ARUBA™
The Mobile Edge Company

White Paper

Location & Tracking on the Mobile Edge

Peter Thornycroft
Product Marketing Manager

Introduction – The mobile edge

The fixed edge of the enterprise network is giving way to the 'mobile edge' – a new way of connecting users to information. The mobile edge transcends the enterprise network perimeter, appearing wherever the user needs access to information – on the campus, in a regional or branch office, at retail outlets, at home, and on the road. The mobile edge makes use of existing high-speed networks – the corporate LAN, the corporate WAN, and the Internet. It does not replace these existing networks, but deploys on top of them as a service overlay, preventing disruptive equipment changes and preserving investment. At the same time, the mobile edge permits a large-scale reduction in cost for the wired network through port consolidation, reduced equipment needs, reduced power requirements, and the elimination of move/add/change costs.

The mobile edge embraces mobility, security and convergence, solving the problems of all three areas simultaneously while delivering dramatic economic savings over traditional fixed networks.

One consequence of mobility is a need to determine the location of users and network assets, or of intruders. Precise location information is very valuable in itself, and it can also be used to enable other location-based services. Uses of location tracking and location-based services on the mobile edge of the enterprise network include locating users, intruders, rogue access points (APs) or other security threats, assets (such as medical equipment in a hospital), and sources of radio frequency (RF) interference.

This paper examines the various technologies used for location and location-based services in WLANs, while comparing their effectiveness and complexity. It will be seen that the mobile edge architecture allows location services to be added to a WLAN economically but without compromising the capacity or performance of the network.

Applications for location information

Location information is valuable in many areas, including security, RF coverage assessment, network management and information for users. Since accurate location information as a service of WLANs is a recent feature, some of the more sophisticated applications are just emerging. The following is a short list of examples:

- Locating 'rogue' APs, which are a security risk
- Locating sources of radio interference so they can be avoided or remedied
- Locating particular clients, either people or equipment
- Locating a Wi-Fi phone making an emergency call (E911)
- Identifying whether a client is inside or outside the building, or a public space
- Providing floor plans or directions for the user's current location
- Providing the network manager with tools to visualize the RF space in a building.

Requirements for location accuracy

When dealing with location tracking in the corporate context, the first question is what degree of accuracy is required. When considering the list of corporate applications above, the practical goal is usually to identify the correct floor of the building, and the correct cubicle – as a goal, the predicted location should be within 10 feet (3 meters) of the true location.

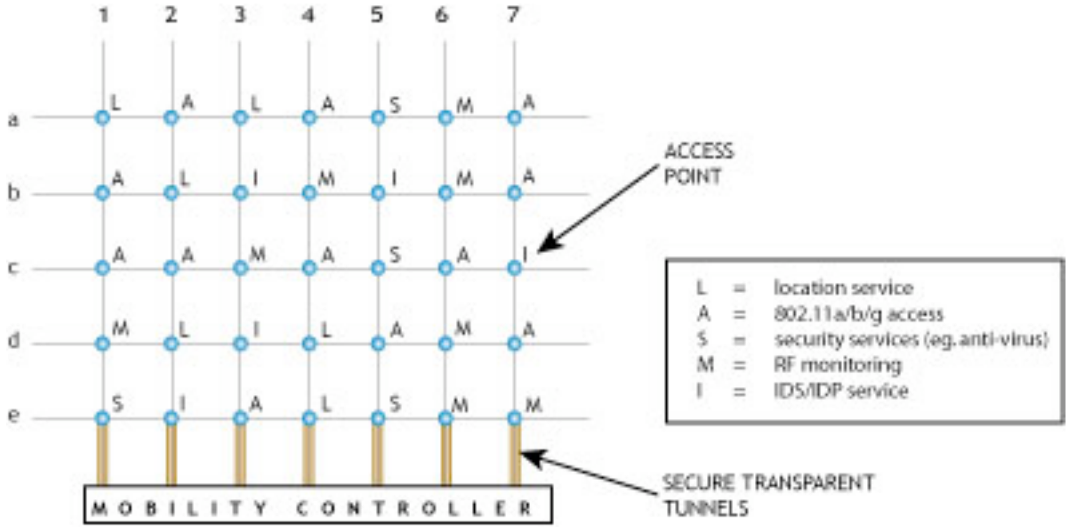

Figure 1

A mobile edge network deployment – the Wireless grid comprises densely-spaced APs delivering services on-demand. Secure tunnels are established transparently over any IP network from each AP to a central mobility controller that dynamically assigns the function of each AP.

Methods for distance and location estimation

This is a summary of available technology for determining client location in wireless networks. There are four methods of interest:

- GPS
- Time Difference of Arrival
- ‘Nearest AP’
- Received signal strength

GPS

A network of satellites transmits signals which the receiver processes to provide a location, in three dimensions, accurate to about 30 feet (10 meters). Since GPS receiver chips are now relatively inexpensive, this method is used in cellular networks, providing data for location-based services. However, GPS is not used in WLANs because satellite reception in buildings is not reliable, and the cost of a GPS chipset is still significant compared to the cost of a Wi-Fi client.

Time Difference of Arrival (TDOA)

This technique sometimes involves putting a special modulation on the radio signal from the tag: other techniques are able to use the Wi-Fi signal without enhancement. However, they all use a special receiver circuit in the infrastructure to measure the relative time-delay of signals arriving at different APs. Since time is proportional to the distance traveled, this allows estimation of the distance to each audible AP and the location of the client can be derived. Some vendors are using this method in outdoor systems, such as ports and airports. However, the standard Wi-Fi chipsets used in APs cannot implement it so this method is not often used with indoor WLAN infrastructure.

Nearest AP

Properly, this should be named ‘associated AP.’ It assumes that a client connected to (‘associated with’) an AP is closer to that AP than to any other: hence it must be located within a radius of that AP (or in its ‘cell’ area). The associated AP in a WLAN is easily determined, but in traditional networks where APs are spaced on the order of 60 feet (20 metres) apart, the range of uncertainty is too wide to be useful for many location-based services.

Received signal strength

It is the received signal strength method, with some improvements, that is used in WLANs today. At its simplest, this method uses the signal strength measured at an AP for the client device to be tracked; the RF power loss between transmitter and receiver is related to the distance traveled. Received signal strength indication (RSSI) results in an estimate of distance but not direction.

To determine the location of a client, at least three APs in the network must be able to detect and measure the client’s signal strength. This is the basic RF triangulation method. Its accuracy is dependent on two factors:

- The accuracy of each measurement. There are many sources of inaccuracy in deriving distance estimates from an RSSI measurement. The RSSI figure itself may be accurate, but especially with mobile clients, the orientation of the antenna, whether there is a human body in the path, the presence of walls, cubicles or metal objects can all have a large effect on the path loss. This translates into uncertainty about the exact distance.

- The longer the distance to be measured, the more inaccurate the measurement. For the purposes of this paper we assume that the accuracy of an RSSI-based distance estimate is proportional to the actual distance from AP to client: doubling the distance from AP to client tends to double the uncertainty of the measurement.

At this point it is useful to consider some practical figures. In a traditional WLAN, APs are spaced at about 60 feet. Thus the typical RSSI-based measurement (AP to client) would be about 50 feet, with a range of uncertainty. Therefore, a measured figure of 50 feet could in fact be anywhere from perhaps 25 to 75 feet. This raw figure does not meet the requirement stated earlier that location should be known to the correct floor of the building and the correct cubicle. This has led vendors to enhance the basic triangulation method, but before considering enhancements, there are some practical difficulties in deriving even basic RSSI triangulation measurements.

Practical difficulties with RSSI triangulation

In WLANs the RF triangulation method works generally as explained above, but there are some complications.

WLANs are cellular, where neighboring APs operate on different RF frequencies (channels) to avoid interference. The Wi-Fi medium access control layer (MAC) allows any station in a basic service set (BSS: we will use the term ‘cell’ in this paper) to transmit at any time. Therefore all stations (including the AP) should be listening on the cell’s RF channel all the time, to avoid missing transmissions.

For an AP to provide an RSSI measurement for a client in its cell is an easy task. However, to measure a client in a neighboring cell, the AP must break off its service, switch channels, listen until it hears a transmission from the desired client, measure it and then return to its own channel to resume service. The AP ‘steals’ time away from its own cell in order to locate a client in an adjacent cell. This has little effect on service when a cell is lightly loaded: a client transmitting while the AP switches channels will not receive an acknowledgement, and will retransmit its data. However, time stealing can cause errors in multi-media traffic (voice and data) that may not be amenable to retransmission, and it seriously reduces the peak traffic level an AP can support (tests indicate a 30 percent reduction, which worsens as the number of location measurements increases).

This method has been adopted by a number of WLAN vendors because they have little alternative, due to another factor: the maximum range of a Wi-Fi connection in a building is on the order of 100 feet. Thus, APs further than that distance from the client will not be able to make a measurement, and as we have seen, the traditional AP spacing is about 60 feet. Since only a few APs are within range of a given location, they will all have to be used for RF measurements to assist with client location.

Some vendors have changed the architecture somewhat by adding software to the client to continuously measure RSSI from the AP, and to report this measurement back to a location platform within the network. This approach has some advantages, such as offloading the processing required from the AP to the client, but has not been generally adopted, as modern WLANs periodically recalculate and change their RF plans to react to environmental changes, so the transmitted signal strengths change over time.

Enhancements to the RSSI triangulation method

Air Monitors

The paragraphs above explained the use of time-stealing APs to monitor other channels while nominally providing coverage of their own cell. An alternative technique is to deploy dedicated RF monitors named 'Air Monitors' (AMs). AMs are identical to APs (the same hardware and software), but they are configured permanently in 'listening' mode. This is a very useful capability, because the AMs contribute not just to location accuracy but also by improving security coverage, detecting RF sources that may be security risks or interferers. The drawback of using dedicated AMs is that they add to the capital costs of the network.

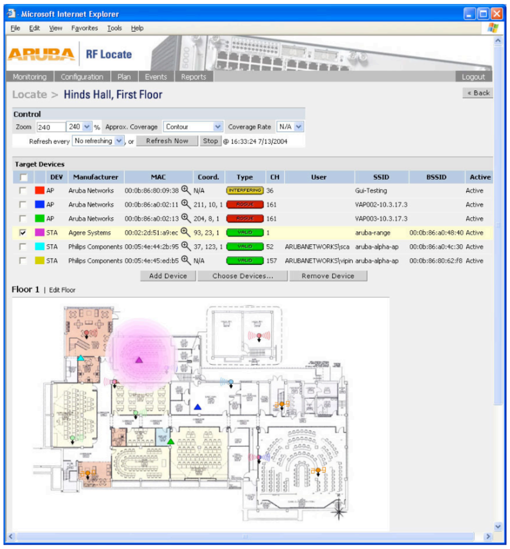

Locate > Hinds Hall, First Floor

Control

Zoom 240 240 % Approx. Coverage Contour Coverage Rate N/A

Refresh every No refreshing, or Refresh Now Stop @ 16:33:24 7/13/2004

Target Devices

□		DEV	Manufacturer	MAC	Coord.	Type	CH	User	SSID	BSSID	Active
□	🟥	AP	Aruba Networks	00:0b:86:80:09:38	N/A	INTERFERING	36		Gui-Testing		Active
□	🟦	AP	Aruba Networks	00:0b:86:a0:02:11	211, 10, 1	ROGUE	161		VAP002-10.3.17.3		Active
□	🟩	AP	Aruba Networks	00:0b:86:a0:02:13	204, 8, 1	ROGUE	161		VAP003-10.3.17.3		Active
☑	🟪	STA	Agere Systems	00:02:2d:51:a9:ec	93, 23, 1	VALID	1		aruba-range	00:0b:86:a0:48:40	Active
□	■	STA	Philips Components	00:05:4e:44:2b:95	37, 123, 1	VALID	52	ARUBANETWORKS\sca	aruba-alpha-ap	00:0b:86:a0:4c:30	Active
□	🟨	STA	Philips Components	00:05:4e:45:ed:b5	N/A	VALID	157	ARUBANETWORKS\vipin	aruba-alpha-ap	00:0b:86:80:62:f8	Active

Floor 1

Figure 2 Multi-point triangulation pinpoints any 802.11 device

Floor-plans and estimation

Another technique involves manual input to assist with the RF model. Since the APs work in 'RF space' and corporations in physical space, it is usual to require the network manager to import or draw a building's floor plan and place the APs in their locations, so that triangulation measurements (which are relative to the APs' positions) can be translated to physical locations. Some vendors take this a step further, requiring the network manager to classify the various cubicle, wall, floor and other obstructions in the building. (For instance, an exterior concrete wall might attenuate a signal three or four times more than an interior wall.) This somewhat improves distance measurement accuracy, but at the expense of the administration time involved.

Walkarounds and RF fingerprinting

Another method for improving location accuracy is a walkaround calibration. Here, an engineer with a special client device walks around the building, stopping at defined points to allow the network to take measurements. This is a different approach from building modeling, as it allows the network to build up a database of actual measurements and relate them to particular locations. A walkaround improves the accuracy of locations, but with a few shortcomings: it can be time-consuming and expensive, it only improves measurements in those locations actively surveyed, and if the WLAN or the building is modified (APs or furniture reconfigured or moved) it becomes invalid. A walkaround might reduce the uncertainty in our location model from 50 feet to 15 feet: close to the target accuracy, but with considerable cost and complexity; and the accuracy will deteriorate over time.

Statistical analysis

Some vendors have developed systems that take raw RSSI measurements and process them by applying probability theory in conjunction with survey measurements. This approach can provide very accurate results under certain conditions, but it requires a walkaround and usually an extra computing platform in the network. These systems are used in special situations, but not usually in WLANs.

The mobile edge architecture and location measurements

The mobile edge architecture and location tracking technology are unique in a number of respects:

Wireless services are delivered through a ‘wireless grid’ of densely-spaced APs. The network, consisting of mobility controllers and APs, contains all the functionality necessary to provide accurate location measurements.

Once the network is installed, the network manager imports a floorplan for the building, and marks the position of each AP on the plan – the only manual step necessary.

Instead of fingerprinting the building with a walk-around client, or requiring the network manager to mark up the floorplan with RF obstacles and construction materials, the wireless grid is utilized for self-calibration. Each AP in turn is put in monitor mode, taking accurate RSSI measurements for all APs within range – of which there will be many. This allows characterization of the in-building environment, and specific calibration of many sections of the building with real-world infrastructure-based figures. The resulting measurements are built up into a database and used to convert raw RSSI measurements to calibrated location determinations. The result is superior to building material modeling and similar to client calibration, but without the drawbacks: no walkarounds are required, and changes in the RF plan are accounted for by the infrastructure.

The mobile edge network uses both dedicated AMs and also time-stealing of APs: both are involved in location of clients. When a client is to be located, RSSI measurements are taken at all APs and AMs in range, except those providing uninterruptible service (e.g. an active voice call) on another RF channel. The density of APs allows many measurements to be made, contributing to more accurate location in three ways.

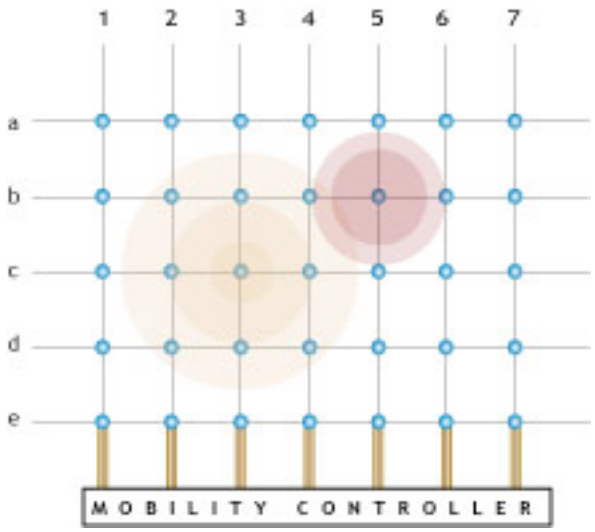

Figure 3

Wireless grid APs self -calibrate to determine RSSI strength, interference sources and location measurements, using APs close to the target. Here C3 transmits while all others calibrate, then B5. Calibration is performed globally as well as locally.

- First, many more APs will be in range of the client, resulting in more measurements and a more accurate location estimate.
- Second, the distance to the nearest APs will be significantly smaller than in a traditional network of sparse APs.
- Third, the many APs that are not at that moment providing wireless service will be constantly in AM mode, enabling long-term, time-averaged RSSI measurement. This contrasts with networks where APs must always steal time-cycles away from providing service in order to measure RSSI from clients on another RF channel. Where time-stealing is the only available method (in the absence of AMs), either the wireless service or the location measurement will suffer from the time-sharing of a single radio resource.

Tools and features for location tracking on the Mobile Edge

The mobile edge is deployed as a grid of densely-spaced APs, enabling the following tools for the network manager:

1. RF coverage visualization. This is a representation of the building showing the different APs and RF channels in use, along with the network software’s estimates of coverage and signal strength.

2. **Client tracking.** The network manager can designate a given client, and the network will locate and track the client across the floorplan. The wireless network can also provide audit information showing historical movements of a particular client and reports of client location and movements.

3. **Location for troubleshooting.** This covers a wide range of activities, which for troubleshooting require knowledge of the location of a problem. If the location is known accurately enough, it may be possible to avoid the cost of an engineer dispatch. If that is not possible, the engineer can save considerable on-site time if the location of the problem is known accurately.

4. **Asset tracking.** Important equipment may be tagged with a commercially-available Wi-Fi tag, allowing it to be tracked like any other client.

5. **Location of rogue APs.** Rogue APs are unauthorized APs, usually brought into the building innocently by employees. However, once connected to the corporate LAN an unprotected rogue AP allows any client access to corporate resources. Rogue APs are readily detected by network AMs and APs. When a rogue AP is detected in the network, an alarm is generated and information about the rogue AP, including its location, is provided to the network manager. The wireless network software also allows a rogue AP to be shut down remotely, either by RF jamming or from the wired side.

6. **Location of hackers.** Malicious hacking attempts are a rare but real threat to a corporate Wi-Fi network, whether 'man-in-the-middle,' 'denial of service' or other attacks. The wireless network has tools that automatically detect and locate such threats, allowing mitigating actions to be directed promptly to the correct location.

7. **Dynamic capacity management.** In conjunction with other RF management algorithms, the mobility controller can detect AP cells with excess load, and dynamically adapt the network to spread the load over adjacent cells.

8. **Location of insecure clients.** Corporate networks generally have strict procedures for maintaining their desktop computing assets. However, with the recent advances in Wi-Fi client security, it is possible that some clients are not using the required levels of security, and thus present risks. The wireless network identifies these clients, and locates them for prompt remediation.

9. **Location of banned clients.** Where a client (or its owner) is forbidden from connecting to the network, such attempts can be tracked and located, assuring good network security.

10. **Location of interferers.** Certain equipment is known to affect WLANs, especially in the 2.4GHz band for 802.11b or 802.11g: some cordless phones and microwave ovens, for instance. Such interference may be intermittent, which makes it difficult to diagnose in traditional WLANs. The wireless network continuously monitors for sources of interference, providing their location to the network manager.

11. Location for emergency phone calls. As Wi-Fi telephony grows, the need for locating these handsets has emerged. Many states have E911 regulations specifying that an emergency phone call must have its location passed with the call to the public safety access point (PSAP).

12. Location for usage classification. The wireless network can deliver location information, allowing a network manager to identify whether clients are inside or outside the building, or in a public reception area rather than a lab area. This information can be used, for instance, to prevent users external to the building from connecting to the network, although they may receive a usable signal.

13. Interfacing to corporate software applications. This is an emerging area, where for instance a wireless PDA client could call up a floorplan of the immediate surroundings to assist in finding a particular conference room. Wi-Fi tagged assets could be automatically tracked by reservations software or an ERM application.

Conclusion

The mobile edge enables connections across a space rather than at the end of a wire, making location of objects within the edge an important new feature. This paper has given an overview of the different methods that can be used for location on the mobile edge, comparing their accuracy

Although many alternatives exist, most successful indoor location techniques are based on the RSSI triangulation method. But basic RSSI triangulation does not provide sufficient accuracy for many of the users of location information.

While techniques such as analysis of building material and walkaround calibration can improve the accuracy of RSSI measurements, they add considerable expense and complexity to the network installation, and the improvement in accuracy erodes over time, as the environment changes.

A recent advance in network deployment – the mobile edge architecture, deployed with a dense 'wireless grid' of APs - allows with self-calibration techniques and enables location of RF sources to an accuracy of the dimension of a cubicle.

Additionally mobility software can provide a rich integrated toolset for managing the network, reducing or eliminating the need for costly ancillary equipment or tools for these purposes.

Location measurements are useful in themselves, and as input to location-based services. The paper concluded with a survey of the ways accurate location information enhances the effectiveness of the wireless LAN.

About Aruba Wireless Networks, Inc.

Aruba Wireless Networks is a fast-growing enterprise infrastructure company enabling the Mobile Edge, an evolutionary new network architecture that addresses three top concerns of IT managers—mobility, security, and convergence. The Mobile Edge extends the reach of enterprise networks, providing secure access to information and voice services anywhere a user needs them, enabling new applications, allowing organizations to compete more effectively, and bringing about dramatic economic benefits. To deliver the Mobile Edge, Aruba manufactures and markets a complete line of fixed and modular mobility controllers, wired and wireless access points, and an advanced mobility software suite. Privately-held and based in Sunnyvale, California, Aruba has operations in the United States, Europe, the Middle East, and Asia Pacific, and employs staff around the world. To learn more, visit Aruba at http://www.arubanetworks.com

10.14.ltme.pt.10.1