

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
11 August 2005 (11.08.2005)

PCT

(10) International Publication Number
WO 2005/074212 A1

(51) International Patent Classification⁷: **H04L 12/56**

(21) International Application Number:
PCT/US2005/003576

(22) International Filing Date: 27 January 2005 (27.01.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/540,504 29 January 2004 (29.01.2004) US

(71) Applicant (for all designated States except US): **QUALCOMM Incorporated** [US/US]; 5775 Morehouse Drive, San Diego, California 92121 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **NANDA, Sanjiv** [US/US]; 16808 Daza Drive, Ramona, California 92065 (US). **WALTON, Jay Rodney** [US/US]; 85 Highwoods Lane, Carlisle, Massachusetts 01741 (US).

(74) Agents: **WADSWORTH, Philip R.** et al.; 5775 Morehouse Drive, San Diego, California 92121 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, 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, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE,

[Continued on next page]

(54) Title: DISTRIBUTED HIERARCHICAL SCHEDULING IN AN AD HOC NETWORK



(57) Abstract: An ad hoc network with distributed hierarchical scheduling is disclosed. In one aspect, stations in a network mesh detect interfering neighbor stations and form interference lists. Stations transmit their interference lists. Scheduling stations schedule allocations for child stations in response to interference lists, received remote allocations, or a combination thereof. Coordination messages are transmitted including frame structure, allocations, and interference lists, among others. In another aspect, an ad hoc mesh network may be organized into a tree topology. In an example wireless backhaul network, this matches traffic flow. Distributed, hierarchical scheduling is provided where parents schedule communication with children while respecting already scheduled transmissions to/from interferers and to/from interferers of their respective children. Procedures to construct interference constraints for distributed, hierarchical scheduling are described, resulting in efficient scheduling and reuse in an ad hoc wireless network, without centralized scheduling. Various other aspects are also disclosed.



EG, 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, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

— as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

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.

DISTRIBUTED HIERARCHICAL SCHEDULING IN AN AD HOC NETWORK

Claim of Priority under 35 U.S.C. §119

[0001] The present Application for Patent claims priority to Provisional Application No. 60/540,504 entitled "Distributed Hierarchical Scheduling" filed January 29, 2004, and assigned to the assignee hereof and hereby expressly incorporated by reference herein.

BACKGROUND

Field

[0002] The present invention relates generally to communications, and more specifically to distributed hierarchical scheduling in an ad hoc network.

Background

[0003] Wireless communication systems are widely deployed to provide various types of communication such as voice and data. A typical wireless data system, or network, provides multiple users access to one or more shared resources. A system may use a variety of multiple access techniques such as Frequency Division Multiplexing (FDM), Time Division Multiplexing (TDM), Code Division Multiplexing (CDM), and others.

[0004] Example wireless networks include cellular-based data systems. The following are several such examples: (1) the "TIA/EIA-95-B Mobile Station-Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular System" (the IS-95 standard), (2) the standard offered by a consortium named "3rd Generation Partnership Project" (3GPP) and embodied in a set of documents including Document Nos. 3G TS 25.211, 3G TS 25.212, 3G TS 25.213, and 3G TS 25.214 (the W-CDMA standard), (3) the standard offered by a consortium named "3rd Generation Partnership Project 2" (3GPP2) and embodied in "TR-45.5 Physical Layer Standard for cdma2000 Spread Spectrum Systems" (the IS-2000 standard), and (4) the high data rate (HDR) system that conforms to the TIA/EIA/IS-856 standard (the IS-856 standard).

[0005] Other examples of wireless systems include Wireless Local Area Networks (WLANs) such as described in the IEEE 802.11 standards (e.g. 802.11 (a), (b), or (g)). Improvements over these networks may be achieved in deploying a Multiple Input

Multiple Output (MIMO) WLAN comprising Orthogonal Frequency Division Multiplexing (OFDM) modulation techniques.

[0006] WLANs are commonly deployed to provide data services to one or more user terminals within an area. A cellular-like deployment of WLAN coverage areas may be deployed with an access point providing coverage for each cell. This nominal cellular architecture assumes that each access point has wired backhaul connectivity (e.g. a T1 backhaul). For example, 802.11 hotspots with a T1 backhaul are commonly deployed in various locations to provide access. A typical cellular-like WLAN deployment is illustrated in FIG. 1. In this example, data links 110A-C connect Internet 102 with Access Points (APs) 104A-C. Each access point 104 has a coverage area, commonly referred to as a cell 130, in which various User Terminals (UTs) 106 may communicate with the access point 104 via a wireless link 120.

[0007] Due to the limited cell radius and dense deployment of APs, as well as the expense, limited bandwidth and limited availability of T1 backhaul, alternate methods of backhaul are of interest. Various ad hoc networks may be formed. For example, consider a multi-hop wireless backhaul architecture, where APs form a peer-to-peer mesh network. In such a hierarchical architecture, only the APs, and not the user terminals (UTs) supported by each AP, participate in the backbone mesh. An earlier, similar approach is the Ricochet model, developed by Metricom Inc., of San Jose, California, now owned by YDI Wireless, Inc., of Falls Church, Virginia. In order to enable a mesh network, protocols are needed to establish communication between nodes in the mesh, also needed are techniques for scheduling reception and transmission to maximize throughput and/or minimize interference between mesh nodes and/or user terminals communicating via the mesh network. There is therefore a need in the art for distributed hierarchical scheduling in an ad hoc network.

SUMMARY

[0008] According to one aspect, an apparatus is described which includes a receiver for receiving one or more signals from one or more remote devices, a message decoder for decoding one or more first interference lists from coordination messages contained in the one or more received signals from respective one or more remote devices, and a memory for storing the one or more first interference lists and a second interference list comprising identifiers associated with the one or more remote devices.

- [0009] According to another aspect, a method is provided for detecting signals received from one or more remote stations, identifying the one or more remote stations, storing an identifier associated with each of the one or more remote stations in a first interference list, and transmitting the first interference list for reception by neighbor remote stations.
- [0010] According to another aspect, an apparatus is described which includes means for detecting signals received from one or more remote stations, means for identifying the one or more remote stations, means for storing an identifier associated with each of the one or more remote stations in a first interference list, and means for transmitting the first interference list for reception by neighbor remote stations.
- [0011] According to another aspect, an apparatus is described which includes a receiver for receiving one or more signals from one or more remote devices, a message decoder for decoding one or more first interference lists and one or more remote allocations from coordination messages contained in the one or more received signals from respective one or more remote devices, and means for allocating a shared resource for receiving and transmitting by one or more remote devices in response to the one or more first interference lists and one or more remote allocations.
- [0012] According to another aspect, a message is disclosed, operable for transmission by a mesh station in a hierarchically distributed mesh network, comprising an interference list field comprising a list of interfering remote stations, a transmit allocation field comprising one or more allocations for transmission by a child remote station on a shared medium, and a receive allocation field comprising one or more allocations for receiving by a child remote station on a shared medium.
- [0013] According to another aspect, a mesh network is disclosed comprising: a first mesh station for establishing a superframe and transmitting a first coordination message, the first coordination message identifying the superframe and comprising one or more allocations; and one or more second mesh stations for receiving the first coordination message, and transmitting or receiving in accordance with the one or more allocations.
- [0014] According to another aspect, computer readable media is disclosed operable to perform detecting signals received from one or more remote stations, identifying the one or more remote stations, storing an identifier associated with each of the one or more remote stations in a first interference list, and transmitting the first interference list for reception by neighbor remote stations.
- [0015] Various other aspects and embodiments are also disclosed.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0016] FIG. 1 is an example of a prior art WLAN cellular-like system;
- [0017] FIG. 2 illustrates an example embodiment of a system deploying a mesh network backhaul;
- [0018] FIG. 3 illustrates a more generalized example of system comprising a generalized BMS mesh;
- [0019] FIG. 4 depicts an example embodiment of a wireless communication device;
- [0020] FIG. 5 depicts an example embodiment of a method for transmitting and receiving on a mesh network;
- [0021] FIG. 6 illustrates an example embodiment of a superframe;
- [0022] FIG. 7 illustrates an example embodiment of a method for adding a BMS within the mesh network tree;
- [0023] FIG. 8 depicts an example embodiment of a method for generating interference lists;
- [0024] FIG. 9 illustrates an illustrative portion of an example embodiment of a BMS;
- [0025] FIG. 10 illustrates an example embodiment of a method for forming transmission and reception schedules at a BMS;
- [0026] FIG. 11 illustrates an example embodiment of a method for forming constraints in response to received scheduled transmissions;
- [0027] FIG. 12 illustrates an example coordination message;
- [0028] FIG. 13 illustrates an example mesh network; and
- [0029] FIG. 14 illustrates an example superframe structure.

DETAILED DESCRIPTION

- [0030] Distributed ad hoc network embodiments utilizing hierarchical scheduling are detailed herein. When used as a supplement to or replacement for traditional wired backhaul “hotspot” topologies, expense and traffic bottlenecks may be avoided. A typical ad hoc network, such as a mesh network of peer stations, may share a single frequency or shared medium. Alternatively, more than one frequency may be used by a wireless mesh network. For example, one or more frequencies may be used for backhaul, and one or more alternate frequencies may be used for access by user terminals. A backhaul may be deployed using any communication format or comply to any standard, including those examples listed above, and may be modified to comply

with one or more techniques detailed herein. Thus, any physical layer and/or Media Access Control (MAC) protocol may be adapted for use with techniques described herein. In one example, an ad hoc wireless mesh network of peer access points may be deployed to provide backhaul, as described below. User terminal to access point communication may be performed between any nodes of the backhaul mesh and user terminals deployed throughout the coverage area using any WLAN standard (e.g. an 802.11 standard, of any variant). In another embodiment, all stations, whether an AP or a UT, may be described as a Backhaul Mesh Station (BMS), and may be part of the backhaul mesh (e.g., similar to AP functionality) or may be solely a leaf node (e.g., similar to UT functionality). A wireless communication device may comprise all or a subset of the described functionality. For example, a powered station with multiple antennas may be well suited to provide connectivity to adjacent nodes, as well as form a portion of a wireless backhaul. A battery powered handheld user terminal, for example, may be deployed solely as a leaf node, or provide a limited amount of backhaul connectivity. These and other embodiments are detailed further below.

[0031] The wireless backhaul mesh aggregates traffic to a level 0 BMS, and distributes traffic from the level 0 BMS to nodes throughout the mesh. The level 0 BMS, detailed further below, may be connected to an external network (such as the Internet, or another intranet, whether wired or wireless; e.g. a network Point-of-Presence POP). The traffic flowing up and down the hierarchical mesh network is scheduled to avoid interference, when possible, to increase throughput.

[0032] Example embodiments are disclosed herein that support highly efficient operation in conjunction with very high bit rate physical layers for an ad hoc wireless LAN (or similar applications that use newly emerging transmission technologies). The example WLAN supports bit rates in excess of 100 Mbps (million bits per second) in bandwidths of 20 MHz. While wireless backhauls of high bandwidth are desirable, the techniques disclosed herein apply with any throughput rates.

[0033] An example embodiment supports high data rate, high bandwidth physical layer transport mechanisms in general, including, but not limited to, those based on OFDM modulation, single carrier modulation techniques, systems using multiple transmit and multiple receive antennas (Multiple Input Multiple Output (MIMO) systems, including Multiple Input Single Output (MISO) systems) for very high bandwidth efficiency operation, systems using multiple transmit and receive antennas in conjunction with spatial multiplexing techniques to transmit data to or from multiple user terminals

during the same time interval, and systems using code division multiple access (CDMA) techniques to allow transmissions for multiple users simultaneously.

[0034] One or more exemplary embodiments described herein are set forth in the context of a wireless data communication system. While use within this context is advantageous, different embodiments of the invention may be incorporated in different environments or configurations. In general, the various systems described herein may be formed using software-controlled processors, integrated circuits, or discrete logic. The data, instructions, commands, information, signals, symbols, and chips that may be referenced throughout the application are advantageously represented by voltages, currents, electromagnetic waves, magnetic fields or particles, optical fields or particles, or a combination thereof. In addition, the blocks shown in each block diagram may represent hardware or method steps. Method steps can be interchanged without departing from the scope of the present invention. The word “exemplary” is used herein to mean “serving as an example, instance, or illustration.” Any embodiment described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other embodiments.

[0035] FIG. 1 is an example of a prior art cellular-like system 100, comprising one or more Access Points (APs) 104 connected to one or more User Terminals (UTs) 106. Each AP and its associated UTs communicate via a Wireless Local Area Network (WLAN) 120. In the example embodiment, WLAN 120 is a high speed MIMO OFDM system. However, WLAN 120 may be any wireless LAN. Access point 104 communicates with any number of external devices or processes via network 102. Network 102 may be the Internet, an intranet, or any other wired, wireless, or optical network. Connection 110 carries the physical layer signals from the network to the access point 104. Devices or processes may be connected to network 102 or as UTs (or via connections therewith) on WLAN 120. Examples of devices that may be connected to either network 102 or WLAN 120 include phones, Personal Digital Assistants (PDAs), computers of various types (laptops, personal computers, workstations, terminals of any type), video devices such as cameras, camcorders, web cams, and virtually any other type of data device. Processes may include voice, video, data communications, etc. Various data streams may have varying transmission requirements, which may be accommodated by using varying Quality of Service (QoS) techniques.

[0036] System 100, as shown in FIG. 1 is deployed with a centralized AP 104. All UTs 106 communicate with the AP in this prior art example. In the example embodiment, this system may be modified to accommodate direct peer-to-peer communication between two UTs, in order to form a mesh network, as illustrated in FIG. 2 below.

[0037] FIG. 2 illustrates an example embodiment of a system 200 deploying a mesh network backhaul for communication between various user terminals 106 as well as connections to devices through wired or wireless backhaul connection 110 (which connects the wireless backhaul to an external network 102 such as the Internet or other external network). In FIG. 2, as with FIG. 1, various access points 104 provide coverage over a geographic region or cell 130 in which various user terminals 120 may communicate with the access point via wireless links 120. In contrast with FIG. 1, an access point may not have a direct connection to an external network such as the Internet 102. In this example, access point 104C is connected to an external network via connection 110. All the other access points 104 are connected through a backhaul link, or mesh link 210. In this example, access point 104C connects with access points 104B and 104D. Access point 104A communicates with access point 104B via mesh link 210A. Thus, user terminals 106, associated with access point 104A in cell 130A may connect to the external network via a path formed through the mesh network using mesh links 210A and 210B. Thus, as will be detailed further below, an access point may communicate with a parent access point as well as one or more child access points. Furthermore, an access point may have user terminals, which may also be considered children. To avoid confusion in terminology, since each AP is just a station in the backbone mesh, when discussing procedures on the backbone mesh, an AP or UT may be referred to as a backbone mesh station (BMS). Various other cells 130 may also comprise one or more user terminals communicating with their respective access points, details not shown. System 200 serves to illustrate one embodiment of a mesh network using an AP for a BMS.

[0038] FIG. 3 illustrates a more generalized example of system 300 comprising a generalized BMS mesh connected to an external network 310 via connection 320. Note that connection 310 or 320 may be wired or wireless, as described above. In this example, each node in the mesh is called a backbone mesh station (BMS). Each BMS may be described as having a level within the hierarchical tree structure formed. For example, level zero BMS, BMS0, is connected to the external network 310 as a point-of-presence (POP). In an example embodiment, the level 0 BMS is a POP in a wired

network. In general, the connection to the external network may be formed using any communication standard, wired or wireless. Each level 1 BMS is one hop away from the level 0 BMS. In this example, BMS1 A-D are connected in the mesh to BMS0 in the mesh via mesh links 330A-D. Each BMS except for the root, (or level 0) BMS has a parent BMS and possibly one or more children. For example, BMS1 A communicates with level 2 BMSs, BMS2 A-B, via mesh links 340A-B respectively.

[0039] Note that the wireless links used within the mesh network may use any wireless technology, known in the art, including techniques such as those described in the aforementioned standards, modified, as detailed herein. Mesh networking standards and technologies are well known for wired networks, for example, the Internet, and those techniques can be applied to wireless mesh networks as well, with appropriate modification (examples include routing, discovery, packet forwarding, domain name services, DHCP, and the like). Additional examples are known, and are disclosed in various Internet protocols, including the Internet Control Message Protocol (ICMP), the ICMP Router-Discovery Protocol (IRDP), the Routing Information Protocol (RIP), and the Open Shortest Path First (OSPF) protocol.

[0040] In the example embodiment, the mesh network employs MIMO-OFDM techniques. Further, a BMS may communicate to one or more associated devices, such as user terminals, on alternate frequencies or other alternate channel, using any type of wireless link and/or communication standard. Or, the user terminal may be simply considered a child or leaf node in the mesh and may communicate on the mesh network itself.

[0041] In one example (as described above with respect to FIG. 2), a BMS is an access point. User terminals would associate with each BMS within that access point's cell for connectivity. In an alternate embodiment, a user terminal may also function as a BMS. The UT may form part of the wireless backhaul, and may also associate with other user terminals to provide access. As will be detailed further below, various BMSs may provide varying levels of support and features. For example, a BMS deployed with typical features of an access point, such as a steady power supply, multiple antennas, etc., may provide for higher power transmission, higher bandwidths, and so forth. A BMS deployed with features common to a mobile user terminal, such as a remote station cellular phone, PDA, etc., and may exhibit features such as fewer antennas, battery power, limited size, and the like, may support a subset of the features available on other BMSs. There are various additional application scenarios for the hierarchical

scheduling scheme. All stations (APs or UTs) that share the same frequency allocation (e.g. use the same band) may become a part of the mesh and may be referred to as BMSs. A mix of APs and UTs may be deployed, or a mesh consisting of UTs only may be deployed. In this case, a particular UT becomes the Level 0 BMS and the remaining UTs organize in a hierarchical tree structure around the Level 0 BMS. In general, a set of BMSs may include APs, UTs or any combination thereof. Any BMS may be a parent, child, or both. Those of skill in the art will readily deploy myriad combinations of BMSs in light of the teaching herein.

[0042] Returning to FIG. 3, a level 1 BMS, BMS1 C communicates to three level 2 BMSs, BMS2 C-E, via mesh links 340C-E. In this example, BMS2 C and BMS2 E are leaf BMSs. BMS2 D further communicates with level 3 BMSs, BMS3 A-B, via mesh links 350A-B, respectively.

[0043] In one embodiment, a primary motivation may be to supply a wireless backbone mesh for a network of APs. Notice that the throughput requirements for backhaul increase rapidly with increased aggregation. Thus, the BMS that provides a fixed network point-of-presence (POP) carries the aggregate throughput of several APs. The high throughput requirement on the backbone can be achieved through the use of MIMO at the BMSs. In terms of size and location, the BMSs may run on line power, may be of larger size, and may be equipped with 4 or more antennas (for example). This allows the BMSs to operate with higher transmit power and may not require sophisticated sleep cycles to conserve battery power. All these features may be combined to allow for high efficiency and high throughput on the mesh backbone.

[0044] In one embodiment, due to larger range requirement for inter-BMS communication, it may be desirable to use a lower frequency band (e.g. the 2.4 GHz unlicensed band or a 2.3 GHz licensed band) for the BMS mesh Frequency Allocation (FA). However, the increased range from the MIMO allows the mesh FA to be in the 5 GHz band (which has lower range than the 2.4 GHz band) but may be preferable due to lower interference. High availability (low outage) may be desirable for backhaul, since the backhaul network supports the aggregate traffic of many users. Note that the mesh itself provides reliability through redundant routes, since a level 2 BMS can find two-hop connectivity through more than one level 1 BMS in case of failure of the parent device.

[0045] Thus, a set of BMSs form a mesh network in which "adjacent" BMSs can communicate with each other. When all BMSs share a common frequency allocation,

communication on inter-BMS links may cause interference at other inter-BMS links. Distributed hierarchical scheduling, described herein, is designed to provide interference avoidance on the mesh.

[0046] A tree topology is well-suited for efficient transport of broadcast and multicast traffic which may be of interest for wide area networking. Broadcast and multicast traffic may be scheduled in the same manner as unicast traffic in the backbone mesh, as detailed further below. In a richly connected mesh of another configuration, additional procedures would need to be defined to construct a tree topology for broadcast and multicast, but in embodiments detailed herein, the tree topology is already available.

[0047] In one embodiment, it may be convenient to use CSMA/CA and the ad hoc networking concepts from the 802.11 distributed coordination function (DCF) for initial access and discovery of new BMSs on the mesh backbone. However, the high throughput requirements for the mesh backbone may make the use of 802.11 DCF inefficient. Although the DCF may be suitable for a mesh network, DCF procedures do not exploit either the hierarchical structure of the backbone, nor the knowledge about mutual interferers that can be made available from other transmissions within the mesh. Distributed hierarchical scheduling allows for these advantages.

[0048] FIG. 4 depicts an example embodiment of a wireless communication device 400, which may be configured as a BMS, such as an access point 104 or user terminal 106. Transceiver 410, used in a level 0 BMS, receives and transmits on connection 320 according to the physical layer requirements of network 310. Data from or to devices or applications connected to network 310 are delivered to processor 420.

[0049] Processor 420 receives data, signals, messages, etc. from WLAN transceiver 440 (or transceiver 410 in a level 0 BMS). Received data may be processed for transmission on another link, such as a mesh link, or on a WLAN (when serving as an AP). Media Access Control (MAC) processing may be performed in processor 420. Internal control and signaling is also communicated between BMSs. MAC Protocol Data Units (MAC PDUs), also referred to as Physical layer (PHY) Protocol Data Units (PPDUs), or frames (in 802.11 parlance) may be delivered to and received from wireless LAN transceiver 240.

[0050] In an example embodiment, the Adaptation layer (ADAP) and Data Link Control layer (DLC) are performed in processor 420. The physical layer (PHY) is performed on wireless LAN transceiver 440, in conjunction with processor 420. Those of skill in the art will recognize that segmentation of the various functions may be made

in any of a variety of configurations. Processor 420 may perform some or all of the processing for the physical layer. A wireless LAN transceiver may include a processor for performing MAC processing, or subparts thereof. Any number of processors, special purpose hardware, or combination thereof may be deployed.

[0051] Processor 420 may be a general-purpose microprocessor, a digital signal processor (DSP), or a special-purpose processor. Processor 420 may be connected with special-purpose hardware to assist in various tasks (details not shown). Various applications may be run on externally connected processors, such as an externally connected computer (or over a network connection), may run on an additional processor within wireless communication device 400 (not shown), or may run on processor 420 itself. Processor 420 is shown connected with memory 430, which may be used for storing data as well as instructions for performing the various procedures and methods described herein. Those of skill in the art will recognize that memory 430 may be comprised of one or more memory components of various types, that may be embedded in whole or in part within processor 420.

[0052] Wireless LAN transceiver 440 may be any type of transceiver (including any type of receiver and/or transmitter). In an example embodiment, wireless LAN transceiver 440 is an OFDM transceiver, which may be operated with a MIMO or MISO interface. OFDM, MIMO, and MISO are known to those of skill in the art. Various example OFDM, MIMO and MISO transceivers are detailed in co-pending U.S. Patent Application Serial No. 10/650,295, entitled "FREQUENCY-INDEPENDENT SPATIAL-PROCESSING FOR WIDEBAND MISO AND MIMO SYSTEMS", filed August 27, 2003, assigned to the assignee of the present invention. Alternate embodiments may include SIMO or SISO systems. An example receiver may comprise RF conversion components, (such as filters, amplifiers, mixers, oscillators, phase-locked loops, analog-digital converters, etc.), demodulators, RAKE receivers, equalizers, decoders, deinterleavers, combiners, and other components well known in the art. Signals received at transceiver 440 may be demodulated according to one or more communication designs or standards. In an example embodiment, a demodulator capable of demodulating MIMO OFDM signals is deployed. In alternate embodiments, alternate standards may be supported, and embodiments may support multiple communication formats.

[0053] Transmitted signals are formatted in transceiver 440 according to one or more wireless system standards or designs, such as those listed above. Examples of

components that may be included in transceiver 440 are amplifiers, filters, digital-to-analog (D/A) converters, radio frequency (RF) converters, and the like. Data and control channels may be formatted for transmission in accordance with a variety of formats. Data for transmission on the forward link data channel may be formatted in a modulator (not shown) according to a rate and modulation format indicated by a scheduling algorithm. Examples of other components that may be incorporated in a transmitter of a transceiver 440 include encoders, interleavers, spreaders, and modulators of various types.

[0054] Wireless LAN transceiver 440 is shown connected with antennas 250 A-N. Any number of antennas may be supported in various embodiments. Antennas 250 may be used to transmit and receive on various WLANs, including mesh links (e.g. 330, 340, 350, etc.).

[0055] Wireless LAN transceiver 440 may comprise a spatial processor connected to each of the one or more antennas 250. The spatial processor may process the data for transmission independently for each antenna or jointly process the received signals on all antennas. Examples of the independent processing may be based on channel estimates, feedback from a remote station, such as a UT or other BMS, channel inversion, or a variety of other techniques known in the art. The processing is performed using any of a variety of spatial processing techniques. Various transceivers of this type may use beam forming, beam steering, eigen-steering, or other spatial techniques to increase throughput to and from a given user terminal. In an example embodiment, in which OFDM symbols are transmitted, the spatial processor may comprise sub-spatial processors for processing each of the OFDM subchannels, or bins.

[0056] In an example system, a first BMS may have N antennas, and a second BMS may have M antennas. There are thus $M \times N$ paths between the antennas of the first BMS and second BMS. A variety of spatial techniques for improving throughput using these multiple paths are known in the art. In a Space Time Transmit Diversity (STTD) system (also referred to herein as "diversity"), transmission data is formatted and encoded and sent across all the antennas as a single stream of data. With M transmit antennas and N receive antennas there may be $\text{MIN}(M, N)$ independent channels that may be formed. Spatial multiplexing exploits these independent paths and may transmit different data on each of the independent paths, to increase the transmission rate.

[0057] Various techniques are known for learning or adapting to the characteristics of the channel between the two BMSs. Unique pilots may be transmitted from each

transmit antenna. The pilots are received at each receive antenna and measured. Channel state information feedback may then be returned to the transmitting device for use in transmission. Eigen decomposition of the measured channel matrix may be performed to determine the channel eigenmodes. An alternate technique, to avoid eigen decomposition of the channel matrix at the receiver, is to use eigen-steering of the pilot and data to simplify spatial processing at the receiver.

[0058] Thus, depending on the current channel conditions, varying data rates may be available for transmission to various user terminals or other BMS throughout the system. The wireless LAN transceiver 440 may determine the supportable rate based on whichever spatial processing is being used for the physical link.

[0059] The number of antennas may be deployed depending on the BMS data needs as well as size and form factor. For example, a high definition video display may comprise, for example, four antennas, due to its high bandwidth requirements, while a PDA may be satisfied with two. An example BMS access point may have four antennas.

[0060] The following distributed scheduling procedures may be deployed for efficient scheduling on a hierarchical architecture, described above, e.g., the mesh backbone. These procedures are based on several observations: Much of the traffic in the mesh backbone flows upstream from the BMSs to the BMS-POP and downstream from the BMS-POP to the BMSs. The aggregate throughput at the BMS POP is typically a large fraction of the total throughput on the backbone mesh. A tree topology is well suited to this traffic flow, in contrast to a general inter-connected mesh network. The BMS-POP is at the root of the tree, as detailed above. Each BMS in the tree structure communicates with its (one) parent BMS and its (one or more) child BMSs.

[0061] FIG. 5 depicts an example embodiment of a method 500 for transmitting and receiving on a mesh network. At 510, a root BMS, or level 0 BMS, sets a superframe. An example superframe is illustrated below with respect to FIG. 6. The superframe duration becomes the superframe duration for the entire mesh. The level 0 BMS begins the superframe transmission with a beacon followed by a coordination message, detailed further below.

[0062] At 520, each parent in the mesh transmits a beacon, as well as transmit and receive schedules for their children. Each parent announces its transmission and reception schedule in its beacon. Each child is assigned a time to receive transmissions from its parent and a time to transmit to the parent. This includes the Level 0 BMS, as

well as all BMSs along the tree structure of the mesh. At 530, each child receives from its parent and/or transmits to its parent according to the schedule transmitted. The process may then return to 520 and repeat indefinitely using the superframe set up at 510. (The root BMS may modify the superframe, as in 510, as necessary.)

[0063] FIG. 6 illustrates an example embodiment of a superframe 610. Superframe 610 comprises a beacon 620, which may comprise various signals and/or messages commonly transmitted in WLAN beacons, well known in the art. A coordination message 630 is also transmitted in superframe 610. In this example, coordination message 630 is transmitted directly subsequent to beacon 620. In alternate embodiments, the coordination message may be transmitted in alternate locations within the superframe. The coordination message may also be included as part of beacon 620, which may include various other signals and/or messages. The coordination message, examples of which are detailed further below, may include system information, the superframe period and start time, a contention period (if applicable), an interference list (described below), and transmit and receive schedules for communication between the BMS and its children. An allocation of superframe 610 for reception from children, identified as RX 640, as well as an allocation for transmission to children, identified as TX 650, are illustrated in FIG. 6. In general, when the allocations within a superframe are identified within a coordination message 630, that coordination message should be transmitted prior to those allocations. In an alternate embodiment, a coordination message may indicate parameters for a subsequent superframe, and the corresponding allocation within the current superframe may have been announced in a coordination message in the previous superframe. As described above, note that a contention period 660 may be allocated within the superframe.

[0064] In one embodiment, the Level 0 BMS schedules its transmissions to/from the Level 1 BMSs and announces the schedule in its coordination message. The Level 0 BMS needs to make sure that this schedule occupies only a fraction of the superframe, the fraction selected to allow for transmission throughout the rest of the mesh. In one embodiment, the permitted fraction is a value between 0.4 and 0.9 and is a function of the number of levels of hierarchy in the mesh. Various scheduling techniques may be used for determining the amount of transmission and/or reception within the superframe at any given level. Example scheduling techniques are illustrated further below.

[0065] A small portion of the superframe may be set aside by the Level 0 BMS as a contention period. The length of the superframe and the position of the contention

period is indicated in each coordination message. During the contention period, in the example embodiment, newly arriving BMSs use the distributed coordination function (DCF) procedures of 802.11 to communicate with an existing BMS in the mesh with a request to join the mesh. Various messages may be used for association with a BMS, and for setting up parent/child relationships within the mesh. Examples are detailed further below. Before the new BMS can make a request in one embodiment, it should listen to the beacons for several superframes and choose the BMS at the highest level of hierarchy (lowest numerical level) whose coordination message it can receive. If the request is accepted by a Level N BMS, the new BMS will become a Level N+1 BMS and follow the procedures described herein.

[0066] The superframe illustrated in FIG. 6 shows beacon 620 at the beginning of superframe. This is consistent a level 0 superframe, in the example embodiment. The superframe defined by the level 0 BMS will be used by all parent BMSs within the mesh. However, the beacon for each respective parent may be in different locations within the superframe, in accordance with allocations for transmission for that BMS, as illustrated further below.

[0067] FIG. 7 illustrates an example embodiment of a method 700 for adding a BMS within the mesh network tree. At 710, a device, either a user terminal, access point, or other device, (generally referred to as a BMS), desiring to become a part of the mesh network, searches for a BMS within the mesh network. As described above, it may listen to beacons for a period of time to determine the BMS at the highest level of hierarchy from which a coordination message is received. Once identified, the device may access that located potential parent BMS, at 720. For example, the contention period may be used for attempting access. The BMS may send a message requesting addition to the tree to the potential parent BMS. This request may be allowed or denied. If the request is allowed, the BMS will become a child of the parent BMS and transmit and receive in accordance with the coordination message from the parent BMS (and may perform other functions, as detailed below). If the access is denied, the BMS may attempt access with another BMS within a mesh (e.g. a node further down the tree) and attempt to establish a connection therewith (details not shown). At 730, a device may transmit a message identifying the new BMS (or may transmit the information in another message). While this step is optional, it may be used to allow the mesh to expand. When a device associates with a BMS, and broadcasts the ID of the new BMS, there may be one or more adjacent stations (perhaps forming another mesh), nearby,

unaware of the new BMS (or its mesh). The adjacent stations may then join the mesh, perhaps as children of the newly joining device.

[0068] Each BMS attempts to receive all available beacons being transmitted in the mesh. If a BMS is able to hear the beacon of another BMS, then the two BMSs are tier 1 interferers (whether or not they have a parent-child relationship). Each BMS maintains a tier 1 interference list (T1IL). The T1IL is constructed locally based on received beacons. Scheduling is determined in accordance with determined constraints. Some constraints are determined with interference lists. There are tier 1 and tier 2 interference list constraints, described below.

[0069] Tier 1 interference list constraints may be generated as follows. An interference list may contain a list of identifiers, each identifier associated with a remote device such as a BMS. An example identifier for inclusion in an interference list is a device ID, which may be transmitted by a device in a beacon (or other message or signal). When any BMS in the T1IL of BMS X transmits, it interferes with reception at BMS X. Therefore, scheduled transmissions at a BMS in the T1IL become constraints on the receive schedule at BMS X. Similarly, when BMS X transmits, it may interfere with reception at each BMS in its T1IL. Therefore, scheduled transmissions to a BMS in the T1IL become constraints on the transmit schedule at BMS X.

[0070] Tier 2 interference list constraints may be generated as follows. Each BMS includes its T1IL in its beacon. The receiving BMS includes the transmitting BMS in its T1IL and stores the received T1IL from that BMS as a tier 2 interference list (T2IL) associated with the transmitting BMS. Since each BMS is only responsible for establishing its transmit and receive schedule to its children, only the T2IL of its children is needed. (However, the T2IL for the other BMSs in the T1IL may be useful to provide fast recovery in case a non-child BMS on the T1IL becomes an orphan and wants to become a child to reestablish connectivity.)

[0071] FIG. 8 depicts an example embodiment of a method 800 for generating interference lists. As detailed further below, interference lists comprise identifiers of various other transmitting stations such as BMSs (including access points, user terminals, and other devices), which provide interference to a particular BMS. Two types of interference lists are described in the example embodiment, a Tier 1 interference list for identifying interferers of a BMS, and Tier 2 interference lists comprising tier 1 interference lists from one or more BMSs included in its tier 1

interference list. These interference lists, as described below, may be used for scheduling to reduce interference and increase throughput in the network mesh.

[0072] At 810, a BMS listens for all BMS beacons. At 820, a list of all interferers is generated, consisting of BMSs for which a beacon is received. The T1IL will be transmitted to neighboring stations, including the parent BMS, if any. For example, the T1IL may be included in a coordination message (e.g. coordination message 630, detailed further below with respect to FIG. 12). At 830, the BMS receives a tier 1 interference list from each beacon in its tier 1 interference list (or the local T1IL). At 840, form the T2IL as the T1IL received from each child BMS in the local T1IL. Note that the T2IL may alternately be formed containing the T1IL from each BMS in the local T1IL (not limited to those from child BMSs). In this alternative, although the BMS need not schedule for non-children, maintaining the T2IL may be used for transitioning an interferer to a child, should the parent of that interferer become unavailable for any reason.

[0073] FIG. 9 illustrates a portion of an example embodiment of a BMS 400. The various functional blocks shown are illustrative only and may be deployed, as described above, as discrete logic, functions or processes of a processor (e.g. processor 420), in association with storage, such as memory 430, or any combination thereof. In this example, signals are received from a transceiver, such as transceiver 440, and beacon detection block 910 may be used to identify local BMSs from which a beacon may be detected. These BMSs form the tier 1 interference list 950 as shown. Beacon detection block 910 may include any device for detecting a beacon signal (e.g., measuring received signal strength, searching, decoding, deinterleaving, filtering, or any other signal processing techniques, well known in the art). The beacon detection will identify the existence of an interferer, and the interferer's identity. Various parameters may be defined for determining the existence of an interferer. For example, a minimum received signal strength may be required to categorize a BMS as an interferer. Note that, in alternate embodiments, a BMS identifier need not be contained in a beacon signal. The BMS may be identified using any variety of techniques (e.g. characteristics of the received signal, such as a code or offset used for or applied to a pilot signal) or included in a signal or message other than the beacon. A beacon detection block 910, in such an alternate embodiment, will decode the respective signal or message, or determine the identifying characteristic of the signal for identifying the interfering BMS, and the identity may be stored in the T1IL 950.

- [0074] Message decoder 920 is used to decode messages received from a remote station, such as a BMS. An example message (such as a coordination message) may include the tier 1 interference list. In the example embodiment, the coordination message and the beacons from these BMSs comprise the tier 1 interference list for those BMSs, and are decoded in message decoder 920 for inclusion in the tier 2 interference list 960, as shown. Message decoder 920 may also receive information within the coordination message from the parent BMS, if any. That information may be used in scheduler 930 to determine transmission and reception schedules, illustrated further below.
- [0075] A message decoder may be deployed using any combination of discrete logic or processing functions, as described above. A message decoder may receive a signal from a receiver or transceiver, and that signal may have been processed using any number of signal processing techniques, such as RF downconversion, amplification, filtering, analog-to-digital conversion, error correction, decoding, deinterleaving, and the like. The message decoder may also include one or more of those functions. The message decoder, examples of which are well known in the art, may be used to extract various fields of information from messages carried on one or more channels (which channels may be of differing formats). The data extracted from the fields may be made available for further processing (by a processor, as described, or by other logic). As described above, the T2IL 960 may comprise the T1IL associated with each child (and need not include those associated with non-child interferers).
- [0076] Scheduler 930 also may access the tier 1 and tier 2 interference lists as shown, for use in scheduling. Scheduler 930 is shown connected to message generator 940, which generates messages for transmission, such as on transceiver 440. A message generator 940 may be deployed using any combination of discrete logic or processing functions, as described above. The message formed may be included, with or without other messages, in one or more channels for transmission according to any transmission scheme, examples of which are detailed above, and of which numerous examples are well known in the art. The message may be an aggregation of various fields of data, examples of which are detailed below (an example of which is a coordination message 630).
- [0077] Various other messages are known in the art, and may be used within the scope of the present invention. For example, messages for requesting and admitting or denying association with a BMS may be deployed. Acknowledgement (ACK or NAK)

messages may be included in various communication formats. A message for identifying a newly detected potential mesh node may be transmitted, as described above. Various messages for establishing or providing status about the mesh, or traffic on the mesh may be formed. For example, a message may include information such as the number of hops to a known level 0 BMS through an identified BMS may be used for maintaining the mesh. The mesh may be adapted as mesh nodes are added or subtracted. Those of skill in the art will readily adapt the principles disclosed herein to any type of communication format and associated messaging format. The tier 1 and tier 2 interference lists are available to message generator 940. (The transmission of the tier 2 interference list may optionally be deployed in certain embodiments.)

[0078] Note that a leaf node need not comprise a scheduler, nor does it need to listen for an interference list from neighbors. A device configured as a leaf may simply monitor beacons to form an interference list, decode messages from the parent BMS to identify the receive and transmit allocations, and other functions as will be apparent to one of skill in the art in light of the teaching herein.

[0079] FIG. 10 illustrates an example embodiment of a method 1000 for forming transmission and reception schedules at a BMS. At 1010, BMS beacons are received. In this context, the term beacon includes the transmission of a coordination message, including interference list and scheduling information. At 1020, the BMS monitors schedule transmissions in the BMS beacons. Each beacon received at a BMS provides information about scheduled transmissions. The scheduling of a transmission or reception may be referred to as an allocation. An allocation may be indicated using any technique, examples of which are well known in the art. For example, an allocation may include a time during which a shared medium is to be accessed, and one or more addresses for identifying one or more transmitting and/or receiving stations (and/or source or destination addresses). The time may be identified by a start time and duration. Alternately, fixed durations may be previously identified, and a time indication may simply identify the duration (e.g. numbered time slots). Those of skill in the art will recognize myriad ways to indicate an allocation. Allocations may be included in a coordination message, examples of which are detailed below, which may be transmitted along with (or alternately as part of) a beacon, or may be transmitted at any other time. A BMS schedules allocations (either receive, transmit, or both) for its children. The allocation may be included in the coordination message.

[0080] When a BMS receives beacons (or, alternately, a coordination message) from a remote BMS (or other station), allocations contained therein may be referred to as remote scheduled allocations. A remote scheduled allocation, in conjunction with interference lists, may be used to generate constraints for use in distributed hierarchical scheduling, an example of which is detailed further below. Beacons received prior to the transmission of the BMSs beacon impose scheduling constraints on the BMS. In determining its transmissions to/from its children the BMS respects constraints from beacons received prior to the transmission of its beacon. At 1030, form constraints in response to receive scheduled transmissions.

[0081] At 1040, construct a transmission and reception schedule in accordance with the constraints. An example scheduling method is illustrated below with respect to FIG. 11. Once this scheduling has been performed, the appropriate coordination message may be transmitted to facilitate transmission and reception by the children of the BMS, as well as to provide information for local BMSs to avoid interfering with the BMS and its children. Except when it is transmitting, a BMS may continue to listen for other beacons to update its T1IL, as well as the T2IL associated with child T1IL BMSs. For efficient utilization, efficient reuse and reduced processing complexity at each BMS, it may be convenient to assume that the transmit and receive schedule announced in the beacon applies to a subsequent superframe (not the current superframe). However, the procedures defined here are also applicable to the case where the beacons schedule the current superframe. In summary, in this example, at each level of the hierarchy, the transmission of the beacon for a Level N BMS is scheduled by the parent (Level N-1 BMS). All beacons received by the Level N BMS allow the BMS to construct its T1IL and to associate a T2IL for each BMS in its T1IL (in particular, for each child). An example of T1IL and T2ILs at an arbitrary BMS is provided in Table 1 below. These lists are used to determine scheduling constraints at BMS k. Notice that the T1IL at BMS k consists of its children (m, n), its parent (p) as well as other BMSs (q) that do not have a parent/child relationship with k. BMS k stores the T2IL associated with BMS m which is the T1IL of BMS m reported in its beacon. All beacons received by the Level N BMS prior to the transmission of its beacon provide constraints to its transmit and receive schedule. Based on these constraints the Level N BMS determines its transmit and receive schedule to/from its (Level N+1 BMS) children and announces it in its beacon, which, in turn, provides the constraints on the schedule of the subsequent BMSs.

Table 1. Example Tier 1 and Tier 2 Interference List

Tier 1 Interference List (T1IL) at BMS k	Tier 2 Interference Lists (T2ILs) at BMS k
BMS m (child)	T2IL = {k, q}. This is the T1IL of BMS m
BMS n (child)	T2IL = {k, p, r, s}. This is the T1IL of BMS n
BMS p (parent)	T2IL = {k, g, n}. This is the T1IL of BMS p
BMS q	T2IL = {k, m}. This is the T1IL of BMS q

[0082] FIG. 11 illustrates an example embodiment of a method for forming constraints in response to received remote scheduled allocations, such as may be used at 1030, described above. At decision block 1110, if there is a remote scheduled transmission to a BMS in the tier 1 interference list of the scheduling BMS (which may be determined by monitoring the remote scheduled allocations, as described in 1020, for example), proceed to 1120, otherwise proceed to 1130. At 1120, place a constraint on the transmit schedule of the scheduling BMS. In other words, the scheduling BMS should not schedule a transmission from the scheduling BMS during such a remote scheduled transmission in order to avoid causing interference at the receiver of that remote scheduled transmission.

[0083] At 1130, if there is a remote scheduled transmission from a terminal in the scheduling BMS's tier 1 interference list, proceed to 1140, otherwise proceed to 1150. At 1140, place a constraint on the receive schedule of the scheduling BMS. In other words, the scheduling BMS should not schedule a reception (e.g. a transmission from one of its children to itself) during such a remote scheduled reception, since the reception of the transmission would interfere with remote scheduled reception from the terminal in its tier 1 interference list.

[0084] At 1150, determine if there is a scheduled transmission to a terminal in the tier 2 interference list of a child. If so, proceed to 1160, otherwise proceed to 1170. At 1160, place a constraint on a transmit schedule from that child. In other words, the tier 2 interference list associated with that child indicates interferers or potential interferers with that child. If a transmission is scheduled to one of those potential interferers, that child should not transmit, to avoid interference.

[0085] At 1170, determine if there is a scheduled reception from a terminal in the tier 2 interference list associated with a child. If so, proceed to 1180, otherwise the process may stop. At 1180, place a constraint on the transmission schedule to that child. In other words, if a terminal in the tier 2 interference list associated with a child (e.g. in a tier 1 interference list of the child) is to perform a transmission, a transmission should not be scheduled to that child, because it may not be received at the child without interference from the scheduled transmission from the other terminal.

[0086] Constraints may be introduced by remote scheduled allocations, as well as by allocations scheduled by the scheduling BMS. For example, an iterative process may be performed by the scheduling BMS, such that once an allocation is made to one child, one or more constraints may be introduced by that allocation (even if the allocation has not yet been transmitted). The scheduling BMS may perform the scheduling in any manner desired. For example, multiple schedule hypotheses may be generated. A schedule hypothesis may be selected based on increasing throughput, meeting service guarantees, or a combination of both. Various heuristic techniques may be deployed (for complicated mesh scheduling), or brute force techniques may be deployed (when the interference lists are relatively simple). The method 1030 (or any of the various equivalent methods, obvious to one of ordinary skill in the art, in light of the teaching herein), illustrated in FIG. 11, may be iterated repeatedly until the allocations for a superframe are generated. In other words, each allocation may be made in accordance with the existing constraints, then, if the allocation is successful within the constraints, that allocation becomes a scheduled transmission to determine additional constraints, if any.

[0087] A variety of beacons, coordination methods, and combinations thereof may be used. An example coordination message 630 is illustrated in FIG. 12. Various alternate embodiments will be apparent to one of ordinary skill in the art, in light of the teaching herein. In the example embodiment, the beacon and coordination message may be combined. The beacon, coordination message, or combination may include various other fields (not shown) or may omit various example fields detailed herein. The example coordination message transmitted by each Level N BMS contains the following information: a system identification (identifying the POP) 1210; a superframe period 1220 and start time 1230 (determined by the Level 0 BMS for the mesh); a contention period field 1240 for identifying the location of a contention period, if any; a Tier 1 Interference List (T1IL) 1250 including all BMSs whose beacons have been

identified (e.g. decoded over the last predetermined time period, for example); a transmit schedule 1260 for transmissions from the BMS to the parent level (N-1) BMS and the child level (N+1) BMSs; and a receive schedule 1270 for transmissions to the parent level (N-1) BMS and the child level (N+1) BMSs. Tx 1260 and Rx 1270 are used to indicate scheduled allocations. Note that, as described above, the scheduled allocation from a scheduling BMS directs the traffic flow between itself and its children, and may also be used by other scheduling BMSs as a remote scheduled allocation in the scheduling by those other BMSs. Any type of allocation indication may be used in coordination message 630.

[0088] In this example, the system identification 1210, superframe period and start time 1220-30 and contention period 1240 propagate down the hierarchy (unchanged) throughout the mesh.

[0089] Occasionally, especially at lower levels of the hierarchy (higher values of N in a Level N BMS), conflicts may arise, e.g. interfering transmissions get scheduled at overlapping periods. In the example embodiment, because the total bandwidth requirements at the lower levels of the mesh hierarchy are relatively low (little or no aggregation of traffic from other APs), this is not a significant problem. When a BMS detects a collision during a scheduled reception from a child, it may modify its receive schedule for subsequent superframes. Otherwise, when a BMS detects a collision during a scheduled reception from its parent, it either notifies the parent about the collision or, if the child is unable to decode the beacon, it does not transmit its scheduled beacon. In either case, the parent modifies the schedule for subsequent superframes.

[0090] The various embodiments described above may be illustrated with an extended example, using a simplified mesh topology as shown in FIG. 13. In this example, a mesh comprising seven BMSs on three levels is illustrated. BMS A 1310 is the level 0 BMS and has as children three level 1 BMSs, BMS B-D, 1320-40, respectively. Level 1 BMS B has one level 2 child, BMS G 1370. Level 1 BMS D 1340 has two children, BMS E-F, 1350-60, respectively.

[0091] An example superframe structure and example procedures for the mesh, illustrated in FIG. 13, are illustrated in FIG. 14. Superframe 610 is illustrated comprising a superframe 1402 and a receive superframe 1404. Transmission and reception during a superframe are then illustrated in corresponding brackets for the seven BMSs A-G, as shown. A key defines different shading patterns to illustrate each

beacon, transmit allocations, receive allocations, constraints and the contention period. The constraints and allocations shown in FIG. 14 are illustrative only. Table 2 illustrates the tier 1 interference lists for the various BMSs shown in FIG. 13.

Table 2. Example T1ILs

BMS	Level	T1IL	Parent
A	0	B, C, D	----
B	1	A, C, D, G	A
C	1	A, B, D, F	A
D	1	A, B, C, E, F	A
E	2	D, G	D
F	2	C, D	D
G	2	B, E	B

[0092] The superframe in FIG. 14 begins with the Level 0 BMS (BMS A) transmitting its beacon 1406. The beacon defines the length of the superframe, the position of the BMS A transmit 1408 and receive 1410 periods to each Level 1 BMS: B, C and D, and the position of the contention period 1412.

[0093] Level 1 BMS B, C and D hear the BMS A beacon 1406 which informs them of their schedule for transmit and receive to BMS A. BMS B receives at 1418 and transmits to BMS A at 1420. BMS C receives at 1416 and transmits to BMS A at 1422. BMS D receives at 1414 and transmits to BMS A at 1424. On receiving the beacon 1406, BMS B, C, D include the constraints in their transmit and receive schedule as shown in FIG. 14 (discussed further below). Note that according to the T1ILs, BMS E, F, G do not receive the beacon of BMS A.

[0094] BMS C has no child, so its schedule consists of communication with BMS A, which has been defined in the beacon of BMS A. So, BMS C requires no further processing. BMS C transmits its beacon 1426, which contains only its transmit and receive periods with its parent BMS A (1422 and 1416).

[0095] The beacon from BMS A 1406 defines transmissions from A to B, C, D and from B, C, D to A, which become receive schedule constraints for BMS D. These constraints are shown on the receive superframe of BMS D (1414 is the transmission from A to D, A TX 1428 identifies the constraint due to the transmission from A to B

and C, C TX 1430 is the constraint due to C transmitting to A, and B TX 1432 is the constraint due to B transmitting to A). Similarly, receptions scheduled at A, B, C and included in the beacon from BMS A 1406 are shown as constraints on the transmit superframe of BMS D (1434 is the C RX constraint due to C receiving from A, 1436 is the B RX constraint due to B receiving from A, 1438 is the A RX constraint due to A receiving from C, and 1440 is the A RX constraint due to A receiving from B). BMS D needs to design its transmit and receive schedule to its children, BMS E and F, respecting these constraints. The transmit 1446 (D to E and F) and receive 1444 (E and F to D) schedule of BMS D to its children E and F is as shown, and is transmitted with beacon 1442.

[0096] BMS B is scheduled to transmit its beacon 1456 next. The transmit and receive periods announced by BMS D in its beacon provide further constraints to the schedule at BMS B. Thus the transmit period at BMS D (to BMS E and BMS F) becomes a receive schedule constraint 1470 at BMS B while the receive period at BMS D (from BMS E and BMS F) becomes a transmit constraint 1458 on BMS B. Notice that BMS D is not the parent of BMS B but its schedule constrains the schedule at BMS B because BMS D gets to transmit its beacon ahead of BMS B. These are additional constraints to those introduced due to the beacon from A (D RX 1448 due to D receiving from A, C RX 1450 due to C receiving from A, A RX 1452 due to A receiving from C, A RX 1454 due to A receiving from D, B transmitting to A 1420, A TX 1462 due to A transmitting to C and D, B receiving from A 1418, C TX 1464 due to C transmitting to A, and D TX 1466 due to D transmitting to A).

[0097] Moreover, BMS B is the parent of G and needs to schedule transmit and receive periods to G. From previously received beacons from G, B is aware that E is on the T1IL of G. This means that the transmit 1474 and receive periods 1475 from E (scheduled by E's parent D), respectively constrain the receive and transmit schedules from G to B. These constraints are shown on the superframe of BMS G in FIG. 14. Therefore B schedules its transmit period 1460 to and receive period 1468 from G to respect these constraints.

[0098] BMS E is a child in this example, and thus does not need to schedule. If that were to change, it would have a D RX constraint 1471 due to D receiving from A, and a D TX constraint 1472 due to D transmitting to A. BMS E transmits its beacon 1473 as shown. As scheduled, BMS E receives at 1475 and transmits at 1474.

[0099] BMS F is also a child in this example, and thus does not need to schedule. If that were to change, BMS F would have a D RX constraint 1476 due to D receiving from A, a C RX constraint 1477 due to C receiving from A, a C TX transmit constraint 1480 due to C transmitting to A, and a D TX constraint 1481 due to D transmitting to A. BMS F transmits its beacon 1478 as shown. As scheduled, BMS F receives at 1482 and transmits at 1479.

[00100] BMS G is also a child in this example, and thus does not need to schedule. If that were to change, BMS G would have a B RX constraint 1484 due to B receiving from A, an E RX constraint 1490 due to E receiving from D, a B TX transmit constraint 1492 due to B transmitting to A, and an E TX constraint 1494 due to E transmitting to D. BMS G transmits its beacon 1486 as shown. As scheduled, BMS G receives at 1496 and transmits at 1488.

[00101] Notice that through this distributed scheduling, overlapping transmissions have been scheduled: (G to B 1488 overlapping with E to D 1474) and (D to F 1446 overlapping with B to G 1460). Recall the interference lists from Table 2. Transmitter G is not in the T1IL for receiver D. Transmitter E is not on the T1IL of receiver B. Therefore these overlapping transmissions do not cause mutual interference. Similarly, transmitter D is not in the T1IL for receiver G. Transmitter B is not on the T1IL of receiver F. Therefore these overlapping transmissions do not cause mutual interference.

[00102] It is also apparent from FIG. 14 that the scheduling and reuse efficiency of the hierarchical scheduling procedure becomes even more significant when the mesh has further levels of hierarchy. In the example, notice the sparse transmit and receive constraints shown on the superframe schedules of BMSs E, F and G. Thus, if each of these BMS were itself a parent, it would have flexibility in scheduling the transmit and receive periods for its children without causing collisions at already scheduled transmissions. If the scheduling is for the same superframe, then that schedule needs to be accommodated in the part of the superframe that follows its beacon and prior to the contention period. However, if the scheduling is always done for subsequent superframes, then transmit and receive periods can be scheduled in the early part of the superframe (prior to the transmission of the beacon by the parent). This permits maximum flexibility and maximum efficiency.

[00103] Another case of interest is a network where the BMSs do not include the T1IL in their beacons. Therefore the parent cannot construct a T2IL for its children. Then, in the example of Figure 13, BMS B is not aware that the scheduled transmit and receive

periods at E place constraints on, respectively, the transmit and receive schedules at G. In such cases, it is possible that the parent (B) may attempt to be aggressive and schedule a transmission to/from its child (G) that results in collisions: e.g. overlapping periods with reception at E and transmission from G or vice versa. Such collisions may reduce efficiency. However, when such a collision is detected, the parent modifies the schedule for the subsequent superframe. By trial and error, a similar level of efficiency can be achieved. However, in some circumstances, inclusion of the T1IL in the beacon and the procedures described using the T2IL may yield a more effective scheme.

[00104] Another case of interest is when the BMS includes its T2IL in the beacon or coordination message. For example, consider the case that BMS D includes its T2IL in its coordination message (received by BMS A), which indicates that BMS G is "adjacent" to BMS E and BMS C is "adjacent" to BMS F. Fast recovery is enabled as follows. If BMS D goes down for any reason, BMS A does not receive its beacon. At this point, BMS A can immediately reorganize the tree by assigning BMS E to be a child of BMS G and BMS F to be child of BMS C. Thus inclusion of the T2IL in the coordination message enables fast recovery in the face of BMS failure.

[00105] Depending on the frequency of network topology changes and traffic changes, it is possible to reduce the overhead of coordination message transmissions. For example, in relatively stable networks, the coordination message need not be transmitted in every beacon period. The schedule may be left unchanged for a number of beacon periods. In highly dynamic networks where the BMS are not in fixed locations, there may be frequent changes in topology. In this case the coordination message may be transmitted in every beacon period and the tree topology and the schedule are updated frequently. In various alternate embodiments, coordination messages may be transmitted more or less frequently, and/or the frequency may be dynamically updated to adapt to changes in the network environment.

[00106] In summary, aspects detailed above yield various benefits, examples of which follow. An ad hoc mesh network may be organized into a tree topology that matches the traffic flow for the case of a wireless backhaul network. Distributed, hierarchical scheduling is provided where each parent schedules communication with its children while respecting already scheduled transmissions to/from its interferers and to/from interferers of its children. Efficient flow and procedures needed to construct interference constraints to accomplish the distributed, hierarchical scheduling are

described. The result is very efficient scheduling and tight reuse in an ad hoc wireless network without the use of centralized scheduling.

[00107] Those of skill in the art would understand that information and signals may be represented using any of a variety of different technologies and techniques. For example, data, instructions, commands, information, signals, bits, symbols, and chips that may be referenced throughout the above description may be represented by voltages, currents, electromagnetic waves, magnetic fields or particles, optical fields or particles, or any combination thereof.

[00108] Those of skill would further appreciate that the various illustrative logical blocks, modules, circuits, and algorithm steps described in connection with the embodiments disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, circuits, and steps have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled artisans may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present invention.

[00109] The various illustrative logical blocks, modules, and circuits described in connection with the embodiments disclosed herein may be implemented or performed with a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general purpose processor may be a microprocessor, but in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration.

[00110] The steps of a method or algorithm described in connection with the embodiments disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module

may reside in RAM memory, flash memory, ROM memory, EPROM memory, EEPROM memory, registers, hard disk, a removable disk, a CD-ROM, or any other form of storage medium known in the art. An exemplary storage medium is coupled to the processor such the processor can read information from, and write information to, the storage medium. In the alternative, the storage medium may be integral to the processor. The processor and the storage medium may reside in an ASIC. The ASIC may reside in a wireless communication device such as a user terminal, an access point, or mesh network node. In the alternative, the processor and the storage medium may reside as discrete components in a wireless communication device.

[00111] Headings are included herein for reference and to aid in locating various sections. These headings are not intended to limit the scope of the concepts described with respect thereto. Such concepts may have applicability throughout the entire specification.

[00112] The previous description of the disclosed embodiments is provided to enable any person skilled in the art to make or use the present invention. Various modifications to these embodiments will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other embodiments without departing from the spirit or scope of the invention. Thus, the present invention is not intended to be limited to the embodiments shown herein but is to be accorded the widest scope consistent with the principles and novel features disclosed herein.

CLAIMS

What is claimed is:

1. An apparatus comprising:
a receiver for receiving one or more signals from one or more remote devices;
a message decoder for decoding one or more first interference lists from coordination messages contained in the one or more received signals from the one or more remote devices; and
a memory for storing the one or more first interference lists and a second interference list comprising identifiers associated with the one or more remote devices.
2. The apparatus of claim 1, wherein the message decoder further decodes one or more remote scheduled allocations from the one or more coordination messages, and further comprising a scheduler for scheduling one or more allocations for one or more remote devices in accordance with the first and second interference lists and the one or more remote scheduled allocations.
3. The apparatus of claim 2, further comprising a message generator for generating a coordination message comprising one or more scheduled allocations.
4. The apparatus of claim 3, further comprising a transmitter for transmitting the coordination message.
5. The apparatus of claim 1, further comprising a detector for detecting a beacon in a received signal for identifying a remote device.
6. The apparatus of claim 2, wherein the receiver receives information from a remote device during a scheduled allocation.
7. The apparatus of claim 4, wherein the transmitter transmits information to a remote device during a scheduled allocation.

8. The apparatus of claim 4, wherein:
the receiver receives a signal from a parent remote device, the signal comprising a coordination message;
the message decoder decodes a transmit allocation, if any;
the message decoder decodes a receive allocation, if any;
the receiver receives information according to the receive allocation; and
the transmitter transmits information according to the transmit allocation.

9. The apparatus of claim 2, wherein the scheduler constrains transmission allocation for the apparatus when a remote scheduled allocation is directed to a remote device identified in the second interference list.

10. The apparatus of claim 2, wherein the scheduler constrains reception allocation for the apparatus when a remote scheduled allocation is directed from a remote device identified in the second interference list.

11. The apparatus of claim 2, wherein the scheduler constrains transmission from a first remote device identified in the second interference list, and for which the apparatus provides scheduling, when a remote scheduled allocation is directed to a second remote device, the second remote device identified in the first interference list associated with the first remote device.

12. The apparatus of claim 2, wherein the scheduler constrains transmission to a first remote device identified in the second interference list, and for which the apparatus provides scheduling, when a remote scheduled allocation is directed from a second remote device, the second remote device identified in the first interference list associated with the first remote device.

13. A method comprising:
detecting signals received from one or more remote stations;
identifying the one or more remote stations;
storing an identifier associated with each of the one or more remote stations in a first interference list; and
transmitting the first interference list for reception by neighbor remote stations.

14. The method of claim 13, further comprising:
receiving one or more remote allocations;
receiving a second interference list from one or more remote stations; and
scheduling allocations of a shared resource in response to the one or more remote allocations, the second interference list, and the first interference list.

15. The method of claim 13, further comprising:
forming a superframe indicated by one or more superframe parameters;
transmitting the superframe parameters.

16. The method of claim 15, further comprising:
receiving one or more superframe parameters;
forming a coordination message comprising the one or more superframe parameters and the first interference list; and
transmitting the coordination message.

17. The method of claim 14, further comprising:
receiving one or more superframe parameters;
forming a coordination message comprising the one or more superframe parameters, the first interference list, and the scheduled allocations; and
transmitting the coordination message.

18. The method of claim 17, further comprising:
receiving a coordination message at a first station from a second station;
decoding an allocation from the coordination message; and
transmitting information from the first station to the second station in accordance with the decoded allocation.

19. The method of claim 17, further comprising:
receiving a coordination message at a first station from a second station;
decoding an allocation from the coordination message; and
receiving information at the first station from the second station in accordance with the decoded allocation.

20. An apparatus comprising:
means for detecting signals received from one or more remote stations;
means for identifying the one or more remote stations;
means for storing an identifier associated with each of the one or more remote stations in a first interference list; and
means for transmitting the first interference list for reception by neighbor remote stations.

21. The apparatus of claim 20, further comprising:
means for receiving one or more remote allocations;
means for receiving a second interference list from one or more remote stations;
and
means for scheduling allocations of a shared resource in response to the one or more remote allocations, the second interference list, and the first interference list.

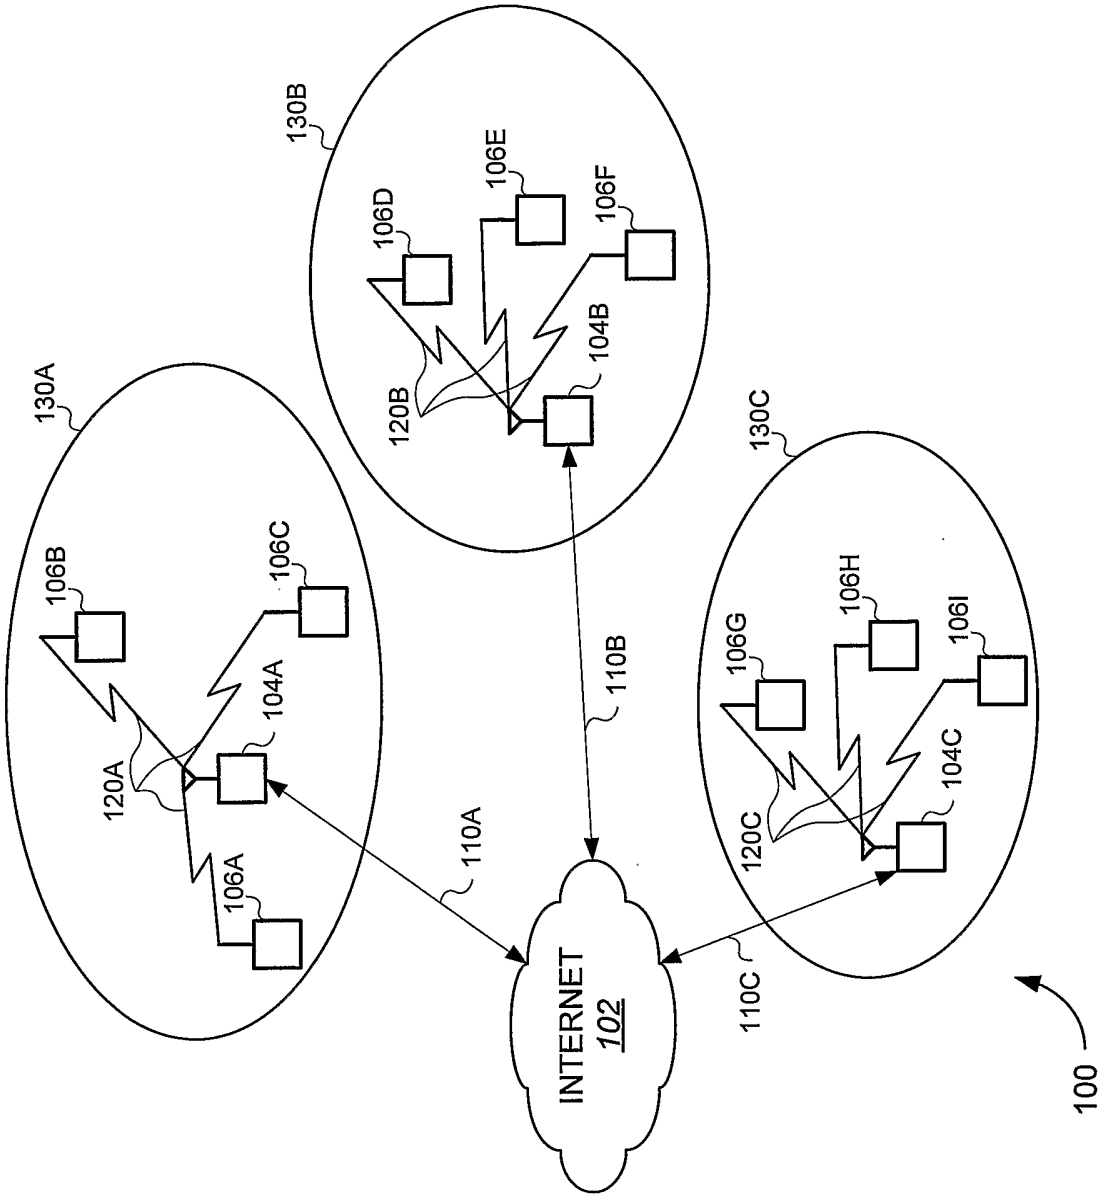
22. An apparatus comprising:
a receiver for receiving one or more signals from one or more remote devices;
a message for decoder decoding one or more first interference lists and one or more remote allocations from coordination messages contained in the one or more received signals from respective one or more remote devices; and
means for allocating a shared resource for receiving and transmitting by one or more remote devices based on the one or more first interference lists and the one or more remote allocations.

23. A message, operable for transmission by a mesh station, comprising:
an interference list field comprising a list of interfering remote stations;
a transmit allocation field comprising one or more allocations for transmission by a child remote station on a shared medium; and
a receive allocation field comprising one or more allocations for receiving by a child remote station on a shared medium.

24. The message of claim 23, further comprising a superframe identification field.

25. The message of claim 23, further comprising a contention period field.
26. The message of claim 23, further comprising a system identification field.
27. A mesh network comprising:
a first mesh station for:
establishing a superframe; and
transmitting a first coordination message, the first coordination message identifying the superframe and comprising one or more allocations; and
one or more second mesh stations for:
receiving the first coordination message; and
transmitting or receiving in accordance with the one or more allocations.
28. The mesh network of claim 27, wherein one or more of the second mesh stations further transmits a second coordination message, the second coordination message identifying the superframe, and comprising one or more second allocations, the second allocations determined to avoid conflicting with allocations in the first coordination message.
29. The mesh network of claim 28, further comprising one or more third mesh stations for receiving the second coordination message and for transmitting or receiving in accordance with one or more allocations contained in the second coordination message.
30. Computer readable media operable to perform:
detecting signals received from one or more remote stations;
identifying the one or more remote stations;
storing an identifier associated with each of the one or more remote stations in a first interference list; and
transmitting the first interference list for reception by neighbor remote stations.
31. The media of claim 30, further operable to perform:
receiving one or more remote allocations;

receiving a second interference list from one or more remote stations; and
scheduling allocations of a shared resource in response to the one or more
remote allocations, the second interference list, and the first interference list.



PRIOR ART
FIG. 1




FIG. 2




FIG. 3


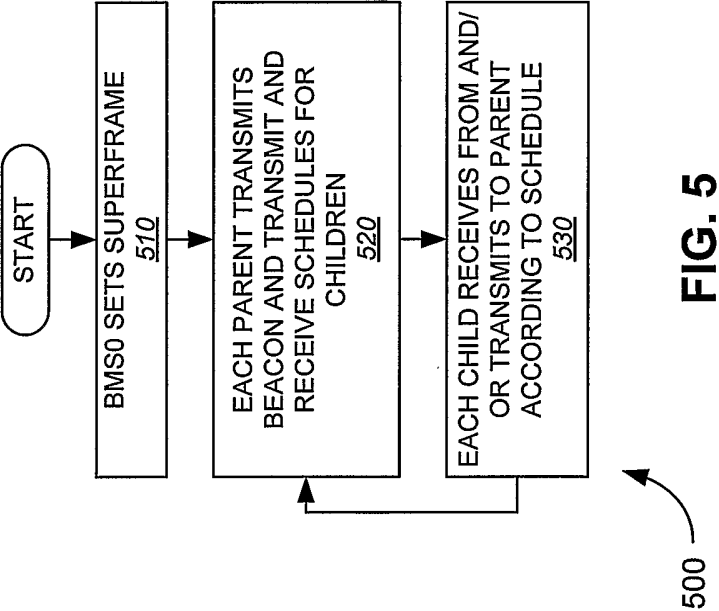


FIG. 4







FIG. 6

7/14



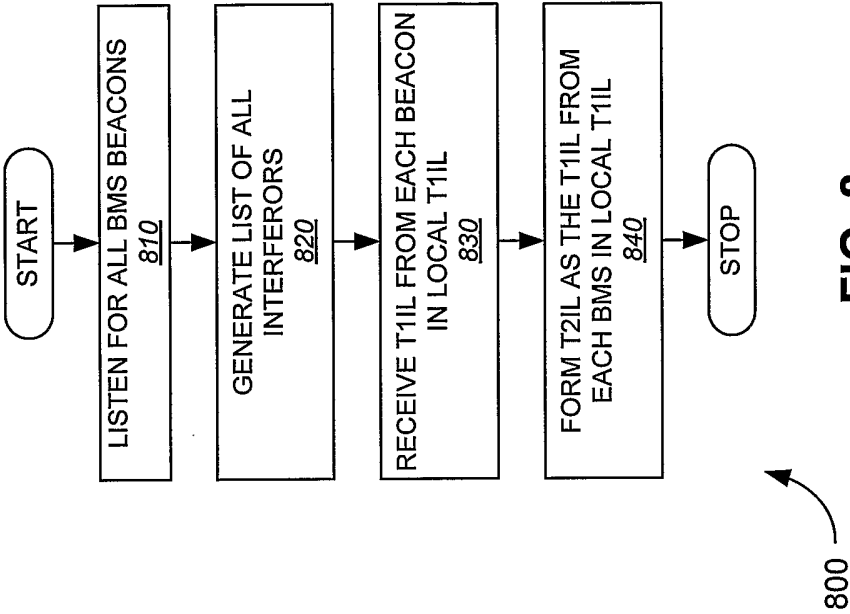


FIG. 8




FIG. 9

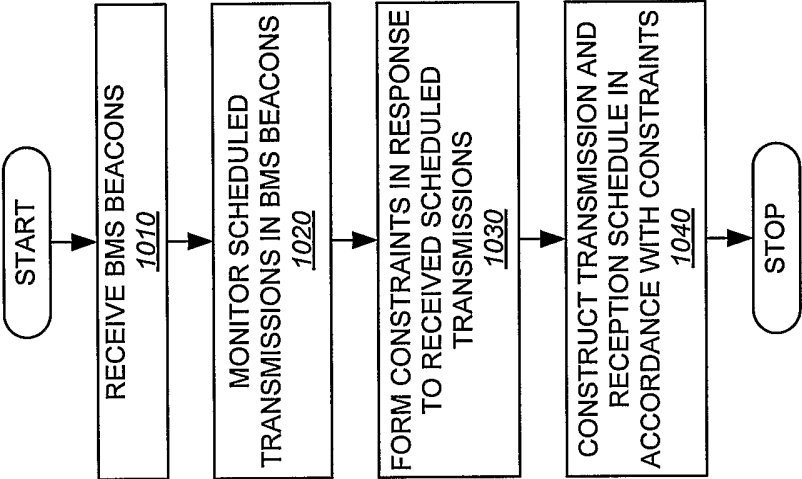


FIG. 10

1000

11/14

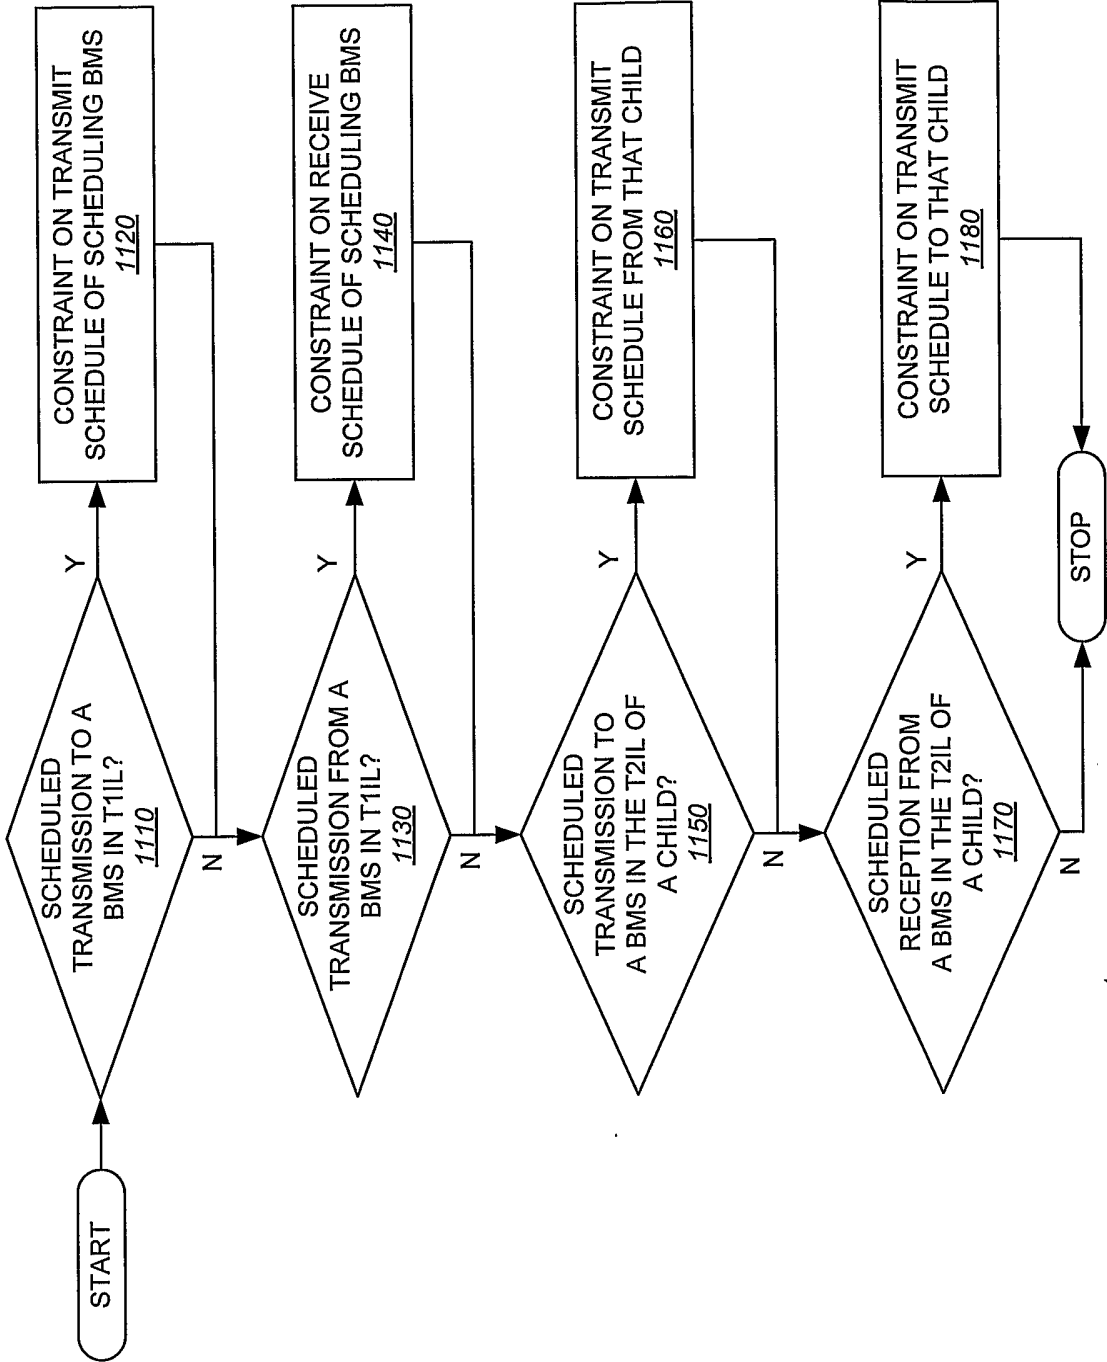


FIG. 11

1030

SYSTEM IDENTIFICATION <u>1210</u>	SUPERFRAME PERIOD <u>1220</u>	SUPERFRAME START TIME <u>1230</u>	CONTENTION PERIOD <u>1240</u>	T1IL <u>1250</u>	TX <u>1260</u>	RX <u>1270</u>
---	-------------------------------------	---	-------------------------------------	---------------------	-------------------	-------------------

FIG. 12

630 ↗

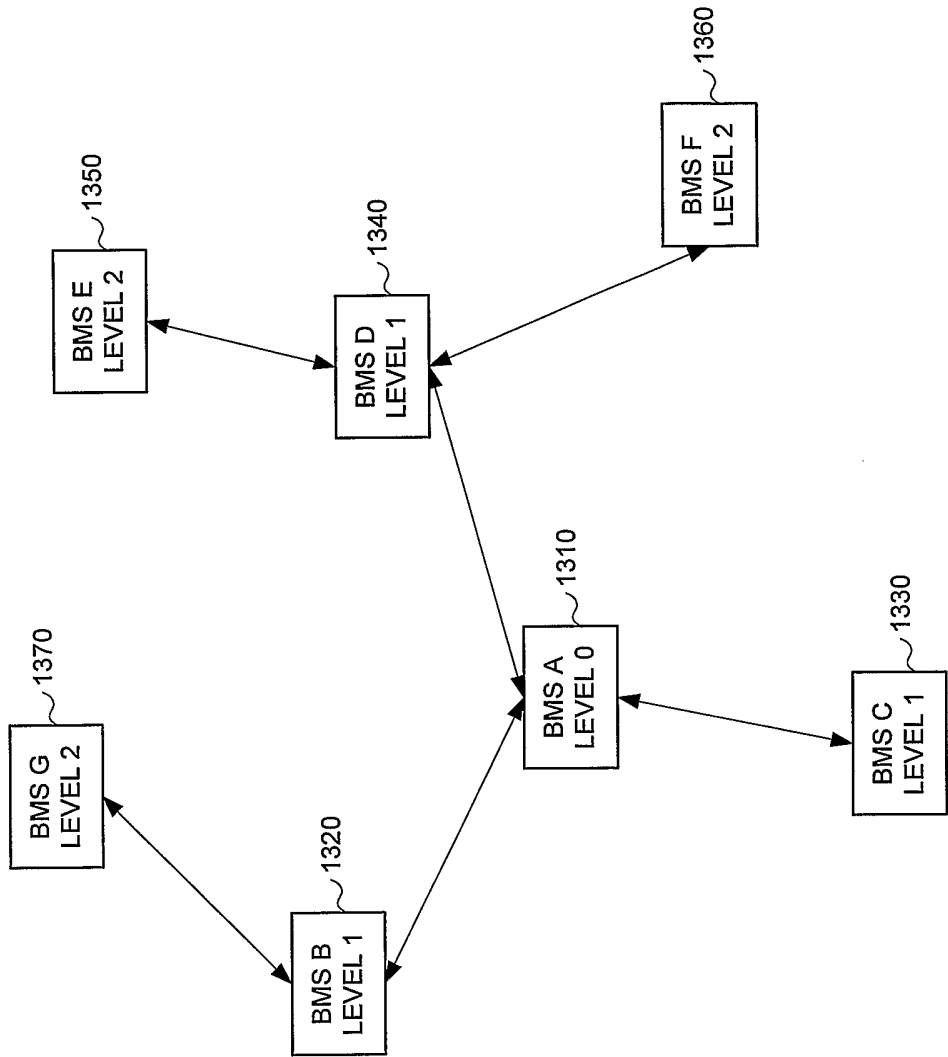


FIG. 13

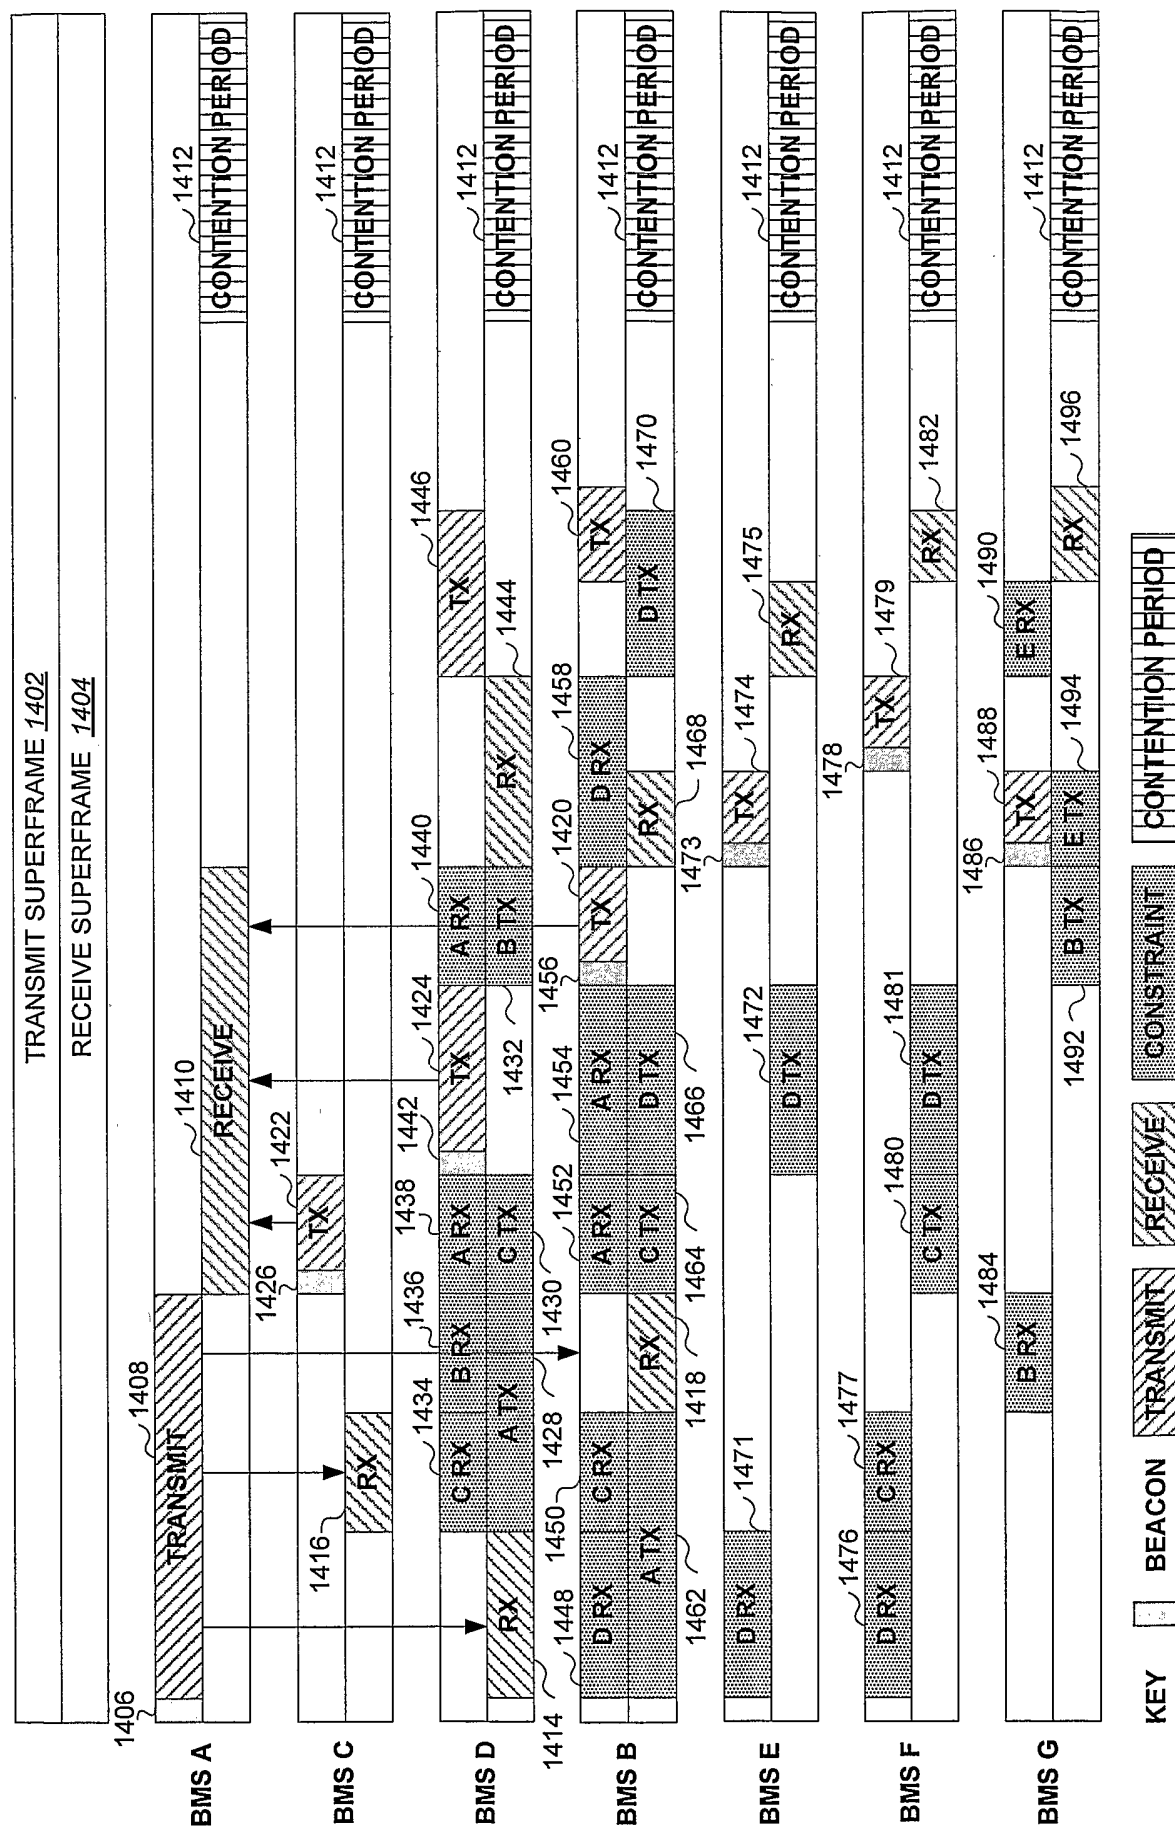


FIG. 14

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2005/003576

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
X	WO 01/28170 A (NOKIA WIRELESS ROUTERS, INC; GARCIA-LUNA-ACEVES, J., JOAQUIN; FULLMER,) 19 April 2001 (2001-04-19) page 8, line 26 - page 11, line 15 page 12, line 25 - page 13, line 4 page 17, line 1 - page 19, line 14 -----	1-31
X	US 2002/067736 A1 (GARCIA-LUNA-ACEVES J.J ET AL) 6 June 2002 (2002-06-06) paragraph '0042! - paragraph '0050! paragraph '0059! - paragraph '0066! -----	1-31
X	US 5 719 868 A (YOUNG ET AL) 17 February 1998 (1998-02-17) column 2, line 15 - column 4, line 56 -----	1-31

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex

* Special categories of cited documents

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- * & * document member of the same patent family

Date of the actual completion of the international search

5 July 2005

Date of mailing of the international search report

28/07/2005

Name and mailing address of the ISA

European Patent Office, P B 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Poort, I

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US2005/003576

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0128170	A	19-04-2001	US 6788702 B1	07-09-2004
			AU 776640 B2	16-09-2004
			AU 2296901 A	23-04-2001
			CA 2390061 A1	19-04-2001
			EP 1249101 A2	16-10-2002
			WO 0128170 A2	19-04-2001
<hr/>				
US 2002067736	A1	06-06-2002	AU 9637801 A	08-04-2002
			WO 0228020 A2	04-04-2002
<hr/>				
US 5719868	A	17-02-1998	NONE	
<hr/>				