1 Optimal Repeated Spectrum Sharing by Delay-Sensitive Users

Yuanzhang Xiao[a] and Mihaela van der Schaar[b]

Department of Electrical Engineering, University of California, Los Angeles

1.1 Introduction

The spectrum is becoming an increasingly scarce resource, due to the emergence of a plethora of bandwidth-intensive and delay-critical applications (e.g. multimedia streaming, video conferencing, and gaming). To achieve the gigabit data rates required by the next-generation wireless systems, we need to efficiently manage the interference among a multitude of wireless devices, most of which have limited computational capability. Central to interference management is the *spectrum sharing* policies, which specify when and at which power level each device should access the spectrum. Given the heterogeneity and the huge number of distributed wireless devices, it is computationally hard to design efficient spectrum sharing policies.

Cloud radio access networks (C-RANs) present a promising network architecture for designing spectrum sharing policies. C-RANs consist of two components, a pool of the baseband processing units (BBUs) and remote radio heads (RRHs), and allocate most demanding computations to the BBU pool (i.e., the “cloud”) [1]–[7]. In this way, C-RANs open up the opportunities of designing efficient (even optimal) spectrum sharing protocols. However, these opportunities come with the following challenges in C-RANs [1]–[7]:

1. How to allocate the computations between the BBU pool and RRHs and minimize the message exchange between them?
2. How to cope with dynamic entry and exit in large networks?
3. How to support delay-sensitive applications that constitute a majority of the traffic in C-RANs?

This chapter presents advances made in the past years on a systematic design methodology for spectrum sharing protocols that are particularly suitable for C-RANs. The spectrum sharing protocols designed by the presented methodology can be naturally implemented in the following two phases:

- the first phase of determining the optimal network operating point, which requires most computation and can be done in the BBU pool; and

[a] Department of Electrical Engineering, UCLA. Email: xyz.xiao@gmail.com
[b] Department of Electrical Engineering, UCLA. Email: mihaela@ee.ucla.edu
[c] This work is supported by NSF CCF Grant No. 1218136.

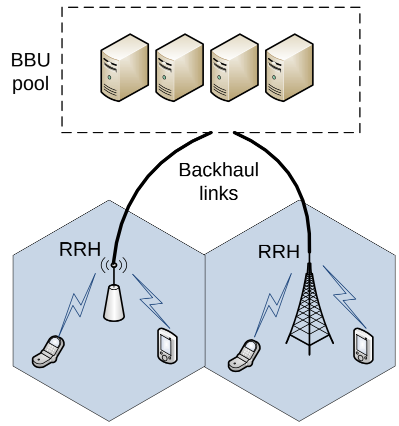

Figure 1.1 Illustration of C-RANs with a BBU pool and distributed RRHs.

- the second phase of distributed implementation by RRHs with very limited computational capability.

Requiring limited message exchange between the BBU pool and the RRHs, the presented methodology results in provably optimal spectrum sharing policies for the C-RANs in interference-limited scenarios. More importantly, the presented methodology is general, and can flexibly reconfigure the BBU pool to compute different optimal operating points in a variety of different CRAN deployment scenarios.

In this chapter, we start with a description of a general model for a C-RAN with the focus on spectrum sharing, and formulate the spectrum sharing design problem. We will present our model for delay-sensitive applications, and illustrate the implication of delay sensitivity on the structure of the optimal spectrum sharing policy. We will give a high-level overview of our design methodology, and discuss the instantiation of the methodology in various deployment scenarios. Finally, we will demonstrate the performance improvement achieved by our design methodology in comparison to state-of-the-art spectrum sharing policies.

1.2 A General Model of Spectrum Sharing in C-RANs

1.2.1 Basic Setup

Consider a C-RAN with a number of cells, which can be either macrocells or small cells (see Figure 1.1 for an illustration). At each time slot and at each frequency channel, there is one wireless device actively served by the base station. Depending on whether the downlink or the uplink is in our consideration, a RRH can be the base station or the active device of a cell. To be general, we will refer to the pair of base station and wireless device as RRH's transmitter and its receiver. Hence, in the downlink (resp. uplink), RRH's transmitter is the base station (resp. the wireless device) and its receiver is the device (resp. the base station).

The channel gain from RRH i's transmitter to RRH j's receiver is g_{ij}. Each RRH i chooses its transmit power level p_i from the set $\mathcal{P}_i \triangleq [0, P_i^{max}]$. Note that the set $\mathcal{P}_i$ contains 0, namely RRH i can choose not to transmit. Denote the joint power profile of all the RRHs by $\boldsymbol{p} = (p_1, \ldots, p_N)$.

Given the power profile, each RRH i obtains a *reward* $r_i(\boldsymbol{p})$. Each RRH i's reward can be any general function that is decreasing in the others power levels. Two representative examples of the reward function are as follows.

Example 1.1 One example of the reward function is the Shannon throughput. Since the RRHs are distributed and cannot decode each other's messages, each RRH treats the interference from the others as noise. Therefore, each RRH i's throughput is

$$r_i(\boldsymbol{p}) = \log_2\left(1 + \frac{p_i g_{ii}}{\sum_{j \neq i} p_j g_{ji} + \sigma_i}\right), \tag{1.1}$$

where σ_i is the noise power at RRH i's receiver.

Example 1.2 Another example of the reward function is the ratio of throughput to power level:

$$r_i(\boldsymbol{p}) = \frac{\log_2\left(1 + \frac{p_i g_{ii}}{\sum_{j \neq i} p_j g_{ji} + \sigma_i}\right)}{p_i}. \tag{1.2}$$

This reward function captures the energy efficiency.

Note that we do not require the RRH's reward to be increasing in its own power level, as in Example 1.2. In other words, we allow very general reward functions. For illustration, we adopt the throughput as the reward in the rest of this chapter.

We define RRH i's local interference temperature I_i as the aggregate interference and noise power level at its receiver, namely

$$I_i = \sum_{j \neq i} p_j g_{ji} + \sigma_i. \tag{1.3}$$

Each RRH i measures the interference temperature with errors. The erroneous estimate is $I_i + \varepsilon_i$, where ε_i is the additive estimation error. Each RRH i quantizes the estimate, and feeds the quantized estimate back to the transmitter. We require each RRH to simply use an unbiased estimator and a two-level quantizer.

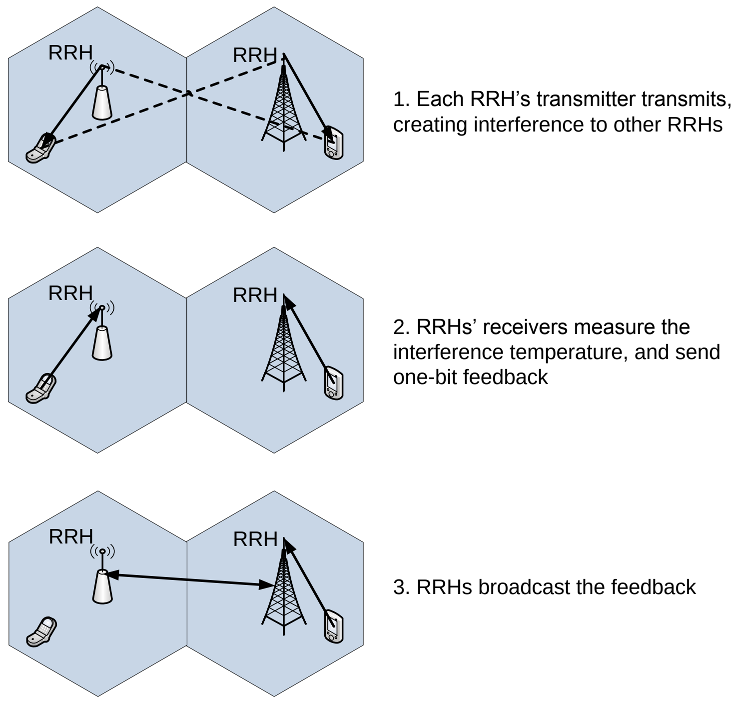

Figure 1.2 Illustration of the spectrum sharing protocol in each time slot.

This results in the following one-bit feedback signal

$$y_i = \begin{cases} 1, & \text{if } I_i + \varepsilon > \text{threshold} \\ 0, & \text{otherwise} \end{cases} . \tag{1.4}$$

1.2.2 Spectrum Sharing Policy

The system is time slotted at $t = 0, 1, 2, \ldots$. At the beginning of time slot t, each RRH i chooses its transmit power p_i^t, and achieves the throughput $r_i(\boldsymbol{p}^t)$. At the end of time slot t, each RRH i broadcasts its feedback signal y_i^t.[1] We define a system distress signal to indicate whether there exists a RRH whose local interference temperature is above the threshold. We denote the (system) distress signal by $y^t \in Y \triangleq \{0, 1\}$, where $y^t = 1$ if there exists a RRH i such that $y_i^t = 1$. See Figure 1.2 for an illustration of the above procedure in one time slot.

The spectrum sharing policy $\boldsymbol{\pi}$ is the collection of all the RRHs' policies, namely $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_N)$. Each RRH i's strategy is a mapping from the history

[1] We could reduce the amount of broadcasting by asking a RRH to broadcast only when its feedback signal is $y_i^t = 1$.

of past distress signals to its set of power levels, which is formally defined as[2]

$$\begin{aligned}\pi_i : \cup_{t=0}^{\infty}(Y)^t &\to P_i,\\ h^t \triangleq \left(y^1, \ldots, y^{t-1}\right) &\mapsto p_i^t,\end{aligned} \tag{1.5}$$

where h^t is the collection of past distress signals at time t.

Note that the above definition of spectrum sharing policies is general and can represent all the existing policies. For instance, it can represent spectrum sharing policies that dictate RRHs to transmit at fixed power levels (i.e., power control policies as in [20][21]). We call such policies "constant policies", and formally define them as follows.

DEFINITION 1.1 (Constant policies) In a constant spectrum sharing policy, each RRH transmits at a fixed power level all the time, namely

$$\pi_i(h^t) = p_i^{const}, \forall i, \forall t, \forall h^t. \tag{1.6}$$

Another class of existing policies of great interest are round-robin time-division multiple access (TDMA) policies [22]. We can use (1.5) to represent these policies as well.

Example 1.3 A simple round-robin policy with N RRHs can be written as

$$\pi_i(h^t) = \left\{ \begin{array}{ll} p_i^{const}, & \text{if } (t \bmod N) = i \\ 0, & \text{otherwise} \end{array} \right. . \tag{1.7}$$

In this case, the RRHs transmit in cycles of N, and each RRH i transmits in the ith time slot in each cycle. Note that the policy $\pi_i(h^t)$ depends only on the time slot t, not the history of distress signals h^t.

DEFINITION 1.2 (Round-robin policies) In a round-robin spectrum sharing policy, the RRHs transmit in cycles and choose the same power levels in each cycle independent of the history. For a round-robin policy with cycle length L, we have

$$\pi_i(h^t) = \pi_i(h^{t+L}), \forall i, \forall t, \forall h^t, \forall h^{t+L}. \tag{1.8}$$

Remark 1.1 Note that our definition of round-robin policies is more general than what people usually think of. Our definition includes Example 1.3 as a special case. It also extends the simple round-robin policy in Example 1.3 by allowing the cycle length to be different from the number of RRHs (i.e., $L \neq N$), allowing each RRH to have different numbers of transmission time slots in each cycle, and so on.

[2] Throughout this chapter, we use a^t to denote a at time t, and $(a)^t$ (resp. $(A)^t$) as the tth power (resp. power set) of real number a (resp. set A).

Remark 1.2 Despite the generality of our definition of round-robin policies, we will show later that the optimal spectrum sharing policy is not round-robin.

DEFINITION 1.3 (Non-stationary policies) Any policy defined in (1.5) that is neither constant nor round-robin is *non-stationary*.

Remark 1.3 Non-stationary policies are a class of general policies. In particular, some of them are TDMA policies without cyclic (or periodic) structures. These non-stationary TDMA policies are optimal for delay-sensitive applications, as we will illustrate in Section 1.3.

1.2.3 Delay Sensitivity

We model the delay sensitivity of the applications run by the RRHs by the discount factor $\delta \in [0, 1)$ [8]-[11][15][16]. Assuming a discount factor $\delta \in [0, 1)$, each RRH i's expected discounted average throughput is

$$R_i(\boldsymbol{\pi}) = \mathbb{E}\left\{(1-\delta)\cdot\sum_{t=0}^{\infty}(\delta)^t \cdot r_i\left[\pi\left(h^t\right)\right]\right\}, \tag{1.9}$$

where the expectation is taken over the history h^t with respect to its distribution induced by the policy and the (random) distress signal. Similarly, each RRH i's expected discounted average energy consumption is

$$P_i(\boldsymbol{\pi}) = \mathbb{E}\left\{(1-\delta)\cdot\sum_{t=0}^{\infty}(\delta)^t \cdot \pi\left(h^t\right)\right\}. \tag{1.10}$$

The discount factor models the delay sensitivity by discounting future rewards. A more delay-sensitive RRH discounts the future throughput more (i.e. has a smaller discount factor), because it has more urgency to transmit.

1.2.4 Problem Formulation

Each RRH i aims to maximize its own average throughput $R_i(\boldsymbol{\pi})$ or minimize its own average energy consumption $P_i(\boldsymbol{\pi})$, while fulfilling a minimum throughput requirement R_i^{min}. Our goal is to design a general methodology for constructing optimal spectrum sharing policies for C-RANs. The design problem can be formulated in a variety of forms, depending on requirements in specific deployment scenarios. One essential feature is that the policy should guarantee some minimum throughput requirements for all the RRHs, which will be introduced as constraints in all the formulations. The objective function can be some energy efficiency criterion $E\left(P_1(\boldsymbol{\pi}), \ldots, P_N(\boldsymbol{\pi})\right)$, (e.g., the weighted average of all the RRHs' energy consumptions, with each RRH's weight reflecting its importance), or some spectrum efficiency criterion $W\left(R_1(\boldsymbol{\pi}), \ldots, R_N(\boldsymbol{\pi})\right)$ (e.g., the weighted

low delay, but unfair for user 4:	1	2	3	**4**	1	2	3	**4**
fair, but high delay for user 1:	**1**	2	3	4	4	3	2	**1**

Figure 1.3 Two simple round-robin schedules with cycle length 8 for 4 RRHs. The first one has the same low delay of 4 for all 4 RRHs, but unfair sharing of transmission opportunities (TXOPs) (i.e., RRH 4 gets later TXOPs). The second one has a fair sharing of TXOPs, but incurs high maximum delay of 7 for RRH 1.

average throughput). Next, we define two policy design problems.

$$\begin{array}{ll} \texttt{MaxPayoff}: & \\ \max_{\boldsymbol{\pi}} & W\left(R_1(\boldsymbol{\pi}), \ldots, R_N(\boldsymbol{\pi})\right) \\ s.t. & R_i(\boldsymbol{\pi}) \geq R_i^{min},\ \forall i = 1, 2, \ldots, N. \end{array} \tag{1.11}$$

$$\begin{array}{ll} \texttt{MinCost}: & \\ \min_{\boldsymbol{\pi}} & E\left(R_1(\boldsymbol{\pi}), \ldots, R_N(\boldsymbol{\pi})\right) \\ s.t. & R_i(\boldsymbol{\pi}) \geq R_i^{min},\ \forall i = 1, 2, \ldots, N. \end{array} \tag{1.12}$$

1.3 The Optimal Spectrum Sharing Policy is Non-stationary

1.3.1 Intuitions

To better illustrate the structure of the optimal spectrum sharing policies, we focus on the case in which the RRHs have strong multi-user interference. The interference is so strong that it is optimal to let only one RRH to be active at each time slot, as in 802.11e MAC wireless networks [17]. Therefore, we will focus on TDMA policies.

All the existing TDMA policies are round-robin policies (including weighted round-robin policies) [17]–[19]. In round-robin policies, time slots are divided into cycles of a fixed predetermined length, and each RRH transmits in fixed predetermined positions within each cycle. The cyclic nature of round-robin policies simplifies the implementation, but imposes restrictions that render round-robin policies inefficient for delay-sensitive applications. The reasons are explained as follows.

For delay-sensitive application, not all the transmission opportunities (i.e. positions) in a cycle are created equal: the earlier transmission opportunities (TXOPs) are more desirable because they result in higher chances to deliver packets before their delay deadlines [17]–[19]. To ensure that the RRH's throughput and delay constraints are met, round-robin policies need a long cycle, and a careful sharing of TXOPs in a cycle.

First, a long cycle is necessary. Suppose that the cycle length is the shortest possible, namely equal to the number of RRHs (as in standard round-robin policies). Then the RRH allocated to the last TXOP suffers severely from delay. We can compensate this RRH for its delay by having a longer cycle and allocating some of the extra TXOPs to it. However, a long cycle results in an exponentially increasing (in the cycle length) number of possible policies to choose from.

Second, a careful sharing of TXOPs is necessary (see Figure 1.3 for an illustration of the following discussion). Suppose that the cycle length is twice the number of RRHs, and that each RRH gets two positions in a cycle. For fairness, no RRH should get two advantageous (i.e. earlier) TXOPs. A possible fair sharing may ensure that the RRH gets both an earlier and a later TXOPs. However, such a schedule is inefficient in worst-case delay: the RRH who gets the first and the last TXOPs in a cycle will experience high delay between consecutive transmissions. As we will illustrate in our motivating example (in the next subsection) and by simulations (in Section 1.6), round-robin policies cannot simultaneously achieve high system performance (e.g. max-min fairness) and fulfill the guarantees in terms of transmission delays required by the delay-sensitive RRHs.

In conclusion, the optimal spectrum sharing policy for delay-sensitive RRHs is non-stationary (i.e., not cyclic) in general.

1.3.2 An Illustrating Example

To further illustrate the performance improvement of optimal non-stationary policies over stationary policies, consider a spectrum sharing scenario in which the RRHs need to determine the transmission schedule and their own transmit power levels. Each RRH seeks to minimize its average energy consumption subject to its minimum throughput requirement. We have proved in [9] that the optimal policy (in the sense of minimizing average energy) has the property that only one RRH transmits at a given time (i.e. the policy is TDMA) and at a fixed power level whenever it transmits. We caution the reader that although the optimal policy is TDMA, it is not round-robin. For a specific numerical example, suppose that there are 3 RRHs all having direct channel gain of 1, cross channel gain of 0.25, noise power of 5 mW, and using the discount factor of 0.6 representing delay-sensitivity. RRHs have the same minimum throughput requirement of 1.5 bits/s/Hz.

We illustrate the policies and their performances in Table 1.1. The power levels are the transmit power levels of the 3 RRHs whenever they transmit. In the optimal constant policy, all 3 RRHs transmit all the time, at the same power level of 186mW.

In round-robin and the proposed TDMA policies, RRHs do not all transmit all the time. For round-robin policies, we compute the optimal policy given the cycle length by determining the optimal (in terms of average long-term energy consumption across RRHs) order of transmission in a cycle and the corresponding power levels. In the optimal round-robin policy with cycle of length 3, RRH

Table 1.1 Illustration of non-stationary vs. stationary policies and their performance.

Policies	Transmit power (mW)	Scheduling	Average energy (mW)
Optimal constant [12][13]	(186,186,186)	all the time	186
Optimal round-robin cycle length 3	(33,144,1432)	123 123 123 ...	108
Optimal round-robin cycle length 4	(43,212,249)	1234 1234 1234 ...	48
Optimal non-stationary (proposed)	(108,108,108)	123323213231...	36

1 transmits first at a low power level (33 mW), RRH 2 transmits after RRH 1 at a higher power level (144 mW) to compensate for having to wait for transmission and RRH 3 transmits last at an even higher power level (1432 mW) to compensate for having to wait even longer. In the cycle of length 4, again RRH 1 transmits at the lowest power level, RRH 2 transmits at a middle power level, and RRH 3 transmits at the highest power level, but the last two power levels are closer together (than in the cycle of length 3) because RRH 3 transmits more often.

In the optimal non-stationary policy, the RRHs all transmit at the same constant power level (108 mW) whenever they transmit; this works because the order in which they transmit is constantly changing. In the last column of Table 1.1, the discounted average energy per RRH per time slot is calculated. Notice that the cycle of length 3 is slightly more efficient than the constant policy, the cycle of length 4 is much more efficient, but the optimal non-stationary policy is more efficient still. Indeed, the optimal policy achieves 80%, 67% and 25% energy savings compared to the optimal constant policy, the optimal round-robin policy with cycle of 3 and with cycle of 4, respectively. Importantly, the energy savings are even more significant when the number of RRHs or the minimum throughput requirement is large (see Section 1.6).

Note that the optimal policy shown in the last row of Table 1.1 is obtained by our proposed distributed online algorithm. The RRHs will determine the schedule online (i.e., determine whether it should be active at the beginning of each time slot) with low complexity.

1.4 New Design Methodology for Spectrum Sharing Policies

Our general methodology can take a variety of forms in different deployment scenarios. Hence, the solutions are dependent on the considered scenarios. However, as illustrated in Figure 1.4, the general methodology has two common key components under all different scenarios: the optimal operating point selection (OOPS) algorithm that is run by the BBU pool to determine the optimal operat-

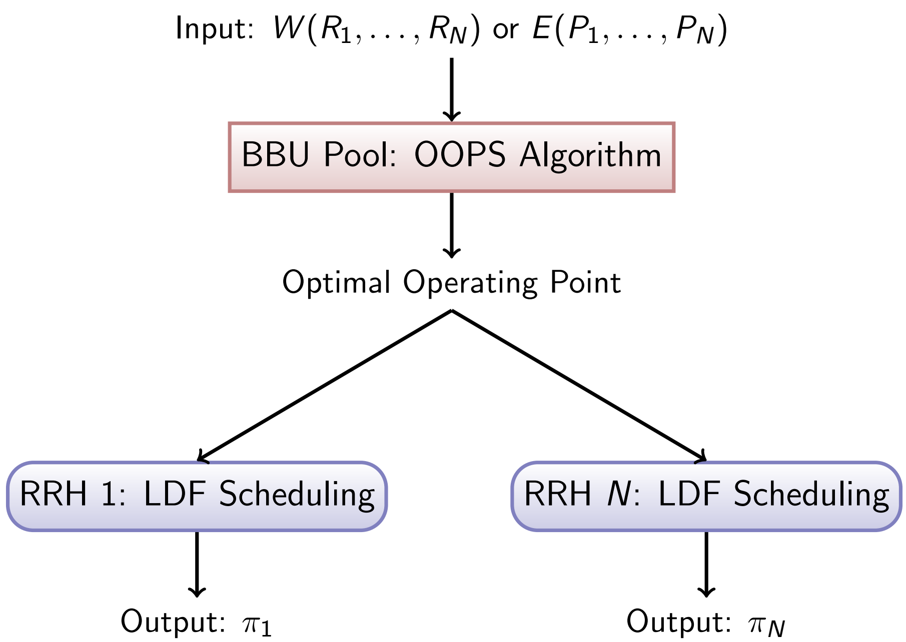

Figure 1.4 The design toolbox for spectrum sharing protocols in C-RANs. Given the performance criterion as the input, the BBU pool runs the optimal operating point selection (OOPS) algorithm, and sends the output to the distributed longest distance first (LDF) scheduling modules at each RRH.

ing point, and the longest distance first (LDF) scheduling algorithm that is run by distributed RRHs to construct the policy. In this section, we introduce the design framework in a baseline scenario [8][9], where the presented methodology is provably optimal and is simple enough for a good understanding of its essence. In this way, we can get the intuition behind the design framework, which can be applied to a variety of other scenarios.

We illustrate the new design methodology for spectrum sharing protocols in CRANs in Figure 1.4. The design toolbox takes as input the performance criterion $E(P_1, , P_N)$ or $W(R_1, , R_N)$ selected by the C-RAN. For example, when the criterion is the weighted sum of energy consumption, the input will be the weights assigned by the designer to each RRH based on its importance. A C-RAN protocol designer can input any desirable performance criterion to the design toolbox, which will then output the optimal spectrum sharing protocol. The design toolbox provides two modules:

- the OOPS (optimal operating point selection) algorithm run by the BBU pool to determine the instantaneous throughput when a RRH accesses the spectrum, as well as the optimal operating point (i.e., the average throughput of each RRH); and

Distributed protocol implementation modules at each user
Optimal operating point selection algorithm:
Input: design criterion, initial “price” (dual variable)
Repeat until desired accuracy reached
Based on “price”, find optimal operating point by Newton’s method
Broadcast the optimal operating point
Update the price using the bisection method
Output: the optimal operating point
LDF (longest distance first) scheduling:
Input: optimal operating point
Calculates “distances from targets” of each user
The user with the largest distance is active
Updates the distances analytically based on the feedback signal
Output: the optimal deviation-proof scheduling

Figure 1.5 The two modules deployed at the BBU pool and at each RRH to determine the optimal spectrum sharing policy.

- the LDF (longest distance first) scheduling run by each RRH to determine whether it should access the spectrum at each time slot.

We give a brief description of the two modules in Figure 1.5.

We briefly discuss the intuition behind the LDF scheduling algorithm. According to our definition of non-stationary policies, the scheduling decision should be made based on the history of distress signals. The central key to and most difficult part of our construction is to prove that it is enough to summarize the history up to each time slot by a particular metric (see [8][9] for the analytical expression of the metric). This metric can be easily computed by each RRH in a completely distributed way, and has a nice interpretation of the “distance from target throughput”. The scheduling decision is then based simply on the metric: the RRH “farthest away” from the target throughput transmits. The way the RRHs update the metric makes sure that the resulting scheduling can indeed achieve the target throughput.

When equipped with these two modules, the RRH can reach the optimal spectrum sharing policy in a distributed manner. As proved in [8][9], the operations performed by both modules converge in logarithmic time to the desired operating point. Importantly, we prove theoretically that the dynamic entry and exit of devices will not affect the convergence speed of the spectrum sharing policy [9].

Hence, the design toolbox is very suitable for CRANs with frequent switch-on and switch-off of RRHs.

In the next few sections, we will present some instantiations of our design methodology in several realistic C-RAN deployment scenarios.

1.5 Applications to Realistic C-RAN Deployment Scenarios

1.5.1 Large-Scale Heterogeneous Small Cell Networks

We consider the first representative deployment scenario of very large-scale heterogeneous small cell networks [10][11]. The unique features of large-scale heterogeneous small cell networks impose the following requirements for efficient spectrum sharing:

- *Deployment of heterogeneous small cell networks:* Existing deployments of small cell networks exhibit significant heterogeneity such as different types of small cells (picocells and femtocells), different cell sizes, different throughput requirements for RRH, etc.
- *Interference avoidance and spatial reuse:* Effective interference management policies should take into account the strong interference among neighboring RRHs, as well as the weak interference among non-neighboring RRHs. Hence, the policies should effectively avoid interference among neighboring RRHs and use spatial reuse to take advantage of the weak interference among non-neighboring RRHs.
- *Scalability to large networks:* Small cells are often deployed over a large scale (e.g., in a city). Effective interference management policies should scale in large networks, namely achieve efficient network performance while maintaining low computational complexity.

In large-scale heterogeneous small cell networks, the design methodology achieves the following:

- A spectrum sharing policy that schedules maximal independent sets (MISs)[3] of the interference graph to transmit in each time slot. In this way, we can avoid strong interference among neighboring RRHs (since neighboring RRHs cannot be in the same MIS), and efficiently exploit the weak interference among RRHs in a MIS by letting them to transmit at the same time.
- A distributed algorithm for the RRHs to determine a subset of MISs. The subset of MISs generated ensures that each RRH belongs to at least one MIS in this subset. Moreover, the subset of MISs can be generated in a

[3] Consider the interference graph of the network, where each vertex is a pair of RRH and its user, and each edge indicates strong interference between the two vertices. An independent set (IS) is a set of vertices in which no pair is connected by an edge. An IS is a MIS if it is not a proper subset of another IS.

distributed manner in logarithmic time (logarithmic in the number of RRHs in the network) for bounded-degree interference graphs. [4] The logarithmic convergence time is significantly faster than the time (linear or quadratic in the number of RRHs) required by existing distributed algorithms for generating subsets of MISs.

- A distributed algorithm for each RRH to determine the optimal fractions of time occupied by the MISs with only local message exchange. The message is exchanged only among neighboring RRHs. The distributed algorithm will output the optimal fractions of time for each MIS such that the given network performance criterion is maximized subject to the minimum throughput requirements.

More importantly, under a wide range of conditions, we analytically characterize the competitive ratio of the proposed distributed policy with respect to the optimal network performance. We prove that the competitive ratio is independent of the network size, which demonstrates the scalability of our proposed policy in large networks. Remarkably, the constant competitive ratio is achieved even though our proposed policy requires only local information, is distributed, and can be computed fast, while the optimal network performance can only be obtained in a centralized manner with global information and NP (non-deterministic polynomial time) complexity.

Through simulations, we demonstrate significant (from 160% to 700%) performance gains over state-of-the-art policies.

1.5.2 C-RANs with Multimedia Applications

We consider the second representative deployment scenario of C-RANs with delay-sensitive multimedia applications [12]–[14]. In this deployment scenario, it is important to provide hard delay guarantees.

Based on the new design methodology, we define a novel quality-of-service (QoS) metric, called continuing QoS (CQoS) guarantees, as follows [14]

$$\textbf{CQoS: } R_i^t(\boldsymbol{\pi}) \geq \gamma_i^{cont} \cdot r_i^{max},\ \forall t = 0, 1, \ldots, \tag{1.13}$$

where $r_i^{max} = \log_2\left(1 + \frac{P_i^{max}}{\sigma_i}\right)$ is the maximum achievable throughput.

CQoS guarantees require a RRH's average throughput starting from *every point in time* to be higher than a threshold. CQoS guarantees are stricter requirements than the conventional QoS guarantees which only guarantee the average throughput starting from the beginning.

A byproduct of the CQoS guarantees is that once they are satisfied, we can also provide upper bounds on the transmission delays of each RRH. First, we define RRH i's transmission delay at any time t as

$$\textbf{Transmission Delay: } d_i^t(\boldsymbol{\pi}) \triangleq \min_{\tau > t} \left\{\tau - t : \pi_i(\tau) > 0\right\}.$$

[4] Bounded-degree graphs are the graphs whose maximum degree can be bounded by a constant independent of the size of the graph.

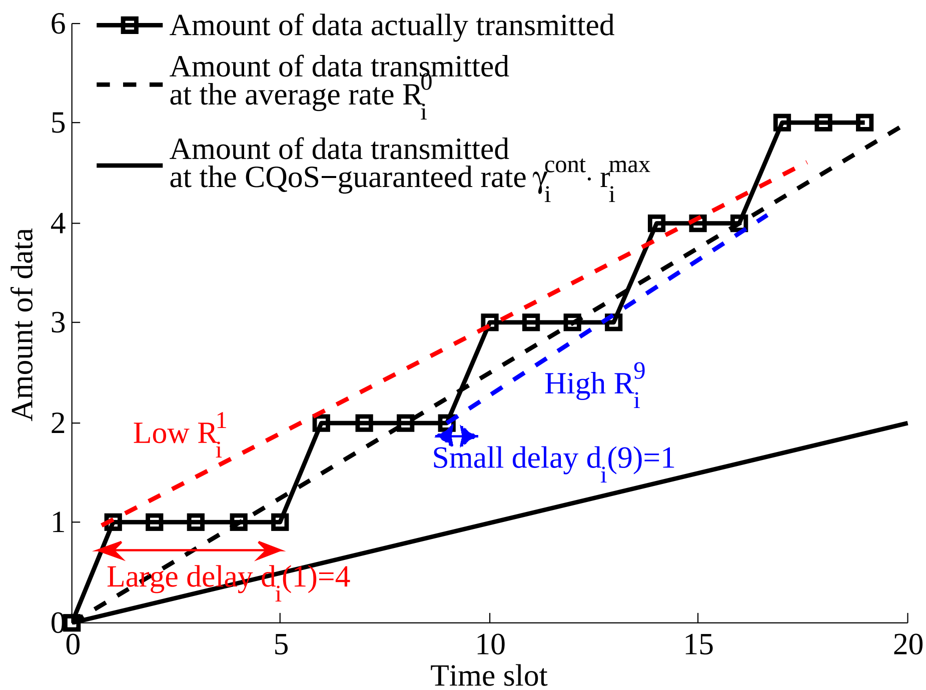

Figure 1.6 Relationship of delay and CQoS guarantees of RRH i. The solid curve with square data points is the amount of data transmitted; each jump in the curve corresponds to a transmission. The two straight lines through the origin are the amount of data transmitted as if the throughput was R_i^0 and $\gamma_i^{cont} \cdot r_i^{max}$, respectively. At each time t, if the continuation throughput R_i^t is higher, the RRH needs to transmit more after time t. Hence, the corresponding delay $d_i(t)$ is lower.

In words, the transmission delay $d_i^t(\boldsymbol{\pi})$ is the minimum wait time until the next transmission. An upper bound on the transmission delays are critical for delay-sensitive applications.

We have proved in [14] that each RRH's CQoS guarantee leads to an upper bound on its maximum delay $\sup_t d_i^t(\boldsymbol{\pi})$. Figure 1.6 illustrate the relationship of delay and CQoS guarantees.

We propose a systematic design methodology, which constructs the optimal TDMA policy that maximizes the system performance (e.g. fairness) subject to the RRHs' CQoS guarantees.

Again, the key feature of the proposed policy is that it is not cyclic as in round-robin policies. Instead, it adaptively determines which RRH should transmit according to the RRHs' remaining amounts of TXOPs needed to achieve the target throughput. We propose a low-complexity distributed algorithm to construct the optimal policy. Simulation results show that our proposed policy significantly outperforms the optimal constant policy and round-robin policies

Table 1.2 Energy efficiency of different policies measured by average energy expenditure (mW). "N/A" means that the policy fails to satisfy the minimum throughput requirements.

Number of cells	7	9	11	13	15
Stationary	20	N/A	N/A	N/A	N/A
Round-robin	19	40	81	13	1320
Proposed	15	28	37	31	106
Energy saving w.r.t. stationary	25%	N/A	N/A	N/A	N/A
Energy saving w.r.t. round-robin	21%	30%	54%	77%	92%

Table 1.3 Spectrum efficiency of different policies measured by average throughput (bits/s/Hz). "N/A" means that the policy fails to satisfy the minimum throughput requirements.

Number of cells	7	9	11	13	15
Stationary	0.9	N/A	N/A	N/A	N/A
Round-robin	2.0	1.8	1.3	1.1	N/A
Proposed	2.7	2.4	2.0	1.7	1.4
Improvement w.r.t. stationary	200%	∞	∞	∞	∞
Improvement w.r.t. round-robin	35%	33%	54%	55%	∞

in peak signal-to-noise ratio (PSNR) for video streaming, by up to 6 dB and 4 dB, respectively.

1.6 Performance Gains

To illustrate the performance gain over existing policies, we consider a network of small cells (e.g. femtocells, picocells). Each small cell serves one device at each time. We randomly place small cell base stations in a 2-dimensional space with an average inter-site distance of 20 m, and randomly place devices with an average distance to their base stations of 5 m. The path loss exponent is 2. The maximum received SNR $\frac{P_i}{\sigma_i}$ is 20 dB. The energy efficiency criterion is the average energy consumption, and the spectrum efficiency criterion is the average throughput. The discount factor is 0.95. The minimum throughput is $R_i^{min} = 1$ bits/s/Hz for all i. All data is the average over the results obtained from 10000 random placements of small cells and devices.

In Tables 1.2 and 1.3, we compare the proposed policies against stationary policies (in which RRHs transmit at fixed powers simultaneously) [12][13] and round-robin TDMA policies [14] in terms of energy efficiency and spectrum efficiency. The minimum throughput requirements and the weights are the same for all the RRHs. The proposed policy significantly improves the spectrum and energy efficiency of existing policies in most scenarios. In particular, when the number of RRHs is large, stationary policies quickly become infeasible (i.e. fail to achieve the minimum throughput requirements), while the proposed policy is feasible. Compared to existing policies, the proposed policy can achieve up to 92% energy saving, and up to 200% spectrum efficiency.

1.7 Related Works

The methodology presented in this chapter is based on two key insights [8]-[15]:

- it is more efficient to access the spectrum in a time-division multiple access (TDMA) fashion, rather than to access the spectrum at the same time (due to interference among RRHs), and
- the optimal way to access the spectrum is not (weighted) round-robin but rather follows a carefully designed non-stationary schedule in which each RRH's transmit power level depends not only on its current state but also on the history of previous states and power levels.

1.7.1 Related Works in Spectrum Sharing

In contrast with these protocols developed based on the new general methodology presented in this chapter, state-of-the-art spectrum sharing protocols [18]-[22] use stationary policies, in which a RRHs transmit power level depends only on its current state. For example, some works [20][21] propose physical-layer power control policies that require the RRHs to transmit simultaneously at fixed power levels over the time horizon in which they interact . Due to strong multi-user interference, stationary power control policies can only achieve low spectrum efficiency and low energy efficiency.

Some works [22] propose stationary medium access control (MAC) layer centric solutions, by modeling the physical layer with collision models and neglecting power control. These solutions adopt contention-free round-robin TDMA [22] protocols, which are suboptimal. The performance loss as compared to the optimal policies is even larger when the RRHs are heterogeneous and experience different channel conditions, have different throughput requirements and demands etc.

Table 1.4 Comparison of the proposed methodology with network utility maximization (NUM), single-user Markov decision process (SU-MDP) and multi-user Markov decision process (MU-MDP), as well as the implications in spectrum sharing scenarios.

	Coupling	Interference regime	Resulting policy
NUM	weak	weak or no interference	stationary
SU-MDP	N/A	N/A	stationary
MU-MDP	weak	no interference	stationary
Proposed	strong	strong interference	non-stationary

Table 1.5 Detailed comparison with Markov decision process (MDP).

	Agents	Action	Value function	Policy
Single-user MDP	single	single action	single-valued	stationary
Multi-user MDP	multiple	action profile	single-valued	stationary
Non-stationary policy design	multiple	action profile	set-valued	non-stationary

1.7.2 Related Theoretical Frameworks

Note that existing theoretical frameworks, such as network utility maximization (NUM) and standard Markov decision process (MDP), are not suitable for designing spectrum sharing policies. This is because they focus on the scenarios with "weakly-coupled" RRHs, namely one RRH's action (e.g., transmit power) does not affect the others' payoffs (e.g., throughput). In spectrum sharing in C-RANs, the RRHs are strongly coupled, since one's transmission may cause strong interference to other RRHs. In addition, standard MDP is mostly used for single-user decision problems, and is often suboptimal when applied to multi-user decision problems. We summarize the key differences from NUM and the MDP theory in Table 1.4. In addition, since both our methodology and MDP result in dynamic spectrum sharing policies, we give a detailed comparison with MDP in Table 1.5.

1.8 Conclusion

In this chapter, we introduce a novel general methodology for designing provably optimal spectrum sharing policies that are particularly suitable for C-RANs. In the protocols designed using this methodology, computationally demanding operations (i.e., determining the optimal operating points) are implemented by the BBU pool, and are separated from simple operations implemented by distributed RRHs with limited computational capability. Moreover, the protocols require limited message exchange between the BBU pool and the RRHs, reducing the burden of the backhaul. The protocols achieve high overall spectrum and energy efficiency, and provide performance guarantees for each individual RRH.

The presented methodology is general and can flexibly configure the cloud to optimize the system performance in different CRAN deployment scenarios, such as large-scale heterogeneous small cell networks, delay-sensitive multimedia communications, and Internet of things. Initial experiment in specific deployment scenarios shows that compared to existing protocols, the proposed protocols can significantly improve the spectrum and energy efficiency.

The design methodology can also be extended to domains other than spectrum sharing. Interested readers are referred to [16] for a treatment in general resource sharing games, to [23] for an application to demand side management in smart grids [23]

References

[1] M. Peng, Y. Li, J. Jiang, J. Li, and C. Wang, "Heterogeneous cloud radio access networks: A new perspective for enhancing spectral and energy efficiencies," *IEEE Wireless Commun. Mag.*, vol. 21, no. 6, pp. 126–135, Dec. 2014.

[2] Z. Ding and H. V. Poor, "The use of spatially random base stations in cloud radio access networks," *IEEE Signal Process. Lett.*, vol. 20, no. 11, pp. 1138–1141, Nov. 2013.

[3] China Mobile. (2011). C-RAN: The road towards green RAN. [White Paper]. Available: http://bit.ly/Ya1zuW

[4] H. Harada, "Cognitive wireless cloud: A network concept to handle heterogeneous and spectrum sharing type radio access networks," in *Proc. IEEE Personal, Indoor and Mobile Radio Commun. (PIMRC)*, 2009, pp. 1–5.

[5] H. T. Dinh, C. Lee, D. Niyato, and P. Wang, "A survey of mobile cloud computing: Architecture, applications, and approaches," *Wireless Communications and Mobile Computing*, vol. 13, no. 18, pp. 1587–1611, Dec. 2013.

[6] J. Zander and P. Mhnen, "Riding the data tsunami in the cloud: Myths and challenges in future wireless access," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 145–151, Mar. 2013.

[7] K. Sundaresan, M. Y. Arslan, S. Singh, S. Rangarajan, and S. V. Krishnamurthy, "FluidNet: A flexible cloud-based radio access network for small cells," in *Proc. 19th Int. Conf. on Mobile Computing & Networking (MobiCom)*, 2013, pp. 99–110.

[8] Y. Xiao and M. van der Schaar, "Dynamic spectrum sharing among repeatedly interacting selfish devices with imperfect monitoring," *IEEE J. Sel. Areas Commun., Special issue on cognitive radio*, vol. 30, no. 10, pp. 1890–1899, 2012.

[9] Y. Xiao and M. van der Schaar, "Energy-efficient nonstationary spectrum sharing," *IEEE Trans. Commun.*, vol. 62, no. 3, pp. 810–821, March 2014.

[10] K. Ahuja, Y. Xiao, and M. van der Schaar, "Efficient interference management policies for femtocell networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 9, pp. 4879–4893, Sep. 2015.

[11] K. Ahuja, Y. Xiao, and M. van der Schaar, "Distributed interference management policies for heterogeneous small cell networks," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 6, pp. 1112–1126, Jun. 2015.

[12] J. Xu, Y. Andreopoulos, Y. Xiao, and M. van der Schaar, "Non-stationary resource allocation policies for delay-constrained video streaming: Application to video over Internet-of-Things-enabled networks," *IEEE J. Sel. Areas Commun., Special issue on adaptive media streaming*, vol. 32, no. 4, pp. 782–794, Apr. 2014.

[13] Y. Xiao, K. Ahuja, and M. van der Schaar, "Spectrum sharing for delay-sensitive applications with continuing QoS guarantees," in *Proc. IEEE Global Commun. Conf. (Globecom)*, 2014, pp. 1265–1270.

[14] Y. Xiao and M. van der Schaar, "Spectrum sharing policies for heterogeneous delay-sensitive users: A novel design framework," in *Proc. Allerton Conf. Commun., Control, and Comput.*, 2013, pp. 85–92.

[15] Y. Xiao and M. van der Schaar, "Optimal foresighted multi-user wireless video," *IEEE J. Sel. Topics Signal Process., Special issue on Visual Signal Processing for Wireless Networks*, vol. 9, no. 1, pp. 89–101, Feb. 2015.

[16] M. van der Schaar, Y. Xiao, and W. Zame, "Efficient outcomes in repeated games with limited monitoring and impatient players," *Economic Theory*, vol. 60, no. 1, pp. 1–34, Sep. 2015.

[17] M. van der Schaar, Y. Andreopoulos, and Z. Hu, "Optimized scalable video streaming over IEEE 802.11 a/e HCCA wireless networks under delay constraints," *IEEE Trans. Mobile Comput.*, vol. 5, no. 6, pp. 755–768, Jun. 2006.

[18] D. Pradas and M. A. Vazquez-Castro, "NUM-based fair rate-delay balancing for layered video multicasting over adaptive satellite networks," *IEEE J. Sel. Areas in Commun.*, vol. 29, no. 5, May 2011.

[19] P. Dutta, A. Seetharam, V. Arya, M. Chetlur, S. Kalyanaraman, and J. Kurose, "On managing quality of experience of multiple video streams in wireless networks," in *Proc. IEEE Infocom*, 2012.

[20] R. D. Yates, "A framework for uplink power control in cellular radio systems," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 7, pp. 1341–1347, Sep. 1995.

[21] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 3, pp. 517–528, Apr. 2007.

[22] M. Fang, D. Malone, K. R. Duffy and D. J. Keith, "Decentralised learning MACs for collision-free access in WLANS," *Wireless Networks*, vol. 19, no. 1, pp. 83–98, Jan. 2013.

[23] L. Song, Y. Xiao, and M. van der Schaar, "Demand side management in smart grids using a repeated game rramework," *IEEE J. Sel. Areas in Commun., Special Issue on Smart Grid Communications*, vol. 32, no. 7, pp. 1412–1424, Jun. 2014.