of their communication. If the link is interrupted for any reason, then communication is also interrupted.

Figure 1 shows a device-to-network link 103 shown as "link 1." Here, an MU, or mobile unit 104, represents the wireless mobile device. An access point, or picocellular system node, shown as a BSU, or base station unit 102, provides an interface to a network (not shown in Figure 1). The terms "mobile unit" and

"MU" are used interchangeably herein, as are "base station unit" and "BSU." A picocellular system node typically refers to a wireless base station with extremely low power designed to cover an extremely small area, such as one floor of an office building. Indeed, short range wireless communication as described herein generally refers to an effective range of 500 meters or less (and with Bluetooth, often less than 100 meters) although some systems may reach up to 1 kilometer depending upon the wireless transmission protocol. Base station units described herein are generally fixed relative to a given location. The location may be any stationary building or geographic space (such as a parking lot or park). Alternatively, the BSU may be fixed relative to a movable object such as a vehicle (e.g., train or plane).

A large circle 107 represents the maximum range of communication between these two devices at a given transmission power class -- the "domain" of the BSU. While the maximum range limits the coverage of a single BSU, a Pocket Mobility Network ("PMN") system employs multiple distributed BSUs to cover a large region. In general, a cellular communications network typically consists of a collection of base stations, where the base stations provide wireless access points for mobile users to obtain a communication link to a wide range of information and communication services. Each base station is set to reside in one stationary "cell." The cells have some geographical overlap and collectively provide coverage to a particular area, typically encompassing many cells.

An important feature of a cellular communications network is that mobile units can "roam" from communicating with one base station to another adjacent base station within the system and not experience any disruption of the communication link while doing this. When a mobile user roams from one base station to another, a procedure called "handoff" is performed, where the link is actually transferred from the currently connected base station to a neighboring base station. To ensure the best performance, the handoff process is desired to be as fast and efficient as possible with little or no loss in data. In addition, data throughput is desired to be as high as possible. Described in detail herein are

methods for synchronizing base stations such that system data throughput and handoff efficiency are optimized.

In Figure 1, a collection of access points, or BSUs 102, 108 and 112-120, are all able to establish a link with the MU 104 and provide access to the network. The large circles or cells around each BSU illustrate the operating coverage or domain of each BSU (e.g., circles 107 and 109 are associated with BSUs 102 and 108, respectively). Typical BSU placement allows for overlapping areas of coverage. The BSU that actually participates in a link is generally the one that has the strongest signal at the MU (but other factors may affect this, as well). When that signal weakens, such as when the MU moves away from that BSU with which it is communicating and toward an adjacent or neighboring BSU, a handoff may occur. A handoff of the link may be made from the BSU 102 to the neighbor BSU 108, creating a new link 105 (link 2).

When the MU 104 first establishes a link with a BSU, it does so using the same procedures that it might use to communicate with any other Bluetooth device. To the MU, the BSU behaves exactly like any other Bluetooth device that operates in conformance with the Bluetooth specification. Thus, it does so while not requiring any modification to MU devices that comply with the Bluetooth specification. In the example of Figure 1, after establishing link 1, the MU 104 can continue to communicate with the network through the BSU 102 for as long as necessary, as long as it stays within the domain of the BSU. While aspects of the invention are described herein as employing the Bluetooth protocol, those skilled in the relevant art will recognize that aspects of the invention are equally applicable with other communication protocols and standards, including IEEE 802.11, Home RF, etc.

When the MU 104 moves, it may move within the range of another BSU (e.g. shown in Figure 1 as the MU moving from location 1 near the first BSU 102 to location 2 near the neighbor BSU 108, where the MU establishes a new link with the neighbor BSU 108). If the MU moves beyond the range of the first BSU 102, the link with the neighbor BSU 108 can be used to maintain uninterrupted communication with the network, provided that a new link is established before

the link with the first BSU is lost. In other words, the MU 104 must establish a link 2 with the BSU 108 in a new domain 109.

Under one embodiment, the PMN is designed in a fashion similar to the General Packet Radio Service ("GPRS") system architecture and thus, may use much of the same terminology. The PMN includes multiple BSUs and a PMN system controller ("PSC") to manage the BSUs. As shown in Figure 2B, each BSU may in turn include one or more base transceiver stations ("BTS"), where each BTS includes Bluetooth hardware and associated software that runs below a Host Controller Interface ("HCI"). A backbone controller to link the BSUs and controller with RF or other wireless links, instead of cabling, may be employed. A network service provider may supply various devices to provide connectivity to networks beyond the immediate premises (or other coverage area) as part of a complete service.

Each BSU may include a Bluetooth RF module, a microprocessor (picocell or baseband controller) with memory (e.g., RAM) and an antenna. The BSUs may mount on a ceiling and provide radio coverage within an approximate 10-meter radius. Each BSU may run both Asynchronous Connectionless Links ("ACL") and Synchronous Connection Oriented ("SCO") services for data and voice connectivity and a control stack, and a voice-over IP software module. Each BSU is based on available Bluetooth chip modules and available Bluetooth protocol stacks.

The BSUs may each be directly connected to the PSC by using appropriate cabling, such as Category 5 cabling. Such cabling is necessary to provide not only a signal path therebetween, but may also supply power to the BSU. Alternatively, a link with the base station controller may be accomplished via wireless techniques. Wiring may be required, however, to provide power to the BSUs.

The PMN may include multiple interfaces (such as Ethernet interfaces), a processor module, a switching module, and interfaces for channelized voice and LAN/WAN or other connectivity (including packetized voice). The PMN software modules may include a voice stack, a data stack, and a control stack. The

control stack handles mobile unit tracking and handoffs, user management, and session management.

A network service provider providing the PMN may also include a channelized (standard) private branch exchange ("PBX"), a voice over IP PBX switch, a direct connection to the public switched telephone network ("PSTN"), a router (for data, voice over IP, or both), a server (for providing various application), a cache, etc. Prepackaged applications for vertical markets, such as hospitals, theme parks, malls, airports, for enterprises and service providers, for private networks in a public space, etc., may also be provided with the PMN.

An example of a PMN is shown in Figures 2A and 2B. Referring to Figure 2A, the base station units are coupled to the Internet or wide area network ("WAN") 202 by way of an edge router 204 and switches 206. As shown in Figure 2A, each BSU is coupled to one switch. Each BSU acts as a bridge between a wired LAN that includes the BSUs, switches 206 and router 204, and the wireless links to the mobile units (*e.g.*, wireless links 103 and 212 to mobile units 104 and 210). The LAN is implemented using Ethernet or an alternative technology. A PMN system controller ("PSC") 208 or "system controller" is coupled to the network 200 or subnet and acts as a systems data communication gateway providing mediation between the PMN and the Internet or WAN 202. The system controller is effectively a collection of functions that may reside or run on one or more computers, such a one or more server computers. While the Internet or WAN 202 are shown, the subnet may be coupled with other networks, such as a LAN, the PSTN, or public land mobile network ("PLMN"). The PSC switches voice and data communications to the appropriate wired or wireless network to which the subnet is coupled.

As shown in Figure 2B, the BSU 120 includes several BTSs. Specifically, the BSU 120 includes BTS 1 220, BTS 2 222 through BTS N 224. Of course, other BSUs likewise may include one or more BTSs. While each BTS is shown in a single block (such as separate radio cards in a single housing forming a BSU), each BTS may be incorporated within a separate housing and connected together to form a single BSU.

The PSC and BSUs may employ the Bluetooth LAN access profile ("LAP"). The LAP uses established networking protocols to enable a computing device or MU to obtain access to a data network. Use of the LAP is analogous to directly connecting to a data network such as via Ethernet. Further details regarding LAP may be found in the Bluetooth protocol noted above.

The network 200 forms the entire domain of the PMN, which is the domain over which a mobile unit can be handed off. If a mobile unit can be handed off between BSUs, then those BSUs are on the same PMN. The router 204 that connects the PMN or subnet both logically and physically to the rest of a corporate network (e.g., WAN 202) is the first router in the data path for MUs communicating with nodes outside the subnet and may be running network address translation ("NAT").

NAT refers to an Internet standard that enables a local area network to use one set of IP addresses for Intranet traffic and a second set of addresses for external traffic. This allows a company to shield internal addresses from the public Internet. This would allow the network 200 to connect to the Internet 202 while not all hosts within the network have globally unique IP addresses. Thus, NAT enables the network 200 to use nonregistered IP addresses to connect to the Internet. The NAT-enabled router translates internal local addresses to globally unique IP addresses before sending packets outside of the network 200. There exist many ways to implement NAT.

Most models of cellular-like systems from General Packet Radio Service ("GPRS") use a non-routed Radio Access Network ("RAN"). In other words, the IP address of the mobile unit is not used to locate the mobile unit once the packet is inside a gateway router. In GPRS, the serving GPRS support node ("SGSN") maintains a mapping of IP addresses to telephone numbers to keep track of locations of mobile units by telephone numbers. As a mobile unit is handed off from one BSC to another, the SGSN tables are updated.

The PMN described in Figure 3 (the multiple subnet architecture) follows a model similar to GPRS. A simplifying factor is that the PMN uses Ethernet within a network or subnet. Once a node (e.g., BSU) makes itself known to the network (described below), the switches and routers 206 that comprise the

subnet will keep track of the correct routing for packets addressed to that node. Another simplifying factor is that the Figure 3 PMN architecture is a fully-routed network.

Multiple subnets may be interconnected to create a larger installation where BSUs and MUs are assigned IP addresses that fall within multiple subnets. As shown in Figure 3, a network 300 includes three subnets 350, 310 and 320, each connected by an MU virtual subnet 330 that includes an MU 331 and the MU 104. As shown, all three subnets 350, 310 and 320 are similar to the previously described network 200 except that none of them contain the subnet to which the MUs belong: the first subnet 350 includes the BSUs 102, 108, 120, 112, 114 and 116, coupled to switches 206, which in turn are coupled to the gateway router 302; the second subnet 310 includes BSUs 312, 314, and 316 coupled to switches 206, which in turn are coupled to a gateway router 302, while the third subnet 320 includes BSUs 322 and 324 coupled to another gateway router by way of switches 206. Between subnets, packet traffic passes through the gateway routers 302 to gain access to a backbone, which can be a backbone router 304. The gateway routers 302 are typically edge routers and may be running, for example, NAT. Alternatively, NAT may be run on the backbone router 304.

Figure 3 illustrates both intrasubnet and intersubnet handoffs within the multiple subnet architecture. An intrasubnet handoff, such as between BSUs 102 and 108, is shown in more detail in Figure 4. Such an intrasubnet handoff is similar to that described above with respect to Figure 1 and is defined as a handoff where the initial or owner BSU and neighbor or target BSU are on the same subnet.

Alternatively, and shown in greater detail in Figure 5, an intersubnet handoff is defined as a handoff where the initial and target BSU are on different subnets, such as the BSU 108 associated with the first subnet 350 and the BSU 316 associated with the second subnet 310. The mobile unit 104 moves from the domain of the BSU 108 to the domain of the BSU 316. The wireless link 105 with the BSU 108 is dropped in favor of a new wireless link 318 established with the BSU 316.

The multiple subnet architecture ("MSA") shown in Figure 5 forms a virtual subnet for the mobile units and set of BSUs together with a point of presence 334 (shown in Figure 3). The point of presence 334 is defined to be the device on the MU virtual subnet 330 that is physically attached to the wired LAN or network (coupled to the backbone router 304 by way of a gateway router 336). The MU subnet 330 is a real subnet in a logical sense, and is logically attached to each of the BSU subnets 350, 310 and 320 by way of the point of presence 334.

Figure 6 shows how the gateway router 336 provides access to the MU subnet 330 and the point of presence 334. The gateway router 336 may be a separate device, or it may be combined with the point of presence 334. The logical route for packets from a host or server 602 to, for example, the MU 104 is over the wired LAN or backbone network to the point of presence 334 and directly to the MU (as shown as the solid arrows in Figure 6). The actual route of packets is not directly from the point of presence to the MU but back through the gateway router 336, backbone router 304, gateway router 302 associated with the subnet 350, through switches 206 and to the BSU 108, all over the backbone network. Mobility under IP, such as via a tunneling protocol, provides communication between the point of presence 334 and the appropriate BSU 108. The point of presence 334 may be part of the PSC 208, or it could be distributed across one or more BSUs. The point of presence and PSC are shown as two separate entities in Figure 3 because they are logically separate. They may be implemented, however, as separate blocks of code on a single server (e.g., a single Linux server).

Slot Alignment

Base stations and network access points in a wireless communication system are typically not synchronized with each other. In frequency hopping ("FH") systems, each base station starts its frequency hopping sequence at an arbitrary time based on its own internal clock. Since time slot durations are also based on the internal clock, slot widths of neighboring base stations vary as well. Consequently, neighboring base stations will most likely not have their slot

edges aligned in time. This situation adversely affects throughput in wireless systems whose base stations have overlapping coverage, particularly in systems where base station-to-base station handoff occurs.

One source of throughput loss occurs when adjacent base stations hop to the same frequency ("channel collisions"). Any mobile units (MUs) that happen to be in the overlap range where they can "hear" both base stations will experience interference when both base stations transmit on the same frequency. Any MU data transmitted during that interference is lost and will have to be retransmitted after one or both base stations have hopped to differing frequencies. Since the slot edges of adjacent base stations are most likely unaligned, a channel collision will most likely interfere with portions of not one but two slots. Thus twice as much data is potentially lost and would have to be retransmitted, thereby decreasing throughput.

This is illustrated by Figure 7A. In the unaligned case, a channel collision is shown when base stations B1 and B2, and a third device all hop to Channel 10. (The third device could be another base station or, in this case, receive time ("Rx") of a mobile "MU1".) B2's transmission time ("Tx") overlaps parts of each of the other two, resulting in interference in parts of slots 5 and 6. Thus the data in both slots 5 and 6 is lost for all of these devices.

In the aligned case shown in Figure 7B, B2's transmission time overlaps only with that for B1, and thus lost data is limited to a single slot (slot 5). Note that the non-overlapping device in this case would be a mobile (MU1) since all the base stations are aligned.

To look at this another way: in the unaligned case there is one chance in 79 that for slot 5, the trailing edge of B1's transmission slot ("Tx") and the leading edge of B2's Tx will both be on channel 10 (for Bluetooth, which employs 79 available channels). There is also one chance in 79 that, for the same slot, the leading edge of B1's Tx and the trailing edge of B2's Tx will both be on channel 10. Thus for any given slot, for a given pair of devices, the probability of a channel collision is $2/79$ for the unaligned case. In the aligned case, the probability is halved ($1/79$).

Consequently, slot edges of neighboring base stations should be aligned so that during a channel collision only a single slot is affected. Channel collisions will occur in such a system, and are largely unavoidable, but synchronizing base station slot edges will mitigate adverse effects by minimizing the number of slot transmissions that experience interference.

Synchronizing base stations also improves base station-to-base station handoff performance, such as in the PMN described above. With unaligned base stations, a mobile unit being handed off from a first base station to a second must resynchronize its time slots to match the second base station. Unproductive time lags result while waiting for portions of one or more time slots to elapse. With synchronized base stations, the second base station is able to make its "first" time slot after a hand-off occur exactly where the first base station's "last t + 1" time slot would have been, thus eliminating the need for the mobile unit to resynchronize after hand-off.

As noted above, the Bluetooth method of TDD causes two types of slots to be present in any system: transmit and receive. If the slot edges are aligned, two different patterns for slot accesses from the perspective of the BSU are possible:

Pattern 1:
0123456789
TRTRTRTRTR

or

Pattern 2:
0123456789
RTRTRTRTRT

where the "T" represents a time slot in which the BSU can start a transmission, where "R" represents a time slot in which the BSU can receive from a slave and where numbers 0-9 simply represent sequential slots. (In the discussion below, an MU is discussed, even though any "slave" may be employed.)

If there are two MUs communicating with a BSU that has only one BTS, the following three situations can occur.

1. Both MUs employ Pattern 1:

The BTS can transmit and receive easily in all time slots

0123456789
 TRTRTRTRTR

2. Both MUs employ Pattern 2:

The BTS can likewise transmit and receive easily in all time slots.

0123456789
 RTRTRTRTRT

3. One MU employs Pattern 1, while the other employs Pattern 2:

Here, however, the BTS "loses" one slot every time the MU is switched.

0123456789
 TR TR TR T
 11 22 11 2

In a first embodiment or "case," when a BSU has more than 1 BTS (such as the BSU 120 of Figure 2B), the MUs with different patterns can be separated, allowing one BTS to run on Pattern 1 and another to run on Pattern 2. This separation works up to a point. If the BTSs which transmit on one pattern are off limits to the MUs on the other pattern, and the BTSs are divided evenly, the BSU could have no more slots for an incoming MU or slave when the BSU is only half full. This would cause dropped calls on a 50% utilized network, which is an inefficient use of resources.

In a second case, if the BSU allocates BTSs a pattern as necessary, once the BSU has allocated all the BTSs, an incoming MU who's pattern is associated with a full BTS, the MU has no slot on which to attach. This is not a major problem when there are 20 BTSs in a base station, as the base station must be over 90% full before this situation could occur. On a base station with only 4 BTSs though, a utilization of just over 75% could yield this condition. This would happen, for example, when 3 BTSs are assigned to Pattern 1 and completely utilized. The fourth BTS has one MU and is assigned to Pattern 2. Another MU on Pattern 1 enters the BSU's area. Since there are no slots available on the 3 BTSs, and the fourth is assigned to Pattern 2, the MU is dropped.

In a third case, the BSU operates in the same way as in the second case. However, when, as in the example above, three BTSs are assigned to Pattern 1 and are completely utilized, while a fourth BTS has only one MU, and is

assigned to Pattern 2. When a second MU on Pattern 2 enters the BSU's area, the BSU allows the MU to join, but accepts the inefficiencies on that channel.

Another resource allocation issue is as follows. Assume, for example, that there are 7 MUs, but they have very minimal traffic, and another 7 MUs that have very high traffic demands. The minimal traffic MUs are on Pattern 1, and the high traffic MUs are on Pattern 2. This situation results in poor resource allocation. The high traffic MUs should be spread to BTSs with low traffic MUs, and allow all the BTSs to serve the MUs to the best of their ability. This would result in the same inefficiencies of the traffic for the third case, where some MUs are on slot pattern 1 and some are on pattern 2.

In a fourth embodiment or case, which makes administration of the BTSs and resource allocation easier, all MUs are set to the same pattern as they enter the network, and all BTSs employ the same pattern. The MUs could be moved from one BTS to another for pure resource reasons, without having to allocate a BTS to a particular pattern, or lose slots to the inefficiency of pattern switching.

A disadvantage of this fourth embodiment is that all BTSs would be transmitting at the same time. In a network where there are only one slot packets, this would be beneficial, because the transmit/receive antennas would not need to be separated. In other words, no transmissions occur during reception slots, and vice versa. In a network of multiple slot packets, this would mean larger amounts of RF energy released into the air while a receive antenna is trying to receive.

Under a fifth case, the BSU may be field configurable by a technician to employ any of the above five alternatives. Thus, a technician can select the best transmission patterns and allocations for BTSs, and change that configuration as circumstances later change after initial configuration.

Modifying Slot Usage Embodiment

Initial base station synchronization can be done at PMN initialization time during base station neighbor discovery. During normal system operation base stations may communicate with each other the "backend network," noted above. Thus, base stations may identify one another, and particularly neighboring base

stations, by employing the backend network. Alternatively, base stations may be able to automatically identify their neighbors as described in greater detail in U.S. Patent Application No. _____, filed _____ (attorney docket no. 34015-8002). In one embodiment, base station synchronization uses dedicated radio transmit slots to communicate directly, *e.g.*, over the air, with other base stations. In one embodiment, a predefined number (*e.g.*, 64) transmit/receive (TX/RX) slot pairs are used, with 128 total slots available for base station synchronization. These slots may also be shared with the base station remote MU monitoring function, where "unconverted Rx slots" per cycle are when the BSU will monitor the signal strength of MUs.

The Bluetooth protocol employs time division duplex ("TDD"), where pairs of transmit and receive slots are repeated. For Bluetooth compatibility, one embodiment uses a repeating cycle of 64 TX/RX slot pairs, or 128 slots total. Normally half of these would be used for transmitting and half for receiving, but one embodiment of the invention converts all except one of the transmit (TX) slots into receive (RX) slots for more effective monitoring. Unconverted RX slots are used for monitoring MU strength and converted TX slots are used for listening for sync information messages from other base stations. The single remaining TX slot is used by that BSU to broadcast its sync information once every 128 slots (*i.e.*, at 80 millisecond intervals). Which slot pair (1-64) contains the single TX slot is what gets decided during sync initialization. This slot will remain assigned to that base station unless it becomes necessary to change it, as described below. Each base station thus moves back and forth between two roles during the ongoing synchronization process: every 80 millisecond, for one slot time, it acts as a "sync server" by providing its synchronization information to any base stations around it who are listening; any base stations within range can detect this special "sync" message, if they have been synchronized, and will use it to synchronize themselves. Then, for nearly half of the time, the base station acts as a "sync client" by listening for sync information from other base stations. For the other half of the time, the base station performs its normal listening for mobile unit signal strength. (During at least every

resynchronization, the BSUs must continue to monitor for MU traffic.) Figure 8 illustrates this method.

Referring to Figure 8, an example of a base station unit that has selected slot 2 during the initial synchronization process is shown. During synchronization, each transmit slot has been converted to a receive slot (e.g., slots 802, 804 and 806) where the base station unit listens for synchronization information from neighboring base stations (thereby acting as a sync client). A single unconverted slot, namely slot 808, is the only slot in 64 pairs that this base station unit transmits its own synchronization information (thereby acting as a sync server). The remaining receive slots (e.g., slots 812, 814 and 816), are when the base station unit listens for mobile unit signal strength.

In one embodiment, each BSU employs various layers of software under the Bluetooth protocol. Each BSU further includes a radio environment monitoring ("REMon") and radio environment management ("REMan") software, that includes a radio environment management component that operates above the host controller interface ("HCI"), and a radio environment monitoring component that operates below the HCI, together with baseband components of the Bluetooth protocol. BSU neighbor discovery and BSU synchronization functions are performed in conjunction with the REMan and REMon components, within the radio environment monitoring and management software. The REMon component functions to receive a set of records from REMan, where the records indicate which MUs to monitor. Each record in the set identifies the MU to monitor, performance indicator types to be measured (e.g., signal strength, signal to noise ratio, error rate, etc.), and the time slot the MU is to be monitored on. The REMon component works with a Bluetooth link controller to convert all except one of the transmit (TX) slots into receive (RX) slots, as described above. Modifying the baseband layer of the Bluetooth protocol stack in the BSUs will permit such a modification in slot assignments to occur. Following synchronization, the BSU returns to a standard Bluetooth TDD, where pairs of transmit and receive slots are repeated. Further details regarding software employed by the BSU are provided below.

Representative Embodiment

During the course of two base stations discovering each other as neighbors they learn whether each other has accomplished slot assignment. If neither has, then they can go no further in the synchronization process, so they record each other as neighbors in their respective neighbor lists and wait until later when one of them does become synchronized, as described below. If both are already synchronized then they just record each other as neighbors. If one has been synchronized and one has not then the unsynchronized base station makes its sync slot selection based on the slot assignments of its neighbors and its neighbors' neighbors, as described below.

The base stations perform synchronization autonomously. The concept here can be visualized by imagining how a water drop generates a ripple in a pond: the waves start at one point and spread outward in all directions. The initial synchronization "water drop" starts at an arbitrary base station in the PMN called the origin, or "seed," base station. After the synchronization propagates to all edges of the network, the entire PCN will have its slot edges aligned. As described more fully below, the BSUs at the same "radius" from the sync origin will be in sync with each other, because they will be giving and taking hand-offs from each other. The system avoids the situation where the entire PMN is in lock-step, from sync origin to outer edge.

The BSUs preferably employ a *re-use distance* of at least 2 "waves" for all sync slot assignments. That is, a sync slot assignment that a base station uses to broadcast its sync information must be unique among all its neighbors and among all its neighbors' neighbors. Beyond that, sync slot numbers can be re-used. In very dense installations 64 available slot assignments may be insufficient for this two-level uniqueness. In that event, the assignable sync slot amount would alternatively be 128 or greater.

Referring to Figures 9A-9C, suitable routines for synchronizing neighboring base stations are shown. Unless described otherwise herein, the functions in the blocks depicted in Figures 9A-9C are well-known or described in detail in the above-noted and cross-referenced provisional patent application.

Indeed, much of the detailed description provided herein is explicitly disclosed in the provisional patent application; most or all of the additional material of aspects of the invention will be recognized by those skilled in the relevant art as being inherent in the detailed description provided in such provisional patent application, or well-known to those skilled in the relevant art.

Referring to Figure 9A, a routine 900 begins in block 902 by selecting a "seed" or initial BSU. A low-complexity random process may be employed to select the system sync starting point or seed BSU. In one embodiment, the system controller 208 broadcasts a message on the backend network to query base stations. The first base station to reply would be designated as the system sync seed. The system controller would command only one base station to assume this role. Alternative embodiments may employ other mechanisms, such as broadcasting messages between BSUs, and eliminating the intervention or need for the system controller to initiate synchronization.

In block 904, the seed BSU chooses a transmission slot. The slot may be chosen arbitrarily. For example, the seed BSU may select the first transmission slot in the series of 64 available slot assignments (*e.g.*, slot 802 of Figure 8).

In block 906, the seed BSU records its sync client information. As described more fully herein, the sync client information includes information regarding wave numbers, BSU addresses, and slot assignments for individual BSUs to which a given BSU recording client sync information is a client. Since the seed BSU is not a client of any other BSU (it is only a server), the seed BSU records some null information, such as "0.0" under block 906.

Under block 908, the seed BSU records its sync server information. As described more fully herein, the sync server information includes a wave number designation, an address of that BSU, and the slot which it has chosen to transmit its sync server information. As described below, the seed BSU may record, for example, "1.B4.1" to indicate that BSU B4 transmits its sync server information in slot 1 during Wave 1.

Under block 910, the seed BSU broadcasts its sync information on its chosen slot. Such transmission may be at, for example, the Bluetooth class 2

power level (4 dBm). Indeed, all BSU transmissions during synchronization may be at the class 2 power level.

An example of BSU synchronization and initialization is shown in Figure 10. Figure 10 shows 23 base stations (represented by designations B1 through B23), which form an overlapping wireless coverage area of a PMN.

Note: the use of hexagons to represent base station cells in Figures 10-15 is for clarity only, and avoids the clutter of overlapping circles. There is no implication that, in a PMN installation, a given base station will always have six or less neighbors. In fact, there will most likely be installations where cells are placed such that BSUs have more extensive overlap, and a given base station may indeed have more than six neighbors. If the PMN installation is in a multistory facility, with base stations on adjacent floors, base stations may also be discovered several floors above and below. These are important so that a user walking up or down a stairway can still experience handoff. There is no limit to the number of neighbors that a base station may have.

In Figure 10, B4 is chosen to initiate the base station system sync process, and thus is the seed BSU. The seed base station is indicated by "0.0" in Figure 10, which shows the state of system sync after one iteration or "wave."

As noted above, each base station's persistent store or memory records at least two pieces of information about its sync. The format of this information is not as important as the content. The recorded information includes at least the following:

1. how the base station gets its synchronization (*i.e.*, its sync client information)
2. how the base station gives its synchronization information to neighbors (*i.e.*, its sync server information)

In each Figure, a base station acting as a sync server is shaded in gray.

In Figure 10, B4's sync client information is "0.0," which designates that B4 is the PMN synchronization seed base station. Only one base station in the system has this property or value. B4's sync server information is "1.B4.1," which indicates that during Wave 1, B4 sends its synchronization information during slot #1. Thus, under routine 900, the seed BSU, B4, in Wave 0:

1. chooses to use, in this example, slot 1 to broadcast its sync information
2. records its sync client information as "0.0"
3. records its sync server information as "1.B4.1" (wave 1, its unique address (shown here simply as "B4"), slot 1)
4. begins broadcasting its synchronization information on its assigned slot

To continue the water analogy, the ripple now begins to spread as BSU's B3, B7, and B8 each simultaneously and asynchronously try to discover their neighbors and get their own synchronization slot assignment. Figure 10 also shows one of several possible outcomes of who synchronizes whom during Wave 1. Assume that B3 is the first non-seed base station to get synchronized, followed by B7, and then by B8. Thus, under Wave 1, the following occurs.

Wave 1, BSU B3:

1. B3 sends B4, and each of B3's neighbors thus far, a "lock request" message so that no other base stations can simultaneously be getting or selecting the same slot. A lock request is similar to a software mutex (*i.e.*, a software lock that indicates to other BSUs that the slot is assigned to that particular requesting BSU, and that other BSUs are locked out of using that slot until it is released by the requesting BSU). The lock request allows only one BSU at a time to perform synchronization, thereby locking out other BSUs from attempting to simultaneously synchronize. Each BSU may, for example, record a table of slots locked by other BSUs in earlier waves. Alternatively, as described herein, each BSU may simply record sync client information that includes slots locked by neighboring BSUs.
2. B4 returns a "lock granted" message to B3 along with any neighbor's neighbors and their synchronization slots (not established yet under this example). B4 (indeed each of B3's neighbors) will receive a lock-request message. Unless it is already engaged in a similar transaction with another BSU, B4 will return a lock-granted

message. In that message it tells B3 all the Tx slots that it (i.e. B4) already knows about as being in use, due to it's (i.e. B4's) neighbors. B3 needs this neighbor's neighbor information to pick a slot that won't conflict. The reasoning is that all of B3's geographic neighbors (i.e. within 17 meters or so) got discovered during Neighbor Discovery, so any slots they use are off limits for B3 to choose as its Tx slot. Then, to cover the case where one BSU's sync broadcast might "carry" farther than normal, the system also disallows any slots in use by any of B4's neighbors' neighbors (which could be as much as 34 meters or so away from B4). Only when it has all this occupied-slot info can B4 pick a slot that will not collide with one of its neighbors or one of its neighbors' neighbors. Another reason is to ensure that each neighbor can receive from each of its neighbors. If any BSU has two neighbors transmitting on the same slot, it will not be able to hear either one. Thus if a neighbor has a neighbor transmitting on slot N already, if you chose slot N, not only would your neighbor not be able to hear you, it would cease to be able to hear its neighbor that it is already listening to.

3. This lock/acknowledge interaction is repeated for all remaining B3 neighbors. If any are already locked by another base station, the BSU waits until it is unlocked, then lock it, or selects another slot.
4. B3 realizes that B4 has slot 1 and that no other neighbors are synchronized yet; therefore, B3 chooses an available slot e.g., slot 3.
5. B3 records its sync client information as "1.B4.1" (copied from what B4 broadcast).
6. B3 records its sync server information as "2.B3.3" (wave 2, B3, slot 3).
7. B3 sends an "unlock request" message to B4, then to any remaining neighbors, in reverse order of locking. Following such unlocking,

other BSUs that have yet to synchronize may then synchronize themselves.

Wave 1, BSU B7:

1. B7 notices that among his discovered neighbors, B4 has a higher "priority" than B3 since it has a lower wave number, so B7 decides to synchronize to B4, as B3 did.
2. B7 sends a "lock request" message to B4 (which has the lowest slot number).
3. B4 returns a "lock granted" message to B7 along with any neighbor's neighbors and their synchronization slots (not established yet in this example).
4. B7 sends a "lock request" message to B3 (which has the next lowest slot number).
5. B3 returns a "lock granted" message to B7 along with any neighbor's neighbors and their synchronization slots (not established yet in this example).
6. B7 performs a lock request/lock grant message exchange with the rest of its neighbors (none in this example).
7. B7 chooses a slot number, e.g., the next lowest available slot number: "5."
8. B7 records its sync client information as "1.B4.1" (copied from what B4 broadcast).
9. B7 records its sync server information as "2.B7.5" (wave 2, B7, slot 5).
10. B7 sends an "unlock request" message to other neighbors (none in this example), then B3, then B4, in reverse order of locking.

Wave 1, BSU B8:

1. B8 notices that among his discovered neighbors, B4 has a higher priority than B7 since it has a lower wave number, so B8 decides to synchronize on B4 as B3 and B7 did.
2. B8 sends a "lock request" message to B4 (which has the lowest slot number).

3. B4 returns a "lock granted" message to B8 along with any neighbor's neighbors and their synchronization slots (not established yet in this example).
4. B8 sends a "lock request" message to B7 (which has next lowest slot number).
5. B7 returns a "lock granted" message to B8 along with any neighbor's neighbors and their synchronization slots (not established yet).
6. B8 performs a lock request/lock grant message exchange with the rest of its neighbors.
7. B8 chooses, e.g., the lowest available slot number: '7'.
8. B8 records its sync client information as "1.B4.1" (copied from what B4 broadcast).
9. B8 records its sync server information as "2.B8.7" (wave 2, B8, slot 7).
10. B8 sends an "unlock request" message to other neighbors (none in this example), then B7, then B4, in reverse order of locking.

At this point, "Wave 1" is considered finished – neighbors at the smallest "radius," or distance from the seed, have been synchronized.

Referring to Figures 9B and 9C, routines illustrating the above functions are shown. As shown in Figure 9B, a routine 920 begins in block 922 where a receiving or "current" BSU receives a sync signal transmitted by a neighboring BSU. In block 924, the current BSU determines whether that neighbor BSU is synchronized. If so, then in block 926, the current BSU determines if itself is synchronized. If so, then the current BSU simply records the neighboring BSU in its list of neighbors. If not, then in block 930, the current BSU performs the slot selection and other processes described below with respect to Figure 9C.

If the current BSU in block 924 determines that the neighbor BSU is not synchronized, then in block 932, the current BSU determines whether it itself is synchronized. If not, then in block 934, the current BSU waits until a synchronization wave is received. If the current BSU is synchronized, then in block 936, it allows the neighbor BSU to perform slot selection and other

processes as shown in Figure 9C. Following blocks 934 and 936, the current BSU records the neighboring BSU in its list.

Referring to Figure 9C, a routine 950 begins in block 952 where the current BSU receives a sync signal from a neighboring BSU. In block 954, the current BSU sends a lock request message to one of its neighboring BSUs. In block 956, the current BSU receives a lock grant message from that neighboring BSU. In block 958, the current BSU determines from its neighbor list whether any more neighboring BSUs exist. If so, then the routine loops back to again perform blocks 954 through 958.

If no additional neighbors exist, then in block 960, the current BSU chooses a transmission slot in which to transmit its sync information. For example, the BSU may choose the next lowest available slot based on slots previously selected by BSUs in this or earlier waves. In blocks 962 and 964, the current BSU records sync client and server information, respectively. In block 966, the current BSU sends an unlock request to the last neighboring BSU that it received a lock granted message from. In block 968, the current BSU determines whether any additional neighbors exist. If so, then the routine loops back to again perform blocks 966 and 968 as the BSU sends unlock requests to neighboring BSUs.

In light of the above details, a discussion of subsequent waves will be described with less detail. Those skilled in the relevant art will recognize that the subsequent waves occur in a manner similar to that described above. Therefore, after Wave 1 ends, Wave 2 commences.

Here, B3, B7, and B8 start acting like sync servers during their designated slot time while base stations at the next radial increment out (*i.e.*, base stations B2, B6, B11, B12 in Figures 11 and 12), who are simultaneously and asynchronously trying to get their own synchronizations established, will now start to be successful at it. The steps are the same as described above. The net results of this wave, with assumptions made for the sake of the example, are below.

In Wave 2:

Base stations B2 and B6 are first synchronized in this example, as shown in Figure 11.

1. In this example, B2 gets its synchronization first, from B3, who is B2's only synchronized neighbor.
2. B6 then has three neighbors who have been synchronized (B2, B3, B7); B6 synchronizes off B3 because B3 has the lowest previous wave slot number (3).

Thereafter, in this example, base stations B11 and B12 get synchronized, as shown in Figure 12.

1. B11 has two neighbors who have been synchronized (base stations B6, B7); B11 synchronizes off of base station B7 has the lowest previous wave slot number (5).
2. B12 has three neighbors who have been synchronized (B11, B7, B8); B12 synchronizes off of B7 since that has the lowest previous wave slot number (5).

Referring to Figure 13, base station B8 has no takers for its synchronization services since all its neighbors (B4, B7, B12) are already synchronized.

At this point, "Wave 2" is considered finished since neighbors at this second radial increment have been synchronized.

Wave 3 now commences. In exactly similar ways as described above, B2, B6, B11, and B12 now begin acting as sync servers and provide synchronization for base stations at the next or third radial increment: B1, B5, B10, B14, B15, B16. The results of this wave are shown in Figure 14. Note that two slot numbers (1, 5) get re-used in this wave, which is acceptable since the *re-use distance* must be at least two waves. (Recall that a base station establishing synchronization must choose a slot number unique among its neighbors and its neighbors' neighbors). The process continues with waves 4 and 5 to synchronize base stations B9, B13, B18, B19 and B20, and B17, B21, B22 and B23, respectively, in the same manner as described above.

Having each base station examine the wave number of already-synchronized neighbors prevents back-propagation that would lead to a never-ending synchronization process. The synchronization "wave front" is kept moving forward by requiring a base station in a current wave to establish synchronization with a base station already synchronized during the previous wave.

After the entire network has gone through the above synchronization process once, each base station is now acting in both sync server and sync client roles. The initial synchronization establishment described above has to be repeated only when a base station's neighbor list changes.

What happens if a base station is added or comes up after initial system synchronization establishment? This situation is illustrated in Figure 15. Assume that B7 comes up after base station neighbor discovery and base station synchronization is completed. It is now in its own unsynchronized pocket. Since B7 wasn't up earlier (e.g., it was added after system initialization), B11 got synchronized one wave later (wave 4) than in the previous scenario. As it discovers its neighbors, B7 notices that all of them are synchronized, so B7 synchronizes off the base station with the earliest wave number and the lowest sync slot number – which is B4 in this scenario. If B4 were down and B7 had to choose among the remaining neighbors, then B7 would choose B3 because of its lower wave 2 slot number. Here, B7 chooses slot 51. B7 also assigns wave 3 to itself since it is between a wave 1 base station and a wave 4 (it could also have chosen wave 2).

If B11 had chosen slot 1 (it could have without violating re-use distance), then there would be a problem because B7 is now coming up between two synchronized base stations both transmitting synchronization information on the same slot. In that case B7 will tell the base station with the later wave number (B11) to choose for itself another slot. Once B11 does so, B7 can proceed to its own synchronization slot selection. B11 having to choose another slot may cause a similar conflict for itself and a neighboring base station. If so then B11 will likewise tell a base station in the next wave further out to choose a new slot, and so on. This reselection may ripple out for an indeterminate number of base

station "layers," but it will eventually get resolved even if it does have to propagate all the way to the edge of the PMN.

If neither of two neighboring BSUs are synchronized (and have not chosen a slot), one of the BSUs may page the other when it does become synchronized. The synchronized one may then ensure it is a Bluetooth master so that the BD_Addr and CLK needed for synchronization is communicated to the unsynchronized BSU, and afterwards the unsynchronized BSU may choose a slot. If the synchronized BSU is not a master, then it performs a master/slave switch with the unsynchronized BSU.

As noted above, BSUs in a given ring or radius are generally, but not necessarily, the sync master to BSUs in the next ring farther away from the seed BSU. Therefore, the BSUs are effectively in a "hierarchical" arrangement. BSUs in one ring are synchronized, but may not necessarily be closely synchronized with BSUs in a distant ring. BSUs repeat synchronization periodically, such as at a repeating cycle of 128 slots (80 milliseconds). The "virtual master" clock value that a BSU receives from its sync master during its assigned slot is passed along to that BSU's neighbors during the same 128 slot cycle, during which that BSUs clock may have little time to drift.

Software and Functionality Details

As noted above, each BSU includes REMon and REMan software. A synchronization manager, performing much of the functions described herein, works with REMan (above the HCI), and REMon (below the HCI). Referring to Figure 16, software operating on two neighboring BSUs is shown. The discussion of software for one BSU is described below; those skilled in the relevant art will recognize that the software is substantially similar to that in the neighboring BSU.

As shown, a synchronization manager 1602 exchanges data and commands with REMan 1604 and REMon 1606. The synchronization manager 1602 functions with a neighbor discovery process 1608 to locate and record the identify of neighbor BSUs. Further details regarding neighbor discovery are found in a co-pending U.S. patent application entitled Wireless Base Station

Neighbor Discovery in a Communications System, Such as a System Employing a Short-Range Frequency Hopping Scheme, as noted herein. The synchronization manager 1602 also provides synchronization information to neighbor BSUs so that those BSUs can initiate their own synchronization (when acting as a sync server) and receive synchronization information from neighbor BSUs and update local synchronization (when acting as a sync client), as described herein.

Synchronization is a concurrent process independent of links with mobile units and handoffs of those links. One constraint may be that links with mobile units cannot be made by any BSU that has not achieved synchronization. The synchronization manager does not do the actual synchronization of BSUs, but instead facilitates the synchronization process by performing inquiries and pages to establish links with client BSUs, as shown in Figure 16. When links with clients are established, the synchronization of internal clocks (such as local clock offsets) are shared across the link. Any client BSU with which a link is established is by definition a neighbor BSU. If the synchronization manager identifies a neighbor BSU, the identify of that BSU is provided to the neighbor discovery process 1608.

Each BSU performs the synchronization process as part of a BSU synchronization profile. This profile is registered as a service with a service discovery protocol ("SDP") 1610, whenever the synchronization manager 1602 is first run (such as at power up). Upon startup of a BSU, the synchronization manager assumes the role of a sync server. However, if based on links with neighboring BSUs it determines it should switch roles and become a client, it then does so. In the sync server role, the synchronization manager 1602 initiates dedicated inquiries as it looks for neighbor BSUs. This includes defining a dedicated inquiry access code under the Bluetooth protocol. Under the Bluetooth protocol, each packet starts with an access code, and a unique dedicated inquiry access code ("DIAC") may be selected during such synchronization operation. When the server BSU receives a response from a BSU, the synchronization manager 1602 of the server BSU initiates paging to

the responding BSU and sets up an ACL link, and then a connection to the synchronization managers of the client BSU.

The server synchronization manager 1602 registers the BSU of the client synchronization manager as a neighbor. This registration information is used to update a Service Availability field on a service record for BSU synchronization. The SDP 1610 contains a service record pertaining to a PNN clock synchronization status of the BSU. In addition to allowing the synchronization manager 1602 to register a service with the SDP 1610, the synchronization manager is able to set and read synchronization information in the service record. Thus, the synchronization status of the BSU on which the service record exists may be updated with a value by the synchronization manager 1602. Under this situation, the server BSU is the "master" on the link, and the client BSU is the "slave." In general, if client synchronization information received from a synchronization manager of the client BSU has a status indicated as "better" than the synchronization status of a server BSU (as reflected in the service record), then a master-slave role reversal is performed according to the Bluetooth protocol. Better may refer, for example, having a lower wave number, closer to the seed, or other indication that the synchronization is closer to the source. In this alternative embodiment, however, the client is never synchronized in the first place. It is only searching for a master because it has not synchronized yet, and thus will never have "better" synchronization.

Inter-BSU timing synchronization is accomplished by communications between peer REMon dedicated radio applications. Synchronization is periodically refreshed to correct for drift between BSUs, as explained herein.

Each baseband controller on a BSU shares a common clock frequency oscillator so that there is no frequency drift between these devices, and no periodic refresh for frequency drift is required. However, depending upon the actual implementation, refreshing inter-BSU timing that results in adjusting the REMon dedicated radio applications on a BSU may require subsequent immediate adjustments to each other baseband controller on that BSU. Every Bluetooth unit has an internal system clock that determines the timing and hopping of the transceiver radios. The Bluetooth clock is derived from a free

running native clock that is generally never adjusted and never turned off. However, in the above embodiments, the native clock is directly modified. Under an alternative embodiment, for synchronization with other units, only temporary offsets are used that, added to the native clock, provide temporary Bluetooth clocks which are mutually synchronized.

For the initial Bluetooth embodiment, there is one time period of interest, namely, a 10 μ sec "window of uncertainty time" within which the slot edge is deemed valid. At a maximum clock drift of 12.5 nanosec per slot, the system may drift 10 μ sec over 800 slots (which takes 500 millisecc). Thus, performing a resynchronization every 80 millisecc is well within the 500 millisecc maximum allowable interval. If a system requires greater than 64 assignable slots (such as in dense traffic areas), then REMon may be increased to provide 128 slot pairs. In this case, the resynchronization interval doubles to 160 milliseconds.

In a fully synchronized system, each base station will be listening for synchronization broadcasts from base stations acting as a sync server. Any received broadcast syncs will be combined according to a weighting process. The weighting process used may give greater or lesser weight to synchronization broadcasts based on their originating base stations' proximity to the seed. On one slot out of every 128 that a client BSU will, in turn, act as a sync server for zero or more BSUs who are at the next radial increment farther away. The waiting process used by a BSU will in one embodiment give greater weight to synchronization information coming from BSUs closer to the seed BSU, and smaller weight to BSUs farther away. This weighting process will be modifiable in software subject to "empirical tuning." Alternatively, signal processing performed may weight feedback constants, where such weights are calculated mathematically.

The "daisy-chaining" inherent in the synchronization process may result in gradually increasing synchronization differences or deltas between the seed BSU and BSUs that are geographically farther away. However, as noted herein, this is not a problem because only localized synchronization between neighbors is required (between neighboring BSUs), because handoff occurs only between neighbors.

Referring to Figure 17A, an example of a synchronization server record 1700 is shown that effectively stores server synchronization information particular to a given BSU. That BSU's unique address is stored in a Bluetooth device address ("BD_ADDR") field 1702, which corresponds to a unique 48 bit address specified under the Bluetooth protocol. The BD_ADDR uniquely identifies each Bluetooth device universally and thus, uniquely identifies a particular BSU associated with a record.

For the Bluetooth embodiment, a piece of information communicated during synchronization is a *network sync address* or neighbor sync master 1704, which is the Bluetooth device address of the (virtualized) synchronization source. It is the address used by every base station during the slots in which it is acting as a sync client or sync server. Knowing the network sync address allows the base station to know the hop sequence with which sync information is being transmitted. The network sync address is used only for this sync operation.

A wave number field 1706 and a transmit slot field 1708 store a wave number during which the BSU transmits its synchronization information when it acts as a sync server, and a transmit slot when it provides such information, respectively.

Another piece of information that may be transmitted as part of "sync information" is the IP address of the system controller. Each base station needs to know this. In installations where the system controller is running a Dynamic Host Configuration Protocol ("DHCP") server, each base station employs this address to request its own IP address. Having the system controller's IP address in the "sync information" will guarantee automatic propagation of this address to all base stations in an efficient manner.

A synchronization server record may also include the BSU's IP address, which is a local IP address 1712 created where the system controller is running DHCP. This IP address may be provided to neighboring base stations during synchronization. The record may include other information (such as system data) to be shared with neighbor BSUs. Indeed any field may be provided, and

include an appropriate flag to indicate that the field data is to be provided to neighbor BSUs during synchronization.

The synchronization server record 1700 may also include information regarding the record itself (which is typically not provided to neighbor BSUs), such as an initialization time field 1714 that identifies the date and time that the record was created (typically upon power up of the BSU). A resynchronization time field 1716 provides an indication as to when the synchronization process should again be performed to refresh such synchronization. This period may be system configurable. The time may be a periodic time at which synchronization is performed (e.g., every 80 milliseconds or hour), or some other value, such as a specific day and time. The server record 1700 may also include a record expiration field 1718 that identifies when the record expires, which may be set such that the record never expires, or set to a specific time (2:00 o'clock a.m.) at which it expires and initiates synchronization again.

Referring to Figure 17B, a neighbor list record that includes client synchronization information is shown as a record 1750. In this depicted embodiment, a separate record is stored and maintained at each BSU, where each record is associated with one neighbor BSU. (Under an alternative embodiment, a single record identifying all neighboring BSUs and their associated information may be employed.) While the server record 1700 defines how the base station provides its synchronization information to neighbors, and the information it provides, the neighbor list record includes information regarding how the base station gets its synchronization. Thus, the neighbor list record 1750 includes Bluetooth device address field 1702, the address corresponds to the neighbor from which it received its server information. The wave number field 1706 and transmit slot field 1708 define the wave and transmit slot from which it received its synchronization information. Likewise the neighbor list may also include the neighbor's IP address in field 1712 when the system controller employs DHCP.

Some Alternative Implementations

Under one alternative embodiment shown in Figure 18, each BSU transmits a continuous synchronization signal over wired media 1802 (such as unused wire pairs in the PMN backend network or perhaps over the power line), instead of over the air. The synchronization signal would have a regular duty cycle, such as a square wave or sine wave. Dedicated circuitry 1804 in base stations extract the synchronization data from the signal. Dedicated source circuitry 1806 in the PMN generates the synchronization signal. A feed-forward system compensates for signal delay, which increases with the length of wire traveled by the signal.

A benefit of this approach is that it reduces complexity in the base station. A potential disadvantage of this embodiment is added circuitry cost, as well as constraints placed on how the system's backend network gets wired. For example, referring to Figure 15, neighboring base stations B4 and B8 could not be wired such that they end up at opposite ends of a network segment, having maximum signal delay between them.

Under another alternative embodiment, instead of utilizing direct base station-to-base station RF resources to communicate and synchronize, a pair of base stations could utilize one or more "common" MUs that they can both "hear." One base station, the "owner," controls the MU but has asked the other base station to monitor the MU's link quality in case a handoff proves necessary. When one BSU communicates with another, or are both communicating with a BSU, their slots are necessarily aligned, and thus they are synchronized. The owner may then synchronize to transmit slots received from the MU, and thereafter synch with neighbors.

Under the Bluetooth standard, only one base station can be the owner (assuming the system does not use scatternets). By virtue of the owner-slave relationship, the owner base station and MU have their slots aligned. The other base station, which does not have a Bluetooth link with the MU but which is monitoring the MU signal strength, gets slot edge alignment information indirectly via its monitoring activity. A drawback to this embodiment is its

reliance on the presence of such MUs in the coverage overlap zones between base stations. If there do not happen to be any such MUs, then synchronization does not occur. Another drawback with this approach is reduced accuracy since both base stations are synchronizing to a third party, the MU, instead of to each other. The MU's own error adds directly to the synchronization delta. Overall synchronization stability could be reduced.

As will be evident from the detailed description provided herein, aspects of the invention minimize the number of packets that must be retransmitted when channel collisions occur. Channel collisions are largely unavoidable, and they normally invalidate data in the overlapping portions of both colliding transmission slots in an unaligned system. By synchronizing base stations such that any colliding transmission slots are fully overlaid, the system reduces the impact of channel collisions to a single slot of data.

Aspects of the invention also provide for slot reassignment to efficiently perform synchronization or other tasks where a BSU acts both as a client and as a server. Aspects of the invention allocate slot patterns to radios controlled by the BSU to permit most efficient resource allocation.

Further, aspects of the invention improve base station-to-base station handoff. A base station monitors the wireless link quality of MUs on behalf of other base stations. For an MU that an owner base station wants a neighbor base station to monitor for RF link quality, embodiments of this invention enable the neighbor base station to very quickly listen for the MU since the receive slot time is known. Only when both base station slot edges are aligned can the neighbor immediately listen without first having to account for slot edge misalignment.

Skipped slots become unnecessary when an MU leaving one base station is handed off to a new base station. Data to the MU does not have to wait for the next slot "leading edge" to arrive at the new base station; the new base station's next slot begins precisely where the previous base station's slot ended.

Handoff setup becomes faster since there is no need to send slot edge information from old to new base stations, because as they are already aligned. There is less setup data to calculate and transmit.

While base station units are generally described herein, aspects of the invention may employ any "Bluetooth switch." Such a switch may have less functionality and be cheaper to implement than a base station unit. Aspects of the invention apply to nodes in a network, such as network access points, stationary nodes in an picocellular communications network, peer-to-peer stationary network access points, and the like.

Those skilled in the relevant art will appreciate that the invention can be practiced with various telecommunications or computer system configurations, including Internet appliances, hand-held devices, wearable computers, palm-top computers, cellular or mobile phones, multi-processor systems, microprocessor-based or programmable consumer electronics, set-top boxes, network PCs, mini-computers, mainframe computers, and the like. Aspects of the invention can be embodied in a special purpose computer or data processor that is specifically programmed, configured, or constructed to perform one or more of the computer-executable instructions explained in detail herein. Indeed, the term "computer", as used generally herein, refers to any of the above devices, as well as to any data processor. Data structures and transmission of data particular to aspects of the invention are also encompassed within the scope of the invention. In general, while hardware platforms such as stationary and mobile devices are described herein, aspects of the invention are equally applicable to nodes on the network having corresponding resource locators to identify such nodes.

Incorporated by reference herein are all above references, patents, or applications and the following U.S. applications, which are assigned to the assignee of this application: Application No. _____, entitled Link Context Mobility, Such as for use in Wireless Networks, (Attorney Docket No. 34015.8001) _____; and Application No. _____, entitled Wireless Base Neighbor Discovery in a Communication System, Such as a System Employing a Short-Range Frequency Hopping Scheme (Attorney Docket No.

34015.8002); Application No. _____, entitled Connection Initiation in Wireless Networks Including Load Balancing (Attorney Docket No. 34015.8004); application no. 60/288,301, entitled Reducing Mutual Channel Interference in Frequency-Hopping Spread Spectrum Wireless Communication Systems, Such as Bluetooth Systems (Attorney Docket No. 34015.8005); Application No. 60/288,300, entitled Method and System for Indicating Link Quality Among Neighboring Wireless Base Stations (attorney docket no. 34015.8006); and application no. 60/311,716, entitled Virtual Bluetooth Devices as a Means of Extending Pairing and Bonding in a Bluetooth Network (Attorney Docket No. 34015.8007). Aspects of the invention can be modified, if necessary, to employ the systems, functions and concepts of the various patents and applications described above to provide yet further embodiments of the invention.

Unless the context clearly requires otherwise, throughout the description and the claims, the words "comprise," "comprising," and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of "including, but not limited to." Words using the singular or plural number also include the plural or singular number respectively. Additionally, the words "herein," "hereunder," "above," "below," and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application.

The above detailed descriptions of embodiments of the invention are not intended to be exhaustive or to limit the invention to the precise form disclosed above. While specific embodiments of, and examples for, the invention are described above for illustrative purposes, various equivalent modifications are possible within the scope of the invention, as those skilled in the relevant art will recognize. For example, while steps are presented in a given order, alternative embodiments may perform routines having steps in a different order. The teachings of the invention provided herein can be applied to other systems, not necessarily the PMN system described herein. The various embodiments described herein can be combined to provide further embodiments. These and other changes can be made to the invention in light of the detailed description.

These and other changes can be made to the invention in light of the above detailed description. In general, the terms used in the following claims, should not be construed to limit the invention to the specific embodiments disclosed in the specification, unless the above detailed description explicitly defines such terms. Accordingly, the actual scope of the invention encompasses the disclosed embodiments and all equivalent ways of practicing or implementing the invention under the claims.

While certain aspects of the invention are presented below in certain claim forms, the inventors contemplate the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as embodied in a computer-readable medium, other aspects may likewise be embodied in a computer-readable medium. Accordingly, the inventors reserve the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the invention.

CLAIMS

I/we claim:

- [c1] 1. A system for wirelessly exchanging communications with at least one mobile unit, the system comprising:
- a first base station unit coupled to a network;
 - a second base station unit coupled to the network, wherein the first and second base station units are configured to communicate wirelessly with the mobile unit under a Bluetooth protocol, wherein the Bluetooth protocol is a short-range, frequency hopping protocol for transmitting and receiving information during time division duplex transmit and receive slots;
- wherein the first and second base station units are further configured for:
- handing-off communications with the mobile unit from the first to the second base station unit; and
 - aligning the transmit and receive slots with respect to the first and second base station units.
- [c2] 2. A system for wirelessly exchanging communications with at least one mobile unit, the system comprising:
- a first base station unit coupled to a network;
 - a second base station unit coupled to the network, wherein the first and second base station units are configured to communicate wirelessly with the mobile unit under a short-range, frequency hopping protocol for transmitting and receiving information during time division duplex having pairs of transmit and receive slots, wherein the pairs of transmit and receive slots repeat over a cycle of n pairs of slots;
- wherein the first and second base station units are further each configured for:

synchronizing the transmit and receive slots with respect to the first and second base station units; and

during synchronizing, transmitting synchronization information during one transmit slot of the n pairs of slots, and converting all other transmit slots in the n pairs of slots to receive slots to thereby receive synchronization information for other base stations.

[c3] 3. The system of claim 1 wherein the first and second base station units are further each configured for:

selecting one of the transmit slots of the n pairs of slots; and

notifying the other base stations that the one transmit slot has been selected.

[c4] 4. In a network, a method for wirelessly exchanging communications with at least one mobile unit, wherein the network includes first and second base stations units coupled to the network, and may include a system controller coupled to the network, the method comprising:

at the first base station unit, receiving a communication signal from the second base station unit, wherein the first and second base station units are configured to employ a wireless communications protocol, and wherein the wireless communications protocol does not provide for handoff of communications links between base station units;

at the first base station unit, determining if the second base station unit has been synchronized based on the communication signal;

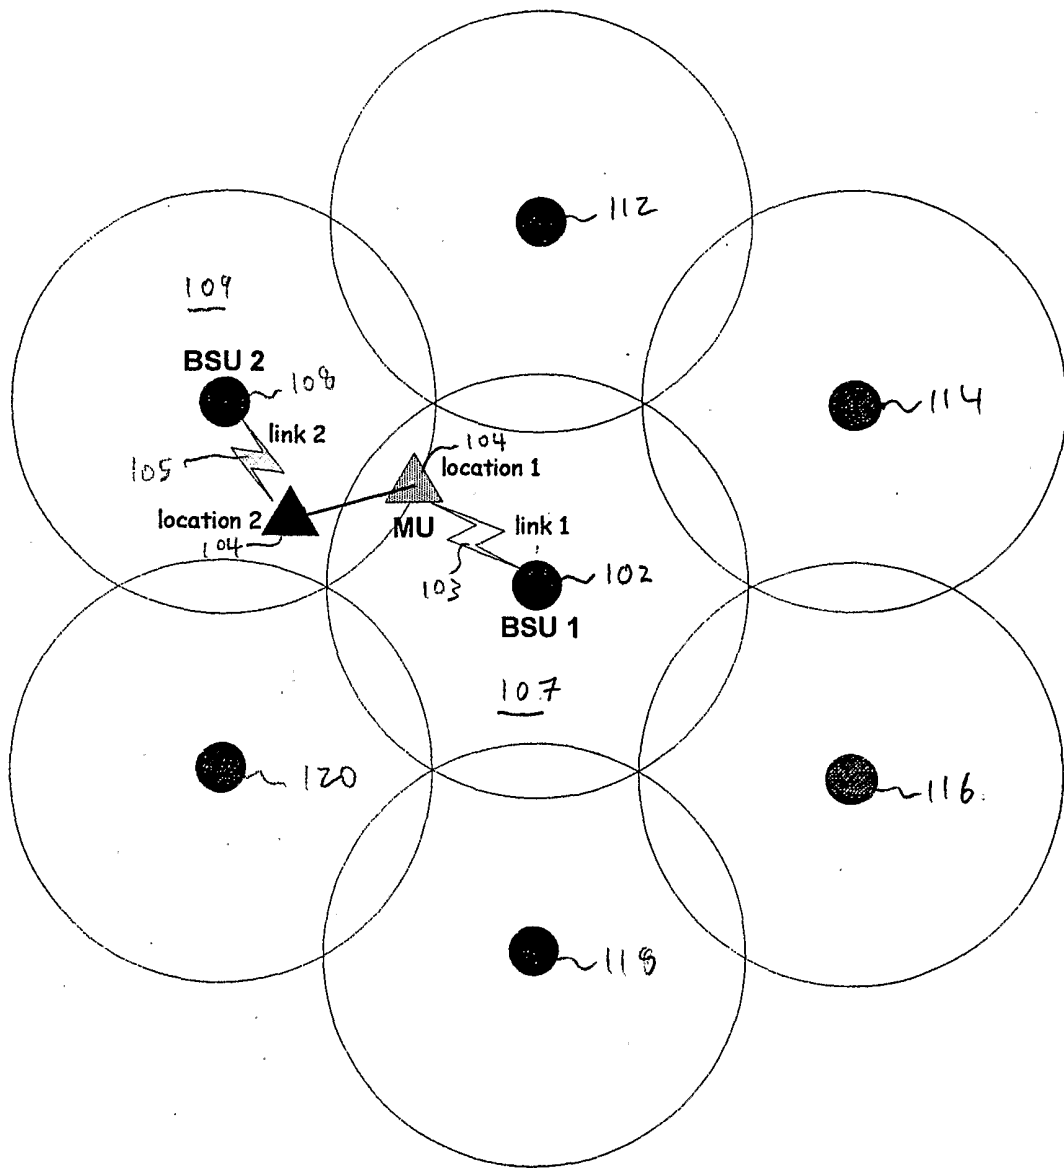
at the first base station unit, if the second base station unit is synchronized, then synchronizing an internal clock based on the synchronized second base station, wherein the synchronizing is performed without assistance from a system controller.

- [c5] 5. The method of claim 4 wherein the a wireless communications protocol is a Bluetooth protocol.
- [c6] 6. The method of claim 4 wherein the communication signal is a wireless communication signal.
- [c7] 7. The method of claim 4 wherein base station units in the network are synchronized in waves based on an initial communication signal from an initial seed base station unit.
- [c8] 8. The method of claim 4, further comprising:
at the first base station unit, if the second base station unit is not synchronized, then waiting until a new communication from another base station unit in the network that is synchronized.
- [c9] 9. In a network, a method for wirelessly exchanging communications with at least one mobile unit, wherein the network includes first and second base stations units coupled to the network, and may include a system controller coupled to the network, the method comprising:
receiving a synch signal from a neighboring base station unit;
locking neighboring base station units from synchronizing;
choosing a transmission slot;
recording synchronization information, and
unlocking neighboring base station units.
- [c10] 10. A computer-readable medium containing a data structure for use by a network access point in a short-range communications network, the data structure comprising:
at least one record stored at the network access point in the short-range communications network, wherein each record comprises:
a unique identifier for the network access point, wherein the network access point may communicate wirelessly under a neighboring network access point, wherein the short-range wireless

protocol permits the neighboring base station unit to communicate wirelessly with an effective range of less than 500 meters;

a synchronization wave number identifying when the network access point is synchronized; and

synchronization transmit slot identifying when the network access point transmits its synchronization information to neighboring network access points.



200

Fig 1



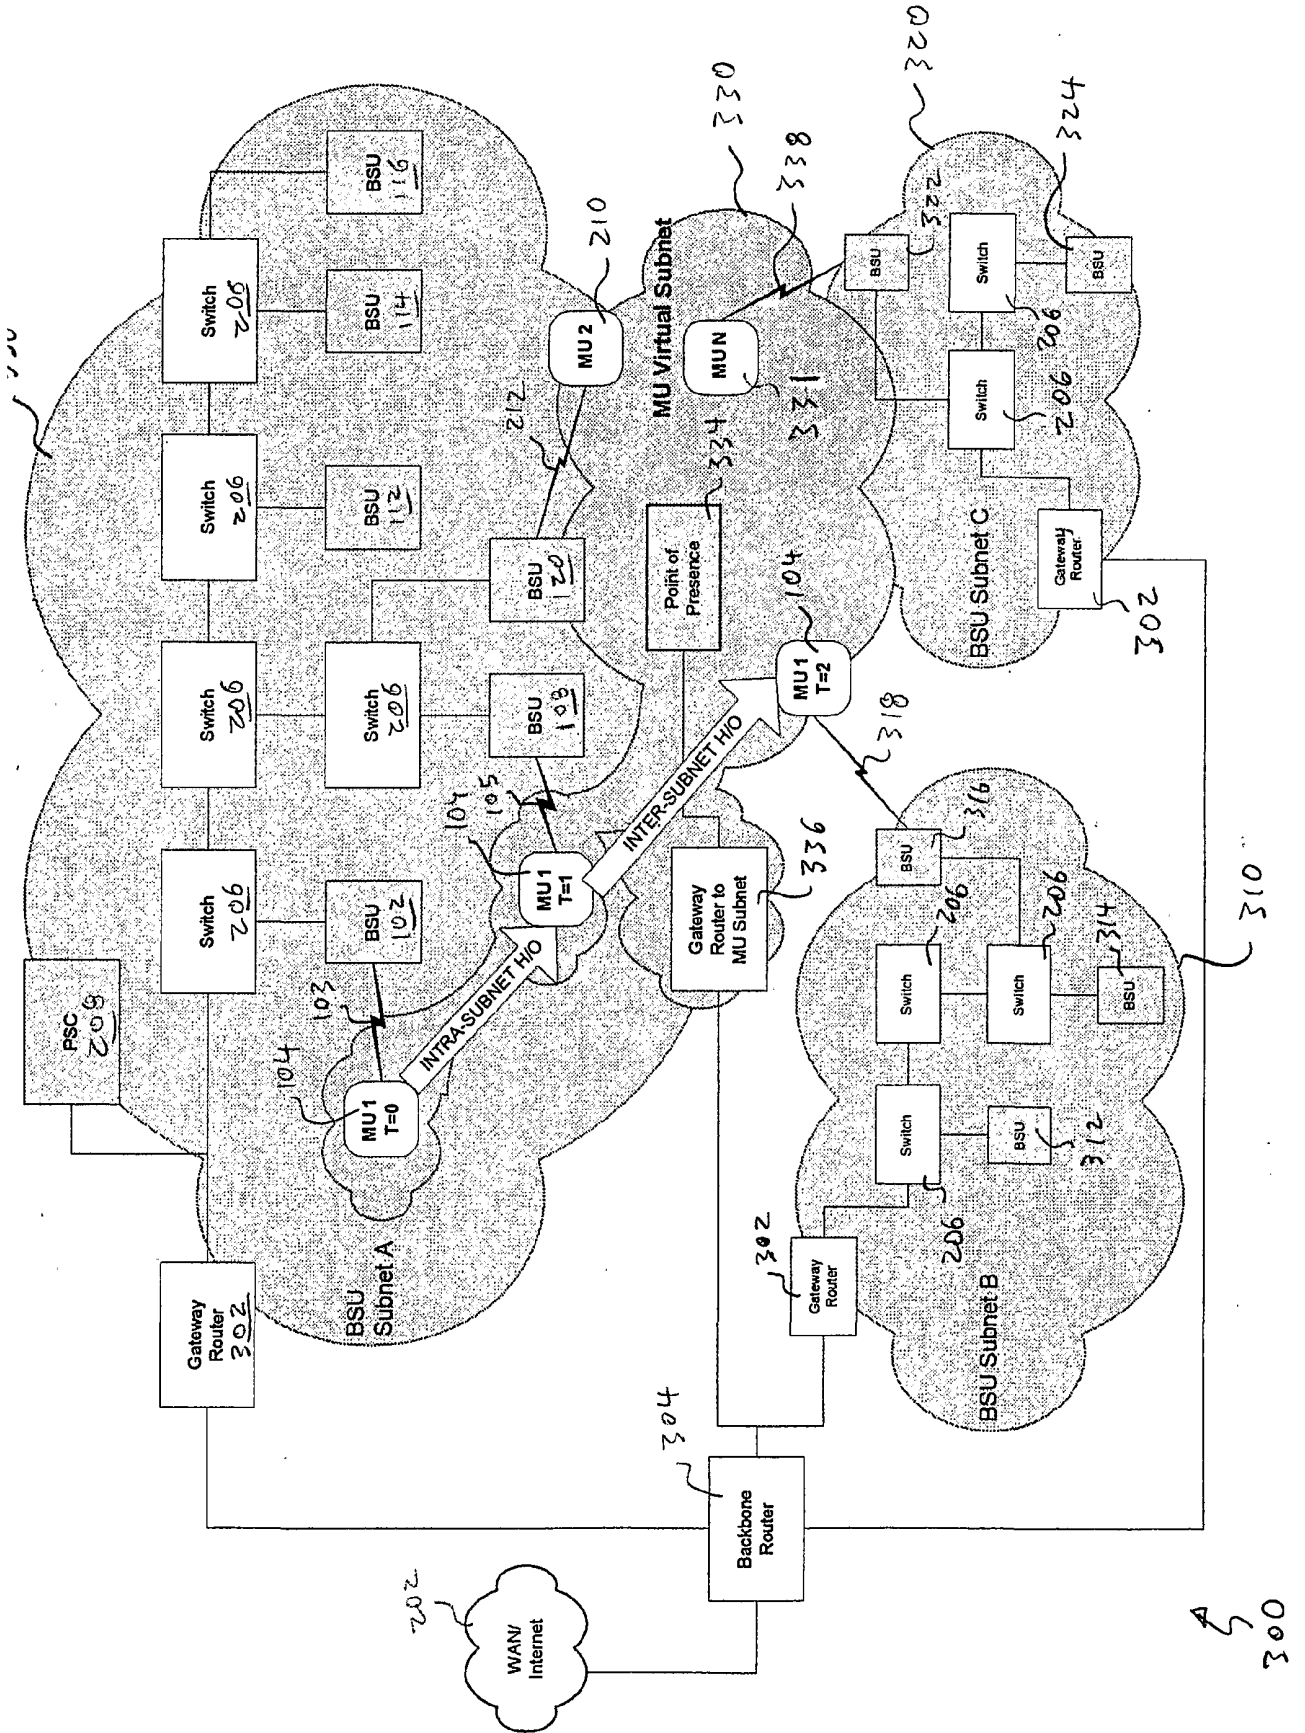


FIG. 2A

FIG. 2B






Figure 4 Intra-subnet Hand-off





Figure 5 Inter-subnet Hand-off



Unaligned case:




FIG. 7A PRIOR ART

Aligned case:





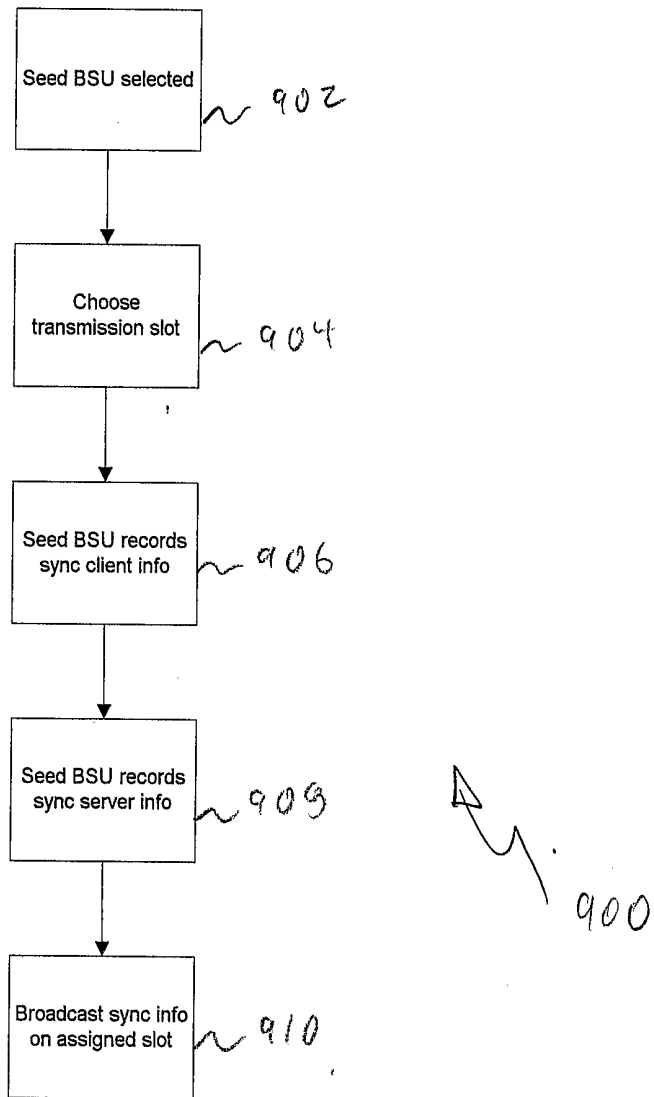
FIG. 7B



Note: **Rx** slots in bold are converted Tx slots

FIG. 8

FIG. 9A







FIG. 9B







FIG. 10







FIG. 12






FIG. 14




FIG. 15




FIG 16

Synchronization Server Record

Bluetooth device address ("BD Addr")	<unique 48 bit address>	1702
Neighbor's sync master	<predetermined BD Addr>	1704
Wave Number	<integer>	1706
Transmit Slot	<integer>	1708
System Controller Address	<IP Address>	1710
BSU's IP Address	<local IP Address>	1712
Initialization Time	<date, time>	1714
Re-synchronization Time	<date, time>	1716
Record Expiration	<date, time>	1718

A

Fig. 17A

1700

Neighbor List Record (including client synchronization information)

Neighbor's Bluetooth device address ("BD Addr")	<unique 48 bit address>	1702
Wave Number	<integer>	1706
Transmit Slot	<integer>	1708
Neighbor's IP Address	<local IP Address>	1712
Initialization Time	<date, time>	1714
Re-synchronization Time	<date, time>	1716
Record Expiration	<date, time>	1718

A

Fig. 17B

1750




FIG. 18