



(22) Date de dépôt/Filing Date: 2004/09/09

(41) Mise à la disp. pub./Open to Public Insp.: 2005/06/09

(30) Priorité/Priority: 2003/12/09 (10/731,348) US

(51) Cl.Int.⁷/Int.Cl.⁷ H04L 29/06


(71) Demandeur/Applicant:
AVAYA TECHNOLOGY CORP., US

(72) Inventeur/Inventor:
BENVENISTE, MATHILDE, US

(74) Agent: MACRAE & CO.

(54) Titre : ARCHITECTURE REPARTIE POUR DEPLOYER DE MULTIPLES RESEAUX LOCAUX SANS FIL

(54) Title: DISTRIBUTED ARCHITECTURE FOR DEPLOYING MULTIPLE WIRELESS LOCAL-AREA NETWORKS



(57) **Abrégé/Abstract:**

A novel distributed architecture for deploying a plurality of wireless local-area networks is disclosed. The illustrative embodiment enables some of the protocols of a protocol stack to be "off-loaded" from wireless access points to a remote central controller. In particular, protocols whose services operate correctly regardless of the presence or magnitude of processing and communication delays are off-loaded to the central controller, while protocols whose services operate correctly only when delays are below a threshold remain in the wireless access points. The illustrative embodiment thus enables the use of "thin" (i.e., reduced-functionality) access points, thereby reducing cost and facilitating maintenance when deploying multiple wireless local-area networks (for example, on a corporate or academic campus).



Abstract

A novel distributed architecture for deploying a plurality of wireless local-area networks is disclosed. The illustrative embodiment enables some of the protocols of a protocol stack to be "off-loaded" from wireless access points to a remote central controller. In particular, protocols whose services operate correctly regardless of the presence or magnitude of processing and communication delays are off-loaded to the central controller, while protocols whose services operate correctly only when delays are below a threshold remain in the wireless access points. The illustrative embodiment thus enables the use of "thin" (*i.e.*, reduced-functionality) access points, thereby reducing cost and facilitating maintenance when deploying multiple wireless local-area networks (for example, on a corporate or academic campus).

DISTRIBUTED ARCHITECTURE FOR DEPLOYING MULTIPLE WIRELESS LOCAL-AREA NETWORKS

Field of the Invention

[0001] The present invention relates to telecommunications in general, and, more particularly, to a novel architecture for deploying a plurality of wireless local-area networks.

Background of the Invention

[0002] Figure 1 depicts a schematic diagram of an exemplary wireless local-area network (LAN) 100 in the prior art comprising access point 101 and stations 102-1 through 102- N , wherein N is a positive integer, interconnected as shown. Each station 102- i , wherein i is a member of the set $\{1, 2, \dots, N\}$, is a device such as a notebook computer, personal digital assistant (PDA), tablet PC, *etc.* that wirelessly transmits signals to and receives signals from other stations in local-area network 100 via access point 101.

[0003] In order for access point 101 and stations 102-1 through 102- N to communicate, they must agree as to the meaning of the signals that they transmit. In particular, they must agree on: who talks when, what constitutes a "0" and a "1," how an error is detected and corrected, *etc.* In the terminology of telecommunications, these agreements are called *protocols*, and a particular set of protocols is known as a *protocol stack*.

[0004] Stations and access points typically employ software, hardware, or a combination of software and hardware to implement a particular protocol stack. Figure 2 depicts a conceptual diagram of two software-based protocol stacks 201 and 202- i that are implemented by access point 101 and station 102- i , respectively, in accordance with the prior art, wherein i is a member of set $\{1, 2, \dots, N\}$.

[0005] Telecommunications engineers design protocols to provide a variety of *services*, such as error control (*i.e.*, detecting and/or correcting bit errors in transmitted data), flow control (*i.e.*, compensating for differences between the rate at which data is transmitted to a device and the rate at which the receiving device processes the data), and medium access control (*i.e.*, ensuring that only one device at a time transmits into a *shared communications channel*). Because telecommunications systems typically involve a plurality of communicating devices that operate independently, protocols must be able to provide

desired behavior (*i.e.*, the protocol services must be *correct*) for different, and often unpredictable, timing relationships between the devices.

[0006] Telecommunications engineers formally describe a protocol via a *protocol specification*. A protocol specification comprises: (i) a service to be provided; (ii) one or more assumptions about the environment in which the protocol executes; (iii) a *vocabulary* of messages used to implement the protocol; (iv) an *encoding* of each message in the vocabulary; and (v) a set of *procedure rules* guarding the consistency of message exchanges. [Design and Validation of Computer Protocols, G. J. Holzmann, Prentice Hall, 1991]

Summary of the Invention

[0007] The present invention enables some of the protocols of a protocol stack to be "off-loaded" from an access point to a remote processor. In particular, the illustrative embodiment employs a remote central controller that can serve a plurality of access points, thus enabling a plurality of wireless local-area networks to be constructed with reduced-functionality access points (*i.e.*, access points that implement a proper subset of a protocol stack) and a single remote controller. By enabling the use of "thin" (*i.e.*, reduced-functionality) access points, the illustrative embodiment offers two benefits: first, it reduces the cost of the access points, which is especially advantageous when deploying a large number of wireless local-area networks (for example, on a corporate or academic campus). Second, the illustrative embodiment can facilitate the maintenance of large numbers of wireless local-area networks, since a substantial number of software updates will often require changes only to the central controller.

[0008] The illustrative embodiment is based on the recognition that some protocols operate correctly regardless of timing relationships between events, while other protocols operate correctly only in accordance with one or more particular timing assumptions, such as a *maximum timing delay* (*e.g.*, a *timeout* for waiting to receive a message, an *inter-frame spacing* for medium access control, *etc.*). In other words, the *correctness* of some protocol services is independent of any delays or timing anomalies, while the correctness of other protocol services is contingent on communication and processing delays not exceeding a particular threshold.

[0009] For the purposes of this specification, a protocol service is defined as "**correct**" if it operates in accordance with the protocol's specification.

[0010] The illustrative embodiment comprises: (a) a plurality of access points, wherein each of the access points performs a first protocol service for a respective network,

and wherein the correctness of the first protocol service is based on a maximum timing delay; and (b) a central controller for: (i) receiving an input signal from each of the plurality of access points, and (ii) transmitting to each of the plurality of access points an output signal based on the input signal from that access point and a second protocol service, wherein the correctness of the second protocol service is independent of the maximum timing delay.

Brief Description of the Drawings

[0011] Figure 1 depicts a schematic diagram of wireless local area network 100 in accordance with the prior art.

[0012] Figure 2 depicts a conceptual diagram of protocol stacks implemented by access point 101 and station 102-*i*, as shown in Figure 1, in accordance with the prior art.

[0013] Figure 3 depicts a schematic diagram of a plurality of local-area networks in accordance with the illustrative embodiment of the present invention.

[0014] Figure 4 depicts a conceptual diagram of a protocol stack implemented by access point 301-*i*, as shown in Figure 3, in accordance with the illustrative embodiment of the present invention.

[0015] Figure 5 depicts a conceptual diagram of a protocol stack implemented by central controller 303, as shown in Figure 3, in accordance with the illustrative embodiment of the present invention.

[0016] Figure 6 depicts a block diagram of the salient components of access point 301-*i*, as shown in Figure 3, in accordance with the illustrative embodiment of the present invention.

[0017] Figure 7 depicts a block diagram of the salient components of central controller 303, as shown in Figure 3, in accordance with the illustrative embodiment of the present invention.

[0018] Figure 8 depicts a flowchart of the operation of access point 301-*i*, as shown in Figure 3, in accordance with the illustrative embodiment of the present invention.

[0019] Figure 9 depicts a flowchart of the operation of central controller 303, as shown in Figure 3, in accordance with the illustrative embodiment of the present invention.

Detailed Description

[0020] Figure 3 depicts a schematic diagram of local-area networks 300-1 through 300-*K*, where *K* is a positive integer, in accordance with the illustrative embodiment of the present invention. As shown in Figure 3, each local-area network 300-*i*, wherein

$i \in \{1, 2, \dots, K\}$, comprises thin wireless access point 301- i , which communicates with wireless stations 102 in network 300- i . Each local-area network 300- i has a respective shared-communications channel; as will be appreciated by those skilled in the art, the separate shared-communications channels might be due to spatial separation, different frequencies, *etc.*

[0021] Each thin wireless access point 301- i is connected to central controller 303 via bi-directional link 304- i . As described in detail below, each thin wireless access point 301- i performs those protocol services whose correctness depends on a maximum timing delay. In addition, when wireless access point 301- i detects a particular condition (*e.g.*, receiving a signal from a station, detecting that the shared-communications channel has been idle for a time interval, *etc.*) that triggers a protocol service whose correctness does not depend on any maximum timing delay, thin wireless access point 301- i sends a signal to central controller 303 via bi-directional link 304- i to perform the protocol service. Thin wireless access point 301- i also receives signals via bi-directional link 304- i from central controller 303 after central controller 303 has performed the protocol service. A detailed description of the operation of thin wireless access point 301- i is disclosed below and with respect to Figures 6 and 8.

[0022] Central controller 303 receives signals from thin wireless access points 301- i through 301- K via links 304-1 through 304- K , performs protocol services whose correctness does not depend on any maximum timing delay, and transmits protocol service outputs to thin wireless access points 301- i through 301- K via links 304-1 through 304- K . In some embodiments, central controller 303 might be connected to an external network (*e.g.*, a wide-area network (WAN), the Internet, *etc.*), thereby functioning as a gateway, as is well-known in the art. In such embodiments, central controller 303 receives messages from the external network and forwards the messages to the appropriate access points for dissemination to the destination stations. In addition, central controller 303 receives from the access points messages that are directed to the external network, and transmits these messages via the external network accordingly. A detailed description of the operation of central controller 303 is disclosed below and with respect to Figures 7 and 9.

[0023] As will be appreciated by those skilled in the art, in some embodiments bi-directional links 304-1 through 304- K might instead be implemented wirelessly. The illustrative embodiment, however, has the advantage of enabling concurrent communications between central controller 303 and thin wireless access points 301- i through 301- K without any potential transmission interference.

[0024] Figure 4 depicts a conceptual diagram of access point 301-*i* implementing protocol stack 401-*i* in accordance with the illustrative embodiment of the present invention. Protocol stack 401-*i* consists of those protocols of stack 201 whose services are correct only if communication and processing delays do not exceed a maximum threshold. For example, in an Institute of Electrical and Electronics Engineers (IEEE) IEEE 802.11b wireless local-area network, protocol services in protocol stack 401-*i* would include: Distributed Coordination Function (DCF) medium-access control, Point Coordination Function (PCF) medium-access control, error control, and flow control. As will be appreciated by those skilled in the art, all of the aforementioned protocol services belong to the Data Link Layer of the Open Systems Interconnect (OSI) seven-layer model. In addition, protocol stack 401-*i* includes protocol services of the OSI Physical Layer (e.g., encoding data in radio signals, etc.).

[0025] Figure 5 depicts a conceptual diagram of central controller 303 implementing protocol stack 501 in accordance with the illustrative embodiment of the present invention. Protocol stack 501 consists of those protocols of stack 201 whose services are correct regardless of the presence or magnitude of any communication or processing delays. For example, in an IEEE 802.11b wireless local-area network, protocol services in protocol stack 501 might include generating a polling list for the Point Coordination Function (PCF) and Wi-Fi Protected Access (WPA) authentication. In a wireless local-area network that supports IEEE 802.11e, an extension of 802.11b proposed by IEEE Task Group E that enables quality-of-service (QoS), protocol services in protocol stack 501 could also include traffic monitoring, determining Enhanced Distributed Coordination Function (EDCF) parameters for different traffic classes, etc.

[0026] Figure 6 depicts a block diagram of the salient components of thin access point 301-*i* in accordance with the illustrative embodiment of the present invention. As shown in Figure 6, thin access point 301-*i* comprises receiver 601-*i*, processor 602-*i*, memory 603-*i*, and transmitter 604-*i*, interconnected as shown.

[0027] Receiver 601-*i* is a circuit that is capable of receiving signals over a wireless shared-communications channel, in well-known fashion, and of forwarding them to processor 602-*i*. In addition, receiver 601-*i* receives signals from central controller 303 over link 304-*i* in well-known fashion. As will be appreciated by those skilled in the art, in some embodiments of the present invention receiver 601-*i* might be a single receiver that can receive signals both wirelessly and via wireline, while in some other embodiments receiver 601-*i* might comprise two physical receivers (for example, a radio and an Ethernet network

interface card). In any case it will be clear to those skilled in the art how to make and use receiver 601-*i*.

[0028] Processor 602-*i* is a general-purpose processor that is capable of executing instructions stored in memory 603-*i*, of reading data from and writing data into memory 603-*i*, and of executing the tasks described below and with respect to Figure 8. In some alternative embodiments of the present invention, processor 602-*i* might be a special-purpose processor. In either case, it will be clear to those skilled in the art how to make and use processor 602-*i*.

[0029] Memory 603-*i* is capable of storing data and programs, including protocol stack 401-*i*, used by processor 602-*i*, as is well-known in the art, and might be any combination of random-access memory (RAM), flash memory, disk drive, etc. It will be clear to those skilled in the art how to make and use memory 603-*i*.

[0030] Transmitter 604-*i* is a circuit that is capable of receiving signals from processor 602-*i*, in well-known fashion, and of transmitting corresponding signals (i) over a wireless shared-communications channel, and (ii) over link 304-*i*, in well-known fashion. As will be appreciated by those skilled in the art, in some embodiments of the present invention transmitter 604-*i* might be a single transmitter that can transmit signals both wirelessly and via wireline, while in some other embodiments transmitter 604-*i* might comprise two physical transmitters (for example, a radio and an Ethernet network interface card). In any case it will be clear to those skilled in the art how to make and use transmitter 604-*i*.

[0031] Figure 7 depicts a block diagram of the salient components of central controller 303 in accordance with the illustrative embodiment of the present invention. As shown in Figure 7, central controller 303 comprises receiver 701, processor 702, memory 703, and transmitter 704, interconnected as shown.

[0032] Receiver 701 is a circuit that is capable of receiving signals over links 304-1 through 304-*K*, in well-known fashion, and of forwarding them to processor 702. In addition, as described above, in some embodiments receiver 701 might also receive messages from an external network (e.g., a wide-area network, the Internet, etc.). It will be clear to those skilled in the art how to make and use receiver 701.

[0033] Processor 702 is a general-purpose processor that is capable of executing instructions stored in memory 703, of reading data from and writing data into memory 703, and of executing the tasks described below and with respect to Figure 9. In some alternative embodiments of the present invention, processor 702 might be a special-purpose

processor. In either case, it will be clear to those skilled in the art how to make and use processor 702.

[0034] Memory 703 is capable of storing data and programs, including protocol stack 501, used by processor 702, as is well-known in the art, and might be any combination of random-access memory (RAM), flash memory, disk drive, *etc.* It will be clear to those skilled in the art how to make and use memory 703.

[0035] Transmitter 704 is a circuit that is capable of receiving signals from processor 702, in well-known fashion, and of transmitting corresponding signals over links 304-1 through 304-K, in well-known fashion. In addition, as described above, in some embodiments transmitter 704 might also transmit messages on to an external network (*e.g.*, a wide-area network, the Internet, *etc.*). It will be clear to those skilled in the art how to make and use transmitter 704.

[0036] Figure 8 depicts a flowchart of the operation of access point 301-*i* in accordance with the illustrative embodiment of the present invention.

[0037] At task 810, access point 301-*i* detects a particular condition (*e.g.*, receiving a protocol service request from a station 102-*j*, detecting that the shared-communications channel has been idle for a time interval, *etc.*) that triggers a protocol service *S*.

[0038] At task 820, access point 301-*i* checks whether protocol service *S* is a service whose correctness depends on a maximum timing delay. If so, execution continues at task 830, otherwise execution proceeds to task 840.

[0039] At task 830, access point 301-*i* performs protocol service *S*, the code for which is in protocol stack 401-*i* in memory 603-*i*. After completion of task 830, execution proceeds to task 860.

[0040] At task 840, access point 301-*i* sends, via transmitter 604-*i*, a request to central controller 330 to perform protocol service *S*, along with any necessary information (*e.g.*, input variables, *etc.*).

[0041] At task 850, access point 301-*i* receives, via receiver 601-*i*, a response from central controller 330 after central controller 330 has performed protocol service *S*.

[0042] At task 860, access point 301-*i* sends, if necessary, a message via transmitter 604-*i* to station 102-*j* that indicates that protocol service *S* has been performed. As will be appreciated by those skilled in the art, such a message, which might include an acknowledgement, an output value, *etc.*, might be appropriate for some protocol services requested by station 102-*j*, but might be unnecessary for other protocol services requested by station 102-*j*, or protocol services triggered by other events (*e.g.*, an idle shared-

communications channel, *etc.*) After completion of task 860, the method of Figure 8 terminates.

[0043] Figure 9 depicts a flowchart of the operation of central controller 303, in accordance with the illustrative embodiment of the present invention.

[0044] At task 910, central controller 303 receives, via receiver 701, a request from access point 301-*i* to perform a protocol service *S*. As described above in task 840, the request might also contain input variables or additional information needed for central controller 303 to perform service *S*.

[0045] At task 920, central controller 303 performs protocol service *S*, the code for which is in protocol stack 501 in memory 703.

[0046] At task 930, central controller 303 sends an acknowledgement message via transmitter 704 to access point 301-*i* that indicates that protocol service *S* has been performed. As will be appreciated by those skilled in the art, for some protocol services the acknowledgement message might also include additional output information (*e.g.*, issued EDCF parameters for an IEEE 802.11e network, *etc.*). After completion of task 930, the method of Figure 9 terminates.

[0047] **Illustrative Embodiment: IEEE 802.11**

[0048] The remainder of the detailed description discloses a particular embodiment of the present invention for IEEE 802.11 wireless local-area networks, including the proposed IEEE 802.11e specification. As will be appreciated by those skilled in the art, this is just one example of an embodiment of the present invention.

[0049] **Introduction**

[0050] Wireless communications involves a radio that communicates wirelessly with the subscribers in a limited area around it and with the larger network through a wired connection. In cellular communications one or more radio(s) reside at what is known as the base station. The coverage area within which subscribers may communicate wirelessly is known as the 'cell'. Analogous notions in wireless LANs (WLANs) are the access point (AP) and the Basic Service Area (BSA), respectively. The collection of clients served by an AP is known as the Basic Service Set (BSS).

[0051] Wireless LANs provide wireless peer-to-peer communication between stations and access to the wired network. A collection of stations communicating directly via the wireless medium is known as an ad hoc network. Stations may communicate with stations in a different WLAN cell with the aid of a wired backbone network, known as the distribution system. An AP is a station that serves as a gateway to the distribution system; it is

analogous to the base station of a cellular communications network. A single-cell wireless LAN is adequate for a home or a small office. Multiple-cell wireless LANs are used to cover larger buildings or spaces.

[0052] Wireless LAN standards specify the protocols for the physical and MAC layers. IEEE 802.11 is the standard dominating the wireless LAN market today. IEEE 802.11g is the newest 802.11 standard - finalized in 2003 - which provides higher speeds than 802.11b in the unlicensed 2.4 GHz ISM band. IEEE 802.11a, the 5GHz extension to 802.11b/g provides high speeds. Other wireless LAN standards include: HiperLAN/2 (the European counterpart to the "American" 802.11a standard).

[0053] WLANs operate in the unlicensed portions of the spectrum, where they provide interference-free simultaneous transmissions on multiple channels; each cell transmits on a single time division duplex (TDD) channel. Channel access mechanisms for asynchronous data transfer fall into two categories: distributed contention based and centralized contention free. Under contention-based access methods, stations access the channel when there is data to transmit, thus risking collision with transmissions attempted by other stations. The distributed random access protocol in 802.11 WLANs is known as the distributed coordination function (DCF). In the QoS-enhanced protocol in the proposed 802.11e draft standard, known as the enhanced distributed coordination function (EDCF), stations access the channel according to the priority of their traffic. Contention-free access methods permit a single station to transmit at a time. With centralized contention-free protocols, a controller - typically the AP - polls stations to send or receive data. The deterministic polling protocol in 802.11 wireless LANs is known as the point-coordination function (PCF). The QoS-enhanced protocol in the 802.11e draft standard, known as the hybrid coordination function (HCF), also employs polling for contention-free access, but with increased channel use efficiency.

[0054] Stations associated with a cell, or a basic service set (BSS), compete for channel access for a variety of reasons. These include the transmission of data packets; the reservation of the channel for the transmission of data packets; or the reservation on the polling list of a deterministic multiple access protocol, like PCF or HCF. The PCF and HCF rely on distributed multi-access methods to claim the channel.

[0055] The DCF employs the Ethernet MAC protocol, which has been adapted to provide the basic access mechanism for 802.11 WLANs. The rules for channel access prohibit a station from attempting transmission of a newly arrived packet if the channel is busy. Carrier sensing is used in order to determine whether the channel is idle. If not idle, transmission is deferred by a randomly selected delay following completion of the current

transmission; this avoids collision with transmission attempts by other stations waiting for the release of the channel. This deferral time is used to set the backoff timer, which is decreased only when the channel remains idle following a transmission for a period equal to the Distributed Inter-Frame Space (DIFS). A station attempts transmission when this timer expires. The AP is allowed to access the medium sooner (after a PIFS) idle time interval.

[0056] EDCF varies, according to the priority of the traffic, the length of the time interval a station waits for an idle time interval before attempting transmission, or before counting down, immediately following a busy period. This time interval, which is known as the Arbitration-Time Inter-Frame Space (AIFS), is shorter for higher priority packets, enabling them to get to the channel faster. Once this time expires, the backoff counter is decreased whenever the channel is idle for a time interval equal to the time slot.

[0057] Centralized MAC Architecture

[0058] In order to reduce the expense and complexity of base stations/APs in wireless communications systems, it is possible to limit their functionality and offer complementary functionality at a single remote controller. Such arrangement reduces the cost and complexity of the access points, which is especially advantageous when deploying a large number of wireless local-area networks (for example, on a corporate or academic campus). Base stations in cellular can consist of RF transmitters/receivers relaying signals back and forth to the MSC for processing.

[0059] The idea of reducing the complexity of an AP applies also to WLANs, but with some limitations. Unlike in cellular systems, transmission in a WLAN relies on random access through tightly timed frame exchanges. In other words, the rules governing whether or not a station attempt transmission are based on the time that elapsed since the medium became idle following a medium busy condition. Introducing a delay in registering the 'end-of-busy' time would put the station – in this case the AP – at a disadvantage relative to other stations competing for access to the wireless medium, as the latter will have a head start in counting down the elapsed idle time. With synchronization based on carrier sensing and a highly interactive MAC protocol, the communication delays between the RF component performing carrier sensing and access control must be minimal. In particular, introducing a transmission delay from the AP to the controller of 5 microsecs (on an 802.11b WLAN) or greater delays the AP's response so that its attempted transmission might collide with a frame from a station. In general, the delayed AP responses will disadvantage it with respect to access to the wireless medium.

[0060] To illustrate the potentially undesirable results of improperly centralizing AP functionality, consider the following example. The components of a WLAN IC chip set

consist of the following: (1) the radio front end, which transmits and receives the signal; (2) the baseband processor, which modulates and demodulates the signal; and (3) the media access controller, or MAC, which handles the protocol translation and control. Suppose that the AP's media access controller resides at a remote gateway/switch, which is separated by a wired Ethernet link from the radio front end and baseband processor, both located at the BSA. It is assumed that the STA has all components co-located, hence the communication delay between the radio front end and media access controller are negligible. Suppose now that the STA and the AP both have frames pending transmission. Consider an AP and several stations operating on an 11 Mbps 802.11b link. A communication delay on the AP side in excess of 5 microsec would cause the AP to have a delayed response to the end of a medium busy period. While the MAC protocol allows the AP to access the medium sooner (after a PIFS) idle time interval, the AP's delayed response would cause it to transmit after an idle time interval of length DIFS or longer, causing possible collisions with stations.

[0061] This example therefore provides guidelines for partitioning the AP functions into two parts, (1) a 'thin' AP and (2) a centralized controller so that the resulting MAC performance does not suffer, and the AP has minimal functionality.

[0062] In this example, the centralized AP controller is referred to as the QoS Cell Manager (QCM). The QCM can co-exist in the Gateway together with other centralized functions, like Self-Configuration (which performs RF planning and radio resources management tasks). The QCM serves a multitude of 'thin' APs that are located in each cell. The 'thin' AP must perform the radio and baseband processing functions, but is not limited to these functions alone. The QCM provides the AP functions complementary to those residing in the 'thin' AP. The medium access controller functions can be split between the QCM and the thin AP. In this paper, we refer to a 'thin' AP as an *Access Port*, which is designated by the lower case abbreviation 'ap'.

[0063] In this example, the QCM is embedded in a gateway (e.g., for connecting to the Internet, etc.) All intelligence may reside at the QCM. The QCM monitors traffic and issues a QoS information element containing the EDCF parameters; i.e. AIFS, CWmin, CWmax, PF, MSDULifetime, etc. In this example, the QCM performs admission control upon receipt of a TSPEC request from a STA. To the extent this request involves a traffic stream going through the network (other than sidelink traffic), acceptance or rejection of a request would be communicated to upper layers. The QCM generates a Polling Schedule and reissues it as needed, based on the demand. Upon receipt of an RR, the QCM supplies a Real-Time Polling schedule. Finally, the QCM is responsible for scheduling downlink traffic.

In doing so it designates the content of TXOPs which is forwarded to the ap for transmission. In performing these functions, the QCM receives from the client information that is transparently routed through the ap. Alternatively, some of these functions may be relegated to the ap.

[0064] The minimal functionality of an ap is the transmission of TXOPs and polling. This involves the following:

- Buffering of MPDU fragments and TXOP contents
- Carrier sensing – determining whether the wireless medium is idle.
- Medium access control (-IFS delay and backoff) – If the channel is idle, the wireless medium may be accessed following a PIFS idle period and any required backoff interval.
- Acknowledgement – Following the successful receipt of a transmission, the ap will send acknowledgment. If after transmitting a frame or a poll, an acknowledgement is not received within a SIFS time interval, the ap will respond accordingly.)
- Polling -- The ap transmits the MSDUs received from the QCM and buffered at the ap. If a TXOP overlaps a polling schedule, the ap will interleave polls with MPDUs according to the specified schedule.
- Reassembly of MPSDUs and transmission to the QCM

[0065] In addition, the ap may perform other functions, which would be otherwise be carried out by the QCM.

[0066] In this example, clients interoperate with the ap-QCM complex via the IEEE 802.11e protocols. The split of the AP functionality is transparent to the clients, which are the 802.11 stations. Communication between the ap and the QCM can be done through a wired or wireless connection, according a proprietary protocol at the MAC layer.

[0067] Network Management Functions

[0068] In a centralized architecture, key AP functions are split between the QCM and the access port. In addition, the QCM may contain intelligent network management functions that not described by standards. In this example, the following network management functions reside in the QCM: TSPEC Admission Control, Traffic Monitoring, and Scheduling.

[0069] TSPEC Admission Control

[0070] The QCM is responsible for admission of TSPEC requests for TXOPs by traffic streams routed either through the AP (and distribution system) or peer to peer within a cell. If a TSPEC request provides a range of service parameters, there will be room for

negotiation of the TXOP schedule, before a final schedule is set. If there is no TSPEC request made and/or a TSPEC is not available, the upper layers will provide a nominal TSPEC request through the Gateway. Ideally, TSPEC admission decisions would be incorporated in call admission control taking place in higher layers. TSPEC admission control will take into account the priority and parameter of the request for TXOPs, the commitments made to TSPEC requests already admitted, the margin for spectrum re-allocation through schedule negotiation of admitted TSPEC requests, and the load on the wireless link in the cell. For example, a TSPEC request of a given priority would be accepted if the WLAN is lightly loaded, but denied if the bandwidth committed to TSPEC requests exceed a corresponding threshold value. These threshold values would increase with priority.

[0071] Traffic Monitor

[0072] The QCM is responsible for setting the parameters of the random access protocol. For instance, in IEEE 802.11e, the QoS information element contains the EDCF parameters. These parameters are adapted by monitoring the wireless channel to determine the number/frequency of collisions and successful transmissions. Information is also provided by the clients on the success and failure rate of transmissions by traffic category. With this information, the QCM will adapt the EDCF parameters for each category in order to meet QoS requirements.

[0073] Scheduler

[0074] The QCM schedules transmissions on the wireless link. It sequences for transmission the downlink frames and packages them into TXOPs that are sent to the ap, one at a time. Transmission of downlink TXOPs may be coupled with polling in CAPs. Selection of the frames placed in a CAP and their sequencing is designed to reflect QoS objectives, such as maximizing throughput and meeting latency and jitter restrictions.

[0075] AP Functions

[0076] In accordance with this example, the AP functions are shared between the QCM and the access port. In theory, the ap functions may range between the minimum functionality described herein and the complete set of AP functions. The following functions of an access point are described below: Downlink Transmission, Polling, and EDCF uplink transmission.

[0077] Downlink Transmission

[0078] The QCM pre-packages a TXOP and sends it to the ap for transmission together with polling instructions, if applicable. Fragmentation of frames received at the MAC sublayer from upper layers – known as the MSDUs (MAC Service Data Units) – may be performed either by the QCM or by the ap. The ap transmits the fragments, MPDUs (MAC

Protocol Data Units) , receives acknowledgements and, if possible, retransmits failed MPDUs. Upon expiration of a TXOP time limit, the ap returns an acknowledgement vector (VACK) to the QCM. The vector acknowledgement indicates whether and which of the TXOP frames failed transmission, as it consists of 0s or 1s indicating whether the individual frames are successfully transmitted (bit set to 0) or not (bit set to 1). VACK is ANDed with a mask of 1s corresponding to the transmitted frames. If the returned VACK=0, all frames have been successfully transmitted. If certain frames failed earlier transmission attempts, the QCM determines whether they should be retransmitted, and if so, will incorporate them into subsequent TXOPs.

[0079] Polling

[0080] The QoS Cell Manager queues traffic and schedules MSDUs to be transmitted within a CAP according to the time-to-live/priority/jitter sensitivity. Use of non-standard queues and queue management methods will be pursued to this end.

[0081] The QCM determines the real-time poll schedule and data/poll interleaving. The real-time poll schedule is more specific than the schedule issued upon TSPEC acceptance. The former specifies the clock time range within which a poll must be sent to a client. The ap buffers (and fragments) any MSDUs included in the CAP, it transmits MPDUs and polls on the wireless channel according to received schedule, returns a vector ACK to the QCM for the downlink MSDUs in the CAP, and sends data received from clients to QCM. It also repeats unanswered polls, provided there is time left in the CAP. If fragmentation takes place in the ap, reassembly of MSDUs will also take place in the ap.

[0082] EDCF uplink transmission

[0083] The ap handles RTS/CTS exchanges as specified for the AP by the 802.11 standard. The ap responds to an RTS with a CTS and acknowledges the receipt of frames from clients. It then forwards the data to the QCM.

[0084] It is to be understood that the above-described embodiments are merely illustrative of the present invention and that many variations of the above-described embodiments can be devised by those skilled in the art without departing from the scope of the invention. It is therefore intended that such variations be included within the scope of the following claims and their equivalents.

[0085] What is claimed is:

1. An apparatus comprising:

(a) a plurality of access points, wherein each of said access points performs a first protocol service for a respective network, and wherein the correctness of said first protocol service is based on a maximum timing delay; and

(b) a central controller for:

(i) receiving an input signal from each of said plurality of access points, and

(ii) transmitting to each of said plurality of access points an output signal based on the input signal from that access point and a second protocol service, wherein the correctness of said second protocol service is independent of said maximum timing delay.

2. The apparatus of claim 1 wherein said first protocol service belongs to a layer selected from the group consisting of: physical layer, and data link layer.

3. The apparatus of claim 2 wherein said first protocol service is selected from the group consisting of: a medium access control service, an error control service, and a flow control service.

4. The apparatus of claim 1 wherein said second protocol service is selected from the group consisting of: an authentication service, an authorization service, a traffic monitoring service, an admission control service, and a polling list maintenance service.

5. The apparatus of claim 1 wherein said central controller is also for:

(iii) receiving a datum via a wide-area network, and

(iv) transmitting said datum to at least one of said access points.

6. The apparatus of claim 1 wherein each of said access points is also for receiving a first datum from said central controller and for transmitting a second datum based on said first datum to at least one station in said respective network.

7. A method comprising:

(a) performing a first protocol service with a first processor, wherein the correctness of said first protocol service is based on a maximum timing delay;

(b) transmitting a first signal to a second processor, wherein said second processor is for performing a second protocol service, and wherein the correctness of said second protocol service is independent of said maximum timing delay; and

(c) receiving from said second processor a second signal based on said second protocol service.

8. The method of claim 7 further comprising (d) detecting a first condition, wherein (a) is in response to (d).

9. The method of claim 8 wherein said first condition comprises the transmission of a signal over a shared-communications channel.

10. The method of claim 8 wherein said first condition comprises an idle time interval for a shared-communications channel.

11. The method of claim 8 further comprising (e) detecting a second condition, wherein (b) is in response to (e).

12. The method of claim 11 wherein said second condition comprises the transmission of a signal over a shared-communications channel.

13. The method of claim 11 wherein said second condition comprises an idle time interval for a shared-communications channel.

14. The method of claim 7 further comprising (d) detecting a condition, wherein (b) is in response to (d).

15. The method of claim 14 wherein said condition comprises the transmission of a signal over a shared-communications channel.

16. The method of claim 14 wherein said condition comprises an idle time interval for a shared-communications channel.

17. The method of claim 7 wherein said first protocol service belongs to a layer selected from the group consisting of: physical layer, and data link layer.

18. The method of claim 17 wherein said first protocol service is selected from the group consisting of: a medium access control service, an error control service, and a flow control service.

19. The method of claim 17 wherein said second protocol service is selected from the group consisting of: an authentication service, an authorization service, a traffic monitoring service, an admission control service, and a polling list maintenance service.

20. A method comprising:

- (a) performing a first protocol service for a first network at a first processor, wherein the correctness of said first protocol service is based on a maximum timing delay;
- (b) performing said first protocol service for a second network at a second processor;
- (c) transmitting a first signal from said first processor to a third processor;

(d) performing a second protocol service for said first network at a third processor, wherein the correctness of said second protocol service is independent of said maximum timing delay;

(e) transmitting a second signal from said third processor to said first processor, wherein said second signal is based on said second protocol service;

(f) transmitting a third signal from said second processor to said third processor;

(g) performing said second protocol service for said second network at said third processor; and

(e) transmitting a fourth signal from said third processor to said second processor, wherein said fourth signal is based on said second protocol service.

21. The method of claim 20 wherein said first protocol service belongs to a layer selected from the group consisting of: physical layer, and data link layer.

22. The method of claim 21 wherein said first protocol service is selected from the group consisting of: a medium access control service, an error control service, and a flow control service.

23. The method of claim 20 wherein said second protocol service is selected from the group consisting of: an authentication service, an authorization service, a traffic monitoring service, an admission control service, and a polling list maintenance service.

FIG. 1 (PRIOR ART)


WIRELESS LOCAL-AREA NETWORK 100

FIG. 2 (PRIOR ART)




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9

