Interference-Aware Sector Design in IEEE 802.11ad Networks

Guillaume Gaillard, Gentian Jakllari

► To cite this version:

Guillaume Gaillard, Gentian Jakllari. Interference-Aware Sector Design in IEEE 802.11ad Networks. 30th International Conference on Computer Communications and Networks (ICCCN 2021), Jul 2021, Athens / Virtual, Greece. hal-03276187

HAL Id: hal-03276187

<https://hal.inria.fr/hal-03276187>

Submitted on 1 Jul 2021

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.

Interference-Aware Sector Design in IEEE 802.11ad Networks

Guillaume Gaillard, Gentian Jakllari

guillaume.gaillard@toulouse-inp.fr, jakllari@enseeiht.fr

RMESS, IRIT, ENSEEIHT, INPT, University of Toulouse, France

Abstract

IEEE 802.11ad is the first standard capable of delivering Gigabit capacities in a WLAN setting by exploiting mmWave frequencies. Its key innovation is the Sector Level Sweep, a procedure for establishing directional links over mmWave by sweeping over a pre-computed set of phased-array antenna sectors. Crucial to the performance of IEEE 802.11ad networks and involving an exponentially-large search space, the sector design is left open to implementation. As a result, several sector designs have been introduced that are aimed at maximizing the SNR of a given direct link while ignoring the interference projected to other links. If highly competitive in small deployments, a sector design that does not take interference into account can lead to poor overall performance in multi-AP deployments.

In this article we present IASSEN, the first approach to Interference-Aware Sector deSign in IEEE 802.11ad Networks. Anchored in an intimate understanding of COTS 802.11ad hardware, IASSEN models the SINR for each network link as function of the low-level parameters driving the sector design. After modeling the problem of interference-aware sector design as an integer nonlinear programming problem, IASSEN introduces a heuristic carefully combining a quick convergence to a local optimum with an approach for avoiding poor local optimum points. A simulation study using measurement-based and realistic sector patterns shows that IASSEN improves overall network performance by as much as 100% when compared to state-of-the-art approaches that do not take interference into account.

1 Introduction

The millimeter wave frequencies, with several GHz of mostly idle and unlicensed spectrum, are touted as a key enabler of 5G-and-beyond networks capable of meeting the exponentially increasing demand for wireless capacity [1]. Nevertheless, this frequency band is not idle by accident – it suffers from strong pathloss and poor penetration of obstacles [2, 3], making it inherently unsuitable for indoor deployments where, paradoxically, demand is highest. To overcome the mmWave handicap, IEEE 802.11ad [4] relies on phased-array antennas [5] capable of directing the main part of their energy toward the desired peer. The so-called beams, oriented directly at a device and indirectly via reflections (e.g. on walls), are constructively combined to maximize the communication gain. However, while phased-array antennas allow the establishment of communication links indoors, they introduce the challenge of beamforming. Communication devices need to identify each other in space so as to direct their beams in the proper direction [6, 7, 8].

Beamforming becomes more challenging when using several access points shared among multiple users in mmWave wireless local networks. [9, 10, 11, 12]. As such networks grow in scale to meet the increasing demand for capacity in indoor places such as classrooms, station halls, etc. interference emerges as a key limiting factor [13].

To elucidate, consider the simple topology in Fig. 1, comprising four devices forming two concurrent bidirectional links in a room. We aim at determining the best eight configurations (two for each device, one for transmitting and one for receiving) of the

Figure 1: An indoor scenario with four devices. The transmission $x \rightarrow y$ causes interference at v because of the non-trivial side-lobe at x in direction of v .

phased-arrays of the devices antennas in order to obtain beam patterns (the directions where the antennas focus their energy) achieving two often contradictory goals:

- Maximize the strength of received signal in both directions of each link.
- Minimize the interference projected on neighboring devices, which are not the intended receiver.

Previous works address this challenge by proper power allocation [14, 15] or a scheduling strategy [16, 17].

BounceNet [18] introduces a centralized approach using best effort scheduling strategies based on measured SNRs, predefined beams and interference estimation. It acquires full knowledge of signal paths between any two devices to deduce the transmission gains in every direction. However, BounceNet does not redefine the antenna beams: the antenna configurations are chosen in a subset that was previously designed. Hence, low-cost devices, presenting multi-directional beams (i.e. side-lobes), will fail to avoid interference.

The standard IEEE 802.11ad introduces the Sector Level Sweep (SLS) procedure to address the beam-forming challenge. A sector refers to a set of weights that impacts the complex gains of each antenna element in the antenna array, with the objective of creating a transmitting or receiving beam in a particular direction. SLS involves exploring all the possible sectors via the exchange of control packets to identify the best sector for a particular destination device [4, 8]. With the number of possible sectors being exponential on the number of antenna elements, however,

802.11ad introduces a limited set of pre-computed sectors, generally called a codebook. Beamforming is thus reduced to selecting one of the codebook sectors. It is shown in [19] that the codebook of the first commercial off-the-shelf (COTS) 802.11ad hardware are designed for maximizing coverage not link performance. A solution, ACO, is introduced in [20] to design sectors that individually optimize a single link throughput based on Signal-to-Noise Ratio (SNR) measurements. While ACO is efficient and significantly improves upon the IEEE 802.11ad COTS sector designs, by focusing solely on SNR it ignores interference. As a result, while it improves the performance a given link, it can increase the interference levels and have a negative impact on other links and the overall network performance.

In this work, we introduce the first Interference-Aware Sector deSign for IEEE 802.11ad Networks (IASSEN). Taking interference into account when designing phased-array antenna sectors introduces a complex tradeoff between maximizing one's performance while minimizing the interference your communication projects to other links in the network. What is more, the potential number of sectors, defined by the different amplitude and phase shift values that can be assigned to each antenna element, is exponential on the number of antenna elements. In short, IASSEN addresses this challenge by a) Modeling the signal-plus-interference-to-noise ratio (SINR) for each link, and relating it with the low-level parameters of the phased-array antennas, b) Formulating the problem of interference-aware sector design as an Integer Nonlinear Programming (INLP), c) Introducing a heuristic carefully combining a deterministic descent toward a local optimum with a controlled random exploration for avoiding poor local optimum points, and d) Anchoring the solution on commercial off-the-shelf (COTS) IEEE 802.11ad hardware.

Our main contributions may be summarized as follows:

- We model the interference-aware sector design as an Integer Nonlinear Programming (§ 3) and instantiate it on TP-Link Talons AD7200 routers, IEEE 802.11ad COTS hardware.

Figure 2: A Sector Level Sweep (SLS) phase in IEEE 802.11ad. The transmitter successively sends one SLS frame per *sector*. The receiver reports back the per-frame Received Signal Strength (RSS) values.

- We introduce IASSEN:
  - a centralized interference-aware sector selection scheme for multi-access point IEEE 802.11ad networks (§ 4). IASSEN carefully combines a quick convergence to a local optimum with an approach for avoiding poor local optimum points.
- We carry out a thorough performance evaluation of IASSEN using simulations with measurement-based and realistic beam patterns (§ 5). Results from multiple scenarios, varying from simple topologies to dense deployments (e.g. 12 access points side-by-side in a room) show that IASSEN significantly improves the network performance when compared to ACO, validating our strategy of interference-aware sector design.

After an overview of related approaches (§ 6), we conclude and give insights into future work (§ 7).

2 A primer on sector selection in IEEE 802.11ad

Two IEEE 802.11ad devices aiming to establish a link (e.g. an Access Point (AP) and a Station (STA)) need to select in a limited time their best antenna configurations, i.e. maximizing the gains in transmission and reception power in the direction of their peer. Known as the beamforming training challenge [21], it is addressed in the standard [4] by the introduction of the Sector Level Sweep (SLS) procedure.

2.1 Sector Level Sweeps

The 802.11ad physical layer defines a set of antenna configurations called *sectors* (Section 2.2) because the gain is focused toward a portion of a sphere.

During the SLS phase (Fig. 2), a transmitter, x , sweeps short frames over all its available sectors. The receiver, y , feeds back a Received Signal Strength (RSS) to the transmitter for each SLS frame:

$$\text{RSS}_{x,y} = P_t(x)G_t(x,y)G_{ch}(x,y)G_r(y,x) \quad (1)$$

where $P_t(x)$ is the transmission power of x , $G_{ch}(x,y)$ is the propagation gain on the link (x,y) and $G_t(x,y)$, $G_r(y,x)$ are the transmit and receive antenna gains, respectively.

The standard introduces a second sweep phase, inverting the transmitter and receiver roles, enabling the refinement of the reception sectors. However, its application is implementation dependent.

2.2 Antenna sectors

IEEE 802.11ad uses phased-arrays, a common directional antenna technology. They combine copies of a signal received or transmitted by a set of n radiating antenna elements. Each antenna element is characterized by its contribution to the main signal: the phase and strength it grants to its copy.

A sector corresponds to a set of *weights* that impacts the complex gains of each antenna element. A weight comprises an amplitude factor (the strength given to an antenna element) and a phase shift (the delay experienced by the copy of the signal). Formally, the weights, i.e. the amplitude factors $w_k^{x_t}$ and phase shifts $\Delta\phi_k^{x_t}$ constitute s_x^t , the transmission *sector* of device x :

$$s_x^t = \{w_k^{x_t} \in \mathbb{R}^+, \Delta\phi_k^{x_t} \in [-\pi, \pi]\}_{k \in [1,n]} \quad (2)$$

Therefore, the gain of a transmitting phased-array antenna as function of its antenna element contributions and weights is:

$$G_t(x,y) = \left| \sum_{k \in [1,n]} w_k^{x_t} a_k^{x,y} e^{j(\phi_k^{x,y} - \Delta\phi_k^{x_t})} \right| \quad (3)$$

where $a_k^{x,y}$ and $\phi_k^{x,y}$ are, respectively, the amplitude and phase of the signal received at y and transmitted by antenna element k of device x .

The reception gains $G_r(x,y)$ and sectors s_x^r are expressed similarly.

Figure 3: Default Codebook running on the Talons. It comprises 36 different transmission sectors and a unique omni-directional reception sector (Rx). The y -axis refers to the antenna element and the x -axis the element's contribution to the sectors. The presence of the diamond symbol indicates the respective antenna element is active on a particular sector, the color indicates the phase shift.

802.11ad only allows devices to choose among a limited set of sectors (max. 128). As a result, the performance of beamforming is defined by the sector design, i.e. the choice of weights and shifts that best combine the elements contributions. The sector design faces two challenges [20]:

1. Estimate the contribution of each antenna element.
2. Deduce the sectors that best satisfy a given goal.

2.3 Beamforming on the Talons

On the TP-Link Talon AD7200 routers (hereafter called Talons), 256 different weights can be assigned to each of the 32 elements of their phased-array antenna. Specifically, 2 bits define 4 discrete phase shifts ($\{0, \pi/2, \pi, -\pi/2\}$), 3 bits define 8 amplitudes and 3 bits per group of 4 elements define 8 distribution amplifiers. The variations of RSS related to different choices in the bit settings of gain and distribution amplifiers not being significant, the amplification can be limited to an on-off model, i.e. $w_k^{x_t} \in \{0, 1\}$ and $w_k^{x_r} \in \{0, 1\}$. Therefore, five different weights may be assigned to an antenna element: $0, 1, e^{j\pi/2}, e^{j\pi}, e^{-j\pi/2}$, leading to 5^{32} significantly different sectors, much more than the default 37 sectors.

A Talon device keeps in its memory a set of pre-

defined sectors, called a *codebook*, that enables it to form different beams. As depicted in Fig. 3, each column in the codebook refers to a transmit sector (the Rx column refers to the reception sector¹). A diamond symbol in the figure indicates the antenna element k (specified in the vertical axis) is *active* (i.e. $w_k^{x_t} = 1$) and its color corresponds to a phase shift.

3 Problem definition

In this following section, we formulate the problem of interference-aware sector design in IEEE 802.11ad networks. Our goal is to identify sectors for every device such that the sum of SINRs (and by extension, throughput) of all the links in a 802.11ad network is maximized.

Let $devs = \{1, 2, \dots, D\}$ denote the set of 802.11ad devices in the network, $\mathcal{L}$ the set of links, and $\mathcal{H}(y)$ the set of neighbors of a device, y , i.e. the devices q with $RSS_{q,y} \neq 0$. For a given link (x, y) , we compute SINR as follows:

$$SINR(x, y) = \frac{RSS_{x,y}}{\sigma_{x,y} + \sum_{q \in \mathcal{H}(y)} p_q RSS_{q,y}} \quad (4)$$

where $RSS_{x,y}$ denotes the strength of the signal received from x at y . The set of real factors $\mathcal{P} = \{p_1, p_2, \dots, p_D\}$ refers to the distribution of traffic load for each device. When a neighbor q of y uses the channel full-time to transmit, its p_q is 1. When no device is transmitting, all $p \in \mathcal{P}$ get close to 0, and the SINR becomes the SNR. Finally, $\sigma_{x,y}$ corresponds to a noise value accumulated on the link (x, y) .

3.1 Interference-Aware Sector Design modeled as an INLP

As mentioned in § 2.2, a sector refers to a set of amplitude and phase shifts weights. Let $\mathcal{W}$ and Φ denote the discrete sets of possible amplitude factors and phase shifts, respectively. We formulate the problem

¹The default implementation includes only one generic (omni-directional) reception sector.

of interference-aware sector as follows:

$$\begin{aligned}
& \text{maximize } \sum_{(x,y) \in \mathcal{L}} \text{SINR}(x,y) \quad \text{subject to :} \\
& \forall (x,y) \in \text{devs}^2, \forall k \in [1, n], \\
& \quad \begin{cases} (w_k^{x_t}, w_k^{x_r}) \in \mathcal{W}^2 \\ (\Delta\phi_k^{x_t}, \Delta\phi_k^{x_r}) \in \Phi^2 \end{cases} \quad (5)
\end{aligned}$$

Pb. 5 is an Integer Non Linear Programming (INLP). It is non-trivial, as increasing the gains and hence the RSS for a given link increases the interference on the other links (Eq. 4).

3.2 Instance of the INLP in the case of the Talons

In the simplest case of Fig. 1 with Talons, we have:

$$\mathcal{W} = \{0, 1\}, \Phi = \{0, \pi/2, \pi, -\pi/2\}, D = 4, n = 32$$

leading to a search space of 5^{256} values. Even for this introductory example, an exhaustive search is not possible in real time with current processing capacities. Besides, the dimension of the problem grows geometrically with D and n . Hence, we chose to design IASSEN as a fast but sub-optimal approximation.

4 IASSEN: toward SINR improvement

In this section, we introduce a scalable solution to the interference-aware design problem formulated in § 3.

We anchor this solution in a realistic context such that we could directly implement it on COTS devices.

4.1 Integer multivariate formulation

Our mechanism inherits from optimization methods: we define a *state* variable X with integer components corresponding to a discrete weight of antenna elements in both link directions. We recast the

interference-aware design problem as the minimization of the objective function, $\varepsilon(X)$, called *energy* associated to each state X :

$$\begin{aligned}
& \text{minimize } \varepsilon(X) = - \sum_{(x,y) \in \mathcal{L}, \mathcal{S}(X)} \text{SINR}(x,y) \\
& \quad \text{subject to :} \\
& \quad \forall i \in [1, 2Dn], X_i \in [0, (\text{Card}(\mathcal{W}) - 1) \times \text{Card}(\Phi)] \\
& \quad \text{Eq. 1, Eq. 3, Eq. 4} \quad (6)
\end{aligned}$$

$\mathcal{S}(X)$ is the sector translation of X , where the n first elements $(X_1, \dots, X_n)$ form the transmission sector of the first device, elements $[n + 1, 2n]$ its reception sector, etc. $X_{2dn+q} = 0$ when antenna element q is off in the transmission sector of d (if $q < n$), otherwise when $q - n$ is off in the reception sector of d . In the case of Fig. 1, X has 256 components taking values in $[0, 4]$ (on/off for the amplitude and four phase shifts).

4.2 IASSEN: A centralized heuristic for interference-aware sector design

Considering the exponentially large space of the interference-aware sector design problem, we resort to a heuristic solution. Our heuristic is a synthesis of two classic approaches in literature – Gibbs Sampling [22], and Simulated Annealing [23] – with the objective of combining their respective strengths. Gibbs Sampling converges fast and monotonically but it can be trapped in a local optimum. In contrast, Simulated Annealing converges randomly toward the objective and is less likely to be trapped in a local optimum but its convergence time can be long.

Our approach is driven by two key intuitions. First, if the Simulated Annealing started from well-selected local optimum it would significantly accelerate its convergence time. Gibbs Sampling is efficient at finding a local optimum. Second, we can accelerate the tail-end of Simulated Annealing by stopping it after a chosen duration and using the intermediate result as a starting point for a procedure that converges quickly to an optimum point – Gibbs sampling. Therefore, we introduce a three-phase heuristic, carefully wrapping Simulated Annealing in two

---

Algorithm 1: IASSEN

Data: $\mathcal{N}_{proc}$, m_{ae} , $\mathcal{N}_c$, K , input of (6)
Result: $x = (x_1, x_2, \dots, x_{2Dn})$ best value of X ,
solution to (6)

```

1 IASSEN()
2   for  $p \in [1, \mathcal{N}_{proc}]$  do
3     Generate a random initial  $x^p$  such that each
       sector has  $m_{ae}$  active antenna elements
       ( $x_i^p \geq 1$ );
4     Gibbs phase 1:  $x^p \leftarrow \text{Gibbs}(x^p)$ 
5     Simulated Annealing:
6        $x^p \leftarrow \text{Simulated\_Annealing}(x^p)$ 
7     Gibbs phase 2:  $x^p \leftarrow \text{Gibbs}(x^p)$ 
8   end
9   return  $x = \underset{p \in [1, \mathcal{N}_{proc}]}{\text{argmin}} \varepsilon(x^p)$ ;

10 Gibbs( $x$ )
11    $r \leftarrow x$ ;
12   Choose a random permutation  $perm$  of the
       components of  $x$ ;
13   for  $c \in [1, \mathcal{N}_c]$  do
14     for  $i \in [perm]$  do
15        $r_i \leftarrow \underset{r_i \in [0, \text{Card}(X_i)-1]}{\text{argmin}} \varepsilon(r_1, \dots, r_i, \dots, r_{2Dn})$ 
16     end
17     if  $r$  didn't change break;
18   end
19   return  $r$ ;

21 SimulatedAnnealing( $x$ )
22    $r \leftarrow x$ ;
23   for  $c \in [1, \mathcal{N}_c]$  do
24     1) Decrease the temperature:  $\mathcal{T}(c) \leftarrow \frac{K}{\log_2(2+c)}$ 
25     2) Build the distribution  $\Pi_c$ :
26        $\forall v \in \text{trans}(r), p(v) = \exp \frac{\varepsilon(r^v) - \varepsilon(r)}{\mathcal{T}(c)}$ 
27        $\Pi_c(v) = \frac{p(v)}{\sum_{t \in \text{trans}(r)} p(t) - p(v)}$ 
28     3) Choose next transition according to  $\Pi_c$ :
29        $r \leftarrow r^v$ ;
       Or Choose  $v$  uniformly if  $\Pi_c$  is not defined;
30   end
31   return  $r$ ;

```

Gibbs Sampling runs:

1. a *Gibbs Sampling phase* reaches a local optimum $\mathcal{O}_i$;
2. a *Simulated Annealing phase* explores around $\mathcal{O}_i$;
3. another *Gibbs Sampling phase*, issued from the exploration, reaches a (new) local optimum.

Figure 4: Evolution of Objective function during a parallel execution of IASSEN. The core Simulated Annealing phase explores the space between two Gibbs phases reaching a local minimum.

4.2.1 Gibbs Sampling phased-array weights

We construct a graph whose nodes are the components of *state* X (Eq. 6). A state transition is a change in a single node component. Although transitions are not independent, this formulation ensures a convergence to a stationary state [24].

Gibbs Sampling starts from a previous state or from a random initial sector (Alg. 1, line 3), with a minimum number of active antenna elements. We choose a random permutation of nodes, (line 12), thus increasing the chances of a process reaching a low local optimum.

A Gibbs cycle consists of choosing, for each ordered node, the state with minimum energy, keeping the others fixed (line 15). For each node i , $\text{Card}(X_i)$ transitions are compared (e.g. 5 for the simple scenario). The process is repeated until no change occurs (convergence to a local optimum, line 17) or a maximum number of cycles is reached (e.g. 30). Depending on the delay constraints, this last restriction can trade off the objective function with the execution time.

4.2.2 Simulated Annealing for a better local optimum

This phase is meant to overcome Gibbs Sampling's issue of being trapped in a poor local optimum. In SA, a *temperature function* decreases according to the number of cycles it has run (line 24), favoring the random process. K is a chosen constant and c is the cycle number.

At each c in SA phase, IASSEN builds a distribution of state transition probabilities Π_c (line 27), according to the gain in SINR of each possible transition and the temperature. We denoted $\text{trans}(r)$ the set of possible transitions from state r at a cycle c , and r^v the state after transition v .

In other words, we construct the distribution Π_c based on the negative gain in energy provided by each possible transition. The distribution only depends on the gains achievable at cycle c . Hence, we obtain a formulation similar to the one in [24].

IASSEN chooses the next transition according to Π_c (line 28). If the SA luckily reaches a local optimum, no further gain is possible. In that case, IASSEN randomly chooses the next transition (line 29).

4.3 IASSEN Illustration

Fig. 4 depicts the evolution of the IASSEN through its three phases as it aims to minimize $\varepsilon(X)$ (right axis). During the first *Gibbs phase*, the *temperature* remains constant ($K = 8$, left vertical axis). A few cycles are needed (4 in this example, reported in the horizontal axis) to reach the local optimum. To escape this local optimum point, IASSEN switches to simulated annealing. The second phase lasts 30 cycles in this example as the temperature decreases, narrowing IASSEN's search for an optimum point. To accelerate the convergence to the new optimum point IASSEN enters its third phase involving a second Gibbs run.

4.4 Obtaining channel component estimates as input

We imagine a scenario where the standard is generalized for a multi-AP deployment as follows:

4.4.1 Sweep phases run for any pair of devices

Obtaining RSS values from any neighbor, not only for the associated STA-AP pairs, is required, like in [18]. This implies a precise protocol implemen-

tation that avoids collisions of feedback frames, overlapping sweeps, and a too long beamforming.

4.4.2 A Signal Quality Indicator enables to estimate noise

SINR values are commonly available (e.g. on the Talons) and noise can be otherwise modeled (e.g. with a constant average).

4.4.3 Antenna elements contributions are individually known

Since this can be obtained by comparing the RSS values of specific sectors [20], our first generalization holds.

Besides, IASSEN requires an estimate of the per-device traffic. Indeed, the interference caused by a neighbor depends on the average portion of time it is actively transmitting.

5 Simulation

To evaluate IASSEN we have implemented a set of Monte Carlo measurement-based simulations. Unlike [25] the scope of our simulator is not the full evaluation of IEEE 802.11ad but the comparison of sector design strategies.

We compare IASSEN against two approaches:

1. the *Default* strategy implemented on Talons using a 36-sector codebook.
2. ACO, the strategy introduced in [20] with the aim of maximizing the SNR of a single link.

The simulation framework enables us to define a topology, where we choose number, positions and orientations of the pair of devices. For each topology instance, a given number of successive sector selections is run, and evaluated against a given number of channel realizations: each *realization* of the simulator provides random samples of noise, attenuation, fading, distortion and traffic, according to different laws of probability. Table 1 summarizes the parameters of our simulations.

Table 1: Simulation parameters

	Parameter	Values	Distrib.
	channel realizations	100	
	ant. elem. amp. chan. atten.	max. 20%	uniform
	ant. elem. chan. phase shift	max. 0.1 rad	uniform
	successive runs of SA	2, 3, 100	
	cycles of first SA run	60	
	cycles of next SA runs	30	
m_ae	min. active ant. elem.	6	
M_ae	max. active ant. elem.	25	
	transmission power	100 mW	
γ	path loss exponent	4.2	
λ	carrier wavelength	0.005 m	
$\mathcal{L}$	loss factor	40	
d_0	reference distance for PL_0	1 m	
χ	log-dist. flat fading, std. dev.	3.92 dB	normal
σ	noise mean, std. dev.	-60 dBm, 4 dB	lognormal
$\mathcal{P}$	traffic mean, std. dev., err.	0.5, 0.25, 0.05	normal

5.1 Channel modelling

In order to reflect a real experiment, we simulate a situation where every device obtains RSS and noise estimates to and from their neighbors in a reasonable time.

We model path loss, noise and traffic as normal variables driving the channel realizations. In contrast with the channel simulation, IASSEN, ACO and Default algorithms receive the mean values of noise and path loss as input to compute the RSS, feed Eq. 4 and compute the sectors. We add a short normal error between the traffic prediction and its realization, leading to a more realistic evaluation of IASSEN's performance. We consider a propagation environment where the Friis model is applicable up to a reference distance of 1 m, and the log-distance model for longer distances (see Table 1).

5.2 Antenna gain considerations

We aim to simulate the behavior of the Talons. Steinmetzer et al. measured their characteristic array factor [19], i.e. the matrix of complex gains for each antenna elements in a wide range of angles. We interpolated these measurements in all directions to integrate them in our simulation. Indeed, there were no measurements for the backside, nor outside the range of elevations $(-30^\circ, +30^\circ)$. We set a value 0 at azimuth 180° and elevations -40° and 40° , and completed the measured ranges and the blanks between effective measurements linearly.

For each direct link and antenna element, we add in the simulation, before each run of the algorithms, uniform variations of the amplitudes and phase shifts (see Table 1) to consider channel perturbations and measurement errors.

Fig. 5 shows a vertical projection of the 3D patterns of a device with three neighbors. Colored zones and black curves are the inter quartile range and mean of values, respectively, reported for all the elevations in one degree of azimuth. We show the neighbors around the beam, the peer (dev. 2) between brackets.

In our simulations, we assume the Default and ACO reception sectors are the same as the selected

transmission sector. We show in Fig. 5 the performance of IASSEN, which selects transmission and reception beams showing intermediate gains to peer dev. 2 w.r.t. Default and ACO, but with a reduced gain to interferers 0 and 3. Note that the specific directional value considered by the sector selection algorithms may be outside the inter quartile range, thus not fully represented in the figure.

5.3 Expected throughput considerations

We make use of IEEE 802.11ad to deduce the expected throughput in a given direction from an arbitrary range of SINR values. Indeed, we matched the recommendations of the standard in terms of sensitivity and choice of modulation (MCS) with SINR bounds. Hence, we adopt the same performance evaluations in terms of throughput as [20], which used this approach to overcome the hardware limitations due to the absence of control over the MCS choices of the Talons.

Table 2 lists the expected bitrate bounds associated with each given MCS value. Note that our choice of SINR values is more demanding than the choice of SNR in [20], so we obtain more realistic throughput achievable with the Talons.

5.4 Evaluation Results

5.4.1 Specific

we first highlight the strength of IASSEN in 2 specific scenarios (Fig 6, 7). In both cases, the 3 algorithms are run for 100 variations of amplitudes and phases, and each resulting selection is evaluated on 100 channel realizations. Traffic is constant at 80% load ($p = 0.8$ for all the devices).

The worst case for ACO occurs when the devices form 2 very close links in the opposite corners of a rectangular volume (15 m, 15 m, 3 m, Fig. 6). We expect the beams maximizing the SNR to also interfere the most. In Fig. 7, IASSEN performs satisfactorily but has limited improvement margin: each device points to the opposite corner and the beams impact less the neighboring traffic.

Figure 5: Beams formed after the three sector selections in a 4-mesh topology, for device 1. In our simulation scenario, Default and ACO choose the same sector for transmission and reception. IASSEN trades off the gain to and from device 2 with the interference to and from device 0 and device 3.

Figure 6: Two long parallel links in proximity (AP 0 $\leftrightarrow$ STA 3, AP 1 $\leftrightarrow$ STA 2), each device points to its peer. ACO's gains in SNR w.r.t. Default do not compensate the interference. IASSEN outperforms ACO in terms of SINR and throughput, except for 10% of cases where the direct gains dominate the SINR.

IASSEN significantly improves both Default and ACO in 90% of the cases in Fig 6. The remaining 10% of cases result from a locally stronger attenuation of the interfering signals, or from the narrower beams of ACO: interference is luckily avoided. ACO and Default present similar SINRs and expected throughput, confirming that default beams are already good candidates for single links in easy line-of-sight situations.

ACO provides stronger and narrower beams than Default (Fig. 7). However, in 40% of the cases, IASSEN reaches expected individual throughputs higher than 2.5Gbps, while only 5% for ACO. **While keeping a good signal quality, IASSEN reduces side-lobes strength and, hence, the interference.**

As expected, the SNR values (obtained without interfering traffic) are the highest for ACO, thus the interference increases. In Scenario 6, IASSEN has even lower SNRs than Default, because Default beams are too strong toward the neighbors.

5.4.2 Mesh topologies

Next, we evaluate scenarios where pairs of devices are uniformly distributed in a volume of (15 m, 15 m, 50 cm). We restrain the height in order to remove elevations where antenna element gains are null (experiments in [26] only reported values between -30° and 30°). Instances with 8 and 24 devices are shown in Fig. 8. APs are indexed with the first half of the IDs, STAs are linked to the AP with complementary ID, so each link gets the same ID's sum (e.g. links 0-23, 1-22, 2-21 in Fig. 8b).

For each instance, the 3 algorithms run 3 times over 100 channel realizations. The traffic $\mathcal{P}$ is normally distributed around mean 0.5 and standard deviation 0.25. Each realization creates a normally distributed traffic error varying 5% each p.

The increase in number of devices is only partially reflected in the sums of SINRs: as more interferers are present, individual values decrease. ACO's gain

Figure 7: Two long crossed links (AP 0 $\leftrightarrow$ STA 3, AP 1 $\leftrightarrow$ STA 2), devices facing their peers at opposite corners. ACO chooses narrower beams than Default ones, but is not designed to reduce interference: IASSEN's process leads to more than 500Mbps gain in individual throughput.

Figure 8: Mesh scenarios: device topology (position and orientation of APs and STAs), distribution of sum of SINRS (dB), distribution of expected throughputs (upload, download). IASSEN is resilient to the diversity and the density of many scenarios.

in individual SINRs over Default generally reduces from 5 to 3 dB while IASSEN overcomes ACO by up to 20 dB for Instances 8a, 8b, because the gains in RSS for the AP-STA pairs increases the sum in the divider of Eq. 4 when many interferers share the space.

The inflections in IASSEN's SINRs sum curves, Instance 8b around 30dB and 8a, 20dB, are coincidental: the uniform distribution of the devices lead to contrasts (e.g. in 8b link 11-12 is short and closely aligned, link 7-16 is long and misaligned).

Instance 8a shows a better performance in up-link than in downlink, because the stations are concentrated in the bottom left corner. Furthermore, IASSEN benefits from the relative isolation of link $6 \rightarrow 1$. Fig 8b shows more homogeneously

IASSEN's performance over the diversity of links: IASSEN outperforms ACO by more than 3 Gbps in upload and 1 Gbps in download.

5.4.3 Influence of expected traffic

We run a third simulation campaign to evaluate the impact of traffic load on the performance of IASSEN. We adopte the same rules as in Scenario 8a to form grid topologies with 8 devices.

We vary the mean of the normally distributed p from 0 to 1 by intervals of 0.2. The standard deviation is set to 0.1, and the prediction error to 0.05. The p values are clipped to 0 and 1, both in algorithm inputs and in the realizations.

Figure 9: A scale of traffic share on a set of 8 mesh topologies. When p increases, the SINR distribution degrades and spatial reuse becomes harder.

Table 2: IEEE 802.11ad sensitivity recommendations and SINR

MCS	Sensitivity	Bitrate	SINR
0	-78 dBm	27.5 Mbps	1 dB
1	-68 dBm	385.0 Mbps	8 dB
2	-66 dBm	770.0 Mbps	12 dB
3	-65 dBm	962.5 Mbps	14 dB
4	-64 dBm	1155.0 Mbps	16 dB
5	-62 dBm	1251.3 Mbps	17 dB
6	-63 dBm	1540.0 Mbps	18 dB
7	-62 dBm	1925.0 Mbps	20 dB
8	-61 dBm	2310.0 Mbps	22 dB
9	-59 dBm	2502.5 Mbps	26 dB
10	-55 dBm	3080.0 Mbps	36 dB
11	-54 dBm	3850.0 Mbps	38 dB
12	-53 dBm	4620.0 Mbps	40 dB

For each mean value of p , we evaluate 10 topology instances, 3 runs of the algorithms and 100 channel realizations. Fig. 9 shows the evolution of performance, the curves representing the mean values while the boxplots show minimum, 25%, median, 75% and maximum.

While the sum of SINRs decreases when the load increases for Default and ACO, it remains stable for IASSEN (around 40dB). This is due to a short but regular gain in individual SINRs, which is close to 0dB in average, i.e. the interference caused by the proximity of the devices is degrading the links.

As a result, the network-wise expected throughputs degrade with each increment of the traffic load. **Even for high traffic load, IASSEN systematically outperforms ACO and Default.** The gain of performance is driven by the last quartile (the mean curves are closer to 75% than to the medians).

5.4.4 On the proximity of Access Points

Finally, we adopt a more realistic topology described in ([18], Fig. 12), generalized without reflections nor obstacles: we place up to 12 APs side-by-side on vertical plane (a wall). The STAs are uniformly distributed on a 10m square.

Fig. 10 shows an instance of the topology with 12 links. The traffic load p is set constant and low (0.05). Our evaluation has harder channel constraints, fixed associations, and represents half-duplex devices forming strong side lobes, thus explaining why we obtain lower throughput than the authors in [18].

The proximity of the APs causes more interference in uplink because the transmission beams of the STAs

Figure 10: A shelf of twelve APs lazily oriented toward STAs distributed in 100 squared meters (topology comparable with [18]) and a low traffic pattern ($p = 0.05$). The proximity of the APs inhibits the interference reduction in upload. **IASSEN reaches up to 10Gbps in download.**

overlap. IASSEN then compensates the overall objective by neatly improving the download, so that **under harsh conditions with close APs, IASSEN still outperforms ACO.**

6 Related work

6.1 Optimized beamforming on the Talons

Adaptive Codebook Optimization (ACO) both extracts the full channel information at antenna element level, and successively improves the SNRs by constructing sectors [20]. ACO repeats *training code-*

Figure 11: A training codebook used in ACO [20]. The four phase shifts of each active element are evaluated sequentially. Beside the elements composing the reference sector (sector 0, issued from a previous cycle), search elements enable the exploration of new sectors.

books (Fig. 11) to retrieve the contribution to the gain of each antenna element, using a Fourier Transform and operations on the SNR measurements. ACO builds the best transmitting sector only considering

amplitude estimates: the discrete phase roundings when aligning the phases lead ACO to choose sub-optimal beams.

6.2 Other SNR-based beamforming

Beam selection mechanisms for a set of devices in a dense space have been proposed centrally [27, 18] or distributively [28]. While efficient at mitigating interference by choosing among few predetermined sectors, directly using classical IEEE 802.11ad sweeps [4] and SNR measurements, they do not scale up for larger search spaces.

Interference analysis in directional wireless networks has been subject of intensive research [29, 14]. Beamforming techniques were evaluated experimentally [12], confirming the importance of the reception sector selection. An interference-aware scheduler was proposed [16], assuming conflicts only occur when directional boresights overlap. In an indoor scenario with reflections and COTS devices, the destructive impact of imperfect beams with side-lobes is not considered.

Imperfect beams with strong side-lobes have been considered in more realistic approaches [13, 17]. A strong impact is observed at short range, and a non-blocking impact otherwise. Scheduling heuristics are provided in the case interference is lower enough, where we consider realistic beams indoors.

Finally, some of the classical phased-array management approaches comprise side-lobe suppression techniques [5]. But these latter are designed for cross-path interference of a unique flow. Hence, they are not applicable to our multi-flow scenario.

6.3 Interference-aware optimization

The approach in [24] is similar to ours, although related to omni-directional Wi-Fi and power control. The authors intend to find the trade-off opposing transmission power and channel assessment threshold that maximizes the network throughput. This opposition corresponds to our beamforming problem. Mhatre et al. reduce the problem to a form they identify as NP-hard, and apply a Gibbs Sampler mechanism and the Simulated Annealing theory to find good candidate solutions.

7 Conclusion and Future Work

IASSEN centrally designs sectors for IEEE 802.11ad COTS devices, aligning beams so interference is reduced, thus throughput is improved. Given estimates of phased-arrays characteristics, it explores the space of solutions by smartly piecing together determinism with a Gibbs sampler, and controlled randomness by using Simulated Annealing.

The simulations in various scenarios prove that IASSEN scales up and beats the state of the art in terms of average SINRs and throughput in the majority of cases, e.g. for high densities of traffic or high proximity of devices.

With a whole implementation of IASSEN (e.g. on Talons, overcoming current limitations [26]), we will evaluate its performance under dynamic radio conditions. We envision a faster update of existing sectors, to react e.g. to a new device, or to a drop in amplitude. Because the sector selection only depends on the neighborhood, we study a distributed version of IASSEN, preserving the global correctness. A fairer resource allocation can be considered by extending IASSEN's INLP.

Acknowledgment

This work was supported in part by the Agence Nationale de la Recherche under the ANR JCJC CiTADEL grant.

References

- [1] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking through Flexible Beams," *ACM SIGMETRICS Performance Evaluation Review*, vol. 43, no. 1, pp. 71–84, 2015.
- [2] H. Xu, V. Kukshya, and T. S. Rappaport, "Spatial and temporal characteristics of 60-GHZ indoor channels," *IEEE JSAC*, 2002.
- [3] T. S. Rappaport, Y. Xing, G. R. MacCartney, A. F. Molisch, E. Mellios, and J. Zhang, "Overview of Millimeter Wave Communications for Fifth-Generation (5G) Wireless Networks- With a Focus on Propagation Models," *IEEE Trans. on Antennas and Propagation*, 2017.
- [4] IEEE, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," in *IEEE 802.11ad*, 2012.
- [5] R. A. Monzingo and T. W. Miller, *Introduction to adaptive arrays*. Scitech publishing, 2004.
- [6] H. Hassanieh, M. Abdelghany, O. Abari, D. Katabi, M. Rodriguez, and P. Indyk, "Fast millimeter wave beam alignment," *SIGCOMM 2018*.
- [7] N. J. Myers, A. Mezghani, and R. W. Heath, "Swift-Link: A Compressive Beam Alignment Algorithm for Practical mmWave Radios," *IEEE Transactions on Signal Processing*, vol. 67, no. 4, pp. 1104–1119, 2019.
- [8] M. Daghani, A.-L. Beylot, and G. Jakllari, "An analytical model for assessing the performance of IEEE 802.11 ad beamforming training," in *2020 IFIP Networking*. IEEE, 2020, pp. 307–315.
- [9] K. Venugopal, M. C. Valenti, and R. W. Heath, "Device-to-Device Millimeter Wave Communications: Interference, Coverage, Rate, and Finite Topologies," *IEEE TWC*, vol. 15, no. 9, pp. 6175–6188, 2016.

- [10] H. Assasa, S. K. Saha, A. Loch, D. Koutsonikolas, and J. Widmer, "Medium Access and Transport Protocol Aspects in Practical 802.11 ad Networks," *IEEE WoWMoM 2018*, pp. 1–11, 2018.
- [11] Y. Ghasempour, M. K. Haider, C. Cordeiro, and E. W. Knightly, "Multi-User Multi-Stream mmWave WLANs with Efficient Path Discovery and Beam Steering," *IEEE JSAC*, pp. 1–1, 2019.
- [12] D. Zhang, P. S. Santhalingam, P. Pathak, and Z. Zheng, "Characterizing Interference Mitigation Techniques in Dense 60 GHz mmWave WLANs," in *International Conference on Computer Communication and Networks (ICCCN)*, 2019.
- [13] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and bane of 60 GHz networks: Practical insights into beamforming, interference, and frame level operation," *CoNEXT*, 2015.
- [14] H. Zhang, S. Huang, C. Jiang, K. Long, V. C. Leung, and H. V. Poor, "Energy Efficient User Association and Power Allocation in Millimeter-Wave-Based Ultra Dense Networks with Energy Harvesting Base Stations," *IEEE JSAC*, vol. 35, no. 9, pp. 1936–1947, 2017.
- [15] L.-H. Shen and K.-T. Feng, "Millimeter Wave Multiuser Beam Clustering and Iterative Power Allocation Schemes," in *IEEE VTC*, 2019.
- [16] X. Qin, X. Yuan, Z. Zhang, F. Tian, Y. T. Hou, and W. Lou, "Joint user-Ap association and resource allocation in multi-Ap 60-GHz wlan," *IEEE Transactions on Vehicular Technology*, 2019.
- [17] E. Arribas, A. Fernandez Anta, D. Kowalski, V. Mancuso, M. Mosteiro, J. Widmer, and P. W. Wong, "Optimizing mmWave Wireless Backhaul Scheduling," *IEEE Transactions on Mobile Computing*, 2019.
- [18] S. Jog, J. Wang, J. Guan, T. Moon, H. Hassanieh, and R. R. Choudhury, "Many-to-Many Beam Alignment in Millimeter Wave Networks," *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2019.
- [19] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in off-the-shelf IEEE 802.11ad devices," *CoNEXT 2017*, pp. 414–425, 2017.
- [20] J. Palacios, D. Steinmetzer, A. Loch, M. Hollick, and J. Widmer, "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking - MobiCom '18*. ACM Press, 2018, pp. 241–255.
- [21] S. Kutty and D. Sen, "Beamforming for millimeter wave communications: An inclusive survey," *IEEE Communications surveys and tutorials*, 2016.
- [22] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1984.
- [23] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [24] V. P. Mhatre, K. Papagiannaki, and F. Baccelli, "Interference mitigation through power control in high density 802.11 w lans," in *IEEE INFOCOM*, 2007, pp. 535–543.
- [25] H. Assasa, J. Widmer, T. Ropitault, and N. Golmie, "Enhancing the ns-3 ieee 802.11ad model fidelity: Beam codebooks, multi-antenna beamforming training, and quasi-deterministic mmwave channel," in *Proceedings of the 2019 Workshop on Ns-3*, ser. WNS3 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 33–40.

- [26] D. Steinmetzer, “Performance and Security Enhancements in Practical Millimeter-Wave Communication Systems,” Ph.D. dissertation, Technischen Universität Darmstadt, 2019.
- [27] G. Bielsa, “Analysis and Performance Improvement of Consumer-Grade Millimeter Wave Wireless Networks,” Ph.D. dissertation, Universidad Carlos III de Madrid, 2019.
- [28] D. Steinmetzer, A. Loch, A. García-García, J. Widmer, and M. Hollick, “Mitigating lateral interference: Adaptive beam switching for robust millimeter-wave networks,” in *mmNets 2017*, oct 2017, pp. 29–34.
- [29] S. Singh, R. Mudumbai, and U. Madhow, “Interference analysis for highly directional 60-GHz mesh networks: The case for rethinking medium access control,” *IEEE/ACM Transactions on Networking*, 2011.