Load Balancing in IEEE 802.11 Networks

Li-Hsing Yen[*], Tse-Tsung Yeh, and Kuang-Hui Chi

Abstract

When an IEEE 802.11 network operates in infrastructured mode, the traffic load collectively given by wireless stations (WSs) can be unevenly shared by all available access points (APs), as WSs independently select APs to camp on. This article briefs the load balancing problem in such environments, and reviews state-of-the-art solutions toward this problem. We categorize load-distribution schemes as WS-based and network-based, depending on which part of the network is in charge of load distribution. In addition, the article presents experimental results using off-the-shelf IEEE 802.11 devices. These results demonstrate the effectiveness of load balancing in increasing overall system throughputs.

Index Terms: IEEE 802.11, load balance

1 Introduction

IEEE 802.11 wireless local area networks have been widely deployed as an infrastructure providing wireless data access services in home, corporate, and public environments. When operating in infrastructured mode, a wireless station (WS) equipped with an IEEE 802.11 interface sends and receives frames through an access point (AP) to a wired infrastructure. The AP, which bridges wireless and wired backbones, typically serves as a link-layer point of attachment to the Internet for WSs.

A WS must be associated with an AP before it can receive frame-forwarding service from the AP. To discover available APs, a WS performs probe processes by an active or a passive scan. In

---

[*] Contact author. E-mail: lhyen@nuk.edu.tw

an active scan, the WS broadcasts in some channel a Probe Request frame specifying a particular Service Set Identifier (SSID). If the SSID of interest matches the one that an AP is configured with, the AP can respond a Probe Response to the WS, which can therefore be aware of the presence of the AP. If the WS uses a passive scan instead, it does not issue any Probe Request but listens to Beacon frames broadcast periodically by APs. Based on the information grabbed from the Probe Response or Beacon frames, the WS selects an AP to camp on. Almost all existing IEEE 802.11 adaptors favor the AP with the strongest received signal strength.

The bandwidth offered by an AP is shared among all WSs associated with it. The throughput of an AP increases in proportion to the amount of packet traffic load added by all associated WSs, provided that the overall traffic load does not exceed the capacity of the AP. When AP's workload exceeds or approaches its capacity, the throughput does not increase further. This phenomenon is called overloading or congestion. A simple way to increase overall system throughput is to deploy additional APs covering the same region, in anticipation that heavy traffic load can be distributed among these APs. Unfortunately, as each WS independently selects an AP to camp on, WSs may be associated with few APs while other APs remain idle. Consequently, the traffic load is not fairly shared by APs. This problem motivates load-balancing protocols for IEEE 802.11 networks. The design goal is to make AP associations load-aware, preventing possible associations with congested APs. The ultimate goal is to increase overall system throughput.

Any approach to load balancing should address two primary issues. The first concerns how to define and measure load-related metrics. The second issue concerns how to balance or, at least, to distribute overall traffic load among all available APs. Depending on which part of the network executes the load-distribution process, we categorize load-distribution schemes as WS-based and network-based. This article first addresses the load metric issue and then reviews prior work with the proposed categorization.

2 Load-Aware AP Selection Metrics

A fundamental issue central to load-aware AP selections is to measure the workload of APs. This is an issue as AP's load does not yet have a well-accepted definition and we may not be able to measure AP's load directly. In intuitive terms, a large population of traffic sources often implies high traffic load. Hence a straightforward load metric of an AP is the number of WSs associated with it [1]. However, the IEEE 802.11 standard does not include WS count as an information element in Probe Response or Beacon frames. Some vendors and researchers have proposed to add such information in these frames. The resultant efforts have come to a part of IEEE 802.11e. In IEEE 802.11e, the QBSS load element contains information on the current WS population and traffic levels in the AP. However, the QBSS Load element is not always present in Beacon or Probe Response frames[1], which somewhat impedes widespread applicability of these schemes.

A network-side management entity may also request the WS count information from every AP to achieve network-wide load balancing. Unfortunately, the information cannot be accessed by using SNMP (Simple Network Management Protocol), as WS count has not been defined in any standardized MIB (Management Information Base) including MIB-II and IEEE802dot11-MIB [2]. Although some vendors have extensions to IEEE802dot11-MIB that include such information, these extensions are only applicable to their own products. An alternative is to extend a standardized protocol such as IAPP (Inter Access-Point Protocol) to include relevant information elements.

The WS count alone can only be a rough estimate of AP's workload, as traffic conditions vary significantly among WSs and may change over time. Some other metrics are used for pragmatic considerations. IEEE 802.11e and US patent 2004/0039817 [3] proposed to use channel utilization, the percentage of time an AP is busy transmitting or receiving data during some interval, as a gauge of channel loading. Channel utilization is not an appropriate AP selection metric, however, as it does not capture transmission capabilities of respective APs: an 80% utilized IEEE 802.11g AP can offer even more bandwidth than a 40% utilized IEEE 802.11b AP.

As overloading often entails a jammed sending queue, frame drop rate of real-time sessions (in AP's transmission queue) is another candidate for load metric [4]. It was also observed that queuing and channel-contention delay, which reflects AP's load level, can be estimated by the delay between scheduled and actual transmission time of periodic Beacon frames [5]. These low-level estimations assume certain implementation conventions and may not apply universally to all products.

---

[1] The information element is present if both MIB attributes dot11QosOptionImplemented and dot11QBSSLoadImplemented for APs are set to true.

Throughput is a common metric for quantifying the effectiveness of load balancing schemes. For this reason, Velayos et al. [6] take link-layer throughput as a direct measure of load. Some AP selection heuristics favor APs that maximize expected throughput [1] or potential bandwidth [5]. However, throughput or bandwidth usage is affected by time-varying channel conditions, which is intrinsic to IEEE 802.11 networks, as well as bursty traffic patterns. Usually we can only take a long-term average (in the scale of seconds or tens of seconds [7, 8]) as an estimate.

Bejerano et al. [7] asserted that the load induced by a WS w on its associated AP a is the time that a takes to provide w one unit of traffic. Accordingly, the load that w imposes on a is inversely proportional to the *effective bit rate* that w experiences. This definition aims at predicting AP's load with the information of effective bit rate. However, the predicted load coincides with observed workload only if w consumes all the bandwidth allocated to it, i.e., w has an infinite backlog of packets to send and receive. We argue that for WSs with bounded or dynamic bandwidth demands, AP's load is better expressed in terms of effective or observed throughput.

Given a load metric, we may gather a baseline for under-loaded AP operations. Accordingly, a load level reaching some predefined threshold can be indicative of overloading or congestion [4]. This is an absolute definition of overloading. The definition of overloading can also be relative. That is, an AP is overloaded if its load level exceeds the current average by certain amount [6]. The balance index introduced in [9] has been used to characterize the degree of load balance among servers. For n APs numbered from 1 to n that serve the same set of WSs, let L_i denote the amount of aggregated load imposed on AP i. The balance index β is defined as [10, 11, 6]

$$\beta = \frac{(\sum L_i)^2}{n \times \sum L_i^2}. \tag{1}$$

The value of β becomes 1 when all APs share equal load, and it approaches $1/n$ in case of extreme imbalance. If L_i in (1) denotes the bandwidth share received by WS i for n contending WSs numbered from 1 to n, the value of β essentially quantifies the fairness of bandwidth share among WSs [1].

All existing load-balancing protocols distribute traffic load by managing AP-WS associations. Depending on which part of the network is in charge of such managements, these protocols could

Table 1: Summary of load-balancing methods in IEEE 802.11 networks

Proposal	Need a server?	Load distribution controlled by	Require changes at the AP side?	Require changes at the WS side?
DLBA [12]	No	WS	Yes	Yes
MLT [1]	No	WS	Yes	Yes
Vasudevan *et al.*'s scheme [5]	No	WS	No	Yes
ALDP [8]	Yes	WS	No	Yes
Virgil [13]	Yes	WS	No	Yes
US patent 2004/0039817 [3]	No	WS	Yes	Yes
LBA [6]	No	AP	Yes	No
MAS [14]	No	AP	Yes	No
Cell Breathing [4]	No	AP	Yes	No
ACS [10]	Yes	Server	Yes	Yes
Online Load Balancing [7]	Yes	Server	No	Yes
US patent 2005/0213579 [15]	No	Switch	Yes	No

be classified into two types: WS-based and network-based. Table 1 compares protocols under discussion.

3 WS-based Load Distribution

In a WS-based approach, WSs learn of APs' load status somehow and, accordingly, select an AP that maximizes their potential benefits (*potential bandwidth* [5], for example). APs act passively in the whole selection process. Many WS-based approaches are not designed to achieve system-wide load balance—WSs select APs simply for their own interests. However, seeking an AP that provides the maximal available bandwidth implicitly implements *least-load-first* AP selection, a widely-used load-balancing heuristic.

The acquisition of APs' load condition may be realized in several ways. A WS may measure channel utilization or the delay between the scheduled and actual transmission time of periodic Beacon frames [5]. Such a measurement requires no assistance from any network-side entity. Alternatively, an AP may assist the measurement by broadcasting its current WS population and/or traffic level in Probe Response or Beacon frames [1, 12], preferably with a QBSS Load element if the AP supports IEEE 802.11e. A dedicated server may also be deployed to assist load measurements. In [8], a WS is first associated with some AP, through which it then accesses load metric (throughputs) from a stand-alone server. The server maintains load states for all APs residing in its administrative domain by using SNMP to periodically poll throughput-related MIB objects from these APs. With additional estimate on its own bandwidth consumption, the WS then decides whether a handoff should be conducted to distribute the load. In [13], a WS briefly connects to each available AP for performance test. For each AP, the WS generates test traffic toward a pre-deployed server in the Internet. The WS then selects the AP that provides the highest performance service for a long stay. The selection is based on expected bandwidth and round-trip time.

The WS-AP association management can proceed in a static or dynamic fashion. In static cases, a WS performs AP selection prior to its association with the target AP and does not reassociate to other APs as long as the association holds. A drawback of static AP selection is the inflexibility to adapt to network dynamics. With dynamic AP selection, on the other hand, a WS may determine to reassociate with another AP even if the current association still holds. Dynamic AP selection is better suited to highly dynamic networking environments. However, it may also lead to unstable WS-AP associations or so-called *ping-pong effects*, the phenomenon of repeated association changes from one AP to another.

One cause of ping-pong effects results from uncoordinated but simultaneous AP switching among WSs. For example, if we suddenly power on an AP beside a congested one, all WSs may detect the presence of and decide to switch to the new AP at almost the same time. Consequently, the new AP immediately becomes overloaded due to bursty migrations and all WSs decide to switch back. Then the same scenario repeats.

To avoid ping-pong effects, either static AP selection should be used or there should be a way to distribute re-associations in the temporal domain. For instance, in [8] a WS periodically searches for the best AP that has the least load. When the best AP found is different from the previous one, the WS does not switch to that AP immediately but generates a random value d. The WS can switch to the best AP only after the AP has been identified the best for d successive times.

Such a back-off scheme suppresses the burst of association migrations.

WS-based approaches have the advantage that off-the-shelf APs can be utilized with little or even no modification [5, 8]. However, as WSs select APs for their own interests, these approaches generally do not lead to a network-wide load balance. An alternative is to distribute load by a network-side entity.

4 Network-based Load Distribution

In a network-based approach, WSs behave passively in modifying AP-WS associations. It is a network-side entity (could be an AP, a switch, or a dedicated server) that controls the distribution of AP's load. There are three basic techniques for APs to control their own load level:

- *Coverage adjustment.* Crowded APs can reduce the transmission power of their beacon signal so that new WSs are less likely to discover them [11]. APs may collaborate in adjusting their radio coverage patterns in a way that lightly-loaded APs cover more area than heavily-loaded ones, and there is no coverage hole to ensure continuous coverage [14, 4].

- *Admission control.* An overloaded AP may simply reject new association requests. A non-overloaded AP decides whether it should grant association requests from WSs based on workload status. The request can be granted only when the predicted load level after the association does not exceed some threshold.

- *Association management.* A crowded AP may send an unsolicited disassociation frame[2] to selected WSs that are already associated with it, hoping that these WSs would re-associate with other lightly-loaded APs. Theoretically, the best disassociation candidate is the one for which the corresponding reassociation balances the load among related APs. However, it is impractical to seek such an optimal solution in a fast-changing networking environment as the optimum holds only for the current state. For a heuristic that finds good candidates, the AP may need to know the load level of neighboring APs and the set of APs that each

---

[2] possibly with status code indicating "Disassociation because the AP is unable to handle all currently associated stations"

associated WS is able to access. The ping-pong effects may still be a problem if we allow a disassociated WS to be disassociated again in the future.

For a global view of load distribution, APs may exchange load status information via a wired backbone. Protocols such as IAPP with slight modifications may be used for this purpose. APs identify themselves as overloaded can then use above-mentioned approaches to relieve load [6, 14]. Alternatively, a dedicated server located in the wired infrastructure can be used to collect load-related information [10]. The server learns of the distribution of traffic load and recommends or instructs designated WSs (by communicating with the peers running at these WSs) to change their AP associations. The ability to estimate the bandwidth actually consumed or potentially demanded by each WS may further facilitate association migration (handoff) decisions, as this information eases the following tasks:

- Determining the set of illegible APs that fulfills the implicit bandwidth demand of a particular WS.
- Predicting load shift between APs for each feasible handoff so that the best handoff target can be determined.

The support of multiple transmission rates, which is common nowadays[3], complicates the estimation. The actual rate in use between an AP and associated WS depends on the underlying time-varying channel conditions. To make things worse, it is known [16] that the throughput of WSs with high transmission rates may suffer from the presence of low-rate WSs. This makes handoff decisions difficult in that the load shift due to association migrations is hard to predict.

Switches may also be used to manage AP associations. US patent 2005/0213579 [15] devised a centralized mechanism where a switch connecting a set of APs provides load-balancing functionality within a wireless LAN[4]. In that patent, APs are required to interact additionally with the switch to make decisions on certain operations on behalf of APs. Therefore the switch and APs need to

---

[3] IEEE 802.11a, for example, supports eight transmission rates ranging from 6 to 54 Mbps.

[4] US patents 2004/0063455 and 2004/0156399 also considered a similar architecture, where a switch connecting a set of APs makes load-related decisions.

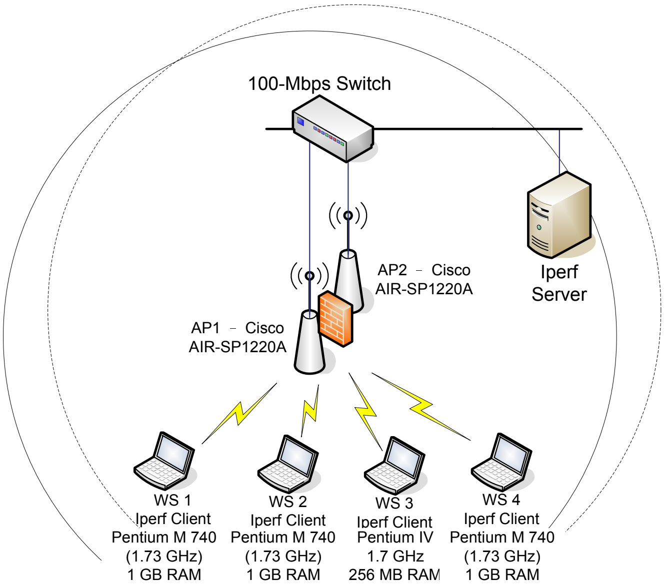

	WS 1	WS 2	WS 3	WS 4
AP1	-52.9	-46.6	-50.9	-48.5
AP2	-63.2	-60.3	-66.1	-60.5

Mean Signal Level (dBm)

Figure 1: Experimental setup. The table in the bottom lists the received signal strengths reported at each WS.

share some (proprietary) protocol; the interoperability with APs from different vendors may be challenging.

5 Experimental Results

We conducted experiments to measure and compare network performance with and without load distribution schemes. In the experiments, we placed two identical IEEE 802.11a APs (Cisco AIR-SP1220A) in a close range with a partition in between. These APs, operating on channels 56 and 64, respectively, were connected through a 100-Mbps switch to a PC running Linux kernel 2.4.20. The Ethernet connecting all these devices was isolated from other LANs, so as to minimize potential influence from any background traffic. A set of notebook PCs were used as WSs, each equipped with an IEEE 802.11 a/b/g PCMCIA interface (D-Link DWL-AG660). These notebook PCs ran Linux kernel 2.4.33 with MADWiFi 0.9.2 driver[5]. Fig. 1 illustrates the experimental setup.

---

[5] http://madwifi.org/

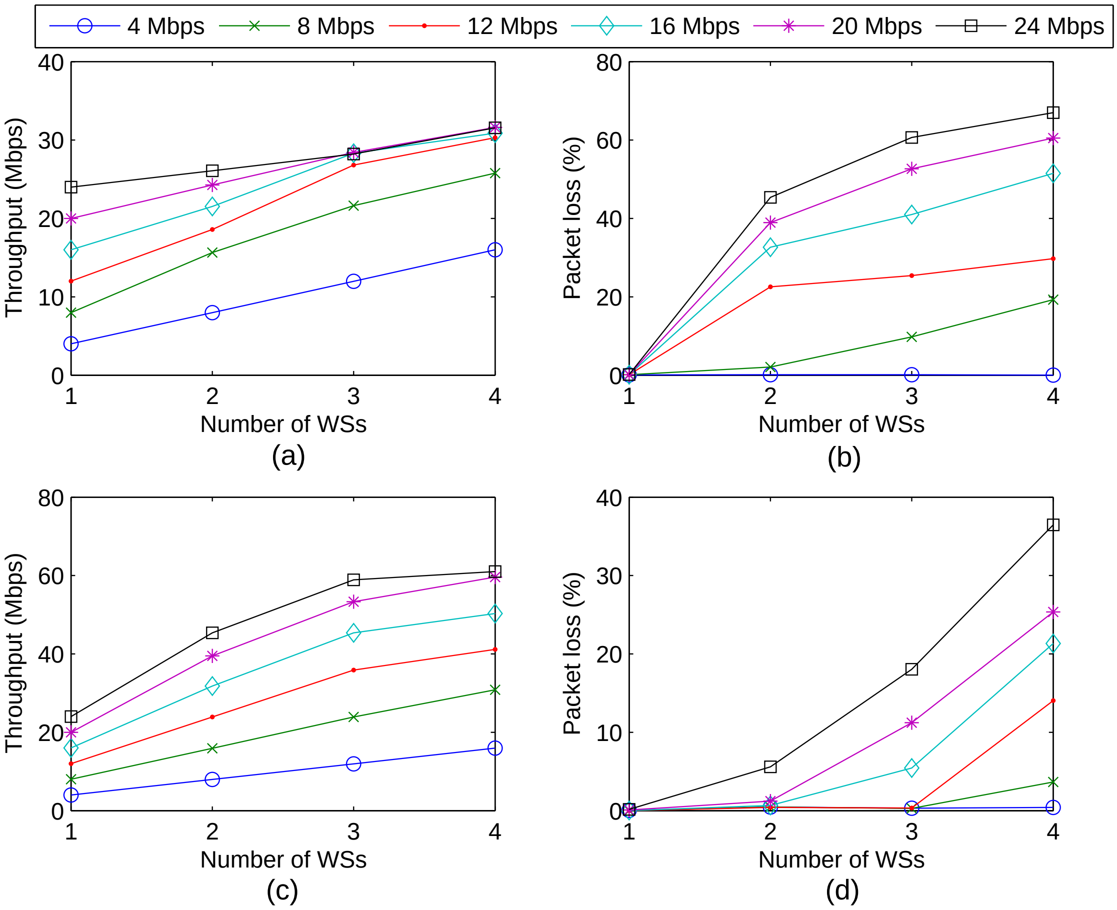

Figure 2: (a) Throughput without load balance. (b) Packet loss rate without load balance. (c) Throughput with load balance. (d) Packet loss rate with load balance.

Our experiments used Iperf[6] to measure throughput and packet loss rate. An Iperf client running at each WS generated fixed-size UDP segments[7] at a constant rate to an Iperf server, the PC located on the Ethernet. The UDP buffer size was set to a default value (64 KByte). The source data rates were all identical and varied from 4 to 24 Mbps. Each test traffic lasted one minute and all statistical results were gathered at the Iperf server.

The first setup did not employ any load-distribution protocol. Four WSs were sequentially added to the network and placed right in front of one AP labeled AP1 (within one meter). These

---

[6] http://dast.nlanr.net/Projects/Iperf/

[7] **Although the whole Internet traffic is mainly composed of TCP traffic, we did not consider TCP traffic because characteristics of the protocol such as slow start and congestion avoidance may influence the validity of our collected results. By getting rid of TCP involvement, collected statistics are fully representative of link-level behavior to the greatest possible extent.**

WSs remain stationary during the test. Consequently, all WSs were associated with AP1 for its stronger signal strength. The aggregate throughput reported at the Iperf server was taken as network throughput, which in this case was contributed by AP1 only. Fig. 2(a) shows how the network throughput relates to the number of associated WSs. It can be seen that the maximum throughput (31.7 Mbps) was upper bounded by one AP's *real* capacity, which is slightly higher than the half of its nominal capacity (54 Mbps). This is consistent with the finding of previous independent studies. When only one WS was associated with the AP, the network throughput completely reflected the given load. When two or more WSs were associated, the throughput no longer increased linearly with the given load, even if the aggregate load did not exceed one AP's real capacity. This can be explained as more WSs raised the degree of channel contention, reducing effective bandwidth each WS received. When the effective bandwidth was inadequate to bear the traffic load the WS generated, considerable outgoing packets were dropped as the sending buffer was mostly full. This argument is supported by measured average packet loss rates, which is shown in Fig. 2(b). The average packet loss rate increased monotonically with the number of contending WSs.

The second setup was identical to the first one, except that a WS-based load-distribution protocol [8] was applied. In this design, a software module called *information collector* running at the PC that also ran Iperf server maintains WS count for each AP and routinely collects throughput-related MIB objects from APs through SNMP. An application running at each WS, called Load Control Entity (LCE), takes charge of association managements. After a WS is associated with some AP (the one with the strongest received signal strength), the LCE requests the information collector for utilization, number of currently associated WSs, and interface speed of each available AP. Based on this information, the AP's normalized residual bandwidth (NRB) after admitting this WS's possible immigration can be computed. The LCE then selects the AP with the highest possible NRB to camp on, which may cause a re-association if needed. Consequently, WSs in this setup were not always associated with AP1, so overall network throughput was contributed by both APs. Fig. 2(c) shows resulting throughputs. Compared with the results of the first setup, the load-distribution scheme significantly improved network throughputs. Particularly, the maximum throughput doubled. The observed improvements were due to effective load distributions among available APs, which can be confirmed by examining the corresponding results of average packet loss rate (Fig. 2(d)). Generally speaking, the throughput-improvement ratios (the ratio of increased throughput to the original) were closely related to the amount of traffic load WSs imposed on APs. When the aggregate traffic load did not exceed 32 Mbps, the improvement ratios ranged from −10.0% to 28.5% with average value 2.5%. In contrast, the ratios ranged from 19.5% to 108.6% with average value 64.6% when traffic load amounted to 32 Mbps or more. For a specific workload, the improvement ratio was higher with fewer contending WSs. For example, consider three cases: (1) two WSs, each generating 24 Mbps traffic, (2) three WSs, each generating 16 Mbps traffic, and (3) four WSs, each generating 12 Mbps traffic. They imposed the same amount of traffic load on APs (48 Mbps in total), but their throughput-improvement ratios were 74.0%, 60.3%, and 35.7%, respectively.

Fig. 3(a) and (b) show respective throughputs of WSs when four WSs were deployed with and without using the load-balancing scheme, respectively. Fig. 3(c) indicates that the load-balancing protocol successfully increased the balance index of AP's load from 0.5 to 1, which implies that AP's load was perfectly balanced (each AP served two WSs and each WS generated identical traffic load.) However, the protocol did not equalize the bandwidth share each WS received, as indicated in Fig. 3(b). For example, the measured throughputs at WSs 1–4 were 18.3, 15.1, 4.84, and 22.8 Mbps, respectively, when each WS generated 24-Mbps traffic load. We believe that WS 3's poor throughput performance was due to its inferior computing/storage capability. While the effect of this slight variation was not significant when all WSs contended for a single AP (Fig. 3(a)), it makes difference when loads were distributed so that only two WSs contended for one AP. The effect was particularly significant when WSs all operated in high data transmission rates. This is clearly revealed by the corresponding balance index values shown in Fig. 3(d).

Such a performance difference would not appear in software-based simulation environment, where every WS acts homogeneously. In practice, a WS might still experience low packet loss rate (or, equivalently, high throughput) even when the serving AP was already congested. Similarly, a load-balanced state which increased the system throughput to a great extent did not necessarily

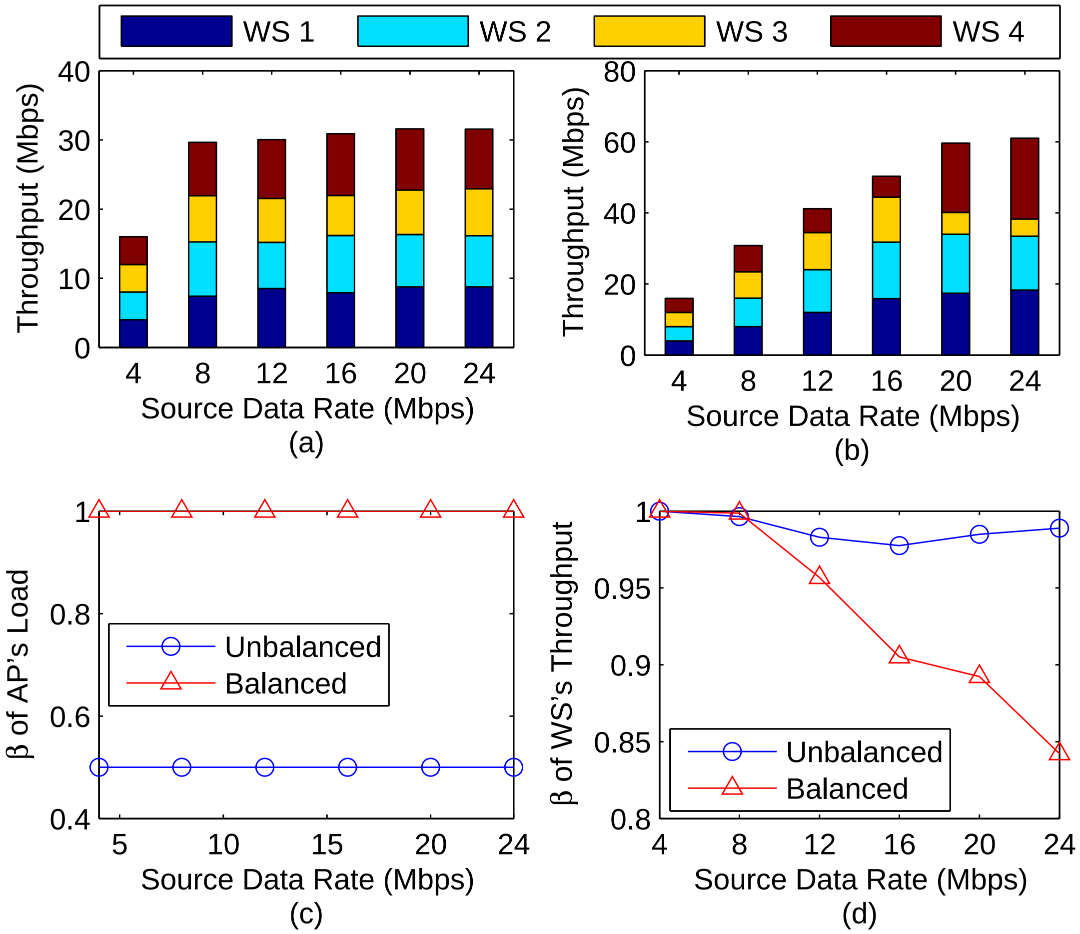

Figure 3: (a) and (b) show throughputs of four contending WSs without and with the load balancing protocol, respectively. (c) and (d) show the balance indices for AP's loads and WS's throughputs, respectively. 'Unbalanced' and 'Balanced' correspond to the results without and with the load balancing protocol, respectively.

lead to uniformly low packet loss rates. The lesson we have learned from these results is that fairness is an independent issue not resolvable merely by load balancing. This is justifiable as the original IEEE 802.11 does not guarantee an equal share of available bandwidth among contending WSs.

6 Conclusions

We have reviewed existing solutions to the load balancing problem in IEEE 802.11 networks. The proposed categorization, concerning which part of the network makes load-distribution decisions, classifies these solutions into WS-based and network-based. We have demonstrated through experiments that overall network throughput can be increased significantly if traffic load can be fairly distributed among available APs.

WS-based approaches are generally not customized to system-wide load balance, but they have the potential advantage of exploiting off-the-shelf APs without much modification. However, care must be taken not to cause ping-pong effects. For this reason, WS-based approaches are suitable for once-only AP selections. Network-based approaches in general do not necessitate modifying WSs. However, either a dedicated server is demanded or certain changes are required at APs/switches. Also information about WS's bandwidth consumption or demand is usually needed for a better reassociation decision. Network-based approaches have the potential of achieving system-wide load balance if designed appropriately.

References

[1] Y. Fukuda and Y. Oie, "Decentralized access point selection architecture for wireless LANs: Deployability and robustness," in *Proc. IEEE VTC2004-Fall*, Sept. 2004, pp. 1103–1107.

[2] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, ANSI/IEEE Std. 802.11, 1999.

[3] M. T. Lee, L.-T. Lai, and D. Lai, "Enhanced algorithm for initial ap selection and roaming," U.S. Patent 0 039 817, Feb., 2004.

[4] O. Brickley, S. Rea, and D. Pesch, "Load balancing for QoS enhancement in IEEE802.11e WLANs using cell breathing techniques," in *IFIP Mobile and Wireless Communication Networks Conf.*, Marrakech, Morocco, Sept. 2005.

[5] S. Vasudevan, K. Papagiannaki, C. Diot, J. Kurose, and D. Towsley, "Facilitating access point selection in IEEE 802.11 wireless networks," in *Proc. ACM SIGCOMM/USENIX Internet Measurement Conference*, Berkeley, CA, USA, Oct. 2005.

[6] H. Velayos, V. Aleo, and G. Karlsson, "Load balancing in overlapping wireless LAN cells," in *Proc. IEEE ICC*, Sept. 2004, pp. 3833–3836.

[7] Y. Bejerano, S.-J. Han, and L. Li, "Fairness and load balancing in wireless LANs using association control," in *Proc. ACM MobiCom'04*, Oct. 2004, pp. 315–329.

[8] L.-H. Yen and T.-T. Yeh, "SNMP-based approach to load distribution in IEEE 802.11 networks," in *Proc. IEEE VTC2006-Spring*, Melbourne, Australia, May 2006.

[9] D.-M. Chiu and R. Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks," *Computer Networks and ISDN Systems*, vol. 17, pp. 1–14, 1989.

[10] A. Balachandran, P. Bahl, and G. M. Voelker, "Hot-spot congestion relief in public-area wireless networks," in *Proc. 4th IEEE Workshop on Mobile Computing Systems and Applications*, 2002.

[11] H. Velayos, "Load distribution in IEEE 802.11 cells," Master's thesis, KTH, Royal Institute of Technology, Sweden, Mar. 2003.

[12] I. Papanikos and M. Logothetis, "A study on dynamic load balance for IEEE 802.11b wireless LAN," in *Proc. 8th Int'l Conf. Advances in Communication and Control: Telecommunications/Signal Processing*, Crete, Greece, June 2001.

[13] A. J. Nicholson, Y. Chawathe, and M. Y. Chen, "Improved access point selection," in *Proc. ACM MibiSys'06*, June 2006.

[14] Y. Wang, L. Cuthbert, and J. Bigham, "Intelligent radio resource management for IEEE 802.11 WLAN," in *Proc. IEEE WCNC*, Atlanta, Georgia, USA, Mar. 2004.

[15] P. J. Iyer, P. Narasimhan, M. Andrade, and J. Taylor, "System and method for centralized station management," U.S. Patent 0 213 579, Sept., 2005.

[16] M. Heusse, F. Rousseu, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM*, Mar. 2003, pp. 836–843.