Collaborative flow control in the DARPA spectrum collaboration challenge

Reference:
Mennes Ruben, Struye Jakob, Donato Carlos, Camelo Miguel, Jabandzic Irfan, Giannoulis Spilios, Moerman Ingrid, Latré Steven.- Collaborative flow control in the DARPA spectrum collaboration challenge
IEEE transactions on network and service management / IEEE [New York, N.Y.]- ISSN 2373-7379 - 17:4(2020), p. 2024-2038
Full text (Publisher's DOI): https://doi.org/10.1109/TNSM.2020.3031078
To cite this reference: https://hdl.handle.net/10067/1733880151162165141

Collaborative Flow Control in the DARPA Spectrum Collaboration Challenge

Ruben Mennes*, Jakob Struye*, Carlos Donato*, Miguel Camelo*, Irfan Jabandžić†
Spilios Giannoulis†, Ingrid Moerman†, and Steven Latré*
*IDLab, University of Antwerp - IMEC, Belgium, [firstname.lastname]@uantwerpen.be
†IDLab, Ghent University - IMEC, Belgium, [firstname.lastname]@ugent.be

***Abstract*—Wireless network technologies are becoming more and more popular. Because of this, important parts of the wireless spectrum become overloaded. Static spectrum allocation, which has been the norm for decades, is not suitable anymore. To maintain the high demand for spectrum and the continuous development of new wireless technologies, there is a need for an intelligent, dynamic spectrum allocation mechanism, where different network technologies collaboratively optimize the spectrum usage. New wireless network paradigms, such as Neutral Host Networks (NHNs) and private 5G, require a smart, spectrum-footprint-aware flow control algorithm to overcome the spectrum scarcity in collaborative way.**

This article presents a strategy, vision and flow control mechanism to implement collaboration in a Quality of Service (QoS)-driven way. The solution in this article is based on policies which may activate depending on its current and neighbor's network states. Through a flow ordering and selection strategy, these policies optimize the spectrum footprint, based on the performance and QoS-requirements of the own and surrounding networks.

The proposed algorithm is tested extensively and validated on a large scale during the DARPA Spectrum Collaboration Challenge (SC2) competition. The results of the SC2 final event and intermediate scrimmages showed that the proposed approach increased the score, indicating increased inter-network collaboration was achieved.

***Index Terms*—Quality of Service, Wireless flow control, Wireless networks and cellular networks, Wireless Spectrum Collaboration**

I. Introduction

COLLABORATION in the wireless spectrum is key to further exploration of the capacity of today's wireless communication. Driven by the increasing need for wireless capacity [1], spectrum scarcity is a familiar phenomenon, especially on the unlicensed industrial, scientific, and medical (ISM) radio bands as operators claim spectrum for peak utilization. In other licensed parts, the spectrum is underutilized. As such, spectrum is wasted in two ways: 1) by introducing collisions and interference in overutilized spectrum bands, and 2) by not using available spectrum resources in the underutilized bands.

Previous work focused on ensuring the coexistence of different technologies, especially for technologies such as Wi-Fi, Bluetooth, and ZigBee. These efforts usually rely on one technology's inherent characteristics, making use of sense-and-avoid techniques or duty cycle changes [2], [3]. Other techniques like Cognitive Radios (CRs) [4], where users can switch from overutilized radio bands to licensed spectrum if its Primary User (PU) is not using it, are another alternative. The downside of both approaches is the focus on its technology. Most of the time, only one specific technology is supported to use the licensed spectrum of the PUs or only one type of technology can be used by the PUs. The collaboration of different technologies and different networks is essential to overcome this spectrum issue.

Two noteworthy new wireless paradigms are private 5G [5], [6], able to operate in very small and specific wireless bands, and Neutral Host Network (NHN) [7], which provide spectrum for other operators. Both paradigms are typically used on crowded places such as enterprises, campuses or large venues (both indoor and outdoor). These new paradigms are applicable to modern emergency response teams following disasters. Often, they need to share the spectrum with legacy wireless networks. In a lot of these scenarios, different flows have different Quality of Service (QoS) requirements and different priorities. Nowadays, every network will try to optimize its performance, often based on link-level metrics such as packet-success-ratio or RSSI. To enable QoS and take the different priorities into account, these new paradigms need to monitor their spectrum footprint and use it in an efficient way, optimizing the overall QoS of the collision domain.

In this work, we present a strategy, vision and flow control mechanism that is able to optimize spectrum footprint based on the priority and QoS of given flows. The strategy and algorithm could be implemented within the context of NHNs or private 5G networks.

The proposed algorithm is validated during the Spectrum Collaboration Challenge (SC2), organized by the Defense Advanced Research Projects Agency (DARPA) [8]. During this three-year competition, teams were challenged to build a wireless radio system that can use the wireless spectrum in a smart, efficient, and collaborative way. Teams combined Artificial Intelligence (AI) with efficient wireless radios to collaboratively optimize the use of the spectrum among radio technologies.

The performance of each team was graded based on successful wireless delivery of traffic flows, which were associated to services with specific QoS requirements. To enforce collaboration, all teams in an ensemble were only awarded the lowest of these grades. We participated in this competition as team SCATTER and reached the sixth position. One of our crucial collaboration features is a smart and collaborative flow control mechanism. In this paper, we present the different flow control strategies implemented by the SCATTER team to support collaboration during the SC2 competition.

The remainder of this paper is structured as follows: in Section II, we discuss the related work on flow control, spectrum optimization, and spectrum collaboration. The problem statement of this work is described formally in Section III. More details about the DARPA Spectrum Collaboration Challenge (SC2) are given in Section IV. In Section V, we briefly describe the architecture of the SCATTER radio system. In Section VI, we describe the proposed algorithm and designed policies to maximize the score. The validation of the proposed algorithm is shown in Section VII. Finally, we conclude our work in Section VIII.

II. Related work

Spectrum collaboration is a broad topic and many different techniques could be applied on different parts of the wireless radio stack. While gain control and other physical layer approaches are essential optimizations towards increasing efficiency, these alone do not suffice to alleviate spectrum scarcity within a given collision domain. To optimize the Quality of Experience (QoE) when wireless devices are operating within a spectrum-scarce environment, higher layers should collaborate. Within this paper, we will not focus on physical layer aspects such as channel gain control [9]. Instead, we will focus on QoE and QoS flow control. In the remainder of this section, we discuss the state-of-the-art concerning spectrum optimization and flow control.

Most flow control algorithms work end-to-end for congestion control [10]. The most famous and widely used flow control mechanism is TCP, which is responsible for end-to-end rate control [11]. Flow control techniques are mainly developed with wired networks in mind. To adapt these flow control mechanisms to wireless environments, different techniques are proposed. One technique that is commonly used in today's wireless networks is network slicing. Network slicing allows different devices to adapt their spectrum footprint based on their needs and requirements by splitting a single physical network into several logical networks, customized for different unique requirements [12]. Slicing is a promising technique, especially for 5G, and different variations are proposed. A lot of these techniques focus on slicing within one operator [13].

To share spectrum among different operators, different solutions are proposed [14]. Salami et al. proposed an algorithm where two Mobile Network Operators (MNOs) could operate in Wideband Code Division Multiple Access networks [15]. In [16], the authors proposed to roam users of adjacent MNOs on each other's networks in a distributed manner without a centralized controller. Other works focus on the protection of incumbents by using the Licensed Shared Access (LSA) concept [17]. LSA would allow an MNO to use spectrum of another operator or incumbent user under a regulator's supervision with predetermined rules and conditions that guarantee operational certainty for the owners. All these techniques are closely related to CR technologies. Within this context, PUs and Secondary Users (SUs) are defined, where SUs can only access the spectrum if the PU is not active in the spectrum [4].

Another phenomenon in wireless systems nowadays is wireless Software Defined Networks (SDNs). One of the most recent wireless SDN orchestration architectures is the 5G-EmPOWER networking framework [18], inspired by the ODIN framework [19], which operates on a Wi-Fi network. The key to the design is the light virtual Access Point (AP) abstraction. The controller manages several physical APs. Different light virtual APs could be deployed on these physical APs and controlled by a centralized controller. Users are again separated into different networks. The spectrum is sliced based on these networks. Our framework differs because it can optimize the network differently for different users in the spectrum. In addition, decisions can be made centralized, distributed, or in a hybrid way.

Like most SDN systems, our framework is flow-based. In contrast with most SDN systems however, no specific rules about routing, firewalls or load balancing are placed in a centralized server. This provides the opportunity of combining the proposed framework with a wireless SDN framework like ORCHESTRA [20], with our framework being responsible for flow control optimization, for informing nodes in the network about each other's states (using the virtual Medium Access Control (MAC) layer), and for applying flow-specific routing rules. In contrast with previously mentioned approaches, the framework developed by the SCATTER team [21] uses the proposed collaborative flow control in this paper to manage and optimize the inter-network QoS and spectrum footprint, using information gathered within our network and received from external ones. Based on this information, the framework creates inter-network slices rather than slices on a per-user basis, as is common in CR systems.

Recent work focuses on the potential interdependence of users' decisions [22]. Vamvakas et al. proposed a dynamic spectrum management scheme for 5G non orthogonal multiple access wireless networks, where they treated the problem as a non-collaborative common pool resource game to support radios that could operate both licensed and unlicensed bands. While the work of Vamvakes et al. is focused on an adoption scheme where all users are encouraged to transmit via an unlicensed band. Based on users' decisions unlicensed spectrum band collapse is minimized by still transmitting under the safer licensed spectrum. Our proposed algorithm focuses on inter-network collaboration to guarantee QoS of accepted flows within the wireless network if the wireless spectrum is close to over-utilization. Note that the proposed algorithm is independent of the underlying wireless optimization, such as Modulation and Coding Scheme (MCS) and gain control or slot scheduling. The proposed algorithm is wireless technology independent. These technologies alone cannot fully address the issue of spectrum scarcity, and higher-level decision-making is vital for coexistence of different wireless technologies within a collision domain [23]. Note that the impact of users' decisions on the algorithms and used technologies is important to understand, but unfortunately, outside the scope of this work.

TABLE I: Overview of used symbols

Symbol	Description
C	Number of channels in the MF-TDMA superframe of the own network. $C \in \mathbb{N}$
T	Number of timeslots in the MF-TDMA superframe of the own network. $T \in \mathbb{N}$
Δ	Duration of the MF-TDMA superframe of the own network in seconds. $\Delta \in \mathbb{R}^{+}$
n	Team ID within ensemble. $n = 0$ represents the own team. $n \in \mathbb{N}$
IM_n	Set of all IMs for a given team n during a given stage.
i	Unique id for an IM.
r_i	Reward for a given IM identified by i. $r \in \mathbb{N}$
h_i	Holding period for a given IM identified by i in seconds. $h \in \mathbb{R}^{+}$
τ_i	Minimum throughput for a given non-file IM identified by i in bits per second. $\tau \in \mathbb{R}^{+}$
l_i	Maximum latency for a given non-file IM identified by i in seconds. $l \in \mathbb{R}^{+}$
δ_i	File transfer deadline for a given file IM identified by i in seconds. $\delta \in \mathbb{R}^{+}$
b_i	Size of a file to transmit for a give file IM identified by i in bits. $b \in \mathbb{N}$
t	Relative collaboration threshold. $t \in [0, 1]$
S_n	The maximum achievable score for a given team n in the ensemble. $S_n \in \mathbb{N}$
s_n	The current individual score for a given team n in the ensemble. $s_n \in \mathbb{N}$
$\hat{s}_n$	The current actual score for a given team n in the ensemble. $\hat{s}_n \in \mathbb{N}$
β_i	Benefit value for a given IM identified by i. $\beta \in \mathbb{R}^{+}$
c_i	Estimated cost for a given IM identified by i. $c \in \mathbb{R}^{+}$
u_i	Estimate of how many bits could be sent within one MF-TDMA slot for a given IM identified by i. $u \in \mathbb{N}$
t_n	Target score in comparison with neighboring team n. $t \in \mathbb{R}^{+}$
a	Aggressiveness factor used in target score formula. $a \in [1, \infty)$
d	Decay factor used in target score formula. $d \in [0, 1]$
ϵ	Minimum difference used in target score formula. $\epsilon \in [0, \infty)$

III. Formal Problem Statement

In this section, we describe the collaborative spectrum footprint problem formally. All symbols used in this work are summarized in Table I.

Consider a wireless environment in which multiple parties operate. Each party, or *team*, desires to achieve a certain level of performance, without restricting others from reaching the same goal. No prior arrangements regarding spectrum assignment or use have been made, and teams are not able to negotiate directly, but can notify others of how well they are performing at any time. How to interpret this performance indicator may be agreed upon by the involved parties beforehand. In our solution, we determine this value as follows.

First, we define the concept of the Individual Mandate (IM). An IM specifies the QoS requirements of a specific flow and associates a *reward* value to it. The reward is a direct measurement of the expected value of achieving the QoS requirements of the IM. We then define IM_n as the set of IMs for a given team n at a certain point in time. We further subdivide IMs into two classes: File mandates and Non-file mandates. Note that more types could be defined based on the different requirements of the IMs.

The non-file mandate is a five-tuple (i, r, h, τ, l). The first member of the five-tuple is a unique ID for IM identification.

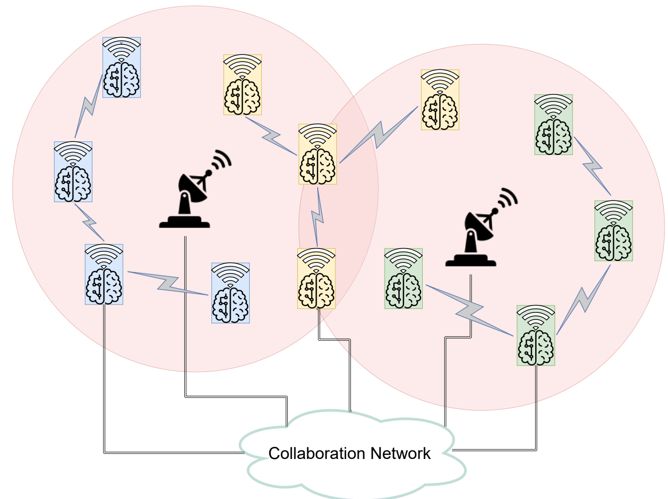

Fig. 1: Network topology with three networks (blue, yellow and green) and two incumbents (black). Each network has one gateway which is connected to a wired backbone and has the ability to communicate with other gateways and collaboration-enabled incumbents.

r represents the reward for accomplishing the requirements of the flow. h defines the holding period of the flow, indicating for how long the flow's requirements must be fulfilled before its reward can be obtained. The minimum throughput and maximum latency for a flow are defined by τ and l, respectively.

The file mandates, on the other hand, are defined by a five-tuple (i, r, h, δ, b). The first three values are the same as defined above. δ describes the file transfer deadline, or the maximum allowed time between generation of the first packet of the file, and successful delivery of the entire file. The size of the file is described by b.

We define the maximum achievable score S_n as the sum of all rewards that can be obtained at a given time. Likewise, the sum of all rewards only for IMs for which the requirements were met continuously over the holding period is s_n, the current score. Based on the actual S_n and s_n and those received from other teams, some strategy facilitating fair, dynamic distribution of the available spectrum can be followed. In the following section, we outline some specifics of the DARPA SC2 competition that influenced the design of our approach. This approach was carefully designed to easily carry over to other environments with different specifics.

IV. DARPA Spectrum Collaboration Challenge

The Spectrum Collaboration Challenge (SC2), organized by the Defense Advanced Research Projects Agency (DARPA), was a three-year competition that started in 2017. The competition aimed at optimizing spectrum usage by using AI in real networks [8]. As team SCATTER, we ended in the sixth position during this competition, out of 35 participants.

During the competition, different teams, each having ten wireless nodes, played in an environment, having to share the available spectrum smartly and efficiently. All teams played numerous matches against four other randomly picked teams, for different scenarios. For each scenario, the mobility scheme, traffic generation, and spectrum properties (e.g., available bandwidth) are fixed and well-defined. Within the ten nodes of each team, one node acts as the gateway. A wired backbone connects this gateway to a collaboration network, through which simple collaboration messages can be exchanged with other teams and, if present, with incumbents. Note that this collaboration protocol is clearly defined, and only a minimum of information can be exchanged. Some information, such as the current score, was mandatory to exchange, while other information, such as nodes' locations, could be shared if so desired.

All matches were executed on a custom-made testbed called SC2 Colosseum. The testbed exists out of 128 Software Defined Radio (SDR) wireless nodes, combined with a complex real-time RF-simulator. Every node is a full-blown server equipped with a Field-Programmable Gate Array (FPGA)-enabled SDR, and a GPU to speed up signal processing or AI-enabled algorithms. DARPA implemented the RF-simulator using the Friis transmission formula and additive white Gaussian noise, taking into account antenna gains, free-space path loss and the impact of random noise.

In the DARPA SC2 competition, all teams were required to use the IM mechanism outlined in Section III. Furthermore, a collaboration threshold t was defined. Teams were expected to ensure that all participants were able to at least achieve a performance indicator t. As such, the final game score of each team was determined by taking a per-second sum of $\hat{s}_n$, where

$$\hat{s}_n = \begin{cases} s_n & \text{if } \forall n' : s_{n'} > t_{n'} \times S_{n'} \\ \min_{n'} s_{n'} & \text{otherwise} \end{cases} \quad (1)$$

By restricting the obtained score when at least one other team does not perform satisfactorily, teams are encouraged to collaborate. Our algorithm thus aims to optimize this $\hat{s}_n$ at all times. Note that the algorithm could easily be used outside the competition, as long as an appropriate, collaboration-enabling scoring function is defined. One example of such a scoring function is the Local Social Welfare Maximizing (LSWM) algorithm, defined by Sinha et al. [24]. In the remainder of this section, we describe the SC2 competition, focusing on the scoring rules.

For each phase of the SC2 competition, multiple rounds were played, each consisting of many matches of a single scenario. An SC2 round works as follows. All scores are reset to zero at the start of each round. The score of one round for a specific team is the sum of the *Match Scores* of each match of the round for the team. Note that each team participated in the same number of matches, and that a Match Score is zero for each team not participating in that match.

Each match is divided into one or more *stages*. Within a stage, each team receives a number of IMs. A team can earn these rewards if they can achieve these QoS requirements for at least h seconds (default 10, specified per IM).

Each second, each team's score is calculated as the sum of the reward values of all IMs that have met their QoS requirements for at least h seconds. For each round, a fixed *collaboration threshold*, between 0 and 1, is defined. If, for

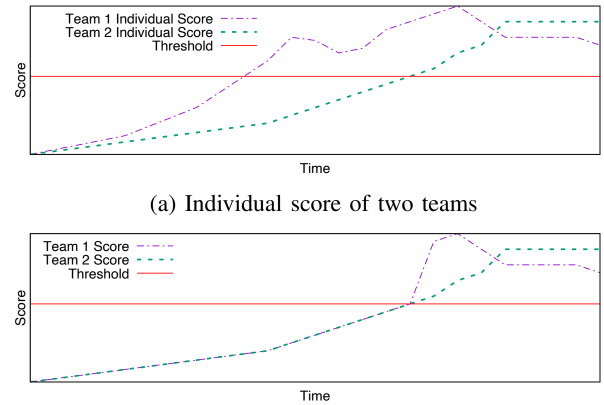

(a) Individual score of two teams

(b) Effective score of two teams

Fig. 2: Example of score where collaboration is intended below the threshold. Note that the maximum score for both teams is equal.

any of the participating teams, the fraction of available rewards scored is lower than this collaboration threshold, every team's score is reduced to that of the lowest scoring team for that second, called the *ensemble score*. As soon as all teams perform above this threshold, each team is awarded their score, as illustrated in Equation 1 and Figure 2. As a result, all 5 teams must perform above the collaboration threshold simultaneously for at least one instant for a match not to end in a 5-way tie.

V. Architecture

In this section, we describe the architecture of the SCATTER system. The SCATTER system is a complete radio stack, supporting optimal control of wireless network communication on all layers. Previous work on the SCATTER system is already available in literature. None of these published works discussed the collaborative flow control mechanisms of the SCATTER system. Overall architecture [25], predicted slot selection [26], physical radio details [27], technology-aware incumbent avoidance [28], [29], and the software architecture details used by the proposed algorithm [21] are discussed in previous work. Note that the discussed architecture could easily be implemented on top of modern wireless radios such as 5G new radio, 4G or Wi-Fi as we discuss later in this section.

A. Overall architecture

As shown in Figure 3 and described in [30], the SCATTER radio system exists out of five independent layers. These layers communicate via a message bus, using ZeroMQ[1], through well-defined messages, using Google's Protobuf[2]. Within the data plane, data can be injected into the network via the network interface. The network interface is responsible for handling IP and performing flow control. This layer forwards the data packets to the MAC layer based on the current

[1]https://zeromq.org

[2]https://developers.google.com/protocol-buffers

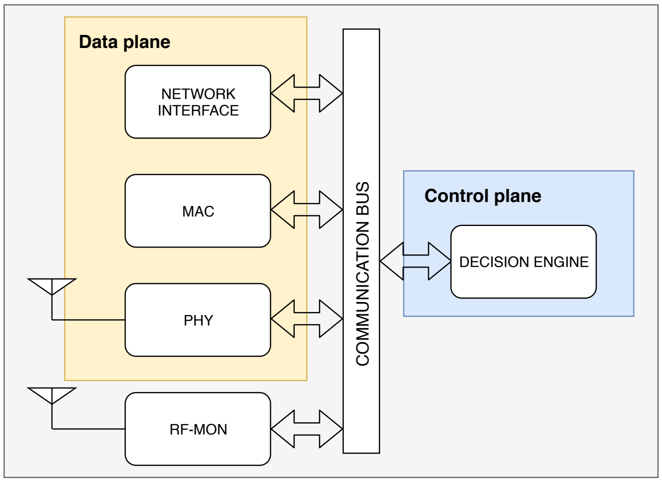

Fig. 3: SCATTER radio system architecture

settings. The MAC layer is responsible for controlling all wireless links on a packet-based level. It forwards, based on its current settings, the data packets to the physical layer. Finally, the physical layer sends out the data packet based on the settings provided by the MAC layer. When the physical layer receives a packet on its wireless interface, this path is followed in reverse.

The control plane, on the other hand, is distributed across all different layers. The decision engine layer receives all possible information about the current state of all flows and layers, and is in charge of changing the settings of the different layers. The Network Interface, MAC layer, and physical layer apply these settings. The RF-monitor layer is an exception. The RF-monitor layer only monitors the spectrum and forwards this information to the decision engine. The decision engine can use this information, together with the information about the flows and the layers, to optimize the current settings. In the next subsection, we take a closer look at the decision engine.

Important to mention is that the physical layer of our radio system is based on LTE and uses a MF-TDMA superframe, as discussed in previous work [25]. Within this work, we assume there are C channels and T timeslots within one superframe. Δ represents the duration of the superframe. Each timeslot exists out of a given number of Transmission Blocks (TBs).

B. Decision Engine Architecture

As shown in Figure 4, the decision engine exists out of two parts: 1) the node part, which contains all modules that are executed on every node; 2) the gateway part, which contains modules that only run on the node that is selected as the gateway. Communication between nodes and gateway is possible by using the own radio stack. Note that this communication is expensive and therefore should be minimized. Each module within the decision making engine could be 1) an optimization module, which can change settings on other layers and will optimize a specific target; 2) a processing module, responsible for processing data so it can be used by other modules; or 3) a communication module, communicating with other layers or other parties. These modules abstract all communication details to exchange information.

Within this paper, we are especially interested in the *Node Mandate Policy* and the *Gateway Policy*. These two modules are responsible for performing flow control and exchanging this information with the Network Interface of each radio. In previous work, we already described the internal details of both modules [21]. Within these modules, there are four basic structures: 1) An abstract representation of an IM or flow, 2) Observers, responsible for observing parts of the system and sharing that information, 3) Policies, responsible for controlling IMs/flows and other parts of the system, 4) Handlers, controlling the policies and observers. We discuss the observers and policies in more detail below.

1) Policies: A policy describes the behavior of handling flows and a set of constraints that describe when a policy can be active. An active policy has (along with other active policies) control over the flows. This is realized by enabling the appropriate handlers. The behavior of the flow could be rule-based, or Machine Learning (ML) could be used to learn policies. Each policy has an implementation both on the node and on the gateway. A policy is first activated on the gateway, after which the gateway notifies all nodes of this activation.

Two classes of policies exist within the system: performance-driven policies and environment-driven policies.

Performance-driven policies are policies that optimize the performance of the system, in the context of the SC2 competition measured by the score. These policies typically try to enable as many flows as possible to increase the performance. Exactly one performance-driven policy can be active at any time.

The environment-driven policies, on the other hand, are policies that limit the active flows proposed by the active performance-driven policy. Based on environmental observations, environment-driven policies could be activated. Such a policy could, for example, restrict the number of active flows to ensure a part of the spectrum is available for other users or incumbents. In comparison with the performance-driven policies, there is no limitation on the number of active environment-driven policies at any time.

2) Observers: Tracking the system and describing the state of subparts of the system is realized by observers. Note that observers can exist both node-side and gateway-side. Observations at a node can differ from those at the gateway. Some information, for example about the status of flows, should be forwarded to the gateway. As this information is shared using the wireless radio, this communication should be minimized.

Within the context of this paper, two observers are essential. First, the Flow Observer (FO) observes all IMs in the system. Node-side, the FO tracks the status of all IMs on that specific node and forwards the necessary information to the gateway FO. This makes it possible to have a complete overview gateway-side and a partial overview node-side. Note that this communication comes with an additional cost. All nodes should forward details about the currently active flows to the gateway. For each IM which has data available, the IM ID, current MCS, and throughput is forwarded every

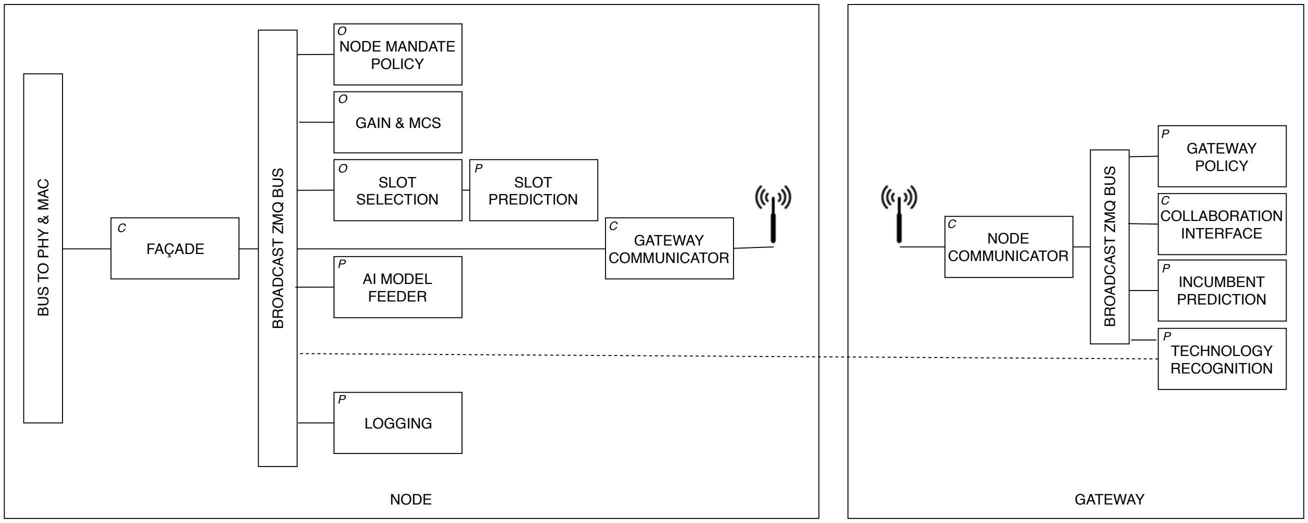

Fig. 4: Architecture of the decision engine of the SCATTER radio system, based on communication modules (annotated with *C*), optimization modules (annotated with *O*), and processing modules (annotated with *P*). The antenna represents a UDP socket that injects data into the own radio stack via the network interface.

observation period. Note that because some IMs are using the same link, MCS statistics can be shared, reducing the communication overhead. The larger the observation period, the less communication is necessary, but the later the gateway can react to the current situation. We propose to use the same observation period as default holding period (in the context of the SC2 competition, 10 seconds). Important to mention is that wireless networks that operate in infrastructure mode (using an AP or base station) already share information about the flows, meaning the additional statistics can be piggybacked there.

Another critical observer for the remainder of this paper is the Neighbor Network Observer (NNO). This observer collects all the information about the neighboring networks. In the context of the SC2 competition, this is that network's team's current score, and the score they need to cross the collaboration threshold.

C. Implementation within state-of-the-art networks

The presented architecture is created in the context of the SC2 competition. The main principles are represented in large state-of-the-art network and radios as well, especially in the context of 5G radio networks. The 4G and 5G Open RAN (Radio Access Network) architecture proposed by O-RAN Alliance [31] has a similar architecture. The RAN Intelligent Controller (RIC) is able to run real-time controlling apps, called xApps, (even from third parties) by using the Application Programming Interface (API) defined by the consortium. The API allows the xApps to interact with the central unit, distributed unit and the radio units of the network infrastructure, even near-real-time [32], [33]. Note that within this RIC, implementation could possibly be simpler than the architecture shown in this work, as the base stations have almost all information already available.

VI. Flow Control Policies

Within this section, we describe our proposed solution to control the spectrum footprint and perform flow control while maximizing our round score. As a basic principle, passing the collaboration threshold with every team is preferable to reaching a high individual score while other teams are not passing the threshold. This strategy promotes collaboration between all the teams of the ensemble.

A. Selecting IMs

All policies need to decide which IMs to enable for achieving an optimal score or for protecting environmental entities, such as incumbents. Besides, as long as not all IMs are to be enabled, policies must not only decide how many IMs to enable, but also which.

In Equation 2, we define β_i as the benefit value for each IM i, where r_i is the reward of IM i as defined above. We define c_i as the estimated cost of IM i. α and g are hyperparameter scalars to tune the benefit function, where $\alpha \in [0, 1]$ and $g \in [1, \infty)$.

$$\beta_i = \frac{r_i^\alpha}{(g \times c_i)^{1-\alpha}} \tag{2}$$

To estimate the cost of an IM, we make a distinction between non-file mandates and file mandates. The cost of non-file mandates is defined in Equation 3. u_i is an estimated number of bits that could be sent in a single TB over the link used for IM i. This value represents the channel conditions of a given link. u_i can be estimated from the current MCS value, computed by an MCS optimization algorithm implemented in the current radio stack. In essence, this delegates channel considerations to the MCS optimizer.

$$c_i = \frac{\tau_i \times \Delta}{u_i} \times \max\left(1, \frac{\Delta}{l_i}\right) \tag{3}$$

With this cost function, we estimate the number of TBs per superframe that are necessary to achieve the requirements of the IM. In the first part of Equation 3, we estimate the number of TBs necessary based on the throughput of the IM. If the maximum latency of an IM is smaller than the length of a superframe, we inflate the cost, in the second part of the equation.

For file mandates, the cost is defined in Equation 4, again estimating the number of TBs necessary to fulfill the mandate.

$$c_i = \frac{b_i \times \Delta}{\delta_i \times u_i} \tag{4}$$

B. Performance-driven policies

We first define the performance-driven policies. Here we do not take into account the additional environmental elements, like protected incumbents. We define two performance-driven policies: the Below-threshold policy and the Above-threshold policy.

1) Below-threshold policy: The Below-threshold policy defines the situation in which networks need to collaborate to increase their score. The policy is active as long there is at least one team not passing the collaboration threshold.

The overall approach is to target a score slightly above the lowest reported other team score. By following this approach, we limit the spectrum footprint for parts where no additional rewards are rewarded. This gives the opportunity for teams with lower score to improve their score. For each neighboring team n, a target score t_n is estimated as defined in Equation 6. We define an aggressiveness factor $a \in [0, \infty)$, a decay factor $d \in [0, 1]$, and a minimum difference $\epsilon \in [0, \infty)$. The formula generates a target score slightly above the other team's score, such that our performance would be at least as high. ϵ ensures a minimum positive difference between our target score and the other team's score. The other parts of the equation are introduced to quickly increase the target if the neighboring team is approaching the collaboration threshold.

$$\bar{S}_n = \max\left(S_n, S_0\right) \tag{5}$$

$$t_n = \left(-\log\left(1 - \frac{s_n}{\bar{S}_n}\right)\bar{S}_n\right)^a s_n^{1-a} \times \max\left(1, \frac{S_0}{S_n}\right)^d + \epsilon \tag{6}$$

The target function, with different aggressiveness factors, is illustrated in Figure 5 for a case where our maximum achievable score equals that of the other team. Clearly, the aggressiveness factor a has an influence on the growth rate of the curve. Our goal is to outperform all other teams as soon as everyone meets the collaboration threshold, as this is the only way to reach the highest match score. Therefore, at the beginning of a stage, the aggressiveness factor a is high. If our attempts to reach a high score seem to prevent the other team from passing the threshold, the aggressiveness factor is decreased.

The decay factor d becomes necessary only if the maximum score of the other team S_n is lower than the maximum score of the own team S_0. The decay factor does not change the curve of the target graph, but instead rotates the curve around point $(0, \epsilon)$, as shown in Figure 6. The target score aims to keep our fraction of achievable rewards scored higher than the other team's. However, as the scoring mechanism selects the lowest *absolute* score as ensemble score, it may be beneficial to sacrifice some of our score in favor of another team's, if that other team has a significantly lower maximum score, even if that means achieving a lower score *relative to our maximum individual score* compared to the other team. While lowering the decay factor can lead to a lower individual score at the moment every team passes the collaboration threshold, or even lead to our team being the last to pass it, it aims to increase the awarded score before this point, hopefully leading to a higher overall match score. While the aggressiveness factor a can increase or decrease how fast the target score grows just before a team is reaching their collaboration threshold, the decay factor d has an impact on the complete curve and can slow down the curve. By slowly decreasing the decay factor from its default value of 1 as long as the ensemble fails to pass the collaboration threshold, the score of the weakest team (in absolute points) may increase, improving the awarded score of all participating teams.

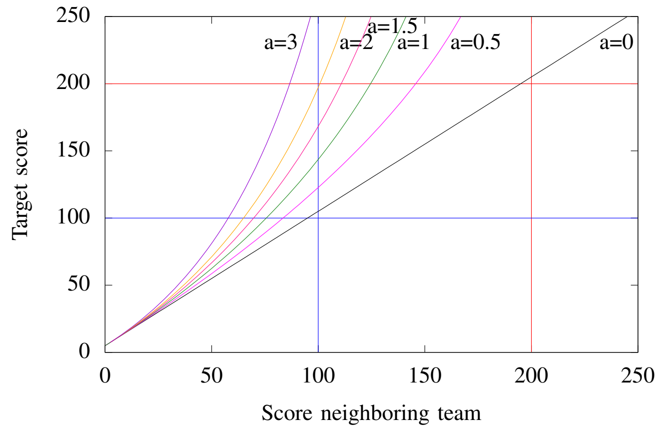

Fig. 5: Example of target functions with different values for aggressiveness factor a. $S_n = S_0 = 200$, $t = 0.5$ and $\epsilon = 5$. The red lines indicates the maximum score possible for the neighboring team (vertical red line) and its maximum score (horizontal red line). The blue lines define the collaboration threshold.

Equation 6 describes the target score in comparison with one other neighboring team. Within the Below-threshold policy, we compute the target score in comparison with all other teams. By default, it is recommended to follow the lowest computed target score. Different strategies can be applied instead, such as using the second-lowest score, or the mean target score. We define the actual target score as $\hat{t}$. This target score $\hat{t}$ is calculated on the gateway. The gateway decides to enable more flows if the current score is too low or to stop flows if the current score is too high. The details of the algorithm are described in Algorithm 1, Algorithm 2, and Algorithm 3. If the current active score is too low, the gateway selects flows to unblock based on the IMs ordering as described in

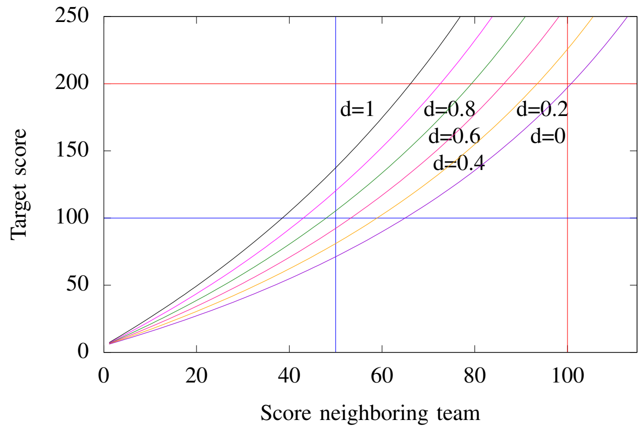

Fig. 6: Example of target functions with different values for decay factor d. $S_n = 100$, $S_0 = 200$, $t = 0.5$, $a = 2$, and $\epsilon = 5$. The red lines indicate the maximum score possible for the neighboring team (vertical red line) and its maximum score (horizontal red line). The blue lines define the collaboration threshold.

Section VI-A and Algorithm 3. If the current active score is too high, IMs are blocked. As described in Algorithm 2, unstable IMs, which do not fulfill their requirements, are blocked first. If the current active score is still not reaching the target score, we block more active IMs based on the ordering described in Section VI-A.

Algorithm 1 Below-threshold Policy

```
stable_IMs = Set of IMs that fulfill their requirements
active_IMs = Set of IMs active but not in stable_IMs
blocked_IMs = Set of IMs not in stable_IMs or active_IMs
s = s_0 + Σ_{i∈active_IMs} r_i
if t̂ < s then
    DECREASE_SCORE_TO_TARGET(t̂)
else if t̂ > s then
    INCREASE_SCORE_TO_TARGET(t̂)
end if
```

Algorithm 2 Below-threshold Policy: Decrease score

```
function DECREASE_SCORE_TO_TARGET(t̂)
    SORT_BY_BENEFIT(stable_IMs, ASC)
    SORT_BY_BENEFIT(active_IMs, ASC)
    for i ∈ active_IMs ∪ stable_IMs do
        if s − r_i > t̂ then
            BLOCK_IM(i)
        end if
        s− = r_i
        if t̂ < s then
            return
        end if
    end for
end function
```

Note that communication for this algorithm should be minimized. As the gateway has an overall view about all the flows, the gateway can communicate with the correct nodes about the changes. Because nodes can control the IMs as well, only the number of IMs should be communicated to the corresponding nodes.

Algorithm 3 Below-threshold Policy: Increase score

```
function INCREASE_SCORE_TO_TARGET(t̂)
    SORT_BY_BENEFIT(blocked_IMs, DESC)
    for i ∈ active_IMs ∪ stable_IMs do
        UNBLOCK_IM(i)
        s+ = r_i
        if t̂ > s then
            return
        end if
    end for
end function
```

2) Above-threshold policy: The Above-threshold policy becomes active as soon as all teams pass the collaboration threshold. At that moment, each team is awarded its individual score. The best strategy here is to get as many stable flows as possible. However, there is one catch; every team within the ensemble needs to stay above the collaboration threshold. To ensure this, we define B as the available budget, as a number of TBs.

When the policy becomes active, the policy estimates how many TBs are needed for all teams to stay above the threshold. This is realized based on the technology recognition framework, described in previous work [28], and the collaboration channels. Together with the number of already used slots by our network, we can estimate a budget B that represents the number of TBs that are safe to claim. Every x seconds, the algorithm releases several IMs based on the available budget B and the estimation c_i of the blocked IMs. The longer the ensemble is above the collaboration threshold, the more risk the network can take to increase the number of flows. Therefore, every x seconds the budget B is increased by several virtual TBs. Every time the ensemble breaks and falls below the threshold, we decrease the growth rate of the budget B.

As the budget B grows gradually, only a few IMs are unblocked at given timestamps. Because of this property, the communication is minimized by transmitting the number of IMs to be unblocked.

C. Environment-driven policies

Environment-driven policies [21] limit the number of active IMs if some specific events change the environment in which the radios are performing. We implement two environmental policies to address changes in the environment due to the presence of protected incumbents, which may be passive (e.g., satellites) or active (e.g., radar). These two policies are discussed below.

Note that no additional communication is necessary to implement these environment-driven policies. Because environment-driven policies can only limit the number of flows and there is always a performance-driven policy active, the same communication mechanism is reused, and no additional communication overhead is introduced.

1) Passive incumbent policy: The Passive incumbent policy is an environment-driven policy that acts in an environment where a passive incumbent is present. The passive incumbent is defined as a static incumbent that only listens to a specific part of the spectrum. The received interference at the passive incumbent due to the transmission power of the nodes of all teams in the ensemble must not exceed a given violation threshold. If a violation occurs, teams are informed through the collaboration channel, and no rewards are awarded as long as the violation continues.

Within this policy, three stages are defined. 1) The incumbent is detected and the policy activates, 2) a violation with the incumbent is detected and the policy prohibits communication interfering with the incumbent, and 3) the recovery stage.

- **Incumbent detection stage**: This stage only generates information about the presence of the incumbent, which is received via the collaboration channel. The policy informs other modules in the system that a passive incumbent is detected. Slot selection and gain modules can react based on this information. The SCATTER system tries to minimize the impact on the incumbent by two means. Firstly, the slot selection module tries to avoid slots within the area of the incumbent (if possible), and secondly the gain adaptation algorithm lowers TX-power.
- **Violation stage**: This stage is triggered when a violation occurs. Here two actions are taken. 1) TX slots at the overlapping frequencies are shut down by setting their duty cycle to zero. 2) All traffic flows on these slots are blocked.
- **Recovery stage**: Once the violation has passed, the Passive incumbent policy allows unblocking of a small number of IMs initiated by the performance-driven policy. This reduces the spectrum footprint after the violation.

Note that although the Passive incumbent policy does not control flows directly, its action of blocking all possible flows over the overlapping frequencies with the incumbent indirectly affects the performance-driven policy, as the state of each involved IM is immediately altered. Once the violation stage ends, the performance-driven policy initiates the unblocking of flows. The passive incumbent policy allows this gradually to avoid new violations due to the spectrum footprint of our radios and the footprint of the other networks.

2) Active incumbent policy: An active incumbent is a protected incumbent that not only senses, as the passive incumbent, but also can use the spectrum by transmitting their waveform. All teams can use the complete spectrum except for the moments the active incumbent is using it. Similarly to the passive incumbent, if collisions happen while the incumbent is actively sending, a violation occurs, and no rewards are awarded as long as the violation continues. This environment-driven policy follows a four stage policy: detection, learning, overlap, and recovery.

- **Incumbent detection stage**: As with the passive incumbent, information about the presence of the incumbent is received and the policy informs the other modules.
- **Learning stage** By combining signal recognition and pattern learning techniques using samples of the spectrum, SCATTER radios recognize and learn the transmission patterns of the active incumbent [29]. Here information from possible violations is also used to enhance the capabilities of the recognition system.
- **Overlap stage**: Once the transmission pattern of the active incumbent is learned, the TX slots that overlap with the incumbent are not used during the transmission of the incumbent.
- **Recovery stage**: When the transmissions of the incumbent do not reach the nodes, all the slots can be used again.

In this policy, the nodes react to this by disabling slots that correspond with the predicted pattern at the time the incumbent is transmitting in the same collision domain as the nodes. Notice that there are no actions to control flows within this policy. However, periodically disabling slots may impact the performance of some IMs. Therefore, if some IMs are not fulfilling their requirements, these flows are the first to be blocked by the Below-threshold policy. In all other cases, we assume that the active flows are necessary to achieve the target score. Other optimization modules are responsible for fulfilling the requirements as well as possible. All details about the behavior and the algorithm of the Active incumbent policy are described in previous work [28], [29].

VII. Validation

To validate the proposed solution, we use the results of the SC2 competition executed on Colosseum, as discussed in Section IV.

A. Performance-driven policies

In this subsection, we present the results of the implementation of both the Below-threshold policy and the Above-threshold policy. The Alleys of Austin scenario is a scenario of the SC2 competition in which five teams play a game of three stages. In each stage, the number of IMs increases, leading to maximum individual scores of 40, 60 and 120 in the three stages.

In Figure 7, we show two games played by the SCATTER team during the final year. In the first example (Figure 7a), the ensemble passed the collaboration threshold during the first stage. Here it is clear to see that the SCATTER radio is trying to achieve the maximum number of points. This is important as, as soon as everyone passes the threshold, teams are scored individually. Near the end of the first stage, one team dropped below the collaboration threshold, and did not recover during the second and third stages, starting at times 300 and 600. In Figure 7b, the ensemble was never able to reach the collaboration threshold. Note that, as the spectrum was fairly saturated, the worst-performing team was unable to achieve their IMs requirements. This is probably because more aggressive teams are more robust and interfere with said team. During these periods, our Below-threshold policy is taking actions.

The target area $[\hat{t}-5, \hat{t}+5]$, annotated with the orange area in Figure 7, becomes important during these below-threshold periods. In both examples, the target area follows the score

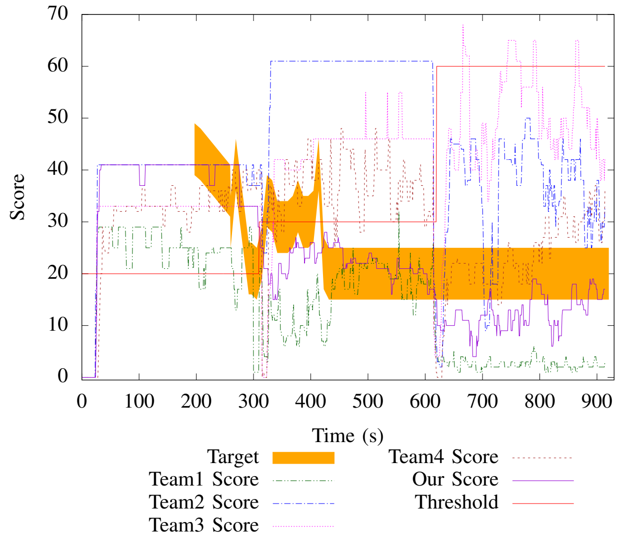

(a) Example 1

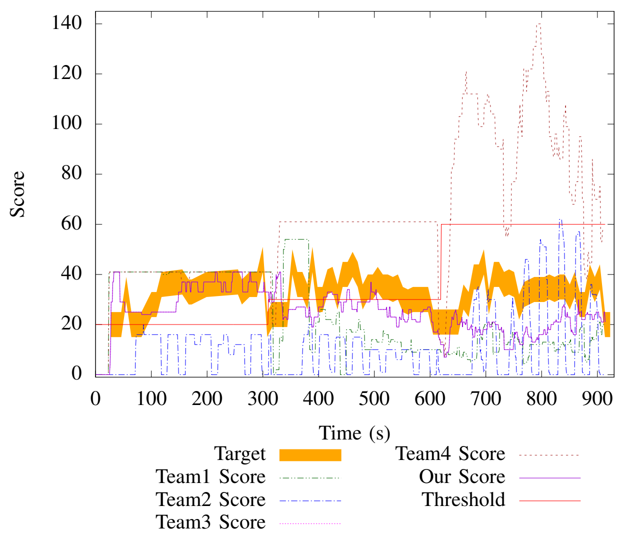

(b) Example 2

Fig. 7: Scoring of the SCATTER system in comparison with other teams during the SC2 competition in a Alleys of Austin scenario

of the weakest team of the ensemble. Note that we are not always able to achieve these scores, as the underlying layers (MAC and physical layer) were not always able to fulfill the requirements of the active IMs. The system still needs to share the spectrum and collisions with other networks occur. This could still have a significant impact on the system.

As mentioned before, the proposed strategy is to increase the score in the long run. Therefore, we discuss the results of the Alleys of Austin scenario in the final event of the SC2 competition, as well as in one of the intermediate events, called *scrimmages*. In Table II, the scores of the Alleys of Austin game during the final event of the DARPA SC2 competition are shown. Here each team was able to play 40 matches with randomly chosen other teams. The collaboration thresholds were 50%, 25%, and 10% for the three stages, respectively.

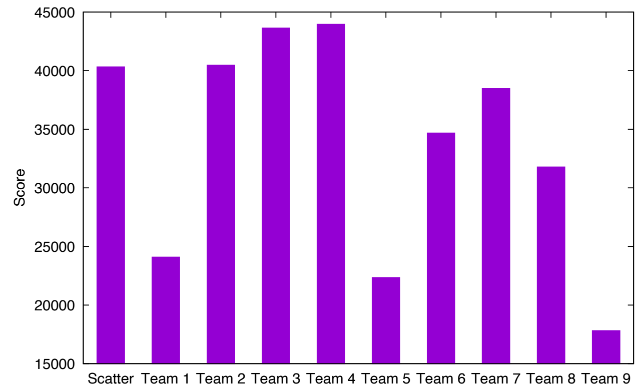

Fig. 8: Total score per team of Alleys Of Austin games during final event of the SC2 competition

In Table III, on the other hand, the results of a scrimmage during the last year of the competition are shown. During these scrimmages, more teams participated, and only a limited number of matches was executed per team (in this case, 5). In contrast with the final event, the collaboration threshold was fixed to 50% for the full match.

During the Alleys of Austin round of the final event of the SC2 competition, team SCATTER ended in the fourth position, as shown in Figure 8 and in Table II. Taking a look at each match reveals that our strategy was able to push 21 ensembles above the collaboration threshold, which is more than any other team. Also, the total number of seconds above the collaboration threshold is the highest of all teams. In terms of average duration of an above-threshold period, SCATTER was only outperformed by Team 4, which reached first place for this round.

This raises the question: why did team SCATTER not rank higher? The answer is two-fold. 1) During the final event of the SC2 competition, the collaboration threshold is lowered for each stage of the Alleys of Austin scenario. This implies that less collaboration is necessary, as the threshold becomes easier to reach. This becomes extra clear when comparing with the scrimmage results in Figure 7, for which the threshold was always high. 2) Teams 2, 3 and 4 are known to have a more robust physical layer. The high spectrum activity during above-threshold periods had a huge impact on the stability of our system, revealing some of its physical limitations.

Note that because the threshold becomes easier to reach, less collaboration is required. More aggressive teams are able to score more points in these matches as the threshold is easier to reach. Nevertheless, the results are still useful. As discussed, the important part is how often and for how long teams were in an above-threshold state. As the strategy of most non-collaborative teams is seemingly to be to be aggressive all the time, they would score highly during periods in which the ensemble reaches the threshold, but these periods would be short and rare.

During the scrimmage however, as shown in Table III, SCATTER reached the third position while only crossing the collaboration threshold during one match. This indicates that the SCATTER approach can push the score during the below-

TABLE II: Alleys of Austin statistics during final event

	Scatter	Team 1	Team 2	Team 3	Team 4	Team 5	Team 6	Team 7	Team 8	Team 9
Number of games	40	40	40	40	40	40	40	40	40	40
Total score	40 352	24 089	40 466	43 630	43 955	22 338	34 677	38 472	31 781	17 814
Ranking	4	8	3	2	1	9	6	5	7	10
Number of matches above threshold	21	13	16	19	18	11	17	18	15	7
Total time above threshold	478 s	80 s	309 s	426 s	470 s	159 s	301 s	396 s	209 s	82 s
Average above-threshold duration	22.76 s	6.15 s	19.31 s	22.42 s	26.11 s	14.45 s	17.71 s	22 s	13.93 s	11.71 s

TABLE III: Alleys of Austin statistics during scrimmage

	Scatter	Team 1	Team 2	Team 3	Team 4	Team 5	Team 6	Team 7	Team 8	Team 9	Team 10	Team 11
Number of games	5	5	5	5	5	5	5	5	5	5	5	5
Total score	31 934	38 339	23 820	18 211	33 414	7613	19 517	16 535	31 689	17 779	27 578	22 800
Ranking	3	1	6	9	2	12	8	11	4	10	5	7
Number of matches above threshold	1	3	2	2	2	0	0	2	3	2	1	2
Total time above threshold	247 s	304 s	250 s	23 s	301 s	0 s	0 s	23 s	321 s	74 s	54 s	23 s
Average above-threshold duration	247.0 s	101.3 s	125.0 s	11.5 s	150.5 s			11.5 s	107.0 s	37.0 s	54.0 s	11.5 s

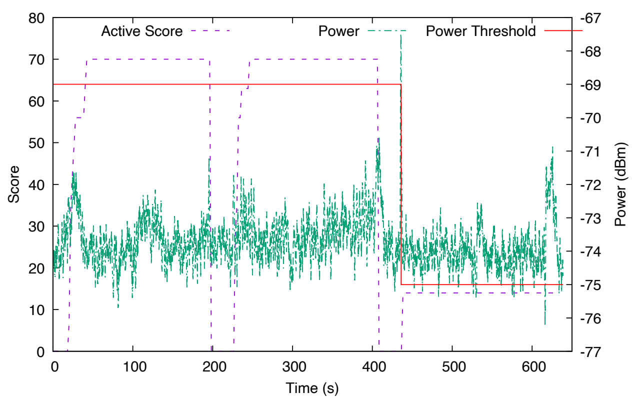

Fig. 9: Active IM score of the SCATTER system in a Passive Incumbent Environment of a Solo game.

threshold phase. Note that only five matches were played. Far from all team combinations were played, meaning there is little to conclude from analyzing the number of matches in which the collaboration threshold was crossed. There was, for example, no match involving the top five teams (SCATTER, Team 1, Team 4, Team 8 and Team 10). Nevertheless, team SCATTER was part of the game where the ensemble was able to stay above the threshold for 247 s. In the other three matches where the ensemble was able to break the threshold, this was achieved for only 53 s, 20 s, and 5 s. Note that, as teams are pseudonymized differently for each scrimmage and the final event, there is no relation between the team numbers in the two tables.

B. Environment-driven policies

If environmental elements are added, environment-driven policies may activate, as discussed in subsubsection V-B1. In Figure 9, a passive incumbent solo match is shown. Here only team SCATTER and a passive incumbent are present. In this match, the incumbent is detected at the start of the match. At that moment, the TX-power of the SCATTER system is reduced. As long as no violation is triggered, only the currently active performance-driven policy controls the active flows. In this match, during the first two stages the Above-threshold policy is active, as clearly visible in Figure 9. When a violation occurs, as is the case at time 400, the power threshold of the incumbent decreases. The Passive incumbent policy reacts to the violation as well, limiting the number of active IMs and informing the other optimization modules in the decision engine about the violation. Note that the performance-driven policy is still pushing to unblock more IMs. This is negated by the environment-driven policy. Unfortunately, we are not able to recover in this match, due to failing wireless control communication in the SCATTER system, required by other layers.

The same behavior is visible in a multi-player passive incumbent match as shown in Figure 10. A violation occurs in the last stage. In Figure 10c, we show that, from the moment the ensemble violates the incumbent, we limit our number of active mandates and reduce the maximum transmission power.

As discussed in subsubsection VI-C2, the Active incumbent policy is implemented to only inform other modules about the behavior of the incumbent. No flows are blocked by the Active incumbent policy itself. Note that the active performance-driven policy still optimizes the policies based on the score of the neighboring teams, and unstable IMs are blocked first if necessary [28], [29].

C. Real-time environment and overhead

As discussed before, the communication overhead is minimal. Policy changes and blocking and unblocking command

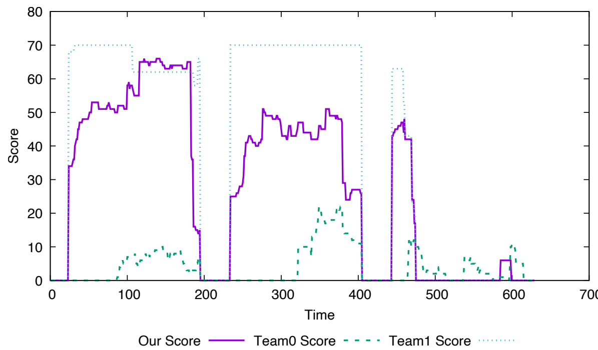

(a) Score of each team

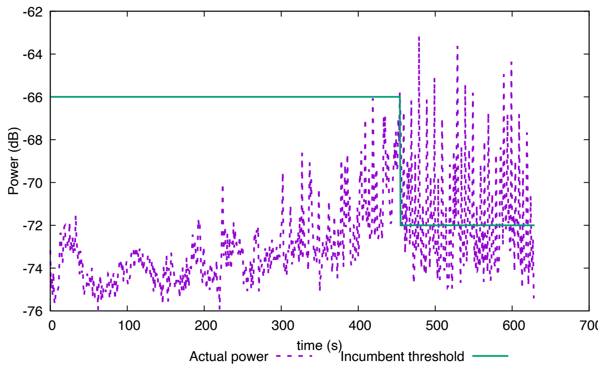

(b) Power measured by incumbent and threshold of incumbent

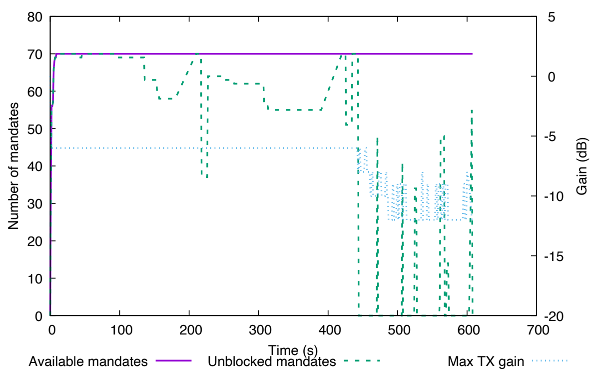

(c) Active flows and max TX power of the SCATTER system

Fig. 10: Example of flow and gain control in a two teams Passive Incumbent Environment

can be easily piggybacked on other messages or could be broadcast to the nodes. The biggest communication overhead by far is the FO, which reports the status of all flows to the gateway. All the nodes need to report the throughput, state and MCS of every IM to the gateway. We use an encoder that was able to encode this information as 48 bit per IM. This additional packet was sent to the gateway every 10 seconds, as this was the default holding and reporting period. Note that within the SC2 competition, the flow status information must be reported by the gateway, and in that sense no additional overhead was created.

In the context of the SC2 competition, the algorithm was evaluated within a large-scale real environment. More than 130 games were played during the final event, and more than 6000 games during only the last year of the competition. Because of the mesh architecture necessary within the SC2 competition, a state update occurred only once every 10 seconds, meaning decisions could be taken based on stale information. However, as many decisions are taken at the gateway, gateway-side information (such as which flows are unblocked) is usually up to date. This 10 second delay may also have advantages. It could give some time to the underlying layers (MAC and PHY) and other modules, to optimize the unblocked IMs. Note that within infrastructure mode, using an AP or base station, the AP or base station already has all information available. No additional communication is necessary, and the state of the network is always accurate and available.

D. Communication validation

In the previous subsections, we focused on the how the proposed framework and implementation of the policies are able to improve the score of all the teams within the ensemble. In this subsection, we show the impact of the proposed and implemented framework in terms of QoS-related metrics.

In Figure 11, Figure 12, and Figure 13, we show some QoS metrics of different IM types (Voice over IP (VoIP), video stream, and static image browsing, respectively) of the same game as shown in Figure 7a. The proposed algorithm reacts to the performance of active, unblocked IMs. During the first 250 seconds, the Above-threshold policy is active, resulting in all active IMs (only VoIP in this phase) are unblocked. Once the Above-threshold policy is deactivated, the proposed algorithm reacts to both the scoring of all teams and the its performance. As shown in Figure 7a we are outperforming one competing team at this point. As such, the algorithm blocks IMs to make spectrum available for that team. It is clear from Figure 12 and Figure 13 that the lowest-performing flows are blocked first. They are not stable (yet) and are not adding points.

It is clear that the proposed algorithm is QoS-aware and that blocking unstable flows has a positive effect on the QoS of the active IMs. Note that we are not able to perform this analysis for IMs of other teams as the required data was not provided by the testbed. However, their scores indicate that the QoS increase applies to other teams' active IMs too.

VIII. Conclusion

In this work, we presented the policies introduced by the SCATTER team to optimize their score during the DARPA SC2 competition. We defined two types of policies: 1) Performance-driven policies, focused on optimizing flow control by maximizing the performance (indicated by the score). 2) Environment-driven policies, limiting the number of active flows if necessary, and providing information to other modules in the system.

We defined two performance-driven policies. Which one of the two is activated, depends on whether all teams are passing some collaboration threshold. The Below-threshold policy is defined to allow the ensemble to collaborate, and to push the overall score by, if needed, reducing the spectrum footprint so that other, weaker, teams can use it to increase their performance. The Above-threshold policy is designed

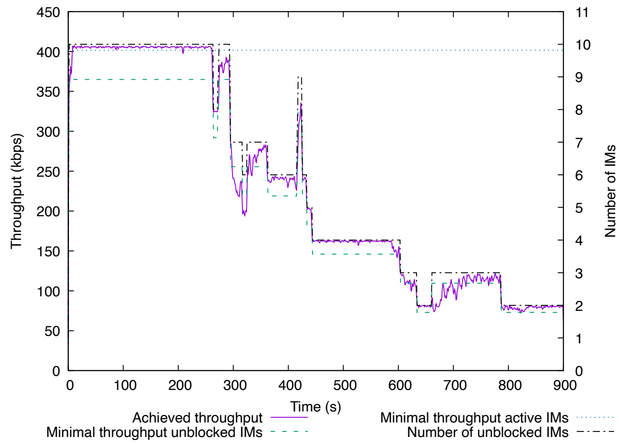

(a) Throughput of VoIP flows

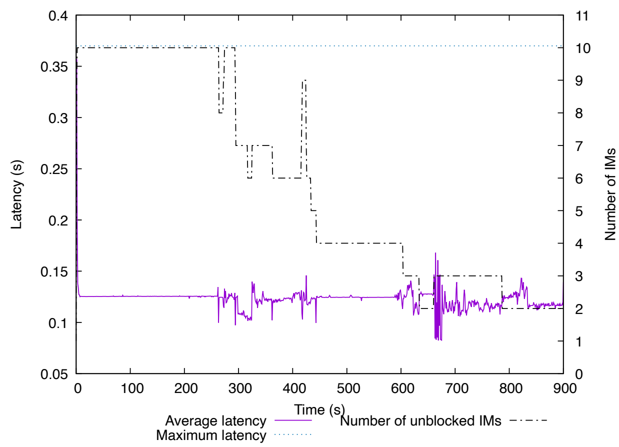

(b) Latency of VoIP flows

Fig. 11: Metrics of the VoIP flows of team SCATTER during the SC2 competition

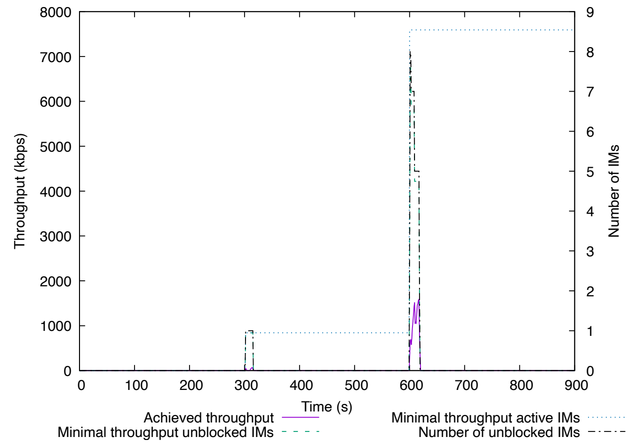

(a) Throughput of video stream flows

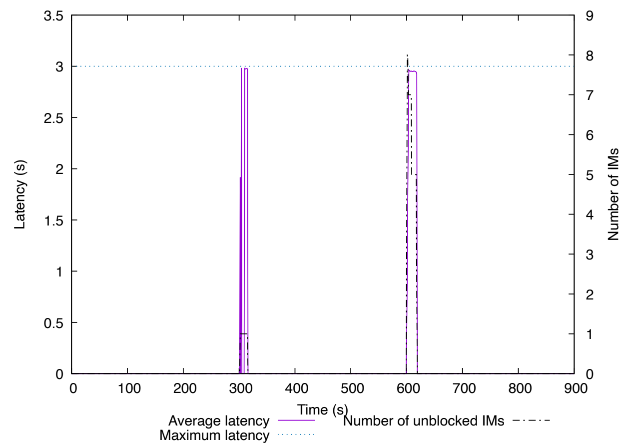

(b) Latency of video stream flows

Fig. 12: Metrics of the video stream flows of team SCATTER during the SC2 competition

to increase the number of active IMs, without pushing other teams below the collaboration threshold for as long as possible. We showed that, by applying these strategies, we could increase the number of matches in which all teams reached the collaboration threshold, as well as the duration of these periods in comparison with other teams during the SC2 competition.

The proposed environment-driven policies are designed to handle environments where (active or passive) incumbents are present. By informing other modules in the system and limiting the number of active IMs, we are able to reduce the spectrum footprint, decreasing the violation with the incumbents. This is accomplished by tuning settings and limiting the number of active IMs, or by using ML modules to predict the behavior of active incumbents to avoid collisions.

This work showed the strategy followed by team SCATTER during the DARPA SC2 competition. All policies are developed in function of this competition, but showed that collaboration with other neighboring networks within the same wireless collision domain could increase the (overall) performance. The entire framework, both type of policies and optimization strategy, could be applied in the context of private 5G networks or NHNs. Future work is necessary to apply the define the necessary policies for these new wireless paradigms as well as how different networks could agree on the priorities on flows.

Acknowledgment

This work has been funded by Research Foundation Flanders (FWO) and the Flemish Government (1S52917N and 1SB0719N).

References

[1] I. Cisco Systems, "Cisco visual networking index: Global mobile data traffic forecast update, 2016-2021 (white paper)," *Cisco Public Information*, pp. 1–35, 2017.

[2] L. Zhang, M. Xiao, G. Wu, M. Alam, Y.-C. Liang, and S. Li, "A survey of advanced techniques for spectrum sharing in 5g networks," *IEEE Wireless Communications*, vol. 24, no. 5, pp. 44–51, 2017.

[3] A. Garcia-Saavedra, P. Patras, V. Valls, X. Costa-Perez, and D. Leith, "Orla/olaa: Orthogonal coexistence of laa and wifi in unlicensed spectrum," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2665–2678, 12 2018.

[4] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "A survey on spectrum management in cognitive radio networks," *IEEE Communications magazine*, vol. 46, no. 4, 2008.

[5] A. Aijaz, "Private 5g: The future of industrial wireless," *arXiv preprint arXiv:2006.01820*, 2020.

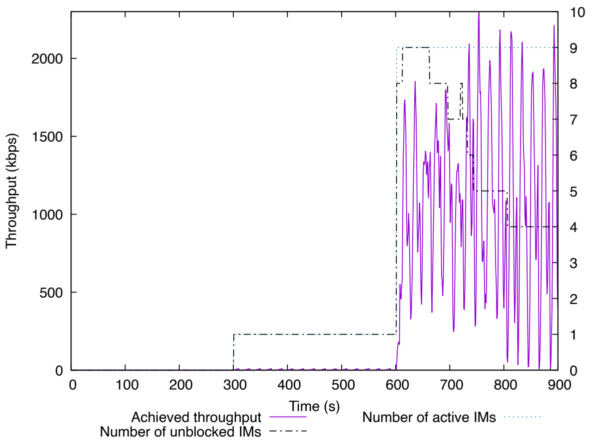

(a) Throughput of file browsing flows

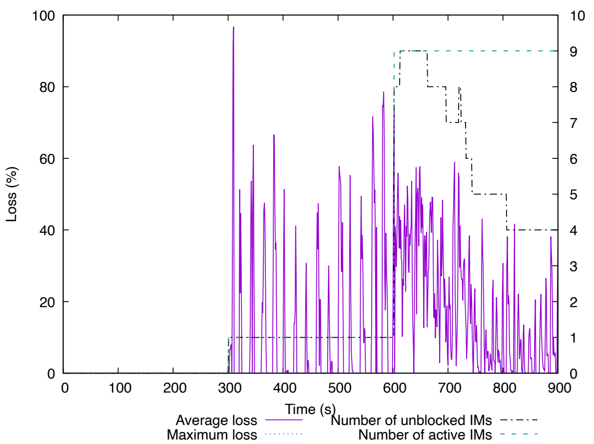

(b) Loss of file browsing flows

Fig. 13: Metrics of the file browsing flows of team SCATTER during the SC2 competition

[6] A. Rostami, "Private 5g networks for vertical industries: Deployment and operation models," in *2019 IEEE 2nd 5G World Forum (5GWF)*. IEEE, 2019, pp. 433–439.

[7] X. Foukas, M. K. Marina, and K. Kontovasilis, "Iris: Deep reinforcement learning driven shared spectrum access architecture for indoor neutral-host small cells," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 8, pp. 1820–1837, 2019.

[8] P. Tilghman, "Ai will rule the airwaves: A darpa grand challenge seeks autonomous radios to manage the wireless spectrum," *IEEE Spectrum*, vol. 56, no. 6, pp. 28–33, 2019.

[9] S. Pandit and G. Singh, "An overview of spectrum sharing techniques in cognitive radio communication system," *Wireless Networks*, vol. 23, no. 2, pp. 497–518, 2017.

[10] M. Chen and A. Zakhor, "Flow control over wireless network and application layer implementation," in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*. IEEE, 2006, pp. 1–12.

[11] V. Jacobson, "Congestion avoidance and control," *ACM SIGCOMM computer communication review*, vol. 18, no. 4, pp. 314–329, 1988.

[12] S. Zhang, "An overview of network slicing for 5g," *IEEE Wireless Communications*, vol. 26, no. 3, pp. 111–117, 2019.

[13] J. Ordonez-Lucena, P. Ameigeiras, D. Lopez, J. J. Ramos-Munoz, J. Lorca, and J. Folgueira, "Network slicing for 5g with sdn/nfv: Concepts, architectures, and challenges," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 80–87, 2017.

[14] R. H. Tehrani, S. Vahid, D. Triantafyllopoulou, H. Lee, and K. Moessner, "Licensed spectrum sharing schemes for mobile operators: A survey and outlook," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2591–2623, 2016.

[15] G. Salami, A. U. Quddus, D. Thilakawardana, and R. Tafazolli, "Non-pool based spectrum sharing for two umts operators in the umts extension band," in *2008 IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications*. IEEE, 2008, pp. 1–5.

[16] A. Yarmohammad and R. Tafazolli, "Decentralized inter-radio access network dynamic spectrum selection scheme," in *2009 International Conference on Wireless Communications & Signal Processing*. IEEE, 2009, pp. 1–5.

[17] M. D. Mueck, S. Srikanteswara, and B. Badic, "Spectrum sharing: Licensed shared access (lsa) and spectrum access system (sas)," *Intel White Paper*, pp. 1–26, 2015.

[18] R. Riggio, M. K. Marina, J. Schulz-Zander, S. Kuklinski, and T. Rasheed, "Programming abstractions for software-defined wireless networks," *IEEE Transactions on Network and Service Management*, vol. 12, no. 2, pp. 146–162, 2015.

[19] L. Sequeira, J. L. de la Cruz, J. Ruiz-Mas, J. Saldana, J. Fernandez-Navajas, and J. Almodovar, "Building an sdn enterprise wlan based on virtual aps," *IEEE Communications Letters*, vol. 21, no. 2, pp. 374–377, 2016.

[20] T. De Schepper, P. Bosch, E. Zeljković, F. Mahfoudhi, J. Haxhibeqiri, J. Hoebeke, J. Famaey, and S. Latré, "Orchestra: enabling inter-technology network management in heterogeneous wireless networks," *IEEE Transactions on Network and Service Management*, vol. 15, no. 4, pp. 1733–1746, 2018.

[21] R. Mennes, J. Struye, C. Donato, and S. Latré, "A dynamic spectrum footprint adaptation framework for collaborative spectrum sharing," in *AnNet 2020*, apr 2020.

[22] P. Vamvakas, E. E. Tsiropoulou, and S. Papavassiliou, "Dynamic spectrum management in 5g wireless networks: A real-life modeling approach," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 2134–2142.

[23] T. Maksymyuk, M. Kyryk, and M. Jo, "Comprehensive spectrum management for heterogeneous networks in lte-u," *IEEE Wireless Communications*, vol. 23, no. 6, pp. 8–15, 2016.

[24] A. Sinha and M. P. Wellman, "Incentivizing collaboration in a competition," in *18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-19)*. Montreal, Canada: Research Collection School Of Information Systems, may 2019, pp. 556–564.

[25] F. A. de Figueiredo, D. Stojadinovic, P. Maddala, R. Mennes, I. Jabandžić, X. Jiao, and I. Moerman, "Scatter phy: A physical layer for the darpa spectrum collaboration challenge," in *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. IEEE, 2019, pp. 1–6.

[26] R. Mennes, F. A. P. De Figueiredo, and S. Latré, "Multi-agent deep learning for multi-channel access in slotted wireless networks," *IEEE Access*, vol. 8, pp. 95 032–95 045, 2020.

[27] F. A. de Figueiredo, D. Stojadinovic, P. Maddala, R. Mennes, I. Jabandzic, X. Jiao, and I. Moerman, "SCATTER PHY: a physical layer for the DARPA spectrum collaboration challenge," in *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN) (IEEE DySPAN 2019)*, Newark, USA, Nov. 2019.

[28] M. Camelo, A. Shahid, J. Fontaine, F. A. De Figueiredo, E. De Poorter, I. Moerman, and S. Latre, "A semi-supervised learning approach towards automatic wireless technology recognition," in *Proceedings of the IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN), Newark, NJ, USA*, 2019, pp. 11–14.

[29] M. Camelo, R. Mennes, A. Shahid, J. Struye, C. Donato, I. Jabandžić, S. Giannoulis, F. Mahfoudhi, P. Maddala, I. Seskar, I. Moerman, and S. Latré, "An ai-based incumbent protection system for collaborative intelligent radio networks," *IEEE Wireless Communications Magazine*, 2020, submitted.

[30] S. D. Giannoulis, C. Donato, R. Mennes, F. A. Pereira de Figueiredo, I. Jabandzic, Y. De Bock, M. Camelo, J. Struye, P. Maddala, M. Mehari, A. Shahid, D. Stojadinovic, M. Claeys, F. Mahfoudhi, W. Liu, I. Seskar, S. Latré, and I. Moerman, "Dynamic and collaborative spectrum sharing: The SCATTER approach," in *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN) (IEEE DySPAN 2019)*, Newark, USA, Nov. 2019.

[31] O. R. Alliance, "O-ran: Towards an open and smart ran," *White Paper, October*, 2018.

[32] S. Niknam, A. Roy, H. S. Dhillon, S. Singh, R. Banerji, J. H. Reed, N. Saxena, and S. Yoon, "Intelligent o-ran for beyond 5g and 6g wireless networks," *arXiv preprint arXiv:2005.08374*, 2020.

[33] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia, "Open, programmable, and virtualized 5g networks: State-of-the-art and the road ahead," *arXiv preprint arXiv:2005.10027*, 2020.

Ruben Mennes obtained a masters degree in computer science from the University of Antwerp, Belgium, in June 2016. In August 2016 he joined the Department of Mathematics and Computer Science at the University of Antwerp and imec where he is active as a Ph.D. student. The focus of his research is mainly on collaborative wireless networks and smart wireless networks, the cross-domain between wireless network technologies and machine learning. He was part of the SCATTER team of the DARPA SC2 competition, where he focused on the structure, design, implementation and research of the decision making and optimisation engine.

Jakob Struye received his master's degree in Computer Science at the University of Antwerp, Belgium in 2017. He then joined the Department of Computer Science at the University of Antwerp, in pursuit of a Ph.D. Initial research efforts focused on time series prediction through neural networks and other brain-like systems. He later joined the SCATTER team, building a collaborative wireless radio participating in the DARPA SC2 competition. Current research interests have shifted towards the wireless spectrum, with a focus on high-throughput low-latency mmWave communication.

Carlos Donato received his B.Sc in Telecommunications Engineering (2013), M.Sc in Telematics Engineering (2015) and Ph.D in Telematics Engineering (2018) from Universidad Carlos III de Madrid, Department of Telematics Engineering. Currently, he holds a post-doc position at imec with IDLab research group in Antwerp. His research interests lie in wireless communications, especially LTE and IEEE 802.11 technologies, resources on demand, energy efficiency, design of network protocols, Software Defined Networking, Cloud-RAN and Artificial Intelligence techniques applied to wireless communications.

Miguel Camelo received the bachelor's degree 1203 in electronic engineering from the University of Ibague, Ibagu/'e, Colombia, in 2006, the master's degree in systems and computer engineering from the University of Los Andes, Bogota, Colombia, in 2010, and the Ph.D. degree in computer engineering from the University of Girona, Girona, Spain, in 2014. He is currently a Researcher with the University of Antwerp—imec, Antwerp, Belgium. He has authored or coauthored several papers in journals and international conferences, and has been involved in several Spanish, Belgian, and European research projects. He also was part of the SCATTER team, a multi-disciplinary research team that participated in the DARPA SC2 competition. His research interests are in the field of applied artificial intelligence and the management of communication networks.

Irfan Jabandžić received his B.Sc. and M.Sc. in automatic control and electronics from Faculty of Electrical Engineering Sarajevo, Bosnia and Herzegovina, in 2012 and 2014, respectively. After working in area of embedded engineering for 2 years, in October 2016 he joined Internet Technology and Data Science Lab. at Ghent University, where he is working towards the Ph.D. degree. He was involved in DARPA Spectrum Collaboration Challenge (SC2) as MAC specialist for SCATTER team. His main research interests are in the field of wireless networks and wireless sensor networks, focusing on flexible and adaptive MAC architectures, heterogeneous wireless technologies coexistence and dynamic spectrum access.

Spilios Giannoulis received his Diploma in Electrical and Computer Engineering (2001) and his PhD (2010) in the area of routing for wireless mobile sensor networks. From 2001 until 2015 he has been a researcher involved in several E.U. and national R&D projects at Industrial systems Institute and the Applied Electronics Laboratory (APEL), University of Patras, Greece. From 2005 until 2015, he was also lecturing at the Technological Educational Institute of Patras, Greece. He is currently a postdoc researcher in imec, Belgium involved in several EU projects in the area of reconfigurable and adaptive wireless networks. He has served as technical leader of the H2020 project WiSHFUL as well as for the Scatter Team that took part in the DARPA Second Spectrum Collaboration Challenge, resulting in winning 1.5 million US dollars prizes in total from DARPA. He is currently involved in the European Space Agency CODYSUN project. His is a member of the IEEE and his main research interests are in the fields of wireless networks, mobile ad-hoc networks, wireless sensor networks, focusing on flexible and adaptive MAC and routing protocols, QoS provisioning, cross-layer and power aware system architecture design, spectrum resource management for public and licensed bands dynamic spectrum access, 4G-5G networks with more than 40 relevant publications in related international conferences and journals.

Ingrid Moerman received her degree in Electrical Engineering (1987) and the Ph.D. degree (1992) from the Ghent University, where she became a part-time professor in 2000. She is a staff member at IDLab, a core research group of imec with research activities embedded in Ghent University and University of Antwerp. Ingrid Moerman is Program Manager of the 'Deterministic Wireless Networks' track at imec. Ingrid Moerman is also coordinating the research activities on mobile and wireless networking at Ghent University, where she is leading a research team of more than 30 members. Her main research interests include: collaborative and cooperative networks, intelligent cognitive radio networks, real-time software defined radio, flexible hardware/software architectures for radio/network control and management, Internet of Things, Next generation wireless networks (5G/6G/…), and experimentally-supported research. Ingrid Moerman has a longstanding experience in running and coordinating national and EU research funded projects. At the European level, Ingrid Moerman is in particular very active in FP7/H2020 programs, where she has coordinated and is coordinating several projects (CREW, WiSHFUL, eWINE, ORCA). Ingrid Moerman was leading team SCATTER, consisting of researchers from IMEC-IDLab and Rutgers University, in the DARPA Spectrum Collaboration Challenge (SC2). The SCATTER team has been awarded with two prizes of 750 000 USD each in Phase 1 and Phase 2 of the DARPA SC2 competition, and was one of the 10 finalist at the DARPA SC2 championship event organized at Mobile World Congress in LA (October 2019). Ingrid Moerman is author or co-author of more than 750 publications in international journals or conference proceedings.

Steven Latré is an associate professor at the University of Antwerp and imec, Belgium. He received a Master of Science degree in computer science from Ghent University, Belgium (2006) and a Ph.D. in Computer Science Engineering from the same university (2011) on autonomic network management. He is leading the IDLab research group at the University of Antwerp, consisting of 80+ researchers. His personal research expertise focuses on deep reinforcement learning and wireless network management. He is author or co-author of more than 100 papers published in international journals or in the proceedings of international conferences. He is the recipient of the IEEE COMSOC award for best PhD in network and service management 2012, the IEEE NOMS Young Professional award 2014, the IEEE COMSOC Young Professional Award 2015 and is a member of the Young Academy Belgium.