



# Network Cooperation for Client-AP Association Optimization

Akash Baid, Michael Schapira, Ivan Seskar, Jennifer Rexford, Dipankar  
Raychaudhuri

► **To cite this version:**

Akash Baid, Michael Schapira, Ivan Seskar, Jennifer Rexford, Dipankar Raychaudhuri. Network Co-operation for Client-AP Association Optimization. WiOpt'12: Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, May 2012, Paderborn, Germany. pp.431-436, 2012. <hal-00766110>

**HAL Id: hal-00766110**

**<https://hal.inria.fr/hal-00766110>**

Submitted on 18 Dec 2012

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Network Cooperation for Client-AP Association Optimization

Akash Baid\*, Michael Schapira†, Ivan Seskar\*, Jennifer Rexford‡, Dipankar Raychaudhuri\*

\*WINLAB, Rutgers University, {baid, seskar, ray}@winlab.rutgers.edu

†Hebrew University of Jerusalem & Google NYC, schapiram@huji.ac.il

‡Princeton University, jrex@cs.princeton.edu

**Abstract**—In a WiFi deployment with multiple access points, optimizing the way each client selects an AP from amongst the available choices, has a significant impact on the realized performance. When two or more such multi-AP networks are deployed in the same region, APs from different networks can cause severe interference to one another. In this paper, we study how *inter-network* interference effects the *intra-network* association optimization and propose a cooperative optimization scheme to mitigate the interference. We model the interference between multiple overlapping WiFi deployments, determine the information that networks need to share, and formulate a non-linear program that each network can solve for optimal proportional-fair association of clients to APs. Assuming a ‘sum of log rates’ utility function, we apply a known  $2 + \epsilon$  approximation algorithm for solving the NP-hard problem in polynomial time. We evaluate the performance gain through large-scale simulations with multiple overlapping networks, each consisting of 15-35 access points and 50-250 clients in a 0.5x0.5 sq.km. urban setting. Results show an average of 150% improvement in random deployments and upto 7x improvements in clustered deployments for the least-performing client throughputs with modest reductions in the mean client throughputs.

## I. INTRODUCTION

To satisfy the exponential growth in mobile data demand [1], mobile operators, broadband Internet providers and stand-alone services such as Boingo and IPass are deploying large-scale WiFi hot-spot networks in urban areas worldwide. This expanding class of WiFi hot-spot access points, expected to grow at more than 350% over the next four years [2], are generally managed through central controllers via solutions such as those from Meraki [3], Ruckus Wireless [4], and BelAir Networks [5]. While performance improvement through centralized control is now standard practise in enterprise WiFi networks [6], [7], [8], [9], the benefits of central control in public locations with multiple overlapping managed networks is unclear due to inter-network interference. In this paper, we study the effect of *inter-network* interference on the *intra-network* optimization problem that each network solves. Further, we propose a scheme for cooperative sharing of channel access information between the different overlapping networks and show the benefits of such cooperation for all the networks involved.

Figure 1 shows a real-world example of overlapping WiFi AP deployments of two leading broadband Internet providers in a  $\sim 1$  sq. km. cross section of the Brooklyn area in New York, USA, compiled using their respective WiFi location

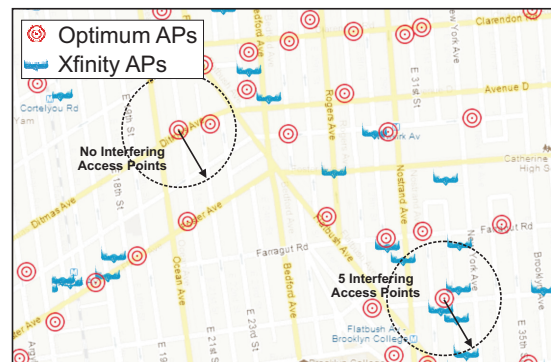


Fig. 1: WiFi AP locations of two providers in Brooklyn, NY

finder services [10], [11]. The exact nature of inter-network interference on the client throughputs in such a scenario depends on the number of co-channel APs, their transmit powers, rate allocation algorithm, and MAC parameter selection. However, inference of these channel access parameters through passive observations is a hard problem and often requires active probing [12]. A key challenge in passive interference estimation is to incorporate the large variation in the number of interferers - for example the number of potential interferers, i.e. Xfinity APs surrounding an Optimum AP in Figure 1 varies from 0 to 5, even in this small scenario. Inference of active interferers becomes even more challenging considering the reality of tens of networks, non-beaconing APs, and dynamic channel selection.

In order to alleviate this inter-network interference, we propose a back-end *operational cooperation* between the networks: each network periodically shares the information about the location and operating channels of its APs with all other networks operating in the same area. Note that clients belonging to one network cannot join other networks in this model. In this paper, we do not focus on the messaging interfaces, which could be implemented either in a distributed fashion, i.e. each AP sends a message to the neighboring APs of other networks or centrally, i.e. the aggregated information is passed through a single interface between networks. Rather, assuming the presence of such side-channel information, we show how each network can optimize client-AP associations to minimize the effects of inter-network interference. Within the scope of the traffic model described in Section III, this form of information exchange followed by *intra-network* optimization




Fig. 2: Qualitative example of why cooperation is needed: (a) Default mode chooses the least distance AP, (b) Intra-network optimization selects least crowded AP if their distance to the client are comparable, (c) Cooperative optimization selects AP with least internal and external interference

is the same as a global optimization considering all APs of all networks as being controlled by a single entity. This follows from the fact that for certain problem formulations, the interference terms in the intra-network problem can be summarized and substituted using the information received from neighboring networks. To the best of our knowledge, such forms of cooperation between multiple managed WiFi networks has received very little attention with only some recent works in the related area of cellular networks [13].

While operational cooperation can be devised for optimizing the channel selection, rate allocation, power control, and back-off windows, we focus on the more tractable case of optimizing client-AP associations: given a set of APs that a client can potentially connect to, selecting the best AP so as to maximize the sum utility of all the clients across all the network. Due to its direct impact on both the client experience (in terms of throughput) as well as the network performance (in terms of traffic load), this problem has been approached through both centralized network utility maximization framework [6], [8] and game-theoretic formulations [14]. In particular, we follow the proportional fairness framework developed in [6] for the basic intra-network optimization and enhance it to incorporate inter-network interference.

**Operational vs. Access Cooperation:** While we propose the operational form of cooperation in this paper, it is important to compare it with a simpler form of cooperation which can be termed *access cooperation*. Through access cooperation between two networks, unlike our assumption, clients of one network can join the other network. While this can increase the coverage area for both the networks, we show that unless the two networks also jointly manage their networks, i.e. solve a global optimization problem, network utility cannot be maximized by only access cooperation. In addition, operational cooperation has three distinct advantages over access cooperation: (i) a network does not have to handle authentication for clients from other networks, (ii) networks do not have to over-provision capacity since they do not have to cater to extra clients and (iii) networks can retain the control of sessions, policy, and billing of their clients.

The rest of the paper is organized as follows. In Sec-

tion II, we describe why cooperation is required in a multi-network hot-spot deployment and outline our system model in Section III. Next, we formulate the optimization problem in Section IV and present simulation results in Section V. Section VI concludes the paper and outlines the future work.

## II. MOTIVATION

Figure 2 shows an illustrative example of cooperation gain. Client  $C1$  is in communication range of three APs of the same network; and the default 802.11 rule as shown in Figure 2(a) is to choose the closest AP (here  $AP1$ ), which gives the highest rate to the client. However, if there is another client  $C2$  attached to  $AP1$ ,  $AP1$  has to divide its downlink transmission time between the two clients, as in Figure 2(b). Assuming proportional fair scheduling, the real throughput that  $C1$  gets from  $AP1$  is only 27 Mbps. Intra-network optimization through a central controller (e.g., Aruba WLAN controllers [15]) can identify this load imbalance and connect  $C1$  to  $AP2$  instead and allow the client to get a throughput of 48 Mbps. In doing so, the network controller assumes that  $AP2$  has sole control of the channel. However in a multi-network setting, a foreign network may have a nearby AP that shares  $AP2$ 's channel. CSMA contention leads to approximately equal time share between the two APs, leading to an actual throughput of only 24 Mbps for  $C1$  if connected to  $AP2$ , as shown in Figure 2(c). Cooperative optimization incorporates the effect of APs of other networks and thus connects  $C1$  to  $AP3$  leading to a throughput of 36 Mbps.

## III. SYSTEM MODEL

We consider a system with  $N$  independently operated WiFi networks with  $U_i$  and  $A_i$  denoting the set of clients and APs in the  $i$ th network respectively. Table I summarizes the notations we use in this paper. Binary variables  $x_{ij}(k)$  indicate the connection state between the  $j$ th client and  $k$ th AP of the  $i$ th network (1 is connected, 0 if not), while  $p_{ij}(k)$  denote the fraction of time provided by the AP to the client. Similarly,  $r_{ij}(k)$  denotes the effective bit rate which includes the retransmission overhead and is assumed to be a step-wise function of the distance between the client and the AP, as specified in the 802.11 standard. Since air time fraction and rate are relevant only for clients connected to an AP,

Symbol	Meaning
$N$	No. of WiFi networks
$U_i$	Set of clients in network $i$
$A_i$	Set of access points in network $i$
$R_{cs}$	Carrier sense radius (equal for all APs)
$R_{int}$	Interference radius (equal for all APs)
$B_{ik}$	Set of co-channel foreign APs within $R_{cs}$ of $k$ th AP of $i$ th network
$C_{ik}$	Set of co-channel foreign APs outside $R_{cs}$ but within $R_{int}$ of $k$ th AP of $i$ th network
$\eta_{ik}$	Number of clients connected to the $k$ th AP of $i$ th network
$r_{ij}(k)$	Wireless PHY rate obtained by the $j$ th client of $i$ th network when connected to the $k$ th AP of that network
$x_{ij}(k)$	Association indicator between the $j$ th client of $i$ th network and its $k$ th AP (value = 0 or 1)
$p_{ij}(k)$	Fraction of time the $j$ th client of $i$ th network gets from its $k$ th AP

TABLE I: Definition of parameters

$p_{ij}(k) = 0$  and  $r_{ij}(k) = 0$  whenever the corresponding  $x_{ij}(k) = 0$ . Thus the  $j$ th client of the  $i$ th network has an effective downlink rate of  $\sum_{k \in A_i} r_{ij}(k)x_{ij}(k)p_{ij}(k)$ .

As is common in commercial WLAN controllers [15], each AP employs a proportional fairness policy. Ignoring the protocol overheads and assuming equal priorities for all clients, proportional fairness translates to equal time share between clients in multi-rate WLAN [16]. Thus for the  $k$ th AP of the  $i$ th network, each of its  $\eta_{ik}$  clients receive a fraction  $1/\eta_{ik}$  the APs airtime. We focus on downlink traffic which forms the majority of WiFi data transmission [17] and assume clients always have pending data requests at the AP. This assumption simplifies the estimation of the client rates significantly and is valid in hot-spot deployments where the number of clients is large enough that each client cannot receive its maximum desired data rate.

In order to account for the inter-network interference, we denote the set of co-channel foreign APs within carrier sense range of the  $k$ th AP of  $i$ th network as  $B_{ik}$  and those outside carrier sense but within interference range (potential hidden nodes) as  $C_{ik}$ . Each AP has to participate in CSMA and thus shares the channel with co-channel APs within its carrier sense radius. We assume that within each network, frequency planning is such that no two APs within carrier sense distance are assigned the same channel. Thus the  $k$ th AP of the  $i$ th network has to share its channel with  $|B_{ik}|$  other APs, bringing its share of the channel access time fraction to approximately  $1/(1+|B_{ik}|)$  [18]. Further we model the hidden node interference (interference from APs outside the carrier sense range but with signals still strong enough to affect ongoing transmissions) by lowering the channel access time further. We introduce a parameter  $\alpha \in [0, 1]$  which captures the average effect of hidden node interference per interferer. The channel access time fraction for the  $k$ th AP of the  $i$ th network is thus also reduced by a factor of  $1/(1+\alpha|C_{ik}|)$ . The choice of the  $\alpha$  parameter can either be made through probe experiments during the deployment stage or be pre-set to the values derived through testbed measurements [19], [20].  $\alpha$  values in the (0.2, 0.6) range satisfy most of our past

experiments on the ORBIT testbed [21].

The objective of the intra-network association optimization, given such a model, is to optimize the set of  $x_{ij}(k)$  variables for maximum utility which we choose to be one which results in proportional fairness. The choice of a proportional fair utility function is a de facto standard in the current EV-DO, 3G cellular systems, as well as in emerging 4G systems based on LTE and WiMAX and has been shown to provide a good balance between resource utilization and fairness of allocation [6], [8], [22]. For cooperative optimization, each network first ascertains the values of  $|B_{ik}|$  and  $|C_{ik}|$  for each of its APs through periodic message exchange with other networks. This information is then used to formulate a similar optimization problem as in the case of intra-network optimization. Note however, that by including the hitherto unknown interference components, the cooperative problem formulation now matches the real interference scenario.

#### IV. PROBLEM FORMULATION AND SOLUTION

##### A. Individual Network Optimization

The intra-network non-cooperative optimization problem formulation is similar to the description in [6]. Since  $x_{ij}(k)$  equals 1 only if client  $j$  is associated with AP  $k$  and channel access time is equally divided between clients connected to an AP, the association optimization within network  $i$  can be denoted by:

$$\begin{aligned}
& \text{Maximize:} && \sum_{j \in U_i} \log \left( \sum_{k \in A_i} r_{ij}(k)x_{ij}(k)p_{ij}(k) \right) \\
& \text{subject to:} && p_{ij}(k) = \frac{1}{\sum_{j' \in U_i} x_{ij'}(k)} \quad \forall k \in A_i, j \in U_i \\
& && \sum_{k \in A_i} x_{ij}(k) = 1 \quad \forall j \in U_i \\
& && x_{ij}(k) \in \{0, 1\} \quad \forall k \in A_i, j \in U_i
\end{aligned} \tag{1}$$

Here the first constraint models the proportional fairness policy of each AP and makes the problem non-linear in  $x_{ij}(k)$  while the second constraint along with the binary constraint restricts each client to connect to exactly one AP. Note that the  $p_{ij}(k)$  in (1) is not the actual time fraction that the client would receive as it does not capture the effect of foreign APs. But without any cooperation, each network has no idea about the number/location of such APs and thus uses this value. Reference [6] shows an efficient approximation algorithm to solve this NP-hard non-linear integer problem for a slightly different problem formulation. This method first requires converting (1) to a relaxed discretized linear program without the integrality constraint on  $x_{ij}(k)$ , i.e., each client is allowed to connect to multiple APs simultaneously. Then the rounding process described by Shmoys and Tardos for the generalized assignment problem [23] is used to arrive at binary values. This polynomial time 2-approximate rounding algorithm thus results in a total utility bounded below by that of the optimal assignment scaled down by a factor of  $2 + \epsilon$ .

## B. Cooperative Optimization

Extending the above formulation based on the assumptions of equal time sharing MAC and availability of  $|B_{ik}|$  and  $|C_{ik}|$  values, the global association optimization problem can be written as:

$$\begin{aligned} \text{Maximize: } & \sum_{i=1}^N \sum_{j \in U_i} \log \left( \sum_{k \in A_i} r_{ij}(k) x_{ij}(k) p_{ij}(k) \right) \\ \text{subject to: } & p_{ij}(k) = \frac{1}{\sum_{j' \in U_i} x_{ij'}(k)} \cdot \frac{1}{(1 + |B_{ik}|)(1 + \alpha |C_{ik}|)} \\ & \forall k \in A_i, j \in U_i, i \in [1, N] \\ & \sum_{k \in A_i} x_{ij}(k) = 1 \quad \forall j \in U_i, i \in [1, N] \\ & x_{ij}(k) \in \{0, 1\} \quad \forall k \in A_i, j \in U_i, i \in [1, N] \end{aligned} \quad (2)$$

The constraints in (2) are a simple extension to those in (1). Note here that the first term in  $p_{ij}(k)$  is directly dependent on the optimization variables  $x_{ij}(k)$ . However  $|B_{ik}|$  and  $|C_{ik}|$  are only dependent on the relative placement of co-channel APs of different networks and are thus constants given a certain topology. So once each network  $i$  knows about the  $|B_{ik}|$  and  $|C_{ik}|$  values for each of its AP  $k$ , it can individually solve the association problem. This joint problem can be solved using the same technique as the individual network optimization.

## V. SIMULATION RESULTS

We compare three cases to quantify gains -

- Least Distance: Each client connects to the closest AP of the same network (benchmark case).
- Intra-Network Optimization: Each network optimizes the association pattern of its clients.
- Cooperative Optimization: All networks share information for optimizing the client association.

Note that in all the three cases we assume that the clients belonging to a network can only connect to APs from that network. The discretized linear program was solved using the open source *lpsolve* solver [24]. All the results presented are averaged over 10 simulation runs. We present results for two deployment scenarios: random deployment and clustered deployment as follows.

### A. Random Deployment

Multiple overlapping networks are considered in a 0.5x0.5 sq. km area, which reflect deployment scenarios in urban hot-spot networks, multi-tenant buildings, or airports. Each network has a variable 15-25 APs placed at uniformly randomly selected points. While there is a minimum separation of 50 meters between two APs of the same network, there is no such restriction for APs of different networks. Reasonable frequency planning is assumed - each AP chooses one of the three orthogonal channels in the 2.4 GHz range to minimize the number of co-channel APs. However due to dense deployment of multiple overlapping networks, choosing a completely isolated channel is seldom possible. Clients are placed at random within the area with the total number of clients of

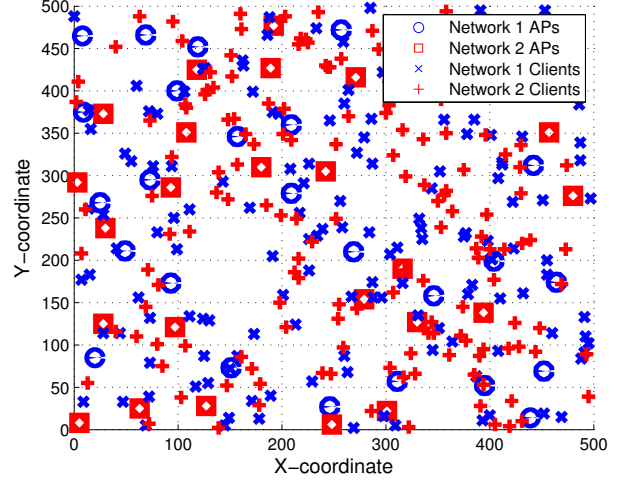


Fig. 3: Topology showing an instance of 2 Networks, 25 APs per network and 150 clients per network at random points

each network set as a parameter. Figure 3 shows an instance of the random AP and client placement. The carrier sense and interference range thresholds of all devices are set to 215 meters and 250 meters respectively as per the specifications in [25]. The physical data rates  $r_{ij}(k)$  are selected based on the distance between the client  $j$  and AP  $k$ , also from [25]. The value of the interference scaling parameter  $\alpha$  is taken as 0.5.

Figure 4 shows the cumulative distribution of the client throughputs for all the clients in the system for the topology shown in Figure 3. The plot shows that while intra-network optimization improves fairness in client throughput, its effect is limited due to the presence of APs of another network. Cooperative optimization more than doubles the 10 percentile throughput from 230 Kbps to 550 Kbps compared to least distance scheme and shows a 77% gain when compared to the same metric in intra-network optimization. Since the cooperative optimization problem (2) decouples into separate problems for each network, utility of each network is individually maximized.

Figure 5 further dissects the comparison between intra-network and cooperative optimization schemes. In this figure, clients are arranged in the increasing order of the throughput they get through intra-network optimization. The key observation here is that almost all lowest throughput clients are better off after cooperative optimization, while the accompanying loss in throughput is inflicted primarily on the clients with high throughputs.

Figure 6 shows the 10 percentile and mean throughput values for simulations with  $N = \{2, 3, 4\}$ , 25 APs, and 150 clients. We note that in each of the cases, the 10 percentile throughputs improve by 140-170% with a small 8-10% decrease in the mean throughput. The achievable mean throughput naturally goes down with increasing  $N$  due to sharing of the spectrum between a larger number of users. Table II shows the effect of variations in the number of APs



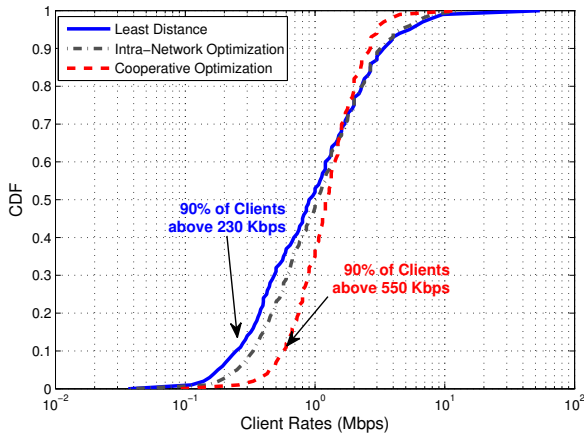


Fig. 4: CDF of client throughputs for all clients.  $N = 2$ ,  $|A_i| = 25$ ,  $|U_i| = 150$




Fig. 5: Per-client comparison between intra-network and cooperative optimization schemes

and clients per network for the case of  $N = 3$ . The key observation here is that the percentage gain brought about due to cooperation increases with AP density, but decreases with client density. The insight from these trends suggests that higher AP densities lead to greater uncertainties that each network has to cope with and thus the information sharing becomes more valuable. However, under a capacity limited regime with large number of users, since all APs are heavily crowded, the relative gain of shifting clients from one AP to another reduces.

### B. Clustered Deployment

Clustered deployments, characterized by a large number of APs placed in a targeted small region are commonly used to serve public places with very high number of peak users, e.g., waiting rooms, mall entrance, etc. In order to study the effects of such topology-specific interference patterns, we considered a clustered topology with two networks. APs of the first network are clustered in three rectangular regions of size 200x200 meters each, while the second network still has a random AP deployment. All other access parameters remain the same as in the random deployment case. Figure 7 shows the CDF of the client throughputs for each network. We observe that since network 1 APs are strongly clustered,




Fig. 6: 10 percentile and mean client throughput values for varying  $N$  with  $|A_i| = 25$ ,  $|U_i| = 150$

$ A_i $	$ U_i $	10 %ile throughput(Mbps)			Mean throughput(Mbps)		
		Least Dist.	Intra Optim.	Coop. Optim.	Least Dist.	Intra Optim.	Coop. Optim.
15	150	0.09	0.13	0.19	0.7	0.66	0.62
25	150	0.1	0.14	0.27	0.78	0.77	0.71
35	150	0.11	0.14	0.31	0.85	0.85	0.77
25	50	0.21	0.33	0.64	1.95	2.17	2
25	150	0.1	0.14	0.27	0.78	0.77	0.71
25	250	0.07	0.09	0.17	0.49	0.47	0.43

TABLE II: 10 percentile and mean client throughput values for varying number of APs and clients with  $N = 3$

the relative effect of network 2 APs on its performance is minimal. Hence cooperative optimization does not improve the client throughputs for this network. Conversely, network 1 clusters strongly effect the performance of network 2, thus cooperating between the two networks leads to large gains for network 2.

### C. Comparison with Access Coordination

A simple alternative cooperation scheme in a multi-network scenario is *access coordination* in which two or more networks agree to allow each others' clients to access their networks. Each client can now connect to the nearest AP of any network. In order to compare the operational cooperation scheme proposed in this paper with an access coordination scheme, we reuse the topology in Figure 3 but allow clients to connect to APs in either network. Figure 8 shows the throughput of each client under the three association schemes with the client indices arranged in the order of increasing throughput. We note that, access cooperation leads to a decrease in the shortest distance between an AP and a client and thus gives higher throughput for almost all clients. However, since access to more APs does not solve the load balancing problem, operational cooperation results in better performance for more than 2/3rd of the lowest throughput clients.

## VI. CONCLUSION

Overlapping managed WiFi networks are increasingly common in urban hot-spot deployments, multi-tenant buildings, office complexes, and airports. While optimizing a single network has been the subject of a number of studies in




Fig. 7: CDF of client throughputs for network 1 (clustered APs) and network 2 (randomly deployed APs).  $|A_i| = 25$ ,  $|U_i| = 150$

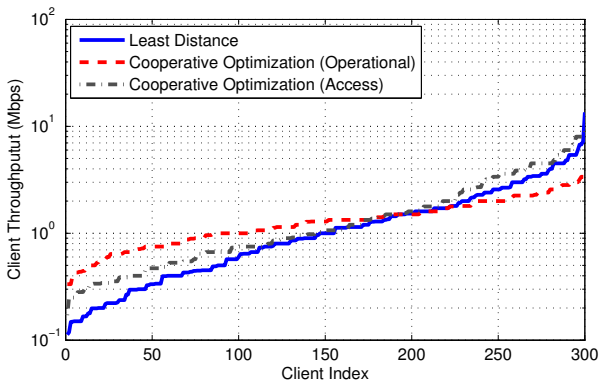


Fig. 8: Throughput of each client for different association schemes (sorted)

the last decade, such inter-network interactions are not well understood. In this paper, we present a first step towards that and study the effect of inter-network interference on the client-AP association optimization problem. Since such planned AP deployments are designed to support a large number of users, balancing the number of clients associated to each AP is important. We show that ignoring the presence of other networks leads to significant throughput degradation, especially for clients at the edge of an APs coverage region. To alleviate this problem, we propose an operational cooperation model under which all networks share the information about the location and operating channel for all their APs. We show that incorporating this information for client-association optimization within a network leads to 140-170% improvements

in the 10 percentile client throughputs when clients and APs are randomly placed. Clustered AP deployments lead to a much higher gain of upto 7x since the value of the shared information increases significantly.

We plan to extend the operational cooperation framework to include channel selection and transmit power optimization. Studying the effects of non-saturated traffic patterns and a mix of uplink and downlink transmissions would also lead to a more pragmatic solution.

## REFERENCES

- [1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011-2016," Cisco White Paper, Feb. 2012.
- [2] PC World News: Number of Wi-Fi Hotspots to Quadruple by 2015, <http://bit.ly/u0x0VP>.
- [3] Meraki, FlexMaster Centralized Wi-Fi Management, [http://www.meraki.com/lib/pdf/meraki\\_whitepaper\\_architecture.pdf](http://www.meraki.com/lib/pdf/meraki_whitepaper_architecture.pdf).
- [4] Ruckus Wireless White Paper: Cloud Controlled Architecture, <http://www.ruckuswireless.com/products/flexmaster>.
- [5] BelAir Networks, Small wireless systems for big networks, <http://www.belairnetworks.com/>.
- [6] L. Li, M. Pal, and Y. Yang, "Proportional Fairness in Multi-Rate Wireless LANs," in *Proceedings of INFOCOM, 2008*, pp. 1004–1012.
- [7] E. Rozner, Y. Mehta, A. Akella, and L. Qiu, "Traffic-aware channel assignment in wireless LANs," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 11, pp. 43–44, April 2007.
- [8] I.-H. Hou and P. Gupta, "Distributed Resource Allocation for Proportional Fairness in Multi-Band Wireless Systems," in *Proceedings of ISIT, 2011*.
- [9] V. P. Mhatre and K. Papagiannaki, "Optimal design of high density 802.11 WLANs," in *Proceedings of ACM CoNEXT, 2006*, pp. 8:1–8:12.
- [10] Optimum WiFi locations, <http://www.optimum.net/WiFi/Find>.
- [11] Xfinity WiFi locations, <http://www.comcast.com/wifi/hotspot.htm>.
- [12] N. Ahmed, U. Ismail, S. Keshav, and K. Papagiannaki, "Online estimation of RF interference," in *Proceedings of ACM CoNEXT, 2008*.
- [13] S. Deb, K. Nagaraj, and V. Srinivasan, "MOTA: Engineering an operator agnostic mobile service," in *Proceedings of MobiCom, 2011*, pp. 133–144.
- [14] O. Ercetin, "Association games in IEEE 802.11 wireless local area networks," *Wireless Communications, IEEE Transactions on*, vol. 7, no. 12, pp. 5136–5143, Dec. 2008.
- [15] Aruba Networks White Paper, ARM Yourself to Increase Enterprise WLAN Data Capacity, <http://bit.ly/5nrjZ>.
- [16] L. B. Jiang and S. C. Liew, "Proportional fairness in wireless LANs and ad hoc networks," in *Proceedings of IEEE WCNC, 2005*, vol. 3, pp. 1551–1556.
- [17] K. Sundaresan and K. Papagiannaki, "The need for cross-layer information in access point selection algorithms," in *Proceedings of ACM SIGCOMM, 2006*, pp. 257–262.
- [18] S. Choi, K. Park, and C.-k. Kim, "On the performance characteristics of WLANs: revisited," in *Proceedings of ACM SIGMETRICS, 2005*, pp. 97–108.
- [19] Z. W. Zhao and Y. C. Tay, "A model for calculating channel share of 802.11 access points with overlapping wireless cells," *Wireless Networking, IEEE Transactions on*, vol. 17, pp. 1581–1593, Oct. 2011.
- [20] S. Maniportsut, B. Landfeldt, and A. Boukerche, "Improving densely deployed wireless network performance in unlicensed spectrum through hidden-node aware channel assignment," *Perform. Eval.*, vol. 68, no. 9, Sep. 2011.
- [21] D. Raychaudhuri *et al.*, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," in *Proceedings of IEEE WCNC, 2005*, vol. 3, pp. 1664–1669.
- [22] Q. Wu and E. Esteves, "The CDMA2000 high rate packet data system," in *Advances in 3G Enhanced Technologies for Wireless Communications, 2002*.
- [23] D. B. Shmoys and E. Tardos, "An approximation algorithm for the generalized assignment problem," *Math. Program.*, vol. 62, pp. 461–474, December 1993.
- [24] "lp\_solve, a Mixed Integer Linear Programming (MILP) solver," <http://lpsolve.sourceforge.net/>.
- [25] Cisco Aironet 1200 Data sheet, <http://bit.ly/dv9Mbl>.