

[0001] SELF-CONFIGURABLE WIRELESS LOCAL
AREA NETWORK NODE

[0002] FIELD OF INVENTION

[0003] The present invention generally relates to wireless local area networks (WLANs), and more particularly, to a method and apparatus for determining and analyzing network topology, configuring network nodes, and resolving functional conflicts that arise in network architectures.

[0004] BACKGROUND

[0005] The term “access point” (AP) as used herein includes, but is not limited to, a base station, an access router (AR), a Node B, a site controller, or other interfacing device in a wireless environment that provides other stations with wireless access to a network with which the AP is associated.

[0006] The term “station” (STA) as used herein includes, but is not limited to, a wireless transmit/receive unit (WTRU), a user equipment, a mobile station, a fixed or mobile subscriber unit, a pager, or any other type of device capable of operating in a wireless environment.

[0007] Typically, a WLAN includes a plurality of APs, wherein each AP is capable of conducting concurrent wireless communications with appropriately configured STAs, as well as multiple appropriately configured APs or ARs, when configured in the “infrastructure mode”. Some STAs may alternatively be configured to conduct wireless communications directly to one another, i.e., without being relayed through a network via an AP. This is commonly known as “peer-to-peer mode” or “ad hoc mode”. Where a STA is configured to communicate directly with other STAs, it may also be configured to function as an AP. STAs can be configured for use in multiple networks, with both network and peer-to-peer communications capabilities.

[0008] In the infrastructure mode architecture, STAs are conventionally connected in a star-type topology to a central AP in order to communicate to each other or to connect to other external networks. Although this architecture has

proven successful in the past, many factors, such as the increasing number of closely located APs, the increasing number of applications for a WLAN, and the fact that APs are restricted to public bands, have resulted in conventional infrastructure mode architectures becoming less desirable. Accordingly, other infrastructure mode topologies have evolved.

[0009] One topology is known as a “mesh” topology, in which WLAN nodes have two or more paths between them which enables the nodes to communicate directly with each other (i.e., as in the ad hoc mode) and to communicate indirectly with each other (via other nodes that relay information). A second topology is known as a “split” architecture, in which one or more access routers (ARs) or access controllers (ACs) are connected via an interconnection to APs present in the network. The ACs provide network-wide monitoring, improve scalability, and facilitate dynamic configurability. The logical interconnection may be a direct connection to the APs, a switched connection, or a routed network connection. The AC and the AP may be collocated in the same physical device.

[0010] In addition to exchanging configuration and control information with the APs, the ACs “split” or share certain functionalities with the APs that are conventionally provided solely by the APs. That is, functions typically provided by standalone or “fat” APs are removed from these APs and are provided by the AC(s). These split-function or reduced-function APs are referred to as “thin” APs. This architecture is similar to a UMTS architecture, where the AC is analogous to a central radio network controller (RNC) and the AP is analogous to a Node B connected to the RNC.

[0011] Figure 1 is a diagram of a network 100 with an infrastructure mode architecture including a plurality of STAs 102a-102n communicating with a fat AP 104. This architecture is often referred to as a fat AP architecture because all of the medium access control (MAC) layer functionalities are located in the AP 104. The STAs 102 communicate with the AP 104, and with one another via the AP 104. The AP 104 incorporates a physical (PHY) layer 106, a real time (RT) MAC layer 108, and a non-real time (NRT) MAC layer 110.

[0012] Figure 2 is a diagram of a network 200 with a split architecture, including a plurality of STAs 202a-202i, several APs 204a-204c, and an access controller (AC) 206. In the split network 200, certain AP functions are split away from the APs 204 and are provided by the AC 206. Although the AP functions may be split in any number of configurations, Figure 2 shows one of the most common arrangements. The APs 204 terminate the infrastructure side of the wireless physical links, provide radio-related management, and provide all RT services to the STAs 202. The AC 206 provides the NRT management functions such as configuration, quality of service (QoS), access control, etc., for all of the APs 204. By sharing functionalities at a higher layer, a better coordinated deployment is possible.

[0013] The AP functional definitions made to support future AC-AP architectures must also be backward compatible to accommodate present-day devices. Since infrastructure mode networks are the present-day convention, it is noted that accommodating hybrid architectures, i.e., those networks with both fat APs and thin APs, will be a significant challenge for future networks.

[0014] An example of a pathological hybrid network 300 with both fat APs and thin APs is shown in Figure 3. The network 300 includes a plurality of STAs 302a-302i; two thin APs, AP1 (304a) and AP2 (304b); a fat AP, AP3 (306); and an AC 308. AP3 306 provides all of its L2 MAC functionalities, including both the RT MAC 310 and the NRT MAC 312. In this deployment, the AC 308 manages all three APs 304a, 304b, 306. Accordingly, there is a conflict with redundancy in the NRT functionalities between the AC (NRT MAC 314) and AP3 306 (NRT MAC 312). This conflict is further aggravated in other network topologies such as, for example, mesh networks, wherein AP functionalities are distributed over the entire mesh network and wherein direct communication between ACs and APs is not always possible.

[0015] Accordingly, it is desirable to provide a method and apparatus to resolve functional conflicts that arise in WLAN architectures.

[0016]

SUMMARY

[0017] The present invention relates to a method and apparatus for configuring network nodes and resolving functional conflicts or redundancies that arise in network architectures. In a WLAN including at least one AC and a plurality of APs, the AC transmits functionality queries to the APs. Upon receiving these queries, the APs transmit query responses, which include the functional capabilities of the APs. The AC then generates a functional map of the functional capabilities available in the network based on the query responses. Conflicting or redundant functional capabilities are identified and are disabled, enabled, or reconfigured by instructions from the AC to the APs having conflicting or redundant functionalities.

[0018] A method for configuring nodes in a WLAN including an AC and an AP begins with the AC sending a functionality inquiry to the AP. The AP responds to the functionality inquiry by sending its functional capabilities to the AC. The AC maps the functional capabilities of the AP. A determination is made whether a capability conflict exists between the capabilities of the AP and the capabilities of the AC, wherein a conflict can include redundant capabilities between the AP and the AC. If there are any capability conflicts, they are resolved.

[0019] An AC for configuring nodes in a WLAN includes a transmitter/receiver, an inquiry device, a capability mapping device, and a capability evaluating device. The inquiry device is in communication with the transmitter/receiver and is configured to send functionality inquiry messages to an AP in the WLAN. The capability mapping device is in communication with the transmitter/receiver and is configured to receive functionality inquiry reply messages from the AP and to map the functionality capabilities of the AP. The capability evaluating device is in communication with the capability mapping device and is configured to determine if any functionality capability conflicts exist between the AP and the AC.

[0020] BRIEF DESCRIPTION OF THE DRAWINGS

[0021] A more detailed understanding of the invention may be had from the following description of a preferred embodiment, given by way of example, and to be understood in conjunction with the accompanying drawings, wherein:

[0022] Figure 1 is a diagram of an infrastructure mode WLAN with a fat AP;

[0023] Figure 2 is a diagram of a split architecture WLAN with thin APs;

[0024] Figure 3 is a diagram of a hybrid WLAN architecture with both fat APs and thin APs;

[0025] Figures 4A, 4B, and 5 are flow diagrams of a method for resolving functionality conflicts or redundancies between an AC and an AP in the architecture shown in Figure 3; and

[0026] Figure 6 is a block diagram of a system including an AC and an AP configured to perform the method shown in Figures 4A, 4B, and 5.

[0027] DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0028] In a preferred embodiment, a WLAN including at least one AC and at least one AP is considered. The AC is preferably configured to split or remove certain L2 MAC functionalities (e.g., NRT MAC functions) from the AP(s) and provide these functionalities to the network. An example of this configuration is shown in Figure 3. Alternatively, the AC may provide all MAC layer functionalities to the network. As noted above, the AC may split and provide any functionality that is typically provided by the APs, including MAC layer functions, PHY layer functions, security methods, management interfaces, etc.

[0029] The present invention is not limited to the network deployment illustrated in Figure 3. Rather, the present invention is applicable to any network deployment in which functional conflicts between network components occur. As previously described in connection with Figure 3, there is a redundancy conflict with the NRT functionalities between the AC 308 and AP3 306, in that both the AC 308 and AP3 306 attempt to provide L2 NRT-MAC functionalities, creating a problem in the network 300.

[0030] Overview

[0031] Utilizing a query-response mechanism, the AC queries all associated APs regarding their respective functional capabilities. Once responses are received from all associated APs, the AC utilizes the responses to generate a functional map of the different functional blocks available in the network. This functional map enables the AC to detect whether any APs are providing duplicate functionalities that may cause an operational conflict.

[0032] If an AP in the system fails to respond or is unable to respond to the AC's query, the AC may assume that the non-responding AP is capable of implementing all functionality modules and that a resulting functionality conflict exists. Failure to respond to the AC's query may indicate one of two possible scenarios: that an AP was implemented according to an alternate standard or that the AP is a legacy AP not capable of responding to the query. Since the AP will be assumed to be a fat AP capable of providing all possible functionalities, the AC stops support of the non-responding AP.

[0033] To resolve a detected conflict or redundancy, the AC generates and transmits a message to the AP ordering that the conflicting module in the AP be disabled or reconfigured to resolve the redundancy. Conflicts among network modules are preferably resolved based on the individual priority of the modules. Modules that reside higher in the network hierarchy (i.e., AC) are given higher priority than nodes that are lower in the hierarchy (i.e., APs). Disabling a module in the AP preferably occurs after a predetermined period of time (e.g., sync-up, future event, etc.), at which point the AC takes control of those functions for the AP and thus resolves the conflict.

[0034] It should be noted that the messaging means described above may also be used to enable functionalities, such as with a functionality enabling message. Additionally, the disabling, reconfiguring, and enabling of functionalities may occur during a system start-up phase or dynamically while the system is operating. Dynamic adjustments may be utilized, for example, to better distribute and re-balance processing power in a system based on functional

maps and node priority. Similarly, dynamically adjusting functions may be preferred in certain architectures, such as mesh networks, wherein the AC functionalities are not situated in a single network node, but are distributed across various nodes.

[0035] It is noted that an AC may split and/or provide some or all MAC layer functions, while allowing the APs to provide some or all PHY layer or multiple PHY layer functions. Further, the approach described herein may be applied to other functionalities such as to multiple PHY layers, multiple radios (RFs), multiple security methods, multiple routing algorithms, different versions of a standard (e.g., 802.11e-WMS/WME/other), management interfaces (e.g., station management entity (SME)), etc., wherein the AC splits and provides these functions.

[0036] AR-AP Capability Signaling

[0037] Figures 4A, 4B, and 5 are flow diagrams of a method 400 for resolving functionality conflicts between an AC 402 and an AP 404. Although the network is typically provided with a plurality of APs, only one AP 404 and its interaction with the AC 402 is shown for purposes of simplicity, it being understood that the AC 402 interrogates all other APs in the network in a like manner. As an initial step, the AP 404 starts up and enables all of its functionalities (step 406).

[0038] The AC 402 sends an inquiry to the AP 404 (step 410). The inquiry can be sent upon initial setup or upon the entry of a new node to the network, such that the system can correctly configure itself as quickly as possible. The inquiry can also be sent on a periodic basis (e.g., once a day as the system should not need to change too frequently) or by being triggered by a particular event (e.g., when a congestion condition arises, the capabilities are gathered as a pretext for load balancing). After the inquiry is sent, the AC 402 sets a timer for receiving a reply from the AP 404 (step 412). The AP 404 receives the inquiry (step 414) and replies to the AC 402 providing its functional capabilities (step 416).

[0039] There is a possibility that the AC 402 may send an inquiry to the AP 404 (step 410) and the AP 404 fails to respond, as indicated by the dashed lines for sending the inquiry and sending the reply. Failure to respond to the AC's query may indicate that an AP was implemented according to an alternate standard or that the AP is a legacy AP not capable of responding to the query.

[0040] After the timer has been set (step 412), a determination is made whether the AC 402 has received a reply from the AP 404 (step 418). If no reply has been received, then a check is made to determine whether the timer has expired (step 420). If the timer has not expired, the method 400 waits for a reply from the AP 404 (step 418).

[0041] If the timer has expired (step 420), this indicates that no reply was received from the AP 404. The AC 402 defaults the AP 404 to having all possible capabilities (step 422). Since the AP 404 is presumed to have all possible capabilities (e.g., that the AP 404 is a fat AP), the AC 402 also halts support of services to the AP 404 (step 424). As applied to the scenario shown in Figure 3, the AC 402 will not attempt to perform NRT-MAC scheduling for the AP 404 because the AP 404 is presumed to have this capability.

[0042] Figure 4B is a flowchart of an alternate method 450 for performing the initial steps of the method 400. The AP 404 starts up and enables all of its functionalities (step 406). The AP 404 attempts to discover the AC 402 (step 452). The AP 404 can attempt to discover the AC 402 through various means. For example, the AC 402 could broadcast packets announcing its presence or a dedicated pre-agreed address on how to join the AC 402 could be known to all APs. After the AP 404 has discovered the AC 402, the AP 404 sends its capabilities to the AC 402 (step 454).

[0043] If the AC 402 has received a reply from the AP 404 (step 418), if the AC 402 has set the AP 404 to having default capabilities (steps 422, 424), or if the AP 404 has sent its capabilities to the AC 402 (step 454), the AC 402 stores and maps the capabilities of all APs in the network (step 502; Figure 5). Next, the AC 402 determines whether the capabilities of any AP conflict or are redundant with the capabilities of the AC 402 or with the capabilities of another

AP (step 504). The capabilities of other APs are checked for situations where there are different levels of APs or where APs have different priorities within the network.

[0044] As one example, assume that the AP 404 has a plurality of functional capabilities, at least one of which is the same as at least one of the capabilities of the AC 402. In order to avoid any conflict within the network, the co-existence of the common function by both the AC 402 and the AP 404 necessitates that this co-existence be resolved by preventing the AP 404 from providing this functionality to other nodes within the network. The AC 402 may utilize an associative memory technique for comparing the AC functionalities with the AP functionalities. However, any other suitable technique for determining a conflict may be employed.

[0045] If there are any capability conflicts or redundancies (step 504), then the AC 402 sends a disable or reconfigure capabilities message to the AP 404 (step 506). In instances where an AP may have priority over another AP or other APs in the network, the AP having priority is permitted to support the functionality in issue and lower priority APs will have the functionality disabled or reconfigured. In situations where none of the APs have a higher priority level but are located at a higher level in the network architecture hierarchy, only the AP having the higher network hierarchy level is instructed to continue support of that functionality. In either case, the AC provides disabling or reconfiguring messages to APs at a lower network hierarchy level.

[0046] Additional factors considered by the AC include functional capabilities derived from cross-vendor components in which conflicts or redundancies arise due to differences in the components or due to one of the APs having superior operational capabilities compared with other APs in the network. This situation may arise where functional capabilities are shared in common with two APs, wherein one of which is a legacy AP, and the AC decides in favor of the more up to date module while disabling or reconfiguring the functional capability of the legacy AP.

[0047] The AP 404 receives the disable or reconfigure capabilities message (step 508) and disables or reconfigures any capabilities that conflict or are redundant with the AC's capabilities (step 510). After adjusting the capabilities according to the message, the AP 404 sends an acknowledgement (ACK) to the AC 402 (step 512). The AC 402 receives the ACK and updates the capabilities map with the current capabilities of the AP 404 (step 514).

[0048] The allocation of capabilities between and among the APs is not limited to MAC layer functions and PHY layer functions and may include security methods, management interfaces, and the like. For example, in mesh networks, the AC can allocate and split functional capabilities over the network in order to provide better balance within the network and to alleviate potential overloading of network nodes. Alternatively, security requirements may be utilized as the overriding factor in allocating and assigning functional capabilities between and among the APs.

[0049] If there are no conflicting capabilities or redundancies (step 504) or after the AC 402 has received the ACK and updated the capabilities map based on disabled or reconfigured capabilities (step 514), a determination is made whether any previously disabled capabilities of the AP 404 need to be enabled (step 516). This scenario may arise, for example, in a load balancing situation where the AC 402 wants to enable or reconfigure functions in the AP 404 that it had previously disabled or reconfigured.

[0050] If there are any capabilities that need to be enabled or reconfigured (step 516), then the AC 402 sends an enable/reconfigure message to the AP 404 (step 518). The AP 404 receives the enable/reconfigure message (step 520) and enables or reconfigures the capabilities listed in the message (step 522). After enabling or reconfiguring the capabilities according to the message, the AP 404 sends an ACK to the AC 402 (step 524). The AC 402 receives the ACK and updates the capabilities map with the current capabilities of the AP 404 (step 526) and the method terminates (step 528).

[0051] If there are no capabilities that need to be enabled or reconfigured (step 516), then the method terminates (step 528).

[0052] Exemplary System

[0053] Figure 6 is a block diagram of a system 600 including an AC 602 and an AP 604 configured to perform the method 400. The AC 602 includes a transmitter/receiver 610, an antenna 612 connected to the transmitter/receiver 610, an inquiry device 614 in communication with the transmitter/receiver 610, a timer 616 in communication with the inquiry device 614, a capability mapping device 618 in communication with the transmitter/receiver 610 and the timer 616, and a capability evaluating device 620 in communication with the capability mapping device 618 and the transmitter/receiver 610.

[0054] The AP includes a transmitter/receiver 630; an antenna 632 connected to the transmitter/receiver 630; a capability determining device 634 in communication with the transmitter/receiver 630; a station management entity (SME) 636 in communication with the capability determining device 634, the SME 636 including a list of the capabilities of the AP 604; and a capability adjusting device 638 in communication with the transmitter/receiver 630 and the SME 636.

[0055] In operation, the system 600 functions as follows. The inquiry device 614 sends an inquiry message to the transmitter/receiver 610 for transmission to the AP 604. When the inquiry message is sent, the inquiry device 614 sets the timer 616. The AP 604 receives the inquiry message via its transmitter/receiver 630. The inquiry message is passed to the capability determining device 634, which accesses the capabilities list in the SME 636 to determine the capabilities of the AP 604. The capability determining device 634 then sends a reply message to the transmitter/receiver 630 for transmission to the AC 602.

[0056] The reply message is received at the AC 602 and is passed to the capability mapping device 618, which maps the capabilities of all APs in communication with the AC 602. If the AC 602 does not receive a reply from the AP 604 and the timer 616 expires, and the capability mapping device 618 defaults the AP 604 to having all possible capabilities (i.e., the AP 604 will be considered to be a fat AP).

[0057] The capability evaluating device 620 examines the capability mapping for all APs and determines which APs have capabilities that conflict with the AC 602. If there are any conflicting capabilities, the capability evaluating device 620 sends a disable or reconfigure message to the transmitter/receiver 610 for transmission to the AP 604. Upon receipt of the disable or reconfigure message by the AP 604, it is forwarded to the capability adjusting device 638 which disables or reconfigures the capabilities specified by the message by signaling the SME 636 which updates the capabilities list accordingly. Once the capabilities are disabled or reconfigured, the capability adjusting device 638 sends an ACK to the AC 602.

[0058] In a similar manner, if the AC 602 instructs the AP 604 to enable or reconfigure any capabilities, the capability adjusting device 638 enables or reconfigures the capabilities specified by the AC 602 by signaling the SME 636 which updates the capabilities list accordingly. Once the capabilities are enabled or reconfigured, the capability adjusting device 638 sends an ACK to the AC 602.

[0059] Although the features and elements of the present invention are described in the preferred embodiments in particular combinations, each feature or element can be used alone without the other features and elements of the preferred embodiments or in various combinations with or without other features and elements of the present invention. Although the various embodiments discussed above are described with reference to certain layers, it should be understood that any of the embodiments can be implemented in any layer or any combination of layers. Further, the features and elements of the present invention may be implemented on a single integrated circuit, such as an application specific integrated circuit (ASIC), multiple ICs, discrete components, or a combination of discrete components and ICs. Moreover, the present invention may be implemented in any type of wireless communication system.

*

*

*

CLAIMS

What is claimed is:

1. A method for configuring nodes in a wireless local area network including an access controller (AC) and an access point (AP), the method comprising the steps of:

 sending a functionality inquiry from the AC to the AP;
 responding to the functionality inquiry by the AP, the response including functional capabilities of the AP;
 mapping the functional capabilities of the AP at the AC;
 determining whether a capability conflict exists between the capabilities of the AP and the capabilities of the AC, wherein a conflict can include redundant capabilities between the AP and the AC; and
 resolving the capability conflict.

2. The method according to claim 1, wherein the mapping step includes storing the capabilities of the AP at the AC.

3. The method according to claim 1, wherein the resolving step includes disabling conflicting or redundant capabilities at the AP.

4. The method according to claim 3, wherein the disabling step includes sending a capabilities disable message from the AC to the AP.

5. The method according to claim 4, further comprising the step of:
 sending an acknowledgement from the AP to the AC upon successfully disabling the conflicting capabilities.

6. The method according to claim 1, wherein the resolving step includes reconfiguring conflicting or redundant capabilities at the AP.

7. The method according to claim 6, wherein the reconfiguring step includes sending a capabilities reconfiguration message from the AC to the AP.

8. The method according to claim 7, further comprising the step of: sending an acknowledgement from the AP to the AC upon successfully reconfiguring the conflicting capabilities.

9. The method according to claim 1, wherein if the AP does not respond functionality inquiry, the AC maps the AP capabilities as if the AP had all possible capabilities.

10. The method according to claim 9, further comprising the steps of: setting a timer, the timer being sent when the functionality inquiry is sent; determining whether the timer has expired; and mapping the AP capabilities as if the AP had all possible capabilities, the mapping being performed after the timer expires.

11. The method according to claim 9, further comprising the step of: halting support of services to the AP.

12. An access controller (AC) for configuring nodes in a wireless local area network, comprising:

a transmitter/receiver;

an inquiry device in communication with said transmitter/receiver, said inquiry device configured to send functionality inquiry messages to an access point (AP) in the network;

a capability mapping device in communication with said transmitter/receiver, said capability mapping device configured to receive functionality inquiry reply messages from the AP and to map the functionality capabilities of the AP; and

a capability evaluating device in communication with said capability mapping device, said capability evaluating device configured to determine if any functionality capability conflicts exist between the AP and the AC.

13. The AC according to claim 12, wherein said capability evaluating device is in communication with said transmitter/receiver and is further configured to send capability disabling or reconfiguring messages to the AP.

14. The AC according to claim 12, further comprising:

a timer in communication with said inquiry device and said capability mapping device, said timer being started when said inquiry device sends a functionality inquiry message, said timer being configured to alert said capability mapping device when said timer expires.

15. The AC according to claim 14, wherein said capability mapping device, upon receiving the alert from said timer, is configured to determine if the reply message has been received from the AP and if no reply message has been received, said capability mapping device is further configured to map the AP capabilities as if the AP had all possible capabilities.