Handover Incentives for WLANs with Overlapping Coverage ⋆

Xenofon Fafoutis and Vasilios A. Siris

Institute of Computer Science (ICS)
Foundation for Research and Technology - Hellas (FORTH)
P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece
{fontas,vsiris}@ics.forth.gr

Abstract. It is well known that in IEEE 802.11 networks, the assignment of low-rate and high-rate users to the same access point significantly degrades the performance of the high-rate users. Our objective is to investigate the implications of the above performance degradation on the incentives for handover between 802.11 wireless local area networks with overlapping coverage. Our focus is on the incentives for supporting handovers, due solely to the improved performance handovers yield for both wireless networks. To study the phenomenon and estimate the potential gain of such handovers, we propose a simple model that predicts the throughput of each access point in different cases. The throughput approximation model can indicate when the handover is expected to be beneficial, and can be used in a handover acceptance policy. Simulation of the proposed procedure suggests that the model is accurate and that there are significant throughput gains for both wireless networks.

Key words: handovers, cooperation incentives, wireless access network

1 Introduction

It is well known that in IEEE 802.11 networks, the assignment of low-rate and high-rate users to the same access point significantly degrades the performance of the high-rate users [1]. This occurs because IEEE 802.11's medium access control protocol gives both high and low-rate nodes equal chances for accessing the shared wireless channel. However, low-rate nodes need more time to send the same amount of data. As a result, high-rate users suffer significant performance degradation, achieving throughput equal to that of low-rate users. When using the term *rate*, we refer to the modulation rate of an 802.11 transmitter.

The objective of the work presented in this paper is to investigate the implications of the above performance degradation on the incentives for handover between 802.11 wireless local area networks with overlapping coverage. In case

⋆ This work was supported in part by the European Commission in the 7th Framework Programme through project EU-MESH (Enhanced, Ubiquitous, and Dependable Broadband Access using MESH Networks), ICT-215320, http://www.eu-mesh.eu

of operator-owned networks, handovers may be supported by cooperation agreements between the operators. The focus of this work is on the incentives for supporting handovers, due solely to the improved performance that handovers yield for both wireless networks, without involving any monetary exchange. Moreover, similar performance incentives arise in the case of overlapping wireless home networks where cooperation agreements are unrealistic. The main assumption of these scenarios is that two or more access points operate in the same channel. Indeed, it is common that there are more than three access points within the range of each other [2]. Hence, the three orthogonal channels available in 802.11b and 802.11g are not sufficient to assign orthogonal channels to different access points. Moreover, as more wireless networks operating in unlicensed bands are deployed over time, the above scenario will be more dominant.

In order to study the phenomenon and measure the potential gain of the aforementioned handovers, we propose a simple model which predicts the throughput of each access point in different context. Similar throughput approximations have been used in [9], [10], [11], [12] and [13]. Our main contribution lies in the application of the above model to predict the impact of handovers and decide whether handing over the nodes that communicate at low rates is beneficial for both access points. Our model suggests that there are cases where it is highly probable that the performance of both wireless networks is significantly improved. For instance, in 802.11b and when the low-rate nodes transmit at 1Mbps, our model suggests that if the number of nodes is uniformly distributed over time there is a 74% probability that the handover is beneficial for both access points, providing, on average, 55% better throughput for the nodes of the low-rate nodes' new host and more than three times better throughput for the nodes of their initial access point. Additionally, based on the necessary conditions for a handover to be beneficial for both access points, we propose a policy that can be used by each access point in order to decide whether to cooperate to the handover.

The rest of this paper is organized as follows. Section 2 summarizes the related work. Section 3 describes the scenario we investigate, and presents the throughput model used in the analysis. Section 4 quantifies the potential gains according to the model. Section 5 compares the analytical model with simulation using NS-2. Section 6 extends the model to include alternative bottlenecks, such as ADSL connections, and investigates their influence. Finally, Section 7 concludes the paper and presents our future work on this subject.

2 Related Work

The handover incentives studied in this paper address the performance degradation problem of 802.11 networks from a new perspective. Related work has approached the same issue in different ways.

One approach to mitigate the problem is to reduce the time low-rate transmissions utilize the shared medium, thus providing time fairness [3][4]. Although such an approach can increase the aggregate throughput of the network, it is un-

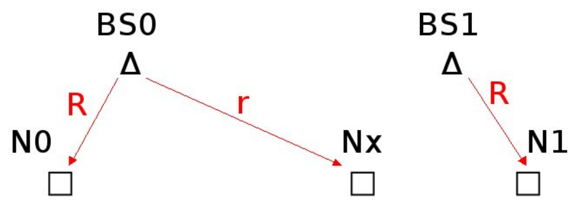

Fig. 1. BS0 serves both high-rate (N0) and low-rate (Nx) users (case a).

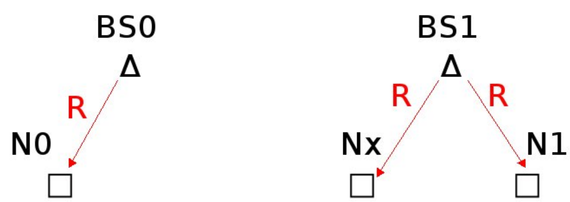

Fig. 2. BS0's low-rate users (Nx) are handed-off to BS1 (case b).

fair to low-rate nodes, as it minimizes their throughput even more. This leads to unsatisfied clients, which should be avoided in an environment with multiple competing providers. Another approach makes use of relay nodes for transmissions that cannot be performed at high rates [5][6][8]. In [5], the authors replace low rate transmissions with a two-hop sequence of shorter range, that enable higher rate transmissions. In [6], the authors propose a system which opportunistically transforms higher-rate 802.11 stations into relays for stations with low data-rates, hence requires the availability of such relay nodes. Another approach is to aggregate the capacity of all the access points and use load balancing mechanisms in order to maximize the network performance [7]. This approach requires all access points to cooperate, which can be assumed in the case where the access points belong to the same operator, but not when they belong to different operators. On the other hand, the focus of this paper is exactly on the cooperation between different operators, and shows that such cooperation can result from performance-oriented incentives even when operators act in their own self-interest.

3 Throughput Model

Consider the case of two access points, BS0 and BS1, Fig. 1. BS0 sends traffic to N0 nodes at high rate R and to Nx nodes at low rate r. Nodes N0 and Nx are the clients of BS0 and its actions target to improve its clients throughput. BS1 sends traffic to nodes N1 at high rate R. Nodes Nx are closer to BS1 and would transmit at a higher rate R, if they were associated to it. This is the base scenario, which we will refer to as case a. Additionally, the following assumptions are made: (a) both access points operate at the same channel, (b) all access points and nodes are in the same contention area and (c) there is at least one node in each of the three node sets. Due to the low-rate transmissions to the set of Nx nodes, the performance for all N0 and N1 nodes degrades.

In the scenario shown in Fig. 2, which we will refer to as case b, the low-rate clients Nx of BS0 are handed-off to BS1. Now, BS1 sends traffic only to the N0 nodes at high rate R, while BS1 sends traffic at high rate R to both its own clients (N1) and the ex-low-rate clients of BS0 (Nx).

The throughput gain of BSi is defined as the ratio of the aggregate throughput of the clients of BSi in case b (Nx clients associated to BS1), over the throughput in case a (Nx clients associated to BS0). This metric will be used to evaluate when the handover of low-rate users (case b) is beneficial. When the gain for both access points is greater than 1, handover improves the aggregate throughput for the clients of both access points. It is important to note that the Nx nodes are clients of BS0 in both cases, even though in the second case they are associated to BS1. Hence, to estimate the gain in both cases, the throughput of the Nx nodes is added to the aggregate throughput of BS0 clients.

The analysis in the following sections focuses on the above simple model, which encompasses the key tradeoffs we want to highlight. The analytical model can be extended to more complex cases, such as more than two access points and the case where additionally to the Nx nodes, some of BS1's clients symmetrically operate at low rates.

3.1 Model

Next we present a model for the throughput in saturated conditions for the downlink direction, i.e., from the access points to the clients. We assume that each access point sends one packet in each round. This is not absolutely true for IEEE 802.11 since the backoff waiting time of each transmission, as defined by the collision avoidance mechanism, is decided probabilistically. However, assuming that the DCF protocol of IEEE 802.11 provides long term fair channel access, the access points will send an equal amount of packets over a long time interval.

If T_0 and T_1 is the time each access point needs to transmit one packet respectively, then the long term throughput in bits per second that each access point will achieve, assuming both access points transmit packets of the same size, is equal to

$$X = \frac{pkt}{T_0 + T_1 + oh} \tag{1}$$

where *pkt* is the packet size in bits and *oh* is the overhead of two transmissions.

According to the DCF protocol of IEEE 802.11, each transmitter needs to sense the medium idle for a time interval equal to DIFS. It then chooses a random number of time slots between zero and the contention window (CW) and waits an additional time interval in order to avoid collisions. The transmission follows on the condition that the medium is idle during the time that it is waiting. After the transmission is completed the receiver waits for a time interval equal to SIFS, during which it switches to transmitter mode, and starts transmitting an acknowledgement at the control rate. The duration of DIFS, SIFS and TimeSlot as well as the initial minimum contention window are defined by the protocol. The time interval of the acknowledgement transmission depends on its size and the transmission rate. Since each transmitter chooses a uniformly random number in the interval between zero and CW, the long-term expected overhead of the collision avoidance mechanism when we have only one transmitter, is equal to CW/2 time slots. On the other hand, if a transmitter senses another transmission while being in backoff, it stops his counter and defers until the transmission is over and the medium is idle again, at which point it continues the backoff procedure at the point the backoff was stopped. When there is more than one contending nodes, their backoff counter runs down simultaneously. Hence, the expected overhead due to the contention avoidance mechanism is again CW/2 time slots, assuming that there are no collisions. In the scenario we investigate, there are two contending transmitters. As a result, there is a collision contention probability equal to 0.0625, assuming the minimum contention window is 15, which is the default value for IEEE 802.11. Based on the above, the overhead of two contenting transmissions is

$$oh = 2(DIFS + SIFS + ACK) + \frac{CW}{2} TimeSlot \; , \tag{2}$$

where the values of $DIFS, SIFS, ACK$, and $TimeSlot$ are defined by the 802.11 standard. This overhead ignores potential collisions; we investigate the accuracy of the model in Section 5.

Case a (no handover): When Nx nodes are assigned to BS0, the expected time interval T_0^a that BS0 needs to transmit a packet depends on the percentage of traffic sent to N0 and Nx nodes, since the duration of the transmission is different due to the different rates. On the other hand, the expected time interval T_1^a that BS1 needs to transmit a packet is independent of the number of its nodes since all operate at the same rate. Hence,

$$T_0^a = \frac{N0}{N0 + Nx} \frac{pkt}{R} + \frac{Nx}{N0 + Nx} \frac{pkt}{r} \; , \qquad T_1^a = \frac{pkt}{R} \; , \tag{3}$$

where $N0$ and Nx are the number of nodes in the N0 and Nx node-set respectively, r and R are the low and high rate, respectively.

The expected throughput of each N0 or Nx node, and each N1 node is

$$X_0^a = X_x^a = \frac{1}{N0 + Nx} X^a \; , \qquad X_1^a = \frac{1}{N1} X^a \; , \tag{4}$$

where X^a is estimated from (1), (2), and (3).

Case b (handover of low rate users): In the case there are no low rate transmissions, hence for both access points the expected duration of a packet transmission is equal to

$$T_0^b = T_1^b = \frac{pkt}{R} \; . \tag{5}$$

The expected throughput of each N0 node, and each Nx or N1 node is

$$X_0^b = \frac{1}{N0} X^b\,, \qquad X_x^b = X_1^b = \frac{1}{Nx + N1} X^b\,. \tag{6}$$

The gain of each access point is calculated using the above expressions. When low-rate nodes are associated to BS0 (case a), throughput is reduced. On the other hand, when the low-rate nodes are associated with BS1 (case b), BS1 shares its share of the wireless channel with the Nx nodes, which are BS0's clients. Case b is always beneficial for BS0, since BS0's clients utilize the wireless channel for more than half of the time in case b, and there are improvements due to removing low-rate transmissions. The inequalities

$$GainBS_0 = \frac{X^b + \frac{Nx}{Nx+N1} X^b}{X^a} > 1 \text{ and } GainBS_1 = \frac{\frac{N1}{Nx+N1} X^b}{X^a} > 1 \tag{7}$$

are necessary conditions for case b to be beneficial for both access points. Since $X^b > X^a$, the $GainBS_0 > 1$ is always satisfied. However, if the assumption that $N0 > 0$ does not hold, then case b is not always beneficial for BS0. The second inequality $GainBS_1 > 1$ is equivalent to

$$\frac{N1}{Nx + N0} > c\,, \text{ where } c = \frac{2 + \frac{ohR}{pkt}}{\frac{R}{r} - 1}\,. \tag{8}$$

This inequality can be used by BS1 to decide if it is beneficial to serve the low-rate nodes of his neighboring access point BS0. We make the following two interesting remarks regarding the above constraint. First, the acceptance constraint does not depend on the ratio of high-rate to low-rate nodes of BS0, but depends only on their sum. Second, for the same R/r ratio a higher value of R yields a higher constraint c. As a result, we expect the constraint to be more restricting for 802.11g than for 802.11b.

By slightly changing the model assuming that the set of Nx nodes consists of two distinct subsets, Ny and Nz, it can be easily shown that given that the handover is beneficial, the gain where both Ny and Nz subsets are served by BS1 is always greater than the gain where the Ny nodes stay with their access point and only the Nz nodes are served by BS1. In fact, the impact of one low-rate node on the throughput is much greater than the impact of additional low-rate nodes. As a result, either all low-rate nodes should be accepted by BS1 or none.

The handover acceptance policy, expressed by (8), requires that the access points know the number of connected nodes and their rates. Assuming that there are no hidden nodes, this information is easy to obtain by sniffing the neighboring traffic. The access point can count unique MAC addresses and extract the rates from their PLCP header. We plan to investigate the implementation details in future work.

4 Model Results

Based on the model presented in the previous section, next we present results to evaluate the number of cases where a handover is advantageous for both access

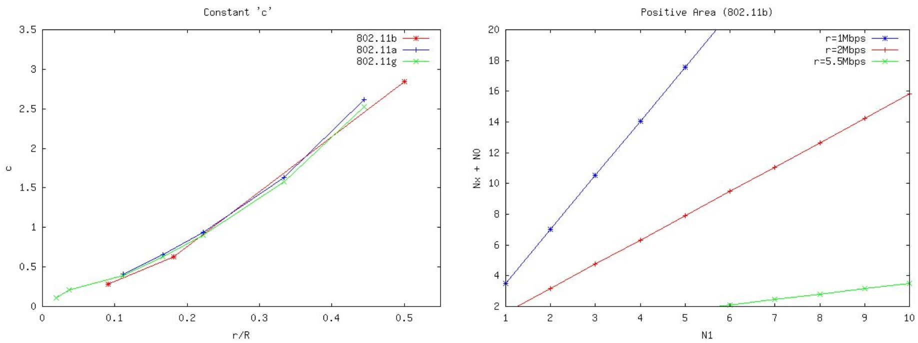

(a) Constant c as a function of r/R. (b) Areas of beneficial handovers.

Fig. 3. Visualization of (8).

points, and quantify the corresponding throughput improvements. We note that we use packet size equal to 1500 bytes for all the following experiments.

Figure 3(a) depicts the value of constant c, as a function of the ratio r/R, for different versions of 802.11, namely 802.11b, 802.11a, and 802.11g; the high rate R is taken to be the highest rate supported by each version. Since for a positive gain scenario the ratio of BS1 to BS0 nodes needs to be greater than this constant, it is obvious that a higher value of c is more restricting. Moreover, since the high rate R is constant for each version, higher values of r make the handover less beneficial. Obviously, the negative impact of low-rate nodes on the throughput is smaller for higher values of r. As expected, the constant is more restricting for 802.11a/g that operate at 54 Mbps. Additionally, 802.11g is slightly less restricting than 802.11a, because the overhead of the latter is slightly higher.

Figure 3(b) depicts the line defined by (8) for 802.11b and its supported rates. The x and y axis shows the number of BS1 and BS0 nodes, respectively. The slope of each line is given by the constant c. Every $N0, Nx, N1$ combination that is below each line can benefit with the handover of low-rate nodes. The area below each line provides a visual estimation of the probability of the appearance of a beneficial scenario, across all possible scenarios that correspond to all $N0, Nx, N1$ combinations. The figure shows that when the low rate is 1 Mbps beneficial scenarios are highly likely. When the low rate is 2 Mbps the beneficial area is cut in half. When the low rate is 5.5 Mbps the beneficial area is very small, and handovers in this case does not improve the throughput for both access points.

Figure 4 depicts the percentage of scenarios where both access points benefit from the handover. We assume that each scenario is a combination of $N0, Nx, N1$, where each set has a uniformly distributed number of nodes in the interval $[1, 10]$. The total amount of different node combinations is equal to 1000. For 802.11b, when the low rate is equal to 1 Mbps the percentage of beneficial scenarios is

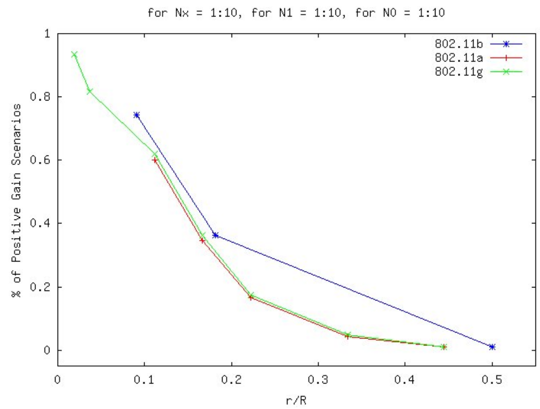

Fig. 4. Percentage of beneficial scenarios.

almost 74%. At 2 Mbps the percentage drops to 40% and at 5.5Mbps the percentage is about 1%. IEEE 802.11a and 802.11g have similar behavior although the percentages are lower, as expected. The figure indicates that when the low rate is less or equal to 12 Mbps there is a significant percentage (more than approximately 35%) of beneficial scenarios. The assumption that number of nodes in each set is uniformly distributed in the interval $[1, 10]$ is, indeed, a simplification. In the real world, the distribution of the number of nodes is expected to be more complicated. However, this simplification provides an estimation on the probability of appearance of a beneficial scenario. Moreover, it can be easily shown that similar conclusions hold for a larger number of nodes.

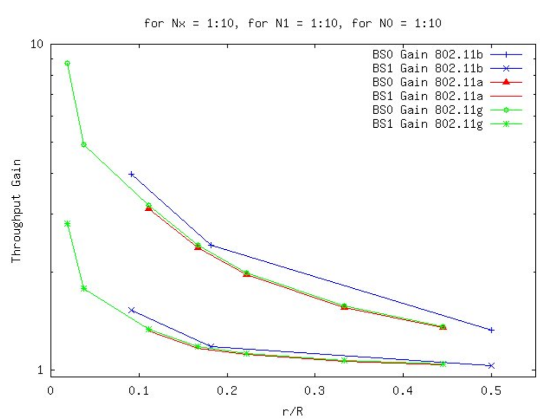

Fig. 5. Throughput gains for BS0 and BS1.

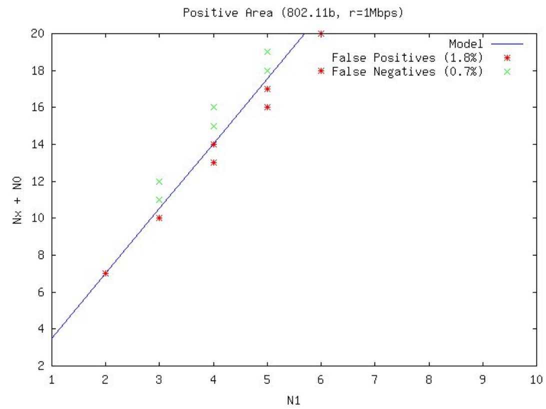

Fig. 6. Comparison of the line with slope c predicted by the model for identifying positive gain scenarios with simulation.

Figure 5 shows that for higher values of the low rate r, the average gain of the beneficial scenarios decreases. The average gain of BS1, for low rate equal to 1 Mbps in 802.11b is about 55%. At 2 Mbps the gain drops to 20%. Observe that the average gain of BS0 is significantly higher than the gain of BS1; this is because BS0's low-rate clients use a part of BS1's channel access time, when they are associated to the latter.

5 Model Evaluation

Next we evaluate the accuracy of the model presented in the previous sections, by comparing its results with those obtained using the NS-2 simulator. The overhead of a transmission is expected to be different in NS-2, since the model does not account for collisions; this simplification will result in lower gains and a lower percentage of beneficial scenarios than those predicted by the model. Note that the collisions do not depend on the number of clients, since we focus on the downlink where the two access points are the sole contenders for channel access.

The default version of NS-2 does not support multiple rates for different receivers and the same transmitter. Hence, we extended NS-2 to support multiple rates, by adapting the rate according to the signal strength of the last received packets. The experiments considered 802.11b, and low rate equal to 1 Mbps.

Figure 6 indicates that with simulation, there is no single threshold that separates the positive scenarios from the negative ones. If the constant c predicted by the model is used for accepting or rejecting handovers, there are 1.8% false positives and 0.7% false negatives, which are shown in the figure as red stars and green crosses, respectively. The false positives and false negatives are very close to the threshold line, and their density is higher for a higher number of nodes.

Figure 7 shows the percentage of false positives and false negatives as a function of the acceptance threshold c. The green line vertical line corresponds to the threshold c predicted by the analytical model, which is equal to 0.285. A threshold equal to 0.287 provides a balance of false positives and false negatives. A more conservative threshold equal to 0.301 yields very few false positives, while a threshold equal to 0.34 yields no false positives.

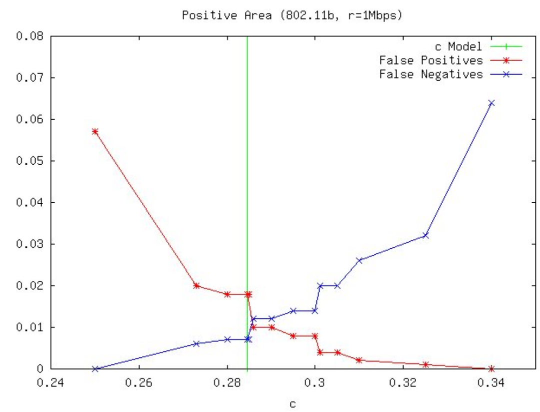

Fig. 7. Percentage of false positives and false negatives for different values of c.

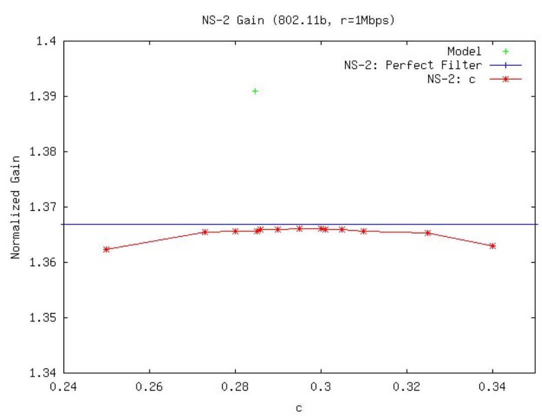

Fig. 8. Normalized gain for different values of c.

Next we consider the normalized gain, which is a function of the acceptance probability (number of accepted scenarios to the total number of scenarios), AP, and the average gain of the accepted scenarios, $AvgGain$:

$$NormGain = AP \cdot AvgGain + (1 - AP) \cdot 1 \tag{9}$$

The normalized gain is a metric that takes into account both the probability of appearance and the gain of a beneficial handover and can be used to evaluate the long term gain of an access point that accepts handovers according to our policy. Figure 8 depicts the impact of false positives and false negatives, for different values of c, on the normalized gain. The straight blue horizontal line is a theoretical optimum filter that perfectly predicts the positive and negative scenarios. This figure shows that the acceptance policy based on the threshold c performs extremely well, giving a normalized gain which is within 0.2% of the maximum gain, which is achieved by the theoretical optimal filter. Moreover, Figure 8 shows that the normalized gain predicted by the model (green '+' at gain of approximately 1.39) differs from the normalized gain estimated using NS-2 by less than 2.5%, indicating that the model is very accurate in predicting the handover gain.

6 ADSL Constraints

Next we extend the model presented in Section 3 to the case where the bottleneck from a wireless user to the fixed network is not the wireless link, but a wired link after the access point, e.g., a user's ADSL connection. For instance, with modern infrastructure in many home wireless networks, which most commonly use IEEE 802.11g, the bottleneck is in the ADSL connection. Although the theoretic maximum throughput of the ADSL technology is 24 Mbps, much lower throughput is common due to various reasons such as signal attenuation.

Assume that C_0 and C_1 is the wired capacity constraint of BS0 and BS1, respectively. In the remainder of this section, due to space limitations, we consider the simple case where both access points have the same wired capacity constraint, $C_1 = C_2 = C$.

Assume XC^a and XC^b is the aggregate throughput of each access point when there is no handover (case a), and when there is handover (case b), respectively, as predicted by the basic model in Section 3. Since case a includes at least one low rate transmission, whereas case b includes none, it is obvious that $XC^a < XC^b$. As a result, there are three zones where the impact of the wired capacity constraint differs. If $C < XC^a$, then the wired capacity is the bottleneck, and the existence of low rate transmitters does not affect performance, hence there are no performance incentives for handover. If $XC^b < C$, then the wired capacity does not affect the performance results presented in the previous sections. The intermediate case where $XC^a < C < XC^b$ is when the capacity constraint has an impact that depends on other factors. When this holds, we find that case b is beneficial for both access points when the following inequality holds:

$$N1 > \frac{Nx(Nx+N0)rR}{C(Nx(r+R)+2N0r+\frac{(N0+Nx)ohrR}{pkt})-(N0+Nx)rR} \,. \tag{10}$$

¿From this equation we see that the number of beneficial scenarios increases, as the wired capacity constraint C increases.

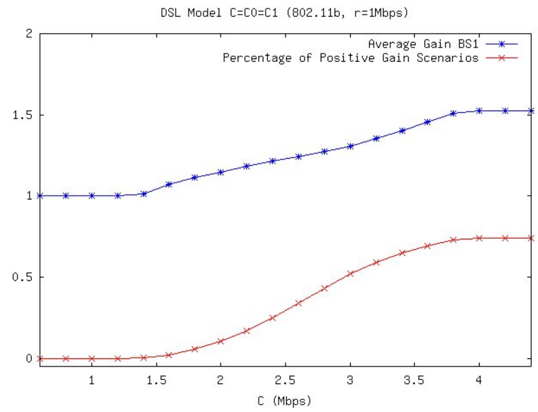

Fig. 9. Effect of the wired capacity constraint on the number of beneficial scenarios and their gain.

Figure 9 depicts, for 802.11b and $r = 1$ Mbps, the percentage of beneficial scenarios and the corresponding average gain from the perspective of BS1, for different values of the wired capacity constraint C. Experiments for all combinations of $N0, N1, Nx$ are performed, where each of these variables obtains values in $[1, 10]$. As expected, the three zones are clearly visible. When C is less than 1.3 Mbps there are no beneficial scenarios. As C increases above this value, the percentage of beneficial scenarios and the corresponding gain increases until 4 Mbps, when the gain is equal to that of the basic model without wired capacity constraints.

An interesting conclusion is that the wired capacity constraint reduces handover gains when its value is less than the half of the maximum effective throughput of the wireless link. Since the most common wireless protocol used in home wireless networks is IEEE 802.11g, which has an effective throughput of about 20-24 Mbps, an ADSL connection which is below 10-12 Mbps would decrease any potential handover gains. In the future, it is likely that we will have a mix of such low rate ADSL connections, together with higher rate connections supported by the deployment of fiber closer to the customers' premises.

7 Conclusion and Future Work

In this paper, we analyzed scenarios where the cooperation between public/home wireless network administrators is motivated solely by the performance improvement for their clients. Using an analytical throughput model with two access points, we computed the gains for both access points, when low rate users are handed-off to the closer access point. Moreover, the analytical model suggests a simple acceptance policy for deciding when handovers are beneficial. The analytical model was compared with simulation results, which verify the accuracy of the model. Finally, we extended the model to include wired capacity constraints, e.g., from ADSL connections.

Ongoing work is investigating the use of priority mechanisms supported by standards such as IEEE 802.11e [14], to control the sharing of the handover gains among the involved access points; such control can also increase the number of beneficial scenarios. We plan to study the details of the implementation of the proposed procedure. We also plan to extend our study with alternative performance metrics, which incorporate notions of fairness in the allocation of throughput among the wireless nodes, and to the case of multiple access points and traffic in the uplink direction. Finally, we plan to investigate the existence of similar performance improvement incentives in the case of wireless mesh networks with overlapping coverage.

References

1. Heusse, M., Rousseau, F., Berger-Sabbatel, G., Duda, A.: Performance anomaly of 802.11b. In: Proc. of IEEE INFOCOM. (2003)
2. Akella, A., Judd, G., Seshan, S., Steenkiste, P.: Self-Management in Chaotic Wireless Deployments. In: WINET Journal. Vol. 13, No. 6, 737–755 (2007)
3. Siris, V.A., Stamatakis, G.: Optimal CWmin Selection for Achieving Proportional Fairness in Multi-rate 802.11e WLANs: Test-bed Implementation and Evaluation. In: Proc. of ACM WiNTECH (2006)
4. Tan, G., Guttag, J.: Time-based Fairness Improves Performance in Multi-Rate WLANs. In: Proc. of USENIX Annual Technical Conference. (2004)
5. Feeney, L.M., Cetin, B., Hollos, D., Kubisch, M., Mengesha, S., Karl, H.: Multi-rate relaying for performance improvement in IEEE 802.11 WLANs. In: Proc. of WWIC. (2007)
6. Bahl, V., Chandra, R., Lee, P.P.C., Misra, V., Padhye, J., Rubenstein, D., Yu, Y.: Opportunistic Use of Client Repeaters to Improve Performance of WLANs. In: Proc. of ACM CoNEXT. (2008)
7. Kandula, S., Ching-Ju Lin, K., Badirkhanli, T., Katabi, D.: FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput. In: Proc. of the 5th USENIX Symposium on Networked Systems Design and Implementation (2008)
8. Liu, P., Tao, Z., Narayanan, S., Korakis, T., Panwar, S.: A Cooperative MAC protocol for Wireless LANs. In: IEEE JSAC. Vol. 25, No. 2 (2007)
9. Kumar, A., Altman, E., Miorandi, D., Goyal, M.: New Insights from a Fixed Point Analysis of Single Cell IEEE 802.11 WLANs. In: Proc. of IEEE INFOCOM. (2005)
10. Kumar, A., Kumar, V.: Optimal Association of Stations and APs in an IEEE 802.11 WLAN. In: Proc. of National Conference on Communications (NCC). (2005)
11. Kasbekar, G., Kuri, J., Nuggehalli, P.: Online Association Policies in IEEE 802.11 WLANs. In: Proc. of 4th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt). (2006)
12. Kauffmann, B., Baccelli, F., Chaintreau, A., Mhatre, V., Papagiannaki, K., Diot, C.: Measurement-based Self Organization of Interfering 802.11 Wireless Access Networks. In: Proc. of IEEE INFOCOM. (2007)
13. Koukoutsidis, I., Siris, V.A.: Access Point Assignment Algorithms in WLANs based on Throughput Objectives. In: Proc. of 6th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt). (2008)
14. Mangold, S., Choi, S., Hiertz, G. R., Klein, O.: Analysis of IEEE 802.11e for QoS Support in Wireless LANs. In: IEEE Wireless Communications. Vol. 10, No. 6, 40–50 (2003)