

CADWAN – A Control Architecture for Dense Wi-Fi Access Networks

Pierluigi Gallo¹, Katarzyna Kosek-Szott², Szymon Szott², Ilenia Tinnirello¹

¹Università di Palermo, Palermo, Italy

²AGH University of Science and Technology, Faculty of Computer Science, Electronics and Telecommunications, Krakow, Poland

pierluigi.gallo@unipa.it, {kosek, szott}@kt.agh.edu.pl, ilenia.tinnirello@unipa.it

Abstract. The growing demands of ubiquitous computing are leading towards the densification of wireless access networks. The challenges of high density deployments can be addressed by network-wide centralized control. To this end we propose CADWAN – a Control Architecture for efficient management of Dense Wi-Fi Access Networks. Its main advantages are: flexibility (it supports software-defined wireless networking), scalability (it uses a three-tier optimization framework), and extendibility (it exploits a unified control interface with support for heterogeneous devices). Furthermore, CADWAN is complementary to ongoing developments in IEEE 802.11, especially 802.11ax.

Keywords. 802.11ax, CAPWAP, dense networks, optimization, SDN, Wi-Fi.

1 Introduction

Wireless access networks need to meet the growing demands of ubiquitous computing. Cisco’s “Visual Networking Index” estimates over 11.6 billion mobile devices with a total of 30.6 exabytes of monthly traffic by 2020. One of the main solutions to the emerging capacity problem is the densification of access networks. For cellular networks, densification implies smaller cells and the coexistence of umbrella cells with heterogeneous macro/micro/femto coverage areas. For Wi-Fi networks, densification comprises both increasing the number of access points (APs) in an area as well as the proliferation of client devices. In both cases densification has several drawbacks:

- increased signal-to-noise ratio, especially for edge users, caused by inter-cell¹ interference,
- degraded performance of random access methods caused by the increased number of contending users (even belonging to independent, but overlapping, cells),
- increased chances of topological situations (e.g., flow-in-the-middle²) in which the carrier sense mechanism does not work well,

¹ We define a Wi-Fi cell as a network consisting of an AP and its associated stations (i.e., user devices).

² The flow-in-the-middle problem occurs when a cell is located between two others which do not sense each other. If the outer cells are under heavy traffic load, the middle cell senses the medium permanently busy and is thus prevented from accessing the channel.

- prominence of capture effect (the correct reception of the strongest frame among concurrent transmissions) which causes uneven capacity distribution.

Therefore, it is not always obvious that densification of Wi-Fi cells increases capacity as expected. The performance improvements could be minimal or, in special cases, completely negligible.

To mitigate these problems, different levels of coordination can be considered. For Long-Term Evolution (LTE), examples include the enhanced intercell interference coordination and coordinated multipoint extensions [1] as well as new standardized control architectures [2]. For Wi-Fi, examples include improving spatial reuse with color coding (802.11ah), or maintaining two network allocation vectors for intra- and inter-cell connections (802.11ax). Additionally, among the proposed solutions is the concept of *centralized control* of such networks. Studies show that for Wi-Fi networks centralized approaches using global network knowledge can outperform channel assignment based on limited local knowledge [3]. However, other Wi-Fi aspects also require centralized coordination including radio configuration (e.g., setting the optimal transmission power [4]) and medium access optimization (e.g., controlling the carrier sensing threshold [5]), which are usually considered on a local (per-AP) scale.

Centralized control of wireless APs has been proposed both in research work and in commercially available products for simplifying the management and security functions of large-scale networks. However, channel access is still typically distributed and randomized, without exploiting the benefits of centralized mechanisms. This is because centralization is harder to implement as a result of the strict time requirements for coordinating channel access operations. There exist solutions that introduce centralized management of the data path. CENTAUR [6] allows configuring non-conflicting operation intervals among the APs by exploiting a central controller for scheduling the activity intervals in each cell. FLUID [4] allows allocating heterogeneous bandwidths among the cells according to the traffic load. These architectures are based on *application-specific solutions*, including customized controllers, signaling mechanisms, and agents at the AP side, which cannot be easily extended for supporting other centralized optimizations. On the other side, *general architectures* based on the paradigm of Software Defined Networking (SDN), which separates the data and control planes, and on the abstraction of the AP programming interfaces have been proposed for dealing with mobility and associations (e.g., Wi-Fi traffic offloading), without explicitly considering the problem of interference management [7][8]. In CloudMAC [9], the virtualization of AP functionalities is extended for processing media access control (MAC) frames remotely in the cloud, where a central controller generates management frames and dissects node replies. CloudMAC allows fine-grained control only of physical layer (PHY) parameters.

Centralization of network functionalities requires the definition of control messages and protocols between Wi-Fi nodes and the central controller. Although many architectures work on customized protocols, a standard solution for centralized Wi-Fi control, called Control And Provisioning of Wireless Access Points (CAPWAP) [10], has been defined well before the introduction of SDN concepts. A limitation of CAPWAP is that the *information elements in the control messages are predefined*. This implies that CAPWAP cannot be used for configuring devices implementing novel standards, without ratifying protocol extensions. Another shortcoming of CAPWAP is that it allows modifying only

predefined parameters, which can be ineffective for re-configuring the MAC operation logic. In parallel, the adoption of the SDN paradigm for Wi-Fi networks has suggested the use of the OpenFlow protocol for configuring the APs as specialized wireless switches. An example of OpenFlow utilization for programming mobility management solutions in Wi-Fi networks is given by OpenRoads [7], while OpenRadio [8] extends the OpenFlow rules by considering additional PHY-related fields for filtering the incoming packets and configuring the forwarding plane. Furthermore, IEEE P1930.1 will define a middleware for controlling APs using the SDN paradigm but it is still in early stages of development. In all cases, however, inter-cell and MAC coordination are not considered.

To fill the gap between CAPWAP capabilities and wireless SDN requirements, we propose CADWAN – a control architecture for efficient management of dense Wi-Fi access networks (Figure 1). It is based on a central controller, a global network view, and APs exposing a programming interface. We have previously shown how this approach can improve the performance of a single Wi-Fi cell [11]. In this paper, we show how multiple cells can be centrally coordinated for performance improvement in dense scenarios. Specifically, CADWAN comprises: (i) an extension of CAPWAP (Section 2) and (ii) an optimization framework (Section 3). The CAPWAP extensions provide a binding for programmable wireless devices, allow station configuration, and ensure full separation of the data and control planes. The optimization framework allows achieving improved efficiency of dense Wi-Fi networks with centralized control (Section 4). We conclude that the proposed approach is flexible, scalable, gives operators greater control over densely deployed networks, and is complementary to existing 802.11-related efforts (Section 5).

2 The Proposed Control Architecture

Coordination among Wi-Fi cells in existing deployments is limited to local coordination (Figure 1). Each AP has its own instances of control algorithms and operating systems and exchanges coordination messages with other APs. Conversely, in CADWAN we propose to manage the complexity of emerging dense Wi-Fi deployments by centralizing the operating system and control algorithms within an SDN controller, indicated as the access controller (AC), while maintaining backward compatibility by allowing legacy coordination messages between APs (Figure 1). The SDN controller can then coordinate the set of APs in the control plane by collecting radio and network measurements and enforcing centralized configuration decisions. Therefore, from the logical point of view, each AP constitutes a managed part of a central system. From a network deployment perspective, this approach allows network operators to coordinate their Wi-Fi cells. APs which are not under the control of the AC represent the constraints to be considered in the optimization process.

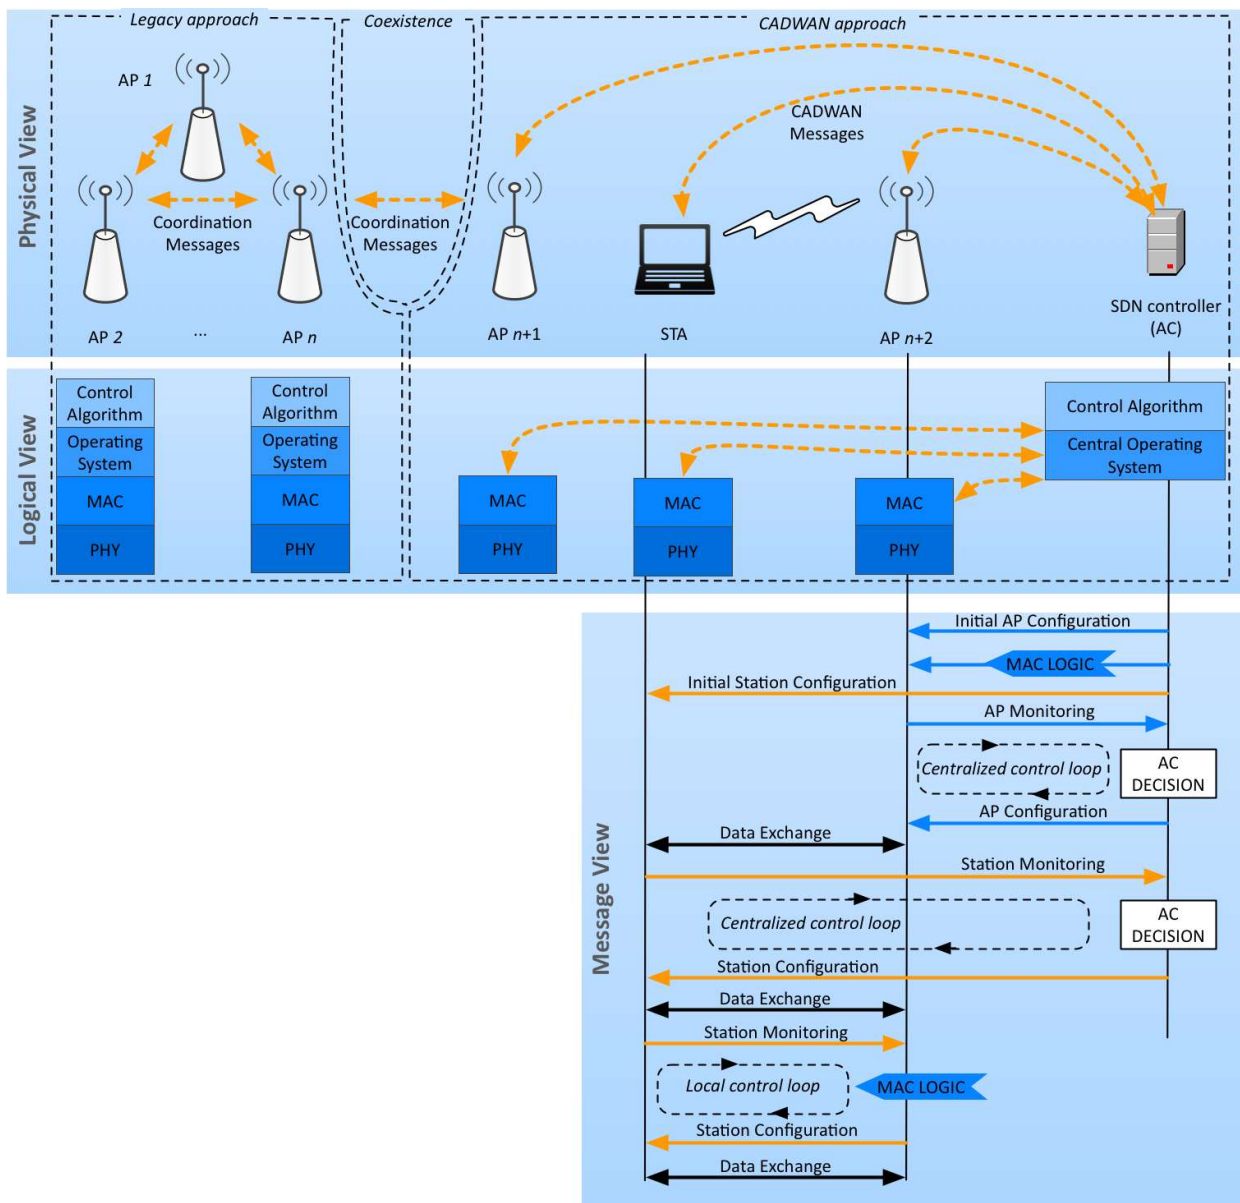


Figure 1. The coordination of Wi-Fi networks using CADWAN, CADWAN coexistence with legacy networks, and CADWAN's message flow.

The second key building block is the programmability of Wi-Fi devices. While in traditional Wi-Fi networks configuration has been mainly limited to frequency and security settings, in the emerging scenario of advanced PHY capabilities and extended operation modes, multiple configuration parameters (transmission power, carrier sense threshold, antenna configuration, etc.) have to be taken into account. Such complexity is destined to grow as soon as wireless devices, based for example on the Wireless MAC Processor (WMP) architecture [12], will allow the programmability of specific MAC features supported by the hardware. In this case, the AC is responsible of defining the MAC operation modes among a set of possible ones or programming a new MAC operation mode by exploiting the new programmable interfaces. The desired MAC operation can be enforced on the controlled APs as soon as

a new network state is detected, as indicated by the injection of the MAC logic into the APs. Alternatively, for improving the reconfiguration latency, the AC can specify a set of MAC operation modes and activation rules to each AP, which in turn can autonomously activate a new mode and notify the associated stations when one of the rules is matched, i.e., the AC sends policies to the AP that is responsible for station configuration (Figure 1). By running context-dependent protocols powerful optimizations of Wi-Fi operation (as described in Section 3) are enabled; however, the network configuration space grows accordingly and the AC is required to be more effective when dealing with increased problem complexity.

While a single network operator administers the AC and its controlled networks, inter-technology coordination, including LTE offloading, Bluetooth/ZigBee coexistence, and inter-operation strategies, may also be considered by the control policies running on the AC, using the ideas proposed in [11].

CADWAN requires a signaling protocol for centralized coordination and policy enforcement. It has three main features, which we describe below. We contrast them with CAPWAP – the currently available signaling protocol for Wi-Fi networks.

First, CADWAN separates the data and control planes. By contrast, in the original CAPWAP proposal, there is an operation mode where all Wi-Fi data and control messages are encapsulated and forwarded to the AC [10]. This approach has certain benefits, but in CADWAN we maintain the operation of the data path completely under the responsibility of the APs, while only the data path parameters that the AC has requested to monitor are forwarded to the global controller. Therefore, the AC is not burdened with processing each frame of all traffic sent over the controlled networks. This ensures SDN-like separation of the data and control planes (Figure 1).

Second, CADWAN supports an enhanced data path configuration through dedicated control messages containing the MAC operation logic, not limited to parametric configurations. CADWAN generalizes the definition of CAPWAP's information elements for coding the instruction set and programming model supported by the programmable devices. For example, for WMP-enabled devices [12] the information element can code the control commands supported by the WMP control interface (loading or activating a specific MAC program) or inject a whole state machine implementing a novel MAC protocol into the device. Whole radio programs can be specified as a set of initialization parameters and the MAC state machine, transported in special control packets [12], as indicated in Figure 1 by the configuration messages.

Third, CADWAN provides flexible control through multi-level control loops as evidenced by the signaling sequence generalized in Figure 1. After initial session establishment, two control message exchanges occur: between the AC and its APs and between each AP and the stations associated with them. Both provide support for programmable devices. This approach enables the configuration of stations by having the AC send policies to the AP, with the latter being an executor: applying these policies locally

and configuring its associated stations accordingly. In the next section, we discuss how messages of the control plane can be structured to optimize network performance.

To summarize, in Table 1 we compare legacy control schemes without SDN (e.g., 802.11aa), existing control schemes with SDN (e.g., [7][8]), and CADWAN.

Table 1. Comparison of the legacy control schemes, existing control schemes with SDN, and the proposed CADWAN architecture.

Feature	Legacy control schemes w/o SDN	Existing control schemes w/ SDN	CADWAN
Architecture type	Decentralized	Centralized	Centralized
Data and control planes	Combined	Separated	Separated
Decision making	Local	Global	Global
Control type	Parametric	Parametric and/of flow-based	Flexible PHY/MAC
Control overhead	Large	Reduced	Reduced
Applicability	General	Application-specific, proprietary solutions	General
Extendibility	Limited	Limited	Improved
Security	Limited	Improved	Improved
Power consumption	High	Reduced	Reduced
Efficiency	Low	Limited	High

3 Optimization Framework

An integral part of CADWAN is its three-step optimization framework, which assures its reliability and scalability. The goal of the framework is to optimize the performance of each Wi-Fi cell in a dense environment. We consider one AC supervising the cells of a single network operator within an area. We

assume that while these cells are fully under CADWAN control, there are also other Wi-Fi cells in the area which are not within the control of the same operator.

The optimization is performed in three steps. First, the SDN controller (AC) attempts to configure cells so that they are maximally distributed in the frequency domain (step 1).³ Then, for cells still interfering with each other due to the limited number of channels, the AC chooses the best channel sharing scheme to coordinate transmissions (step 2). Finally, within each cell, the AP performs local parameter optimization to improve performance (step 3).

The three steps represent different aspects of inter-network coordination. Step 1 is the traditional frequency planning problem. Step 2 provides time-based spectrum allocation. Finally, step 3 has the most innovative aspect: for a given allocated spectrum portion (frequency, time, code) we also consider configuring and tuning the spectrum access rules.

In each step, the AC and the APs exchange information using control loops: the AP provides monitoring data while the controller sends configuration instructions. The details of the information exchange are specific to each step. In general, the controller obtains full area knowledge based on measurements and control information received from the APs which perform periodic background scanning. This information includes

- channel quality estimation:
 - intra- and inter-channel interference,
 - measured energy from non-IEEE-802.11 devices and from other Wi-Fi networks,
- traffic load and medium access performance indicators:
 - number of associated stations,
 - airtime use,
 - number of collisions,
 - QoS traffic-related reports such as the QLoad reports of 802.11aa,
- frequency band support and other network configuration information.

While the first step can occur during the planning and deployment stage, all steps should be periodically performed (with decreasing time scales for each step) during the network's operation. This is necessary because the wireless environment changes over time due to mutating radio link conditions, channel selection by legacy APs in the area and node mobility. In the following, we present the goals and methods for each of the three steps.

3.1 Step 1: Radio Configuration

In the first step, the SDN controller (AC) configures the radio parameters of the APs under its control to optimize network-wide performance taking into account the current radio context including interference from networks outside of its control. This includes APs of other network operators and

³ This is somewhat similar to the 802.11aa management procedure for overlapping networks, which is done locally at each AP.

private APs, which serve as an external constraint for the optimization problem. The main goal is to avoid the overlapping of networks in the frequency domain, because of the small number of available channels and the high-density of deployed APs. In these conditions, it is likely that multiple interfering cells are configured on the same channel and one station senses several APs. Although the carrier sense mechanism allows the coexistence of interfering cells, their capacity can be unevenly allocated. Additionally, stations can be impaired by hidden station transmissions or flow-in-the-middle problems which are difficult to predict from a local perspective. Therefore, this step can be considered a generalized form of frequency planning which, in some cases, will also impact the logical topology of the network. The following methods can be used in this step:

- **Frequency band and channel selection** – the AC performs frequency planning (including channel width configuration) taking into account traffic QoS requirements and device constraints.
- **Transmission power and data rate control** – adjusted for each AP to improve network coverage and minimize interference [13].
- **Station reallocation** – the controller may force clients to associate to lightly loaded APs to equalize the load and increase fairness among the networks.
- **Load balancing** – the AC may encourage clients to associate to more than one AP by maintaining simultaneous connections.
- **Deactivating APs** – the controller may decide to disable selected APs in a dense deployment due to power saving, network scaling, as well as application and load-specific requirements.

3.2 Step 2: Channel Sharing

In the second step, the AC addresses cells which are still interfering with each other (i.e., in close proximity and on the same channel) after the procedures performed in the first step. The SDN controller performs decisions based on the feedback from the APs related to their traffic load and MAC performance indicators. Based on this information, the controller chooses the best sharing scheme to coordinate transmissions in these overlapping Wi-Fi networks. Categories of such channel sharing schemes include:

- **Carrier-sensing-based** – using legacy 802.11 to provide a best effort approach.
- **Imposing traffic limits** – 802.11aa sharing policies (proportional and on demand) can be used for defining traffic allocation limits.
- **Fairness balancing** – selected networks receive higher priority in channel access to mitigate the negative effects of interference and ensure inter-network fairness.
- **Time-division multiplexing (TDM)** – networks take turns in accessing the channel, i.e., periods of time are assigned to each network in which its corresponding transmissions take place without interfering with other networks' transmissions (Section 4.2).
- **Conditional frame decoding** – stations distinguish inter-cell and intra-cell transmissions to improve the probability of decoding overlapping frames [5].

3.3 Step 3: Local Performance Optimization

The last step consists of performance optimization within each cell, as mandated by the AC, according to locally available information. Whereas in the previous steps the APs provided measurement data to the AC, in this step such information can be used locally. For example, if the AC allows stream reservations,

the AP can apply this policy when Voice-over-IP traffic exceeds a given percentage of the total traffic. Examples of solutions that can be adopted in this step include:

- **Grouping stations for MAC policy adoption** – a specific policy can be applied to a subset of stations, e.g., to take part in the coordinated multi-user transmissions envisaged in 802.11ax [5].
- **Parameter tuning** – optimization can be performed with respect to PHY and MAC parameters such as contention window size and carrier sensing threshold [5].
- **Configuring antennas** – spatial reuse can be improved locally by mandating the use of novel antenna techniques such as beamforming.

Since in this step decisions and enforcement are applied locally, the optimization and context adaptation can run faster than in the previous steps.

4 CADWAN Example Use Case

We now describe how the CADWAN architecture can control high-density Wi-Fi cells belonging to a single administrative domain, as in the case of access networks deployed in a multi-floor building⁴, by applying the three-step optimization procedure. We focus on three domains: frequency (step 1), temporal (step 2), and protocol (step 3). We provide an example of CADWAN’s use by assuming the implementation of already existing solutions, hitherto proposed separately but now intelligently combined by the AC to achieve improved Wi-Fi efficiency. This proves that our proposal could be applied in the near future, in real Wi-Fi deployments. The tuning of the individual algorithms used in each particular step is out of the scope of this tutorial.

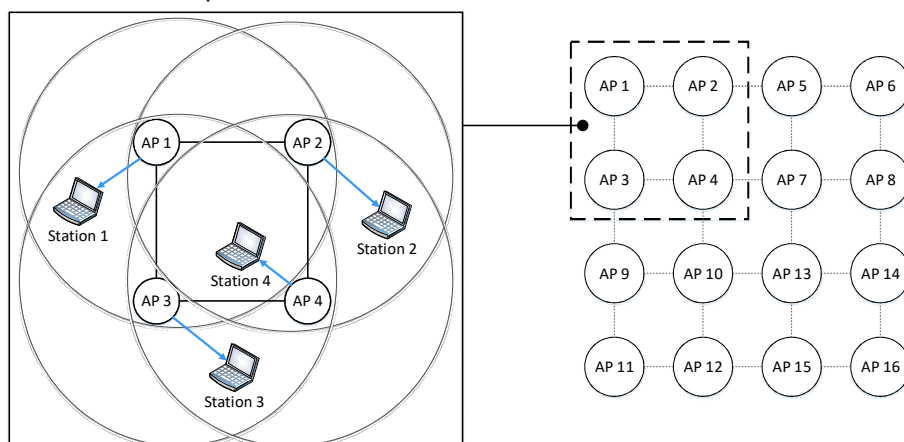


Figure 2. Grid area with overlapping networks: one tile of the big picture (on the left) and the big picture itself (on the right). APs have the same coverage radius so that they hear their cardinal but not their diagonal neighbors.

For the sake of simplifying the description, and without loss of generality, consider a regular grid topology, as shown in Figure 2. Additionally, due to the high density of neighboring Wi-Fi deployments

⁴ This scenario applies also to domestic APs which are close to each other and, despite being owned by different neighbors, are governed by the same network operator and controllable through CADWAN.

and to provide a comprehensible example, we assume that only two unused non-overlapping channels remain available for allocation.⁵

Given the regularity of the topology, we can zoom in on a single tile containing four APs, to explain several interesting phenomena, which are then targeted by CADWAN optimization. For the clarity of Figure 2, we illustrate a basic configuration, in which each AP serves one station; however, later on we analyze configurations with up to 8 APs operating on one channel and up to 30 stations per AP.

4.1 Step 1: Radio Configuration in a Grid Network

For a simple topology (left side of Figure 2), under our assumption that only two channels are available, it is possible that APs 1 and 4 (being out of range) select the same operating channel, while APs 2 and 3 select the orthogonal channel. Under this hypothesis, Station 4 may experience low throughput, because AP 1 is hidden from AP 4 and their transmissions can collide with high probability. Better performance may be achieved if AP 1 is paired with AP 3 and AP 2 with AP 4 on orthogonal channels, since collisions could be avoided. It follows that even in the presented simple scenario it is difficult to find the optimal planning by only considering the APs' visibility, because throughput results also depend on the station placement.

To make these considerations more explicit, we quantify (by ns-2 simulations) the first step of our optimization framework. We measure the throughput performance in two configurations: four downlinks (Figure- 3a) and 16 downlinks (Figure- 3b) under the assumption of two available empty channels.

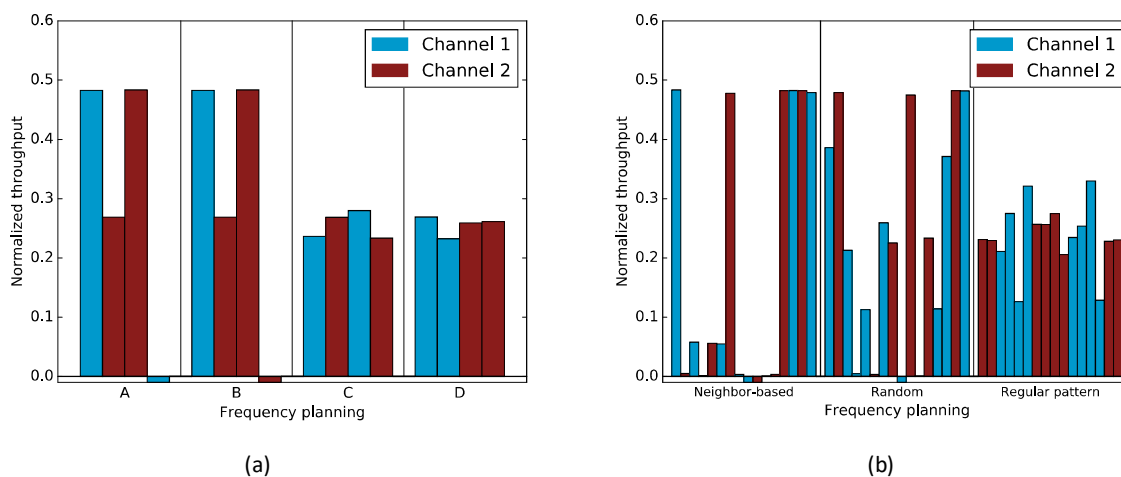


Figure 3. Throughput performance under different frequency plannings: (a) four partially overlapping cells (left part of Figure 2) with uneven traffic loads and (b) a regular grid of 16 APs (right part of Figure 2). Within each planning, cells are presented in their natural order. Negative bars indicate operating channel when zero throughput is achieved.

⁵ In real deployments, free channels are rarely available. Therefore, appropriate channel selection algorithms are required to choose, e.g., the least loaded channels (Section 3).

We consider the effects of an uneven load distribution in the four AP topology (Figure 3a), by assuming that traffic flows directed to the stations associated to AP 2 and AP 4 require 0.27 normalized throughput, while the other flows are saturated. When AP 4 selects the same channel as AP 1, it achieves zero throughput because of the hidden node phenomenon (planning A). When it selects the same channel as AP 2 and AP 3, it again receives zero throughput because of the flow in the middle phenomenon (planning B). Conversely, the position of stations associated to AP 2 and AP 3 allow these APs to use the same channel without throughput impairments: station 3 achieves the cell saturation throughput, while station 2 achieves its offered load. As shown in planning C, since the traffic load of cells 2 and 4 are almost equal to half the saturation throughput, the two cells can select the same operating channel. The higher collision rate of station 2, due to its position, results in a slightly lower throughput than of station 4. Finally, the optimal planning D shows that better throughput fairness can be achieved than in planning C.

Frequency planning based on local interference can be subpar as shown in the more complex topology with 16 APs (right part of Figure 2). For this topology, two channels are enough for tuning each AP on a channel other than its neighbors by repeating planning D. However, by randomly placing even a single station associated to each AP (the simplest approach), such a planning may result in several cells with zero throughput, as shown in the neighbor-based planning in Figure 3b. This figure also shows, for the same topology, the results obtained with a random planning of APs and with a regular pattern based on the results from the 4 AP case. In summary, CADWAN's first step amounts to centralized assignment of the available frequency channels based on the AP and station placement (i.e., topology-specific interference).

4.2 Step 2: TDM-based Channel Sharing

The second step of our optimization framework requires the AC to coordinate the APs of overlapping cells to share the channel. An exemplary channel sharing solution that can be implemented is the following TDM-based approach.

A TDM slot is assigned to each overlapping cell by the AC and these slots are then scheduled so that they interleave. Each slot begins with a beacon transmission, specifying its length, which can depend on network load or priority. Each station is allowed to transmit only within its cell's slot. In the simplest TDM channel sharing approach, the contention among stations in a cell (i.e., within a single slot) is realized using 802.11's basic distributed coordination function (DCF). The AC is able to schedule the slot of each AP by postponing its beacon. If DCF is used within a slot, stations belonging to a given network can transmit only in a predefined time interval after beacon reception. TDM channel sharing may exhibit low performance under non-saturation, so the AC should only switch to such a scheme based on reports of high collision rates and revert to normal medium access if the airtime use drops below a given threshold.

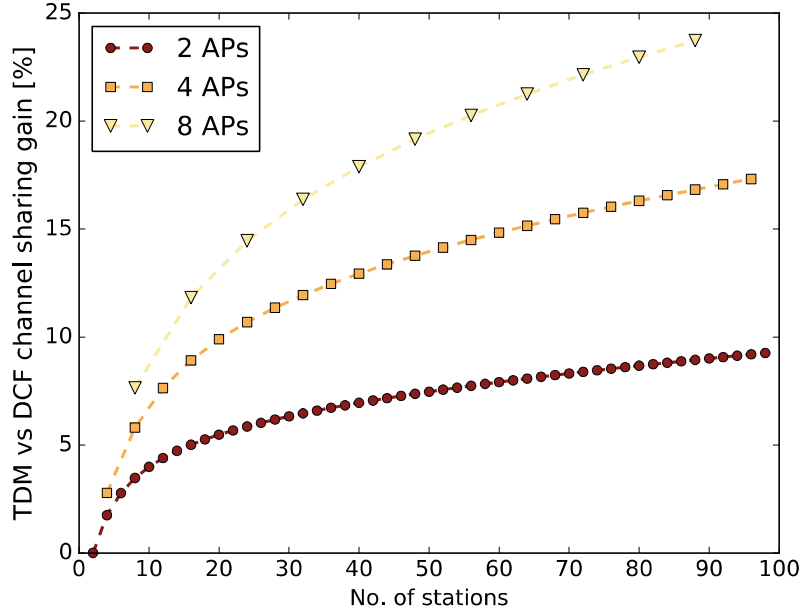


Figure 4. Exemplary channel sharing solution: TDM throughput gain in comparison to legacy DCF.

Figure 4 presents the per-station throughput gain of using TDM in comparison to legacy DCF (for sharing the channel among cells) with the assumption of homogenous slots. The throughput values were calculated according to [14]. We compare the results for $y = \{2, 4, 8\}$ saturated, overlapping cells for a maximum of 100 total stations. There are y APs serving a total of x stations with $\frac{x}{y}$ stations per AP. The TDM throughput gain for y fully overlapping cells is calculated as

$$\left(\frac{T_{\frac{x}{y}+1} \times \frac{1}{y} - T_{x+y}}{T_{x+y}} \right) \times 100 \quad , \quad (1)$$

where T_i is the per-station normalized saturation throughput for a Wi-Fi network comprising i stations. T_{x+y} corresponds to DCF channel sharing and $T_{\frac{x}{y}+1} \times \frac{1}{y}$ to TDM channel sharing.

The TDM gain increases with station density: both for an increasing number of stations per AP and for an increasing number of APs. This gain is a direct result of the decreased contention among the networks. It is therefore possible to improve the efficiency of channel use in the case of densely deployed networks with the centralized CADWAN approach.

4.3 Step 3: MAC-based Optimization

We exemplify CADWAN flexibility in the last step using a classical MAC-based optimization (i.e., *contention window tuning*) in which the AC provides policies for more efficient channel use. The performance of DCF can be improved by using a theoretically calculated optimal contention window value proportional to the number of contending stations [15]. Using this approach, the aggregate

throughput remains almost constant, contrary to DCF, for a wide range of contending stations. Calculating the optimal contention window requires centralized knowledge regarding the number of contending stations (belonging to overlapping cells) and the contention time, obtained from measured traffic parameters. This centralized knowledge and decision making can leverage CADWAN’s control capabilities.

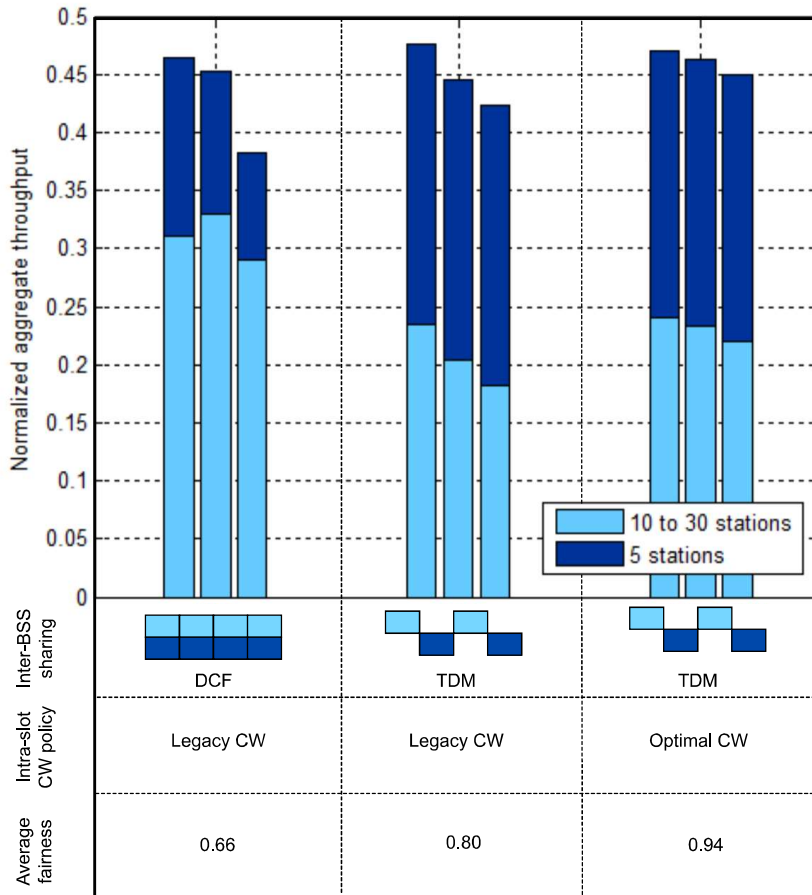


Figure 5. Contention window tuning: normalized aggregate throughput for DCF, TDM, and TDM with optimal CW, where the three bars for each MAC configuration represent 5 stations in one network overlapping with 10, 20, and 30 stations in the second network, respectively.

In our ns-2 simulations we consider two overlapping networks sharing the same channel – one with 5, the other with 10, 20, or 30 stations. Figure 5 presents the aggregate throughput for the various MAC configurations:

- Legacy DCF (after step 1) – stations belonging to both networks contend during 100% of the airtime and use standard CWs,
- TDM channel sharing (after step 2) – the AP instructs the two groups of stations to contend separately, each during 50% of the airtime and using standard CWs in each TDM slot,

- TDM channel sharing with optimal CW (after step 3) – the AP instructs the two groups of stations to contend separately using DCF with optimal CW in each slot.

Switching from legacy DCF to TDM channel sharing increases throughput and improves inter-network fairness. Additionally, applying the optimal CW to TDM can slightly increase throughput, but more significantly – increase intra-network (per-station) fairness. While the use of optimized contention windows is always beneficial, its effects are even more evident when applied jointly with TDM-based channel sharing. We conclude that, using the third step to provide *medium access optimization*, a central controller may improve network performance by changing MAC policies and parameters.

5 Summary

CADWAN improves performance of high-density Wi-Fi networks by managing them in a flexible, scalable, and extensible manner while being complementary to ongoing Wi-Fi developments. For example, the improved spatial reuse with cell “coloring” proposed in 802.11ax can be mapped to CADWAN’s step 2. However, while standards such as 802.11ax define new functionalities, they do not specify when to enable them nor how they should be configured. Thus, a solution such as CADWAN is necessary for Wi-Fi optimization in dense deployments.

Acknowledgements

This work has been partially funded by the WiSHFUL EU H2020 Project (GA No. 258301), the Canaletto project “ADHOC: ADvanced wifi networks for HOme Customized applications”) and by Polish budgetary funds for science for 2013-2015 (project No. IP2012 016672). K. Kosek-Szott and S. Szott are supported by the Foundation for Polish Science (FNP).

Bibliography

- [1] K. I. Pedersen, Y. Wang, S. Strzyz and F. Frederiksen, "Enhanced inter-cell interference coordination in co-channel multi-layer LTE-advanced networks," IEEE Wireless Communications, vol. 20, no. 3, pp. 120-127, 2013.
- [2] C. J. Bernardos et al., "An architecture for software defined wireless networking," in IEEE Wireless Communications, vol. 21, no. 3, pp. 52-61, 2014.
- [3] A. Baid, D. Raychaudhuri, "Understanding channel selection dynamics in dense Wi-Fi networks," IEEE Communications Magazine, vol. 53, no. 1, pp. 110–117, 2015.
- [4] S. Rayanchu, V. Shrivastava, S. Banerjee, R. Chandra, “FLUID: improving throughputs in enterprise wireless lans through flexible channelization”, Proc. of the 17th annual international conference on Mobile computing and networking (MobiCom '11), New York, NY, USA, September 19-23, 2011.
- [5] E. Khorov, A. Kiryanov and A. Lyakhov, "IEEE 802.11ax: How to Build High Efficiency WLANs," Proc. of 2015 International Conference on Engineering and Telecommunication (EnT), Moscow, Russia, 2015.

- [6] V. Shrivastava, N. Ahmed, S. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, A. Mishra, "CENTAUR: Realizing the full potential of centralized WLANs through a hybrid data path," Proc. of 15th Annual International Conference on Mobile Computing and Networking (MobiCom'09), Beijing, China, 2009.
- [7] K.-K. Yap, M. Kobayashi, R. Sherwood, T.-Y. Huang, M. Chan, N. Handigol, and N. McKeown, "OpenRoads: empowering research in mobile networks", ACM SIGCOMM Computer Communication Review, vol. 40, no. 1, pp. 125-126, 2010.
- [8] M. Bansal, J. Mehlman, S. Katti, and P. Levis, "OpenRadio: a programmable wireless dataplane", Proc. of the First workshop on Hot topics in software defined networks (HotSDN '12), New York, NY, USA, 2012
- [9] J. Vestin, P. Dely, A. Kassler, N. Bayer, H. Einsiedler, and C. Peylo, "CloudMAC: towards software defined WLANs." ACM SIGMOBILE Mobile Computing and Communications Review, vol. 16, no. 4, pp. 42-45, 2013.
- [10] P. Calhoun, M. Montemurro, and D. Stanley "Control and Provisioning of Wireless Access Points (CAPWAP) Protocol Specification," IETF RFC 5415, 2009.
- [11] P. Gallo, K. Kosek-Szott, S. Szott, I. Tinnirello, "SDN@home: A Method for Controlling Future Wireless Home Networks", IEEE Communications Magazine, vol. 54, no. 5, pp. 123-131, 2016.
- [12] P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, I. Tinnirello, and G. Bianchi "Wireless MAC Processor Networking: A Control Architecture for Expressing and Implementing High-Level Adaptation Policies in WLANs", IEEE Vehicular Technology Magazine, vol. 8, no. 4, pp. 81-89, 2013.
- [13] O. Oteri, P. Xia, F. LaSita, and R. Olesen, "Advanced power control techniques for interference mitigation in dense 802.11 networks," Proc. of IEEE 16th International Symposium on Wireless Personal Multimedia Communications (WPMC), Atlantic City, NJ, USA, 2013.
- [14] M. Gas, K. Kosek-Szott, M. Natkaniec, A. R. Pach, "3D Markov chain-based saturation throughput model of IEEE 802.11 EDCA," Electronics Letters, vol. 47, no. 14, pp. 826-827, 2011.
- [15] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function", IEEE Journal on Selected Areas in Communications, vol. 18, no. 3, pp. 535-547, 2000.