Self-Management of Wireless Base Stations

Kai Zimmermann, Lars Eggert and Marcus Brunner
NEC Europe Ltd., Network Laboratories, Kurfürstenanlage 36, 69115 Heidelberg, Germany
{zimmermann|eggert|brunner}@netlab.nec.de

***Abstract*— This paper presents an autonomous, self-organizing and decentralized configuration and management system for a group of base stations in wireless networks. Compared to existing systems, where a central node computes and disseminates management information, the system's decentralized operation improves reliability by eliminating central points of failure and can decrease convergence times for large installations by enabling localized reconfiguration. A second novel feature is the integration of external, third-party input into the distributed configuration algorithm, improving the quality of the configuration result and convergence times. The paper describes the decentralized management approach and its prototype implementation. It also presents preliminary experimental results that illustrate the system's scalability properties.**

I. Introduction and Motivation

The Internet has included wireless links almost since its beginning. Satellite-based communication [5], early radio networks such as ALOHANET [1] or packet radio networks [10] provided connectivity without wires in the early Internet. However, these systems were not available or affordable by many users and consequently not very prevalent.

Wireless local area networks (WLANs) based on the IEEE 802.11 family of standards [8] started to provide mass-market wireless connectivity ten years ago and are still becoming increasingly more popular. Today, third-generation cellular networks are becoming a second alternative for wireless access across longer distances. In the future, other wireless access technologies, such as WiMax, various ultrawideband technologies or future-generation cellular networks [11], will provide even more users with a variety of different wireless access technologies. Wireless connectivity is becoming ubiquitous.

Providing wireless connectivity to a larger geographic area requires deployment of multiple base stations, each of which covers a fraction of the total region. This is independent of the specific network technology used to provide this connectivity. The most popular method of deploying these base stations is as access networks that extend a wired core network by a single wireless hop. Deployment of multi-hop wireless access networks is also possible, but less popular due to the intrinsic complexities of this approach, for example, self-interference when forwarding across wireless links.

Once deployed, a group of wireless base stations requires continuous management to provide a uniform service environment, recover from faults or maximize overall performance. Manual management of each base station is only possible for very small groups. As the group of deployed base stations grows, automated management becomes a necessity.

Very few wireless network technologies include adequate management mechanisms. Even if they do, these systems typically only focus on physical or link-specific characteristics and do not manage higher-layer properties. For example, WLAN networks do not include any management functions. Each base station provides an area of wireless coverage that is completely isolated from its neighbors and each must be managed independently of its neighbors.

Existing approaches to managing WLAN access networks that consist of multiple base stations are primarily centralized. A central master system periodically computes a global configuration for the whole network based on available information. It pushes this configuration out to the individual base stations in a piecemeal fashion or they pull their respective configurations in from the master. A centralized approach has several disadvantages. First, it creates a central point of failure. Failure of the master can make the whole system unusable. Second, a central master limits scalability due to processing and communication overheads, especially in environments that require frequent configuration changes. Third, it complicates the system, because this approach introduces additional infrastructure, *i.e.*, the central master.

This paper presents a decentralized approach for management of a group of collaborating base stations. The individual base stations aggregate and share network information. They implement a distributed algorithm that computes a local configuration at each base station based on the shared information such that the overall network-wide configuration is consistent. Although the current prototype described in this paper focuses on managing a WLAN access network, the general mechanism is applicable to other wireless and wired access technologies.

A decentralized approach is inherently more resilient to failure. Because each base station computes a local configuration based on exchanging information with its neighbors, it can react locally to changes in its local environment without involving a central master. Furthermore, a decentralized system allows the group of base stations affected by a local change in their environment to react locally. This can improve scalability, convergence time and communication overhead. A decentralized approach, however, also introduces challenges, for example, to guarantee convergence, establish system-wide consistency of the configuration, as well as trust issues between otherwise independent nodes.

The following Section II briefly describes existing approaches for wireless base station management. Section III presents the decentralized solution that is the key contribution of this paper and its prototype implementation and Section IV descri-

bes the inclusion of external information into the management system. Finally, Section V briefly illustrates some preliminary scalability properties of the prototype implementation and discusses future work.

II. Related Work

Two different paradigms exist in managing wireless networks. Centralized systems use a single master device to configure a group of base stations or a small group of cooperating master devices for very large networks. The second approach is decentralized. Here, the individual base stations are autonomous entities that collaborate as peers to arrive at a consistent, system-wide configuration. This section describes existing approaches in both areas and also briefly describe a third, hybrid approach.

Several companies provide centralized management solutions for groups of wireless base stations [3] [4]. The majority of these systems implement link-layer "wireless switches" that connect base stations that act as wireless bridges to a switched wired network. The link-layer switch implements the management component. This centralized, link-layer approach offers traffic and channel management, policy, bandwidth and access control. Additionally, this solution provides intrinsic roaming, because the management device can handle client movement at the link layer.

Centralized link-layer solutions also have drawbacks. Link-layer broadcast domains cannot arbitrarily grow due to the scalability issues associated with broadcast traffic. Additionally, the topology of the wired network may not allow direct connection of the management system to the base stations. Centralized network-layer solutions address this shortcoming.

Decentralized management solutions are popular to configure mobile *ad hoc* networks. These management systems typically focus on the challenging task of enabling peer-to-peer communication in highly dynamic, mobile environments [2] [9]. In contrast, the decentralized solution presented in this paper focuses on configuring a stationary wireless access network for mobile clients, with the goal of improving efficiency and performance.

Besides centralized and decentralized approaches, hybrid approaches exist [15]. These systems push some functionality from a central system into the base stations, which are therefore slightly more complex than the simple wireless bridges of centralized approaches. Although hybrid systems offer minor scalability increases, they do not completely address the drawbacks of centralized systems.

III. Decentralized Base Station Management

This section describes the decentralized solution for wireless network management, including assumptions, the basic idea, security issues and several example applications.

A base station in the decentralized management system has to fulfill several requirements. Each base station is a full-fledged IP router for its delegated IP subnet, able to operate stand-alone. It needs at least two network interfaces; one to provide wireless services to its clients and a second interface (wired or wireless) for uplink connectivity. Additional interfaces, when present, can act as probe interfaces, provide multi-homed uplink connectivity or offer additional client connectivity on different channels or link protocols.

Base stations automatically distribute the available address space among themselves, configure subnets for client connectivity on their wireless interfaces and configure the addressing of their wired uplinks. In the current prototype, IP auto-configuration occurs through a separate mechanism that was an earlier research effort [12]; the next revision will integrate this process.

The current system uses X.509 certificates signed by a common certification authority to establish trust between base stations and to indicate access network membership. Consequently, each base station has an individual certificate. This certificate has a second function: a hash of the certificate provides a statistically unique identifier for each base station.

A. Basic Operation

Although each base station is able to operate in a stand-alone fashion, integrated management of a group of base stations that provide connectivity to a geographic region requires collaboration. This collaboration occurs through periodic information exchange across the uplink interfaces, which allows each individual base station to adapt its local configuration consistently with its peers.

When a base station starts up, it performs a probing phase – after a brief randomized de-synchronization delay – before configuring itself to provide connectivity to clients. During the first part of this probing phase, it auto-configures its IP connectivity, *i.e.*, it obtains a subnet delegation to serve clients and configures its uplink interface [12]. The base station then performs a channel scan to detect other base stations in its immediate neighborhood and determines their identifiers. It then contacts these neighbor stations over its wired uplink interface and, after successful authentication and authorization, integrates itself into the network-wide information exchange.

Once configured, the base station periodically performs a channel scan to detect changes in its environment. Because client connectivity is disrupted during the scan, the base station performs this scan less frequently when it provides uplink connectivity to clients and more frequently when it does not. It could also use a dedicated probe interface for this purpose, if available. The base station also starts to participate in global and local information exchanges with its peers.

The system uses different kinds of exchange mechanisms for different kinds of information. Information that is globally important, such as encryption parameters or attack status, is disseminated throughout the network using an epidemic communication mechanism. Information that is of local significance only, such as radio frequencies, transmit power or link utilization, is only disseminated locally among the affected neighboring base stations. This differentiation by information type improves the scalability of the system.

The information that each station maintains falls into three different categories. *Private* information, such as logs, is never

disseminated. A base station disseminates *local* information, such as its current channel, transmit power or utilization, to its neighbors, *i.e.*, other base stations within wireless range. This allows a group of neighbors to adapt their configurations in response to local events. A base station periodically disseminates updates about its local state to its neighbors every few seconds and likewise receives their updates.

A third kind of information requires *global* dissemination to all cooperating base stations. System-wide parameters, such as wireless protocol, security parameters or attack status are examples of such global information. The system disseminates global information using epidemic communication [6]. Instead of broadcasting or multicasting such updates to global state, they are piggy-backed onto the periodic local information exchanges between neighbors. This technique prevents broadcast storms when global state updates are frequent.

Disseminating a global configuration change throughout the network in a consistent manner requires transactional semantics. This is a well-known challenge in distributed networks and a wide variety of approaches exist [14]. The current system implements a very simple method of guaranteeing global consistency – election of a central locking service. Future revisions may replace this method with a more scalable variant.

B. Management System Functionality

The current management system specifically targets WLAN access networks. It coordinates radio properties, such as channel use or transmission power, among a group of neighboring base stations. It also implements system-wide functions, such as load balancing. By exchanging utilization information, neighboring base stations can distribute client load by raising or lowering transmission power or link speeds. Overloaded base stations, for example, can push clients at the edge of their range onto other base stations by lowering their transmission power,

A second example of a system-wide management function is self-protection through rogue access point detection. Rogue access points are base stations located within range of the managed WLAN access network that are not part of it. They are potential security threats, because they may attempt to spoof clients into associating with them instead of the actual access network and then intercept their traffic. The current management system detects rogue access points and disseminates their presence throughout the system, preventing any legitimate base station from communicating with these nodes.

A third system function provides a means to obtain a global view of the system, *i.e.*, retrieve local information from all participating base stations of the system, for logging, administrative and monitoring purposes. The decentralized measurement system can support this functionality without the need of an explicit logging function. Instead, a *virtual neighbor* can be created by disseminating its ID throughout the system and inserting it into each base station's neighbor list. The virtual neighbor will receive the local information disseminated by each base station as if it was in radio range of every individual base station simultaneously. The virtual neighbor can aggregate and export this system-wide information for a variety of uses.

C. Security Considerations

A decentralized management system must fulfill several security objectives. First, it must protect sensitive information against unauthorized access. Second, it must protect the distributed configuration algorithm from attacks. Third, it must prevent management functionality to be used as an attack tool, *e.g.*, for flooding attacks. These security aspects are similar to those found in *ad hoc* networks [13].

The use of X.509 certificates and two-way authentication addresses all these security objectives. Traffic encryption protects sensitive information; digital signatures allow verification of the authenticity of management communication, protecting the operation of the distributed algorithm and consequently mitigate the use of management functions for attacks.

Installation of base station certificates and the corresponding certification authority certificates still requires one-time manual configuration of base stations. However, methods for semi-automated certificate configuration – such as physically connecting to a mobile certification authority that auto-installs the required certificates on first boot – can significantly shorten the configuration process. The specifics of such approaches are outside the scope of this paper.

IV. Inclusion of External Information

One challenge for automated configuration of wireless access networks is base stations with overlapping coverage areas that are unable to detect this occurrence because none is within the area of overlap. Such base stations should become neighbors and coordinate their configurations, but fail to detect the each other's presence during the probing phase. Consequently, their configurations will not be coordinated, leading to an inconsistent overall network configuration.

One approach to this problem involves manual configuration, forcing the base stations to treat each other as virtual neighbors (see Section III-B). Obviously, this approach does not fit with the goal of complete self-management.

Another solution is the inclusion of *external* information into the configuration process. This external information does not originate at base stations but is contributed by other nodes into the configuration algorithm. This paper assumes that these other nodes are user-operated clients of the WLAN access network, but they could also be specialized probe nodes under control of the WLAN operator [7].

The use of external information can address the overlap problem. If WLAN clients periodically notify their base station of other clients and base stations within their radio range, the management system can update the neighbor relation when a client enters an overlap area, eliminating or at least significantly reducing the overlap problem.

External information can improve the self-healing and self-optimization functions of a decentralized management system

in other ways. It enables detection of interferences or holes in coverage, can identify rogue access points outside the range of the base stations themselves, aid location tracking. Passively mobile clients – *i.e.*, users-carried devices – can already significantly aid the management system. Actively mobile nodes – *e.g.*, self-propelled robots under control of the management system – are even more useful, because the system can use them to obtain targeted information.

The inclusion of external information also has drawbacks. First, it requires additional software to be present on client nodes. Furthermore, the system must verify the trustworthiness of external information carefully before acting on it.

V. Preliminary Evaluation

This section briefly evaluates a prototype implementation of the management system. Note that the implementation and evaluation is currently ongoing and these results investigate the base functionality and are preliminary. Several aspects of the system, such as the impact of including external information (Section IV), have not yet been experimentally analyzed.

The current system prototype is a Perl daemon that is capable of operating on physical hardware, *i.e.*, a Linux PC equipped with IEEE 802.11b WLAN interfaces and a group of such machines within radio range will self-configure in interaction with one another. However, an analysis of the scalability properties using physical devices is impractical. Therefore, the prototype offers a simulation mode, where multiple copies of the same code execute on a single PC inside a simulated topology. The preliminary measurements in this paper analyze this simulation mode.

The preliminary scalability analysis investigates the convergence time of the group of base stations if all start up within a few seconds of one another, *i.e.*, the time of the initial self-organization, such as after a power failure. Clients are not present. The experiments measure the convergence times of 50 repetitions and calculate mean performance and standard deviations. Each experiment uses a randomly generated, connected base station topology, *i.e.*, the aggregate coverage area of the base station group is not geographically partitioned. This is arguably a common deployment case; the usefulness of integrated management of a group of base stations that cover geographically separate regions is unclear. The number of base stations is a parameter of the experiment and varies from 1 to 100 in increments of 10, with two additional group sizes of 5 and 15 to investigate behavior for small groups.

Figure 1 shows the performance. For smaller groups of 1-20 base stations, the mean initial self-organization time quickly increases from 17 to approximately 20 seconds. For larger groups of 20-100 base stations, the mean initial self-organization time remains between 20 and 25 seconds.

A single base station self-configures in approximately 17 seconds. This is due to the startup behaviour. After the initial randomised 0-10 second de-synchronization delay, a base station initiates a probing phase to detect and contact its neighbors. This operation takes approximately 2 seconds. When no mobile nodes are associated with a base station, it

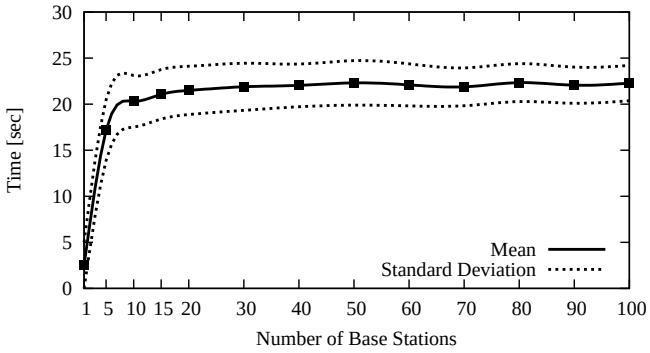

Fig. 1. Initial self-organization time of a group of base stations.

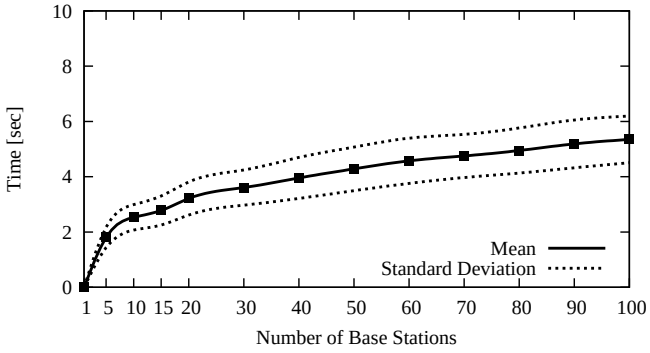

Fig. 2. System convergence time after a change to the global configuration information.

repeats the probing phase. The second probing phase already starts after 10 seconds, in order to allow a base station to detect additional neighbors that were still channel scanning during its initial probing phase. It will then initiate communication with these neighbors and start the information exchange.

The current prototype implements a probing mechanism that can only detect neighbors that are already offering client connectivity, *i.e.*, have switched to infrastructure mode. A future revision may extend this behaviour to detect neighbors that are themselves still channel scanning. This means that with the current prototype, some base stations do not detect all their neighbors during their initial channel scan.

Figure 2 shows the system convergence times if the global configuration information changes at one base station, for example, after an administrative update. The experiments measure the convergence times of 500 repetitions and calculate mean performance and standard deviations. before, the number of base stations is a parameter of the experiment and varies from 1 to 100 in increments of 10, with two additional group sizes of 5 and 15 to investigate behaviour for small groups.

Figure 2 shows a logarithmically growing convergence time up to 5 seconds for 100 base stations. This behavior illustrates the scalability properties of the epidemic update mechanism. Each base station informs its neighbors about changes in periodic intervals. In the current prototype the default is set to 1 second. That means that a change at the global configuration

set at one base station is forwarded after a second at maximum. Shorter periods result in faster convergence times, but increase load on the network.

Note that although this decentralized self-organization is costly for small groups of base stations, it adapts well to larger groups, as indicated by almost constant convergence times for increasing group sizes. Future experiments will verify if this scalability trend holds for groups of several thousand base stations. They will also investigate the dissemination times of changes to global state, and analyze the effects of including external information into the configuration algorithm.

VI. Conclusion

This paper introduced and motivated a distributed approach for wireless base station management and configuration and compared it to existing and mostly centralized solutions. One novel feature of the presented system is the inclusion of external information into the distributed management process to improve the quality of the configuration result. The paper outlined the current system and prototype implementation and presented a preliminary scalability analysis that indicated promising behavior.

A detailed performance and scalability analysis of a more complete system implementation is currently ongoing. It will investigate the performance of additional system functions such as improved channel allocation, load balancing, rogue detection or location tracking and also quantify the quality improvement obtainable by the inclusion of external information. It will also extend the scalability analysis to larger groups of base stations. A more detailed description of the management system and more detailed experimental analysis are available in [16].

Although the current base station configuration system specifically targets WLAN networks, the general idea of decentralized management certainly applies to other wireless and wired networks. The current system provides a decentralized management "middleware" on top of generic methods for information dissemination that can adapt to other network technologies in the future.

VII. Acknowledgement

The authors would like to thank Simon Schütz, Stefan Schmid and Jürgen Quittek for their feedback during the design of the system.

This document is a byproduct of the *Ambient Networks* project, partially funded by the European Commission under its Sixth Framework Programme. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the *Ambient Networks* project or the European Commission.

References

[1] N. Abramson. Development of the ALOHANET. *IEEE Transactions on Information Theory*, Vol. 31, No. 2, March 1985, pp. 119-123.

[2] Advanced Cybernetics Group and Meshdynamics. Why Structured Mesh is Different. *White Paper*, 2004.

[3] Airespace Corporation. Putting the Air Space to Work. *White Paper*, October 2003.

[4] Aruba Wireless Networks. Getting a Grip on Wireless LANs. *White Paper*, January 2003.

[5] V. Cerf. Packet Satellite Radio Reference Sources. *RFC 829*, November 1982.

[6] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart and D. Terry. Epidemic algorithms for replicated database maintenance. Proc. *6th Annual ACM Symposium on Principles of Distributed Computing*, Vancouver, British Columbia, Canada, 1987, pp. 1-12.

[7] S. Felis, J. Quittek and L. Eggert. Measurement-Based Wireless LAN Troubleshooting. Proc. *First Workshop on Wireless Network Measurements (WiNMee 2005)*, Riva del Garda, Trentino, Italy, April 3, 2005.

[8] IEEE-SA Standards Board. ANSI/IEEE Standard 802.11. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *ANSI/IEEE Standard*, 1999.

[9] L. Ji, J. Agre, T. Iwao and N.Fujino. On Providing Secure and Portable Wireless Data Networking Services: Architecture and Data Forwarding Mechanisms. Proc. *International Conference on Mobile Computing and Ubiquitous Networking (ICMU 2004)*, January 2004.

[10] J. Jubin and J.D. Tornow. The DARPA Packet Radio Network Protocols. *Proceedings of the IEEE*, Vol. 75, No. 1, January 1987, pp. 21-32.

[11] N. Niebert, M. Prytz, A. Schieder, L. Eggert, F. Pittmann, N. Papadoglou and C. Prehofer. Ambient Networks: a Framework for Future Wireless Internetworking. To appear Proc. *IEEE 61st Semiannual Vehicular Technology Conference (VTC 2005 Spring)*, Stockholm, Sweden, May 30 - June 1, 2005.

[12] J.J. Silva Tobella, M. Stiemerling and M. Brunner. Towards Self-Configuration of IPv6 Networks. *Proc. Poster Session of IEEE/IFIP Network Operations and Management Symposium (NOMS'04)*, Seoul, Korea, 2004.

[13] F. Stajano and R. Anderson. The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks. Proc. *7th International Workshop on Security Protocols*, Cambridge, UK, April 1999, In *Lecture Notes in Computer Science (LNCS)*, Vol. 1796, Springer Verlag, Heidelberg, Germany, pp. 172-182.

[14] A. Tanenbaum and M. van Steen. Distributed Systems, Principles and Paradigms. Prentice Hall Inc., NJ, USA, 2002.

[15] Trapeze Networks. Defining An Integrated Access Point. *White Paper*, 2004.

[16] K. Zimmermann. An Autonomic Approach for Self-Organising Access Points. *M.S. Thesis*, University of Ulm, Germany, March 2005.