

Design of High Capacity Wireless LANs based on 802.11b Technology

Steven Gordon and Arek Dadej

Institute for Telecommunications Research, University of South Australia

Abstract: The availability and low cost of 802.11b wireless networking products have encouraged a rapid growth in deployments of wireless LANs. As the use of powerful portable computing devices (laptops, PDAs) with WLAN adapters becomes a common expectation, the ability of wireless LANs to cater for large numbers of users with high bandwidth requirements is an increasingly important design consideration. To date, most of the deployed large WLANs were designed for continuous radio coverage, not for high densities of bandwidth-hungry users. In this paper, we report on the results of an experimental and analytical study specifically addressing the problem of designing high density WLANs based on the 802.11b technology. We conclude that in the design of high-density WLANs, due to the limited control over transmission ranges of the wireless devices, particular attention must be paid to means by which the effect of exposed terminals in the network can be reduced.

Key words: Wireless LAN; 802.11; WLAN Capacity; Medium Access Control; RTS/CTS; Transmission Range; Sensing Range; Hidden Terminal; Exposed Terminal.

1. INTRODUCTION

Over the recent years, the availability of low-cost 802.11b wireless networking equipment has seen a rapid growth in deployments of wireless LANs based on this technology. To date, the major design considerations in the deployment of 802.11b WLANs have been the completeness and continuity of radio access coverage, as well as (more recently) the WLAN security. Little attention has been paid to the ability of WLANs to cater for large numbers of users with high bandwidth requirements, densely packed within a well-defined coverage area. As a result, the available technical

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35673-0_28](https://doi.org/10.1007/978-0-387-35673-0_28)

C. McDonald (ed.), *Converged Networking*

© IFIP International Federation for Information Processing 2003

literature, including manufacturer's technical notes, offers little advice on the design of high-density wireless LANs. The design for high density/capacity will inevitably become one of the major considerations in the deployment of large WLANs, rendering the trial and error approach to WLAN deployment prevailing today inadequate.

As part of R&D activities supported by the m.Net Corporation (<http://www.mnetcorporation.com>), we have undertaken an experimental and analytical study aiming at capacity design of large WLANs based on 802.11b technology. Observations from this study can be converted into advice to manufacturers of 802.11b equipment and, more importantly, into engineering guidelines for the design of WLANs with controlled capacity.

The 802.11b technology allows for 3 non-overlapping frequency channels to be used by the access points. It is therefore natural to think that in order to maximise the network capacity, one may try the principle of spatial frequency re-use commonly exploited in cellular networks. With three non-overlapping channels, the ability to form a semi-cellular pattern of access points is somewhat limited; nevertheless some capacity gains should be possible. Our study was therefore designed to answer questions relevant to the design of semi-cellular networks of 802.11b access points.

In the design of a semi-cellular 802.11b WLAN, the following questions must be answered:

1. *How much control does the designer have over the transmission, hearing and sensing ranges of the wireless devices?* The answer is critical in dimensioning the cells so that minimum overlap of coverage (thus interference) occurs between adjacent cells.
2. *Are there any significant limitations on the number of clients that can be supported by one access point (cell)?* A number of known studies have addressed the question of WLAN MAC protocols capacity [7, 10, 1, 2, 3]. The results relevant to our question can be summarised as follows: 1) there are no significant MAC-related limitations on the number of users supported by a cell, and 2) the capacity of a WLAN cell degrades slightly with the growing number of users in the cell, because the settings of the 802.11b back-off algorithm are optimised for small numbers of users, e.g. 10 to 20. In our study, we have adopted the capacity model from [1]. The resulting cell capacities are shown in *Figure 1*. We have also investigated (including experimentally) possible implementation-driven limitations on the number of clients supported by an AP, but none have been identified.




Figure 1. Capacity of 802.11b calculated using Bianchi's model

3. *What is the effect of cell overlap on the total network capacity?* Some overlap between cells is unavoidable, and so the impact of overlap on capacity must be determined. We need to consider not only the simple sharing of radio resource between cells in the areas of overlapping coverage, but also the impact of such phenomena as hidden and exposed terminals.
4. *Given the answers to the previous three questions, what is the optimum layout of cells in respect to network capacity?* An answer to this question leads to recommendations for designing high-density wireless LANs.

Although individual aspects of WLAN have been extensively studied over a number of years, perhaps the only published study of the issues in large WLANs can be found in [4]. Based on the experiences with design and deployment of a large University WLAN, the author presents guidelines for the layout and configuration of cells. However, the semi-cellular approach is not discussed and the only recommendation for high-density coverage is to use multiple channels in the same coverage area and to increase the receiver threshold settings. The latter reduces the coverage area of a cell. However, as discussed later in our paper, this introduces problems due to mismatch between the access point and client transmission ranges. We present analysis of a multi-cell WLAN in Section 4.

2. IEEE 802.11 WIRELESS LANS

The core components of IEEE 802.11, a standard for WLANs, were completed in 1997 [5]. Works on various extensions to 802.11 are still in progress. This section describes aspects of 802.11 necessary as a background to our analysis.

2.1 Physical (PHY) Layer

The majority of WLAN products currently on the market implement 802.11b. The main characteristics of the 802.11b PHY layer are data rates of 1, 2, 5.5 and 11MB/s and the operating frequency range of 2.4GHz to 2.4835 GHz (ISM band). There are 11 operating channels available for 802.11b devices in Australia, three of which are non-overlapping.

2.2 Medium Access Control (MAC) Layer

IEEE 802.11 defines two mechanisms for sharing access to the radio resource: the Distributed Coordination Function (DCF) and the Point Coordination Function (PCF). Most 802.11b products available today implement and use DCF, and hence in this paper concerned with practical aspects of design and deployment of WLANs, we ignore the PCF.

DCF uses Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). *Figure 2* shows an example sequence of MAC frames to illustrate the basic access method of DCF.

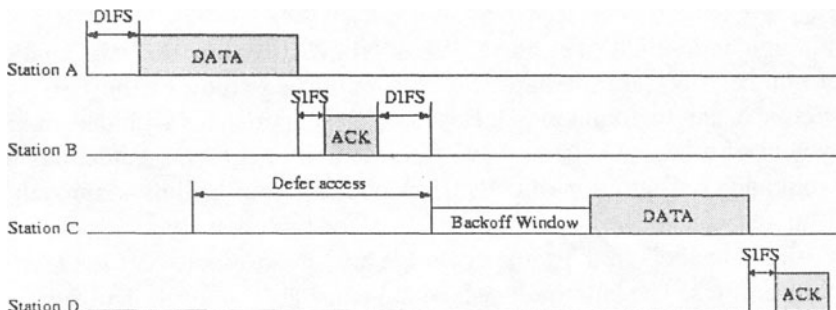


Figure 2. Timing diagram for the 802.11 basic access mechanism

Stations with DATA frames to transmit must first sense the medium (for a period of DIFS, DCF Inter-Frame Space) to ensure that no other stations are transmitting (Station A). If the medium is sensed free, the DATA frame

is sent, the receiver (Station B) waits for a time period of SIFS (Short IFS) and responds with an acknowledgement (ACK) frame. If the medium is not free, access is deferred until it is sensed free (Station C), after which the station waits for a Backoff Window time before transmitting. The Backoff Window length is selected randomly from values between 0 and CW (Contention Window). CW is initially 32, and doubles after every attempt by a station to retransmit the frame (up to a max of 1024). The random backoff mechanism prevents collisions between multiple stations awaiting access.

DCF also defines an optional access mechanism that requires stations to first advertise to all stations within the hearing distance the intention to send (Request To Send, RTS) and then wait for permission from the intended recipient (Clear To Send, CTS). All stations that hear either the RTS or CTS frame defer access until the transmission is over. The RTS/CTS significantly reduces probability of collisions in the presence of hidden terminals. The hidden terminal problem occurs when two (or more) stations outside the hearing range of each other transmit to the same station within the hearing range of both, causing a collision. The presence of hidden terminals results in a significant loss of capacity [7, 6, 9].

In another scenario, if two stations within hearing (sensing) distance of each other are ready to transmit to different destinations, one will defer from sending (to avoid collision). However, if the destinations can only hear their respective source station, not both, then both transmissions could occur in parallel without a danger of collision. This is called the exposed terminal problem. The RTS/CTS mechanism can amplify the loss of capacity due to exposed terminals.

3. COVERAGE AREA OF A STATION

To determine how well the coverage area of access points and mobile stations can be controlled by the user-configurable parameters (e.g. transmit power, data rates), we have performed simple coverage area measurements.

3.1 Experiments

An experimental network was built that consisted of a WLAN segment (APs) and a 100Mb/s Ethernet distribution network. The measurements were made for Cisco Aironet 350 APs and WLAN NICs, and Compaq WL110 WLAN NICs. Details of the experiments can be found in [8].

Three different propagation environments were considered:

1. Open: outdoors with minimal obstructions in the line of sight.

2. Closed: indoor, with up to two walls (plaster board on steel frame) in the line of sight; concrete elements of the building contributing multipath effects.
3. Semi-closed: indoor, in a corridor/hall with multipath effects but no obstructions in the line of sight.

A number of measurements for different data rates and power level settings were carried out. Example results, sufficient to illustrate problems relevant to our study, are summarised in *Table 1*.

Table 1. Maximum transmission range at 11Mb/s in different environments (AP power 1mW)

	Open (m)	Closed (m)	Semi-Closed (m)
Cisco 1mW	45	15	25
Cisco 30mW	50	20	30
Compaq 30mW	60	20	35

The Cisco NICs have a minimum power level setting of 1mW, whereas the Compaq (and many other) NICs have just one power setting of approximately 30mW.

3.2 Discussion

From the measurements, we conclude that the transmission range with the lowest power level setting (minimum cell diameter of 50 m, equivalent to nearly 2000 square metres coverage) is too large to enable effective design of high-density WLANs. Further reduction of transmission range is necessary. However, the required level of power control is not available in current 802.11b hardware.

Consider a simple example. If there are 1000 wireless users in a large convention hall of 1000m² in area, then with a transmission range of 25 metres one cell can cover the entire area. With three APs collocated and operating on the three available non-overlapping frequencies, and practically achievable data rate per channel of 5.5Mb/s, the available data rate per user will be only 16.5kb/s.

One available approach to reduction of the AP transmission range is to attenuate the antenna signal or use directional antennae. These techniques are already popular among wireless LAN designers and allow shaping of cells as required. However, there are no equivalent practical means for limiting the range of wireless NICs (clients). This leads to a mismatch between the transmission ranges of access points and clients, which results in a significant loss of capacity due to exposed terminal effect. The exposed terminal problem affects mainly the uplink traffic and is more pronounced

for large degrees of mismatch between the respective transmission ranges of APs and clients.

Consider example in *Figure 3*, where the transmission range of client 1 is shown as a dashed circle. When client 1 is transmitting, client 2 can hear client 1, and therefore must defer its own transmission. However, if both clients transmitted at the same time, all frames would be successfully received by their respective APs. The large transmission range of client 1 has prevented other clients from sending to their own APs while, in fact, they could do so without colliding. If the client and AP ranges were the same, then both clients in *Figure 3* could transmit at the same time.




Figure 3. Exposed terminal phenomenon due to mismatch between AP and client ranges

In a WLAN with semi-cellular layout of APs, the exposed terminal problem affects only uplink traffic. In *Figure 3*, both A1 and A2 can transmit downlink at the same time, as they are outside each other's hearing range.

As a result of the mismatch between client and AP transmission ranges, capacity calculations for a multi-cell network are far more complex than in simple, symmetric cases. For a cellular layout, there is no longer a direct, linear capacity gain resulting from spatial re-use of radio resources. The gain now depends on the client and AP ranges relative to each other, and on the ratio of upload to download traffic in the network.

An added complexity results from the choice of basic access versus RTS/CTS mechanisms. The RTS/CTS mechanism counteracts the hidden terminal problem, but can amplify the effect of exposed terminal. For example, a client responding to its own AP's RTS will silence all clients in its transmission range for the duration of AP's transmission, thus extending the capacity loss resulting from exposed terminal problem to downlink traffic. Tradeoffs between the effective reduction in hidden terminal problem and the amplification of exposed terminal problem, need to be explored in each specific case of capacity calculations for multi-cell WLAN.

4. ANALYSIS OF MULTI-CELL WIRELESS LANS

In order to investigate the effects resulting from the mismatch between client and AP transmission ranges, we developed a simple analytical model of multi-cell WLANs. The model accounts for the effects discussed in Section 3.2 and allows calculations of total WLAN capacity for different high-density network scenarios.

The capacity of a multi-cell WLAN is calculated by finding the capacity of a single cell, and then evaluating the capacity reduction caused by interference from stations in other cells (due to cell coverage overlap, as well as hidden and exposed terminals). The model can be used to calculate the WLAN capacity for different layouts of APs (co-located APs or cellular layout - see *Figure 4*) and for multiple combinations of user-configurable parameters. The model, implemented in Matlab, starts from finding the capacity of one cell (the *local cell*) and, based on the distances between APs and clients, the transmission ranges, and the average ratio of downlink to uplink traffic, calculates the capacity loss due to interference from stations in surrounding (*interfering*) cells. This is done through calculation of the percentages of stations in the local cell that receive their full share of radio resource, stations that share radio resource with stations in interfering cells, stations that are hidden terminals (in which case, we assume the capacity is reduced by a factor of H), and stations affected by the various modes of exposed terminal phenomenon.




Figure 4. Co-located (left) and cellular (right) layout for multi-cell WLAN. There are 3 non-overlapping channels available: 1, 6 and 11.

The main assumptions used in the construction of the multi-cell WLAN model are:

- The WLAN spans over a large area, and therefore capacity is identical for all cells in the area (i.e. edge effects in the layout of APs are ignored).

- A cell is circular, and the radius of the cell is equal to the transmission range of an AP.
- The sensing range of a station is 1.5 times its transmission range.
- Users are distributed throughout the area in a uniform pattern, and have the same bandwidth requirements.
- Only DATA frames may cause collisions (collisions of ACK frames are ignored).

For further details on the model and the full results, refer to the appropriate sections of [8]. Here, we will only present results for one example case.

In order to analyse the capacity issues in a large WLAN with frequency re-use, we have considered a scenario with 2000 wireless users (e.g. attendees of a large convention) in a coverage area of 10,000 square metres (the example venue here is slightly larger than the exhibition halls in the Adelaide Convention Centre). We assume that the ratio of downlink to uplink traffic can vary from 1:1 up to 5:1. *Table 2* shows other parameters relevant to the scenario.

Table 2. Model parameter values

Parameter	Units	Description	Value
Ra	Metres	AP transmission range	-
Rc	Metres	Client transmission range	25
Sc	Metres	Client sensing range	37.5
H	-	Hidden terminal factor	2
Users	-	Assumed number of users per AP	30
Bianchi_basic	Percent	Cell capacity with basic access	50.31
Bianchi_rts	Percent	Cell capacity with RTS/CTS	51.55

For the scenario under consideration, the maximum user bandwidth (for a single cell, without multi-cell interference) is 184kb/s with the basic access method (189kb/s with RTS/CTS), i.e. 30 users sharing approximately 5.5Mb/s.

Taking into account the interference from users in other cells, *Figure 5* shows the user bandwidth for both (co-located and semi-cellular) layouts of a multi-cell WLAN for the described scenario. The major observations are:

1. The capacity gain from using a cellular layout (over a collocated AP layout) is between 20% and 30%. The capacity is strongly limited by the asymmetry between transmission ranges of clients and access points, and by the long sensing range (1.5 x transmission range) of the clients; the resulting exposed terminal problem leads to significant reduction in capacity.
2. The only way to achieve capacity gains by means of frequency re-use in the scenario characterised by asymmetry of transmission and sensing

ranges is to exploit the asymmetry between the downlink and uplink components of data traffic. This is possible because the shorter sensing range of APs prevents them from sensing clients associated with other APs, thus protecting the downlink transmissions from exposed terminal.

3. The basic access method has significant performance advantages over the RTS/CTS scheme when the downlink/uplink traffic ratio is more than 2:1. This is due to the RTS/CTS scheme amplifying the exposed terminal problem by extending its capacity loss effect to downlink traffic (as explained in Section 3.2.)




Figure 5. User bandwidth using co-located and cellular layouts when each multi-cell network has 2000 users. RTS/CTS results are shown as the dashed lines.

Even though the resulting data rates per user shown in the *Figure 5* do not appear to be overwhelming, the actual capacity gains achieved thanks to frequency reuse are significant. In a WLAN without reuse, all radio resources equivalent to 16.5 Mbps data rate would be shared among 2000 users, resulting in no more than 8.25 kbps per user. In a network with carefully controlled layout and transmission ranges of APs, it is possible to achieve much higher data rates per user (between 25 and 90 kbps), especially when the effects of hidden and exposed terminals are minimised by careful consideration of network operating conditions (e.g. asymmetry of uplink and downlink traffic).

We have to emphasise that the capacity gains shown in *Figure 5* for the cases of basic access and high ratios of downlink to uplink traffic can only be achieved if the resource re-use is maximised by careful control of the layout and coverage areas of APs. Without such careful, systematic design,

the excessive (uncontrolled) overlap between cells will cause significant degradation in total capacity.

5. CONCLUSIONS

We have investigated capacity gains that can be achieved by using a semi-cellular layout of APs in a WLAN. The ability to achieve such gain, and thus to deploy high-density WLANs, is significantly limited by the lack of sufficient control of the transmission ranges in current 802.11b devices.

The radio range of access points can be limited by artificial means (e.g. antenna attenuators and/or directional antennae) which allows dense packing of the APs, but this results in asymmetry between the transmission ranges of APs and clients, and amplifies the effects of the exposed terminal problem.

As the RTS/CTS scheme, often seen only as a means for eliminating collisions caused by hidden terminals thus universally recommended for certain packet and network sizes (see e.g. [1]), amplifies the effects of exposed terminal even further, we strongly recommend that this feature of 802.11 is used with caution, as a result of careful network design. It is very likely that in most practical designs of large 802.11b-based WLANs the use of RTS/CTS will not be advisable.

Finally, it is evident that without proper consideration of all factors discussed here (e.g. cell locations and sizes, RTS/CTS) in the design stage, deployment of high-density WLANs using current 802.11b technology carries significant risk of poor, well below expectations, performance.

ACKNOWLEDGEMENTS

This research was supported by m.Net Corporation. For more information, please visit <http://www.mnetcorporation.com>.

REFERENCES

- [1] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE JSAC*, 18(3):535-547, March 2000.
- [2] F. Cali, M. Conti and E. Gregori. Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit. *IEEE/ACM Trans. Networking*, 8(6):785-799, Dec. 2000.
- [3] K. Chen. Medium access control of wireless LANs for mobile computing. *IEEE Network*, Sep/Oct 1994.
- [4] A. Hills. Large-scale wireless LAN design. *IEEE Comm Mag.*, 39(11): 98-104, Nov. 2001.
- [5] IEEE/ANSI. Wireless LAN MAC and PHY Layer Specifications. IEEE Std 802-11-1999. 1999. See also IEEE Std 802-11b-1999.

- [6] K. Huang and K. Chen. Interference analysis of non-persistent CSMA with hidden terminals in multicell wireless data networks. In *Proc. 6th PIMRC*, pages 907-911, Toronto, Canada, Sep. 1995.
- [7] L. Kleinrock and F. Tobagi. Packet switching in radio channels: Part I-Carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Tran. Comm.*, 23(12):1400-1416, Dec. 1975. (see also Part II in same issue)
- [8] m.Net Corporation. *Design Guidelines for High Density WLANs*. Final Report Volume I and II. 17 April 2002.
- [9] A. Zahedi and K. Pahlavan. Throughput of a wireless LAN access point in presence of natural hidden terminals and capture effects. In *Proc. 7th PIMRC*, pages 397-401, Taipei, Taiwan, Oct. 1996.
- [10] A. Zahedi and K. Pahlavan. Capacity of a wireless LAN with voice and data services. *IEEE Tran. Comm.*, 48(7):1160-1170, July 2000.