Aerohive NETWORKS

solution brief

Radio Resource Management in HiveOS

Introduction

As enterprise wireless LANs have grown from a curiosity providing access in conference rooms into the primary method of connectivity, the number of users supported by those networks has increased. Real-time radio management supports high user densities by balancing user traffic loads between several APs and reserving airtime for what users judge a network experience by: the ability to handle their traffic.

Radio conditions change rapidly, and responding to a constantly-shifting environment requires that a network have the agility to react quickly. Aerohive's cooperative control architecture was designed to enable distributed algorithms to support real-time processing, and this capability is used extensively within the radio management subsystem of HiveOS, the Aerohive operating system.

Design Goals and Benefits

Several features enable HiveOS to serve large networks with many users. At the core, these high-density deployment features promote efficient use of airtime for transmission and attempt to maximize overall network throughput.

Each radio across the entire network optimizes use of radio spectrum by suppressing excess management transmissions and steering clients to the less-loaded 5 GHz radio band. Real-time communications between HiveAPs enable the network to direct client devices towards the least-loaded HiveAP so that total network throughput is maximized.

In addition to client management, network throughput is maximized by the Aerohive Channel Selection Protocol (ACSP), a distributed algorithm for minimizing interference between APs.

High Density Network Support for Client Devices

Early wireless networks were a luxury that used a minimal number of APs to provide network access over a large area. As more users have discovered the benefits of mobile connectivity, the load on networks has risen dramatically. High anticipated network usage means that networks are being designed around user throughput per unit of area; instead of covering square feet, the network planner must now focus on the number of megabits per second per square foot that can be expected. To meet the challenge of building a high-density wireless LAN, HiveOS offers several capabilities that reduce demands on the available airtime and spread the load on the airtime across as many APs as possible.

Airtime Efficiency

The most critical resource in a wireless LAN is the available airtime for transmission, and many of the components of the HiveOS radio management subsystem are designed to reduce contention for airtime.

Band steering. Moving user traffic to the 5 GHz radio band (802.11a and 802.11na) is a long-standing technique to increase total throughput on an 802.11 network. Total area throughput is limited by the number of channels that can be saturated, so the greater the number of non-overlapping channels are available, the greater the possible throughput. Furthermore, the noise floor on 5 GHz channels is often lower due to the smaller number of unlicensed devices. In a dense deployment, the 5 GHz noise floor can be 10 dB lower than the 2.4 GHz noise floor, which provides an immediate 10 dB boost to client SNR. The higher resulting speeds mean that transmissions may require significantly less airtime for the same amount of payload data.

To move user traffic to the 5 GHz band, the network must discover devices that are capable of 5 GHz operation and encourage its use. Aerohive's cooperative control architecture allows the network to build a list of devices capable of 5 GHz operation that is shared among all HiveAPs in a network. When a HiveAP receives a Probe Request from a device on an 802.11a/na channel, that client's record in the radio management database is marked as "5 GHz-capable." When a HiveAP receives a Probe Request from a 5 GHz-capable client device on an 802.11b/g band, it will hold the response to encourage the client to move to a 5 GHz channel.

Coverage area management. Obtaining high throughput in high density environments requires dividing the physical space into as many small radio coverage areas as possible. Smaller coverage areas assist with capacity by reusing frequencies and by enabling client devices to transmit at high data rates that use small amounts of airtime. In high-density environments, HiveAPs transmit management traffic at the highest data mandatory data rate in the wireless network. High data rates for management traffic, especially Beacon frames, shrinks the effective coverage area of each AP and supports the use of high data rates.

Probe Management. Extensive use of Probe Request and Probe Response frames uses airtime that would be better devoted to handling user data. HiveOS has several operations to reduce the amount of airtime used by Probe frames.

Client devices use 802.11 Probe Request frames to build a list of available networks and maintain information used to determine candidate neighbor APs for transition. To build the list of available networks, client devices will use a mix of broadcast and directed Probe Requests. APs typically respond to all applicable client Probe Requests. In a dense environment, a client that sends several broadcast Probe Requests in quick succession may elicit a small flood of Probe Responses from all APs in the area. When APs are configured for multi-BSSID operation, a Probe Response is generated from each BSSID on the AP.

To reduce the airtime devoted to Probe Request and Probe Response transmissions, Aerohive has implemented enhanced processing of Probe Requests in HiveOS. When configured for high-density operation, a HiveAP receiving a broadcast Probe Request will respond for only one of its BSSIDs instead of responding for each of them. In an environment configured for four-BSSID operation, this reduces Probe Response airtime utilization by three-fourths.

Client devices often send out quick blasts of Probe Requests in succession, obligating the AP to send repeated identical responses. When a client repeats its Probe Request transmissions to a HiveAP in a high-density environment, the HiveAP will not respond if it has sent an identical Probe Response within the last Beacon interval (typically 100 ms).

A link budget threshold augments HiveOS high-density Probe Request processing. In high-density environments, a broadcast Probe Request may be received by several APs at varying distances and thus varying signal qualities. To ensure that clients connect to nearby APs and use as little airtime as possible for transmissions, the network should send responses preferentially from APs that will serve the client with a good signal-to-noise ratio. Aerohive's cooperative control architecture allows the network to coordinate probe responses so that Probe Responses are generated from APs with the largest SNR.

Distributed Client Optimization

In a high-density network, the infrastructure must ensure that client loads are distributed evenly across network elements. Early approaches to load management worked to ensure that client loads were spread across radio bands by steering devices to the relatively less loaded 5 GHz band.

In a crowded environment, however, band steering is insufficient. In high-density environments, a client device may be served equally well by several APs with similar link quality. Many client devices use implicit link quality measurements such as collisions, CRC errors, and frame losses as inputs into the roaming algorithm. Crowded radio spectrum can cause poor link parameters, and clients may interpret low link quality as an implicit signal to transition to adjacent APs. In a crowded environment, however, adjacent APs are likely to have similar loads and a transition to a neighboring AP will do very little to improve performance for user traffic. The central problem of load balancing is to distribute the load among a group of APs and remove incentives for clients to roam to similarly-loaded neighbors.

Basic Client Optimization Operations

To build a load balancing algorithm that works across multiple APs and is extensible to all radio bands, HiveOS uses two major operations that allow individual HiveAPs to assert ownership of a client for a limited duration based on the infrastructure view of a client.

The first operation "anchors" a client to a HiveAP for a limited period of time. By using real-time Probe Request processing, a HiveAP takes ownership of a client session. Upon receipt of a Probe Request frame, a HiveAP compares signal quality with its neighbors to decide which AP should transmit a Probe Response. After a client has associated, a HiveAP can take control of a client session by informing its neighbors that it will manage network access for the client device. As long as a session is owned by a HiveAP, neighboring HiveAPs will not respond to Probe Request frames or Association Request frames.

Ownership is released periodically to allow clients to test connections to neighboring APs. As clients move around, it may be that the best HiveAP changes, and periodic releases allow a client to compare several Probe Responses. Clients are also released when the signal-to-noise ratio (SNR) falls below a threshold that indicates that another HiveAP may offer better service. Ownership may also be released as part of a distributed load-balancing algorithm.

Applications of Distributed Client Optimization

First-generation approaches to client management actively disconnect client devices from the network with Deauthentication frames. Not all clients respond to active disconnections by seeking out alternative access points, and many clients enter a “death spiral” where the network attempts to push them to an alternative access point and the clients refuse to move.

The basic operations for asserting real-time ownership of a session and releasing ownership of a session can be used to optimize client connectivity in two ways.

Association and admission control. When a HiveAP receives a request for association, it will accept the association only when no other HiveAPs “own” the session. Requests for association are accepted only when a client device can be served with an acceptable SNR.

Load balancing. Mechanisms to control client associations by managing associations can provide load balancing as well. When a client attempts to associate with a HiveAP, it will only be accepted if the client is not owned by its current HiveAP, it will have acceptable SNR for reasonable operating speed, and the new HiveAP has a low enough load to accept new clients. All HiveAPs have real-time information on the load information of its neighbors, and will avoid responding to requests that can be served from more lightly loaded neighbors. Clients are also released when a HiveAP has become overloaded because it allows other clients to roam and the distributed load balancing algorithm to spread clients among neighboring APs.

Automatic Radio Management Capabilities

Obtaining maximum throughput in an 802.11 network requires efficient re-use of radio spectrum. Aerohive has designed a distributed radio management protocol to enable effective channel re-use in dense environments called the Aerohive Channel Selection Protocol (ACSP). ACSP was designed to offer a stable network experience for client devices while minimizing overlapping channel assignments that sap a network of throughput. To successfully maintain a stable network, a channel selection algorithm should work without imposing an undue computational burden on the network elements, offer repeatable results, and converge quickly to a channel plan. The network should react only to persistent interference rather than transient spikes.

ACSP defines the “best” channel as the channel with the lowest interference because minimizing interference maximizes throughput. To reduce interference, ACSP holds transmission power to the minimum requirement for coverage to an AP’s neighbor and monitors an extensive list of channel parameters to identify the clearest channel for communication with user devices.

Channel cost. At the core of interference minimization is calculation of channel cost. HiveAPs select the channel with the lowest cost, and the cost of a channel is updated in real time based on radio link performance learned from monitoring and background scans. HiveAPs select a channel based on a probability algorithm. This white paper converts costs into approximate changes in the probability that a channel will be selected.

Avoiding co-channel interference is a major goal of automatic channel selection. When a HiveAP assembles information about its radio environment, the cost of a channel is increased when other APs are detected. Each detected neighbor AP increases the cost, but a diminishing scale is used to reflect that in a dense environment with many neighbors, the difference of a single additional neighbor AP is not something that affects throughput as dramatically as adding the first neighbor.

Two static factors are taken into account before running the calculations. Channels in the non-overlapping channel set are given a preference that biases selection in favor of a non-overlapping channel of approximately 2%. Within the 5 GHz band, some channels have more restricted transmission power. To provide headroom for adjusting power, channels with high allowed transmission power are given a bonus ranging from 0.5% to 1.2%.

Dynamic measurement. Channel cost is adjusted based on radio measurements by the HiveAP. Maximum throughput is obtained by minimizing the time of all transmissions, even those of adjacent APs. Interference reduces throughput by making transmission time unavailable. The major radio measurements used in ACSP are channel utilization, the utilization of overlapping APs, and the CRC error rate. HiveAPs directly measure the available airtime at each radio, and select channels based on finding the maximum capacity.

Each factor is assessed a penalty that reduces the likelihood of selecting the channel of up to 3.5%. Thus, if all three factors are poor, the total penalty will be over 10%. Lightly-loaded channels with low error rates are not penalized. Penalties for poor performance increase linearly until reaching a maximum value. The following table shows the highest acceptable values to avoid additional channel costs, as well as the beginning of the maximum penalty band for each factor.

	Channel utilization	**Overlapping utilization**	**CRC rate**
No cost (no penalty)	< 10%	< 5%	< 5%
Maximum cost (-3.5%)	> 50%	> 20%	> 20%

Operational details. When a HiveAP is powered on and initializes channel selection procedures, it enters the Scan state. In the Scan state, each radio tunes to each available channel and measures the radio quality to compute a channel cost. Other than during start-up operation, the Scan state is rarely used. After completing the scan, if it is the only HiveAP, it selects the lowest cost channel and enters the Run state, which is the normal operating state.

If a newly-started HiveAP discovers other HiveAPs, it will instead enter the Listen state. In the Listen state, the HiveAP tries to confirm its channel selection before entering the Run state for normal operation. Each HiveAP in the Listen state will advertise its neighbor count, and the AP with the most neighbors wins the right to use the channel. Multiple HiveAPs can boot simultaneously and enter the Listen state, but only one will move from Listen to Run. All others return to the Scan state to find a new channel.

To allow new HiveAPs to smoothly join an existing network, an Aerohive network selects one exiting HiveAP in the network to be the *arbiter* and places the arbiter in charge of allowing frequency re-use. When multiple APs are in the Listen state and wish to enter the Run state and begin accepting client sessions, they must first find the arbiter and request the right to enter the Run state on a particular channel.

Arbiters may receive multiple requests for channel allocation. When several messages are received, the arbiter will grant only one request. Each channel request includes the channel cost as perceived by the AP wishing to join the network. The arbiter grants the request with the highest cost from the APs trying to join the network to minimize cost over the entire network. All other APs are sent back to the Scan state to find a new channel.
.

Summary

Adaptation of a wireless network to a mercurial radio environment requires that APs react in real-time to changing conditions. Protocols must be carefully designed for stability in many types of environments, and need to offer stable and predictable results when used in increasingly dense deployments. Aerohive's distributed protocols enable a network to make rapid decisions on client requests to enhance airtime efficiency and provide fluid load-balancing, while maximizing throughput to large numbers of clients by selecting the best-performing channel layout.

About Aerohive

Aerohive Networks reduces the cost and complexity of today's networks with cloud-enabled, distributed Wi-Fi and routing solutions for enterprises and medium sized companies including branch offices and teleworkers. Aerohive's award-winning cooperative control Wi-Fi architecture, public or private cloud-enabled network management, routing and VPN solutions eliminate costly controllers and single points of failure. This gives its customers mission critical reliability with granular security and policy enforcement and the ability to start small and expand without limitations. Aerohive was founded in 2006 and is headquartered in Sunnyvale, Calif. The company's investors include Kleiner Perkins Caufield & Byers, Lightspeed Venture Partners, Northern Light Venture Capital and New Enterprise Associates, Inc. (NEA).

Aerohive™ NETWORKS

Corporate Headquarters
Aerohive Networks, Inc.
330 Gibraltar Drive
Sunnyvale, California 94089 USA
Phone: 408.510.6100
Toll Free: 1.866.918.9918
Fax: 408.510.6199
info@aerohive.com
www.aerohive.com

EMEA Headquarters
Aerohive Networks Europe LTD
Sequel House
The Hart
Surrey, UK GU9 7HW
+44 (0)1252 736590
Fax: +44 (0)1252711901

SB1101905