Fifty-first Annual Allerton Conference
Allerton House, UIUC, Illinois, USA
October 2 - 3, 2013

Spectrum Sharing Policies for Heterogeneous Delay-Sensitive Users: A Novel Design Framework

Yuanzhang Xiao and Mihaela van der Schaar[1]

***Abstract*— We develop a novel design framework for spectrum sharing among distributed users with heterogeneous delay-sensitivity (e.g. users with video streaming that requires low delay, and users with video conferencing that requires very low delay). Most existing spectrum sharing policies are *stationary*, i.e. users transmit at constant power levels simultaneously. Under stationary policies, the users have low throughput due to the strong interference from each other. *Nonstationary* spectrum sharing policies, which allow users to transmit at time-varying power levels, can significantly improve the spectrum efficiency. The most well-known and simple nonstationary policy is the round-robin TDMA (time-division multiple access) policy, in which the users access the spectrum in turn. Although the round-robin TDMA policy increases the spectrum efficiency by eliminating multi-user interference, it is suboptimal in terms of quality of experience for delay-sensitive users, especially when they have heterogeneous delay-sensitivity. This is because the round-robin TDMA policy allocates the users' transmission opportunities in a predetermined order such that they have (roughly) the same amount of transmission opportunities in any duration of time. However, some users may have earlier deadlines and need more transmission opportunities early on, while some can wait until later. This heterogeneity in delay-sensitivity is not considered in the round-robin TDMA policy.**

In this paper, we propose nonstationary policies that allocate the transmission opportunities based on the users' delay-sensitivity and their past deadline-abiding transmissions. As we will see, the optimal policy is not cyclic at all as is the round-robin TDMA policy. We also propose a low-complexity algorithm, which can be run by each user in a distributed manner, to construct the optimal nonstationary policy. Simulation results validate our analytical results and quantify the performance gains enabled by the proposed policies.

I. INTRODUCTION

A plethora of bandwidth-intensive and delay-sensitive applications, such as multimedia streaming, video conferencing, and gaming, have emerged and are now increasingly deployed over wireless networks. The proliferation of such applications imposes huge challenges in the design of wireless networks, in which the users share the common spectrum and cause interference to each other. Hence, when designing spectrum sharing policies, it is crucial to consider both spectrum efficiency and delay sensitivity.

Most existing spectrum sharing policies are *stationary*, namely they require the users to transmit at *constant* power levels over the time horizon in which they interact[1] [1]–[8]. Stationary policies are inefficient in many spectrum sharing scenarios where the interference among the users is strong. Under strong multi-user interference, increasing one user's power level significantly degrades the other users' throughput, which results in a low spectrum efficiency. In the extreme cases modeled by the collision medium access control (MAC) model, simultaneous transmission leads to packet loss where no user receives anything.

One way to improve spectrum efficiency is to use TDMA (time-division multiple access) policies, where only one user transmits at each time slot. TDMA policies are *nonstationary*, in the sense that the users transmit at *time-varying* power levels. The most commonly-used nonstationary policy is the simple round-robin TDMA policy, where the users transmit in turn in a predetermined order. The key feature of round-robin TDMA policies is that the users have roughly the same amount of transmission opportunities up to any point in time. In other words, round-robin policies do not optimize the positions at which each user's transmission opportunities are allocated based on their delay deadlines. If some users have earlier deadlines while others have later deadlines, round-robin policies will still allocate the same amount of transmission opportunities for all the users. The consequence is that the users with earlier deadlines cannot transmit all their packets before their deadlines, because the transmission opportunities are occupied by the users with later deadlines, who can actually delay their transmissions to make room for more delay-sensitive users.

In this paper, we propose a design framework of nonstationary spectrum sharing policies for delay-sensitive users. The proposed nonstationary policy allocates the transmission opportunities based on the users' delay-sensitivity, as well as their past delay-abiding transmissions. The resulting transmission scheduling is not cyclic as in round-robin policies. Instead, it adaptively determines which user should transmit according to the users' remaining amounts of transmission opportunities before the deadlines (which can be inferred from their past delay-abiding transmissions). We propose a low-complexity algorithm, which can be run by each user in a distributed manner, to construct the optimal policy. Simulation results show that the proposed policy can achieve performance improvement in orders of magnitudes compared to stationary policies [1]–[8], and achieve up to 100% performance gain over the round-robin TDMA policy.

Note that in our previous works [13]–[14], we proposed design frameworks of nonstationary spectrum sharing policies for users with *homogeneous* delay-sensitivity. We model the user's delay-sensitivity using the discount factor, namely the rate at which a user discounts future transmission opportunities. Specifically, each user discounts the future

[1]The authors are with the Electrical Engineering Department, UCLA, Los Angeles, CA 90095, USA. Email: {yxiao,mihaela}@ee.ucla.edu

[1]Although some spectrum sharing policies go through a transient period of adjusting the power levels before the convergence to the optimal power levels, the users maintain constant power levels after the convergence.

978-1-4799-3410-2/13/$31.00 ©2013 IEEE

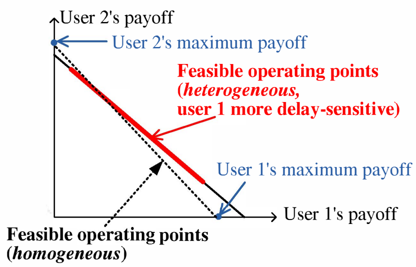

Fig. 1. Feasible operating points when the users have homogeneous delay-sensitivity and when they have heterogeneous delay-sensitivity (with user 1 being more delay-sensitive).

TABLE I
COMPARISON WITH RELATED WORKS.

	Spectrum sharing	Delay sensitivity
[1]–[8]	Stationary	None
[9]–[12]	Stationary	Homogeneous
[13]–[14]	Nonstationary	Homogeneous
Proposed	Nonstationary	Heterogenous

transmissions by a geometric series. The discount factor is the ratio of successive terms in the geometric series. A more delay-sensitive user has a smaller discount factor, because it has more urgency to transmit early and discounts the future transmissions more. In [13]–[14], we assume that the users have the *same delay-sensitivity* (i.e. the same discount factor). Hence, the design frameworks in [13]–[14] are very different from the design framework in this paper, which is aimed for users with *heterogeneous* delay sensitivity.

To illustrate the differences from the design frameworks for users with homogeneous delay-sensitivity, we show the feasible operating points when the users have homogeneous and heterogeneous delay-sensitivity in Fig. 1. When the users have the same delay-sensitivity, the feasible operating points, achievable by time sharing in TDMA policies, is simply the line connecting their maximum payoffs. When they have different delay-sensitivity (assuming user 1 to be more delay-sensitive), TDMA policies will result in feasible operating points lying in another line. More specifically, the intersection of the line with the x-axis is larger than user 1's maximum payoff, and the intersection with the y-axis is smaller than user 2's maximum payoff. Of course, the feasible operating points cannot be the entire line, since user 1's payoff cannot exceed its maximum payoff. The direction of the line is different from the case with homogeneous delay-sensitivity, because when user 2's payoff increases due to more transmission opportunities, the more delay-sensitive user 1 suffers from a larger decrease in its payoff, compared to the case when it has the same delay-sensitivity as user 2. Hence, the case with heterogeneous delay-sensitivity is more challenging to analyze, because we need to determine the direction of the line, and which portion of the line is feasible.

Finally, we summarize the comparison of our work with the existing works in spectrum sharing in Table I. We distinguish our work from existing works in the following categories: the spectrum sharing policy is stationary or non-stationary, whether the users are delay-sensitive or not, and if they are, whether they have homogeneous or heterogeneous delay sensitivity.

The rest of the paper is organized as follows. In Section II, we describe the system model and formulate the policy design problem. Then in Section III, we motivate our proposed policy by showing the inefficiency of round-robin policies in a simple example. We solve the policy design problem in Section IV. Simulation results are presented in Section V. Finally, Section VI concludes the paper.

II. SYSTEM MODEL AND PROBLEM FORMULATION

A. System Model

1) Network model: We consider a wireless network with N users (see Fig 2 for an illustrating example of a system with two users). The set of users is denoted by $\mathcal{N} \triangleq \{1, 2, \ldots, N\}$. Each user has a transmitter and a receiver. The channel gain from user i's transmitter to user j's receiver is g_{ij}. Each user i chooses a power level p_i from a compact set $\hat{\mathcal{P}}_i$. We assume that $0 \in \hat{\mathcal{P}}_i$, namely user i can choose not to transmit. We also assume that the users need to comply with some interference temperature constraints (ITCs) measured at K locations in the network. Depending on different scenarios, the ITCs can be imposed by primary users in a cognitive radio network or the base station in a femtocell network. The channel gain from user i's transmitter to the kth location is g_{i0_k}. Each user i should know the channel gain $\{g_{i0_k}\}_{k=1}^K$ to each measurement location and the interference temperature limit $\{I_k\}_{k=1}^K$ at each location. Hence, each user i's set of *admissible* power levels is

$$\mathcal{P}_i = \{p_i \in \hat{\mathcal{P}}_i : g_{i0_k} \cdot p_i \leq I_k, \forall k = 1, \ldots, K\}. \quad (1)$$

For convenience, we also define user i's maximum admissible power level as $P_i^{\max} \triangleq \max_{p_i \in \mathcal{P}_i} p_i$. Now we can already see one advantage of using TDMA spectrum sharing policies: each user only needs to know the channel gains from its own transmitter to the measurement locations in order to meet the ITCs. In stationary spectrum sharing policies, since users transmit simultaneously, they need to know additional information of channel gains from the other users' transmitters to the measurement locations [1]–[5].

Our system model is general enough to model many wireless communication networks. It can model wireless ad hoc networks where N users transmit in the unlicensed spectrum (e.g. the 2.4 GHz frequency band) without ITCs (K=0). It can also model the uplink (the receivers are co-located) and the downlink (the transmitters are co-located) of a cellular network with possible ITCs imposed by base stations in nearby cells. It can also model cognitive radio networks with N secondary users sharing the spectrum with K primary users, each of which imposes an ITC at its receiver. Similarly, it can model femtocell networks with N femtocells sharing the spectrum with $K = 1$ base station, who imposes an ITC for the femtocells.

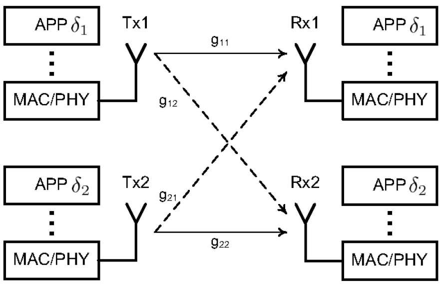

Fig. 2. An example system model with two users. Each user is a pair of transmitter (Tx) and receiver (Rx). Each user has an application (APP) layer (parameterized by the delay sensitivity of the application δ_i) and a medium-access-control/physical (MAC/PHY) layer.

Similar to [1]–[8], we assume that the system parameters, such as the number of users and the channel gains, remain fixed during the considered time horizon.

2) The MAC/PHY layer: We describe the users' medium-access-control/physical (MAC/PHY) layers. We denote the joint power profile of all the users by $\boldsymbol{p} = (p_1, \ldots, p_N)$. Assuming that the users cannot jointly decode their messages and that they treat other users' interference as noise, each user i's instantaneous throughput under the joint power profile $\boldsymbol{p}$ is

$$r_i(\boldsymbol{p}) = \log_2\left(1 + \frac{p_i g_{ii}}{\sum_{j \in \mathcal{N}, j \neq i} p_j g_{ji} + \sigma_i}\right), \tag{2}$$

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

The system is time slotted at $t = 0, 1, \ldots$. We assume that the users are synchronized as in [1]–[8] (e.g. by using a global clock from the global positioning system (GPS)). Since we focus on TDMA policies, we can write the transmission schedule as $\pi : \mathbb{N}_+ \rightarrow \mathcal{N}$, where $\pi(t)$ is the index of the user who should transmit at time t. The simple round-robin TDMA policy is then $\pi(t) = (t \mod N)$.

The spectrum sharing policy specifies the transmission schedule and the users' transmit power levels when they transmit. To maximize its own throughput, each user i will transmit at the maximum admissible power level $P_i^{\max}$. Hence, we only need to design the transmission schedule.

3) The application layer: There exists two different ways to model the application layer. One way is to use a detailed model that specifies the delay deadline and the distortion impact[2] of each packet (see [15] for an example). This is the most accurate model for the application layer. However, it is difficult to use this accurate model in practice. This is because there are too many parameters in this model (e.g. a delay deadline and a distortion impact for each packet). In order to coordinate, the users need to exchange these parameters with each other, which imposes a huge overhead of information exchange. Even in the centralized scenario, it is hard to determine the transmission schedule based on this application layer model. Suppose that the users transmit to the base station in an uplink. The users need to inform the base station of all the parameters before the base station determines the transmission schedule, which is impractical.

Another way to model the application layer is to use a single parameter, namely the discount factor, to represent the delay sensitivity of the application [9]–[14]. In this model, each user i's application layer is abstracted by a discount factor $\delta_i \in [0, 1)$. User i will discount the future throughput by δ_i. A more delay-sensitive user will discount the future throughput more (i.e. have a smaller discount factor), because it has more urgency to transmit now. This abstract model is less accurate, but is amenable to practical implementation, because of the low overhead in exchanging the information on the users' delay sensitivity.

Without loss of generality, we assume that the users with smaller indices is more delay-sensitive, namely $\delta_1 \leq \delta_2 \leq \ldots \leq \delta_N$. Given the transmission schedule, we can determine user i's discounted average throughput as

$$R_i(\pi) = (1 - \delta_i) \sum_{t=0}^{\infty} \mathbf{1}_{\pi(t)=i} \cdot \delta_i^t \cdot r_i(P_i^{\max}, \boldsymbol{p}_{-i} = \mathbf{0}), \tag{3}$$

where $\mathbf{1}_{\pi(t)=i}$ is the indicator function, and $r_i(P_i^{\max}, \boldsymbol{p}_{-i} = \mathbf{0})$ is user i's instantaneous throughput when it transmits at the maximum admissible power level $P_i^{\max}$ and the other users do not transmit $\boldsymbol{p}_{-i} = \mathbf{0}$.

Each user i's overall payoff v_i is then a function of the discounted average throughput. For video applications, we can use PSNR (peak signal-to-noise ratio) as the payoff, whose dependence on the throughput has been well studied. In this paper, we simply use the discounted average throughput as the payoff (i.e. $v_i = R_i$) for a fair comparison with existing works [1]–[8]. We also define user i's maximum achievable payoff as $\bar{v}_i \triangleq r_i(P_i^{\max}, \boldsymbol{p}_{-i} = \mathbf{0})$, which is achievable only when user i transmits all the time.

B. The Policy Design Problem

We want to maximize a social welfare function defined either on the users' normalized payoffs, $W(\frac{v_1}{\bar{v}_1}, \ldots, \frac{v_N}{\bar{v}_N})$, or on the users' absolute payoffs $W(v_1, \ldots, v_N)$. This definition of the welfare function is general enough to include the objective functions deployed in many existing works [1]–[13] as special cases. Example welfare functions include the average throughput $\sum_{i=1}^N \frac{v_i}{\bar{v}_i}$ and the max-min fairness $\min_i v_i$. At the maximum of the welfare function, some users may have extremely low payoffs. To avoid this, each user i can impose a normalized minimum payoff guarantee $\gamma_i \in [0, 1)$. To sum up, we can formally define the policy design problem as follows

$$\begin{aligned} \max_{\pi} \quad & W(\frac{v_1(\pi)}{\bar{v}_1}, \ldots, \frac{v_N(\pi)}{\bar{v}_N}) \text{ or } W(v_1(\pi), \ldots, v_N(\pi)) \\ s.t. \quad & \frac{v_i(\pi)}{\bar{v}_i} \geq \gamma_i, \ \forall i \in \mathcal{N}. \end{aligned} \tag{4}$$

III. A Motivating Example

Before introducing the design framework, we provide a motivating example to show the differences between

[2]The distortion impact of a packet measures how much this packet contributes to the overall video quality.

the cases with homogeneous and heterogeneous delay-sensitivity, and the advantage of the proposed policies over round-robin TDMA policies.

Consider a simple wireless ad hoc network with two symmetric users. For simplicity, we assume that each user can achieve a maximum throughput of 1 bits/s/Hz, which is obtained when the other user does not transmit. First consider the case when the users have the same delay-sensitivity δ. Without loss of generality, we assume that user 1 transmits first in the round-robin TDMA policy. Then under the round-robin TDMA policy, the two users overall payoffs are

$$v_1^{\mathrm{RR}} = (1-\delta)\cdot\sum_{t=0}^{\infty}\delta^{2t}\cdot 1 = \frac{1}{1+\delta},$$
$$v_2^{\mathrm{RR}} = (1-\delta)\cdot\sum_{t=0}^{\infty}\delta^{2t+1}\cdot 1 = \frac{\delta}{1+\delta}.$$

The round-robin TDMA policy achieves the maximum sum payoff, namely 1 bits/s/Hz. However, it is suboptimal in terms of fairness, because the user who transmits later always achieves a smaller payoff. We formally define the fairness criterion as the minimum payoff between the users, namely $\min\{v_1, v_2\}$. Hence, the payoff profile with the optimal fairness is $(0.5,\ 0.5)$, which cannot be achieved by round-robin TDMA policies. Actually, the transmission schedule that achieves the optimal fairness is not cyclic at all. When the discount factor is 0.9, the schedule in the first few time slots is "122121122112". We will show how to construct the optimal nonstationary policy in the case of homogeneous delay-sensitivity in Sec. IV-E.

Now we consider the case when the users have different delay-sensitivity $\delta_1 < \delta_2$. Then the round-robin TDMA policy that achieves the maximum sum payoff should let the more delay-sensitive user (i.e. user 1) to transmit first, which results in the following overall payoffs:

$$v_1^{\mathrm{RR}} = (1-\delta_1)\cdot\sum_{t=0}^{\infty}\delta_1^{2t}\cdot 1 = \frac{1}{1+\delta_1},$$
$$v_2^{\mathrm{RR}} = (1-\delta_2)\cdot\sum_{t=0}^{\infty}\delta_2^{2t+1}\cdot 1 = \frac{\delta_2}{1+\delta_2}.$$

The above payoff profile achieves a higher sum payoff than that in the case of homogeneous delay-sensitivity, because $\frac{1}{1+\delta_1}+\frac{\delta_2}{1+\delta_2} > \frac{1}{1+\delta_2}+\frac{\delta_2}{1+\delta_2} = 1$. In other words, it achieves a higher sum payoff to let the more delay-sensitive user to transmit first. This is also illustrated in Fig. 1, where some feasible operating points under heterogeneous delay-sensitivity lies beyond the dashed line that represents the feasible operating points under homogeneous delay-sensitivity. Simple calculation also tells us that it achieves a higher fairness to let the more delay-sensitive user 1 to transmit first.

From the theoretical results in Sec. IV, we know that the optimal fairness achieved by the proposed policy is

$$\frac{1-2(1/\delta_1-1/\delta_2)}{1+\frac{1/\delta_2-1}{1/\delta_1-1}}.$$

TABLE II

COMPARISON OF FAIRNESS IN THE SIMPLE EXAMPLE ($\delta_2 = 0.99$).

δ_1	0.83	0.86	0.89	0.92	0.95
Fairness, Round-robin	0.498	0.498	0.498	0.498	0.498
Fairness, Proposed	0.582	0.654	0.715	0.758	0.771
Performance improvement	17%	31%	44%	52%	55%

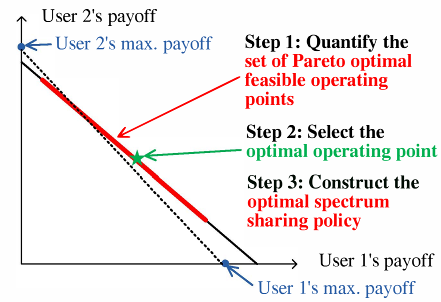

Fig. 3. The procedure of solving the design problem.

To illustrate the performance gain, we fix user 2's delay-sensitivity to $\delta_2 = 0.99$, and change user 1's delay-sensitivity. We show the fairness achieved by the round-robin policy and the proposed policy, as well as the performance improvement over the round-robin policy, in Table II.

Note that in Sec. V, we will evaluate the performance of different policies in more realistic scenarios (e.g. with more users), and show that the performance improvement is great (up to 100%).

IV. Solving The Policy Design Problem

In this section, we solve the policy design problem (4) following the procedure outlined in Fig. 3. We first quantify the set of operating points achievable by nonstationary policies, then select the optimal operating point based on the welfare function, and finally construct the policy to achieve the optimal operating point.

A. *Quantify The Set of Operating Points*

The first step in solving the design problem (4) is to characterize the set of operating points achievable by nonstationary spectrum sharing policies. We write the set of achievable operating points as $\mathcal{B}(\boldsymbol{\delta})$. The following theorem analytically quantify the set $\mathcal{B}(\boldsymbol{\delta})$ given the delay sensitivities $\boldsymbol{\delta}$.

Theorem 1: Given the delay sensitivities $\boldsymbol{\delta}$, if $\delta_1 \geq \frac{N}{N+1}$, we can achieve the following set of operating points using nonstationary spectrum sharing policies:

$$\mathcal{B}(\boldsymbol{\delta}) = \left\{ \boldsymbol{v} : \sum_{i\in\mathcal{N}} \frac{v_i}{c_i\cdot\bar{v}_i} = 1, v_i \geq 0, \forall i\in\mathcal{N} \right\}, \quad (5)$$

where $c_i = \frac{1/\delta_i-1}{1/\delta_1-1}$ for all $i \in \mathcal{N}$.

Proof: See Appendix I. ■

According to Theorem 1, as long as the difference between the users' delay sensitivity is not very large (i.e. $\delta_1 \geq \frac{N}{N+1}$), we can achieve a large set of operating points using nonstationary policies. The achievable operating points lie on

TABLE III

THE ALGORITHM RUN BY USER i.

```
Input: Normalized target payoffs {v_i^⋆/v̄_i}_{i∈N}
Initialization: Set t = 0, v'_j(0) = v_j^⋆/v̄_j, ∀j,
repeat
  Find the users with large "distances":
        J = {j ∈ N : v'_j(t) ≥ 1 − δ_j}
  if J ≠ ∅
    Find the most delay-sensitive one: i* = min_{j∈J} j
    if i = i*
      Transmits at the power level P_i^max
    end if
    v'_{i*}(t+1) = v'_{i*}(t)/δ_{i*} − (1/δ_{i*} − 1), v'_j(t+1) = v'_j(t)/δ_j, ∀j ≠ i*
  else
    v'_j(t+1) = (1/δ_j) · v'_j(t), ∀j ∈ N
  end if
  t ← t + 1
until ∅
```

TABLE IV

THE ALGORITHM RUN BY USER i (HOMOGENEOUS DELAY SENSITIVITY).

```
Input: Normalized target payoffs {v_i^⋆/v̄_i}_{i∈N}
Initialization: Set t = 0, v'_j(0) = v_j^⋆/v̄_j for all j ∈ N.
repeat
  Finds the user with the largest distance: i* ≜ arg max_{j∈N} v'_j(t)
  if i = i* then
    Transmits at the power level P_i^max
  end if
  Updates v'_j(t+1) for all j ∈ N as follows:
  v'_{i*}(t+1) = (1/δ) · v'_{i*}(t) − (1/δ − 1)
  v'_j(t+1) = (1/δ) · v'_j(t)
  t ← t + 1
until ∅
```

a hyperplane. In addition, the payoff achieved by the least delay-sensitive user is no smaller than a threshold. In other words, there are upper bounds on the payoffs achievable by more delay-sensitive users.

Note that the set $\mathcal{B}(\boldsymbol{\delta})$ determined in Theorem 1 may not be the set of *Pareto optimal* operating points. Determining the Pareto optimal operating points in the case of heterogeneous delay sensitivities is still an open problem, and will be an interesting future research topic.

B. Select The Optimal Operating Point

Since we have identified the set of achievable operating points, the problem of selecting the optimal operating point $\boldsymbol{v}^\star$ can be written as

$$\boldsymbol{v}^\star = \arg\max_{\boldsymbol{v}} W, \ s.t. \ \boldsymbol{v} \in \mathcal{B}(\boldsymbol{\delta}), \ \frac{v_i}{\bar{v}_i} \geq \gamma_i, \forall i. \qquad (6)$$

Since the constraints in the above problem are a linear equality and linear inequalities, the optimization problem (6) is convex and easy to solve when W is a convex function in $(v_1, \ldots, v_N)$.

C. Construct The Spectrum Sharing Policy

Given the optimal operating point $\boldsymbol{v}^\star$, we can construct a nonstationary policy to achieve it. The policy can be implemented by each user in a distributed manner. The algorithm run by user i is described in Table III. We can see from Table III that although the optimal nonstationary policy may be complicated, the algorithm to construct it is very simple. The algorithm schedules the transmission based on the users' delay sensitivities and the "distances" from the target payoff, where the distance from the target payoff is the normalized continuation payoff that a user needs to achieve after the current period. Starting from the most delay-sensitive user, the algorithm checks each user's distance. Among all the users with distances larger than some thresholds, the most delay-sensitive one will transmit in the current period. After this, the algorithm updates all the users' distances, which are used to determine who should transmit in the next period.

Theorem 2 ensures that if all the users run the algorithm in Table III locally, they will achieve the optimal operating point $\boldsymbol{v}^\star$.

Theorem 2: If each user i runs the algorithm in Table III, then each user i achieves the optimal payoff $v_i^\star$.

Proof: See Appendix II. ■

D. The Case of Homogeneous Delay Sensitivity

We describe the design framework when all the users have the same delay sensitivity δ. Although most results follow as the special cases of the corresponding results when the users have different delay sensitivities, we will discuss the differences from the case of heterogeneous delay sensitivities as we describe the results. Note that having the same delay sensitivity does not mean that the round-robin TDMA policy is optimal. First, we quantify the set of achievable operating points as follows.

Corollary 1: Given the delay sensitivity δ, if $\delta \geq 1 - \frac{1}{N+1}$, we can achieve the following set of operating points:

$$\mathcal{B}(\delta) = \left\{ \boldsymbol{v} : \sum_{i \in \mathcal{N}} \frac{v_i}{\bar{v}_i} = 1, v_i \geq 0, \forall i \in \mathcal{N} \right\}. \qquad (7)$$

Proof: When all the users have the same delay sensitivity δ, we have $c_i = 1$ for all $i \in \mathcal{N}$. Then this corollary follows as the special case of Theorem 1. ■

Corollary 1 says that as long as the users are not very delay-sensitive (i.e. $\delta \geq 1 - \frac{1}{N+1}$), we can achieve any Pareto optimal operating point by some nonstationary policy. The difference from the case of heterogeneous delay sensitivities is that we can achieve *all* the *Pareto optimal* operating points when the users have the same delay sensitivity. In contrast, when the users have different delay sensitivities, determining Pareto optimal operating points is still an open problem.

Given the set of achievable operating points, the users can solve for the optimal operating point following the same procedure as before. Then to construct the optimal nonstationary policy, each user runs the algorithm in Table IV, which is a simpler version of the algorithm in the case of heterogeneous delay sensitivities. Since the users have the same delay sensitivity, the algorithm determines who should transmit solely based on the users' distances from the target payoff.

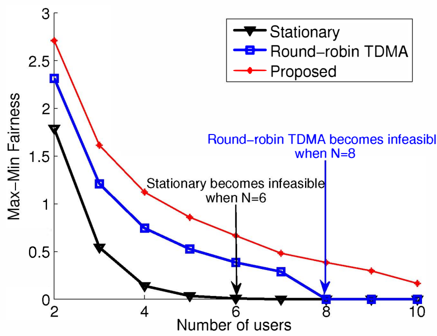

Fig. 4. Comparison of the max-min fairness achieved by the optimal stationary policy, the round-robin TDMA policy , and the proposed policy under different numbers of users (averaged over 1000 channel realizations).

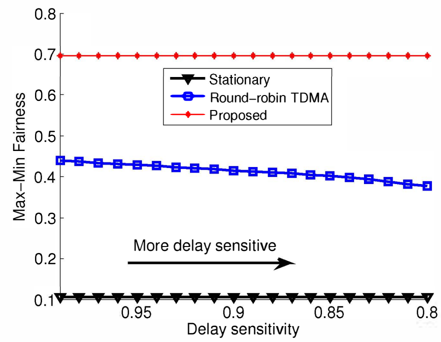

Fig. 5. Comparison of the max-min fairness achieved by the optimal stationary policy, the round-robin TDMA policy , and the proposed policy under different delay sensitivity (averaged over 1000 channel realizations).

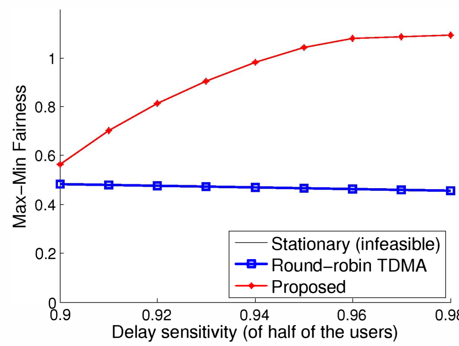

Fig. 6. Comparison of the max-min fairness achieved by the optimal stationary policy, the round-robin TDMA policy , and the proposed policy under different delay sensitivity heterogeneity (averaged over 1000 channel realizations).

V. SIMULATION RESULTS

In this section, we demonstrate the performance gain of our spectrum sharing policy over existing policies. Throughout this section, we use the following system parameters. The noise powers at all the users' receivers are normalized as 0 dB. The maximum transmit powers of all the users are 20 dB. For simplicity, we assume that the direct channel gains have the same distribution $g_{ii} \sim \mathcal{CN}(0,1), \forall i$, and the cross channel gains have the same distribution $g_{ij} \sim \mathcal{CN}(0,0.5), \forall i \neq j$. The performance criterion is the max-min fairness $\min_i v_i$, namely we maximize the payoff of the user that achieves the lowest throughput. In this way, we can achieve both high spectrum efficiency and fairness among the users. Each user's minimum payoff guarantee is 10% of its maximum achievable payoff, i.e. $\gamma_i = 0.1, \forall i$.

We compare against the optimal stationary policies [1]–[8] and the round-robin TDMA policy. In round-robin TDMA policy, the more delay-sensitive user transmits earlier. When they have the same delay sensitivity, the user with a higher minimum payoff guarantee transmits earlier.

Number of users: In Fig. 4, we compare the max-min fairness achieved by the proposed policy and the other two policies under different numbers of users. We fix the delay sensitivity to be $\delta = 0.9$ for all the users. We randomly generate 1000 channel realizations, calculate the max-min fairness under each channel realization, and take the average. First, observe that the optimal stationary policy and the round-robin TDMA policy are infeasible (i.e. fail to achieve the minimum payoff guarantee) when the number of users is large, which is reflected by the zero max-min fairness shown in the figure. In the considered scenario, when there are more than 5 users, the stationary policy becomes infeasible due to strong interference. When there are more than 7 users, the round-robin TDMA policy becomes infeasible due to the delay experienced by the users who transmit later. On the other hand, the proposed policy remains feasible even when the other policies are not. Moreover, when the number of users is large, the proposed policy doubles the max-min fairness achieved by the round-robin TDMA policy (e.g. when $N = 6$), and achieves performance improvement in the orders of magnitude over the stationary policy.

Delay sensitivity: In Fig. 5, we compare the max-min fairness achieved by the three policies under different delay sensitivities. We assume that the users have the same delay sensitivity. Since the stationary and round-robin TDMA policies are infeasible when the number of users is large, we do the simulation for $N = 5$. We can see that although the delay sensitivity does not affect the performance of the stationary policy, the stationary policy has a much lower performance due to multi-user interference. The performance of the round-robin TDMA policy degrades as the users become more delay-sensitive, because the users who transmit later suffer from more performance loss when they are more delay-sensitive. In contrast, the performance of the proposed policy does not degrade when the users are more delay-sensitive. In particular, when the delay sensitivity is $\delta = 0.8$, the proposed policy achieves a performance improvement of 100%.

Heterogeneity in delay sensitivity: In Fig. 6, we compare the social welfare achieved by the three policies, when the users have different delay sensitivity. We assume that there are $N = 6$ users. Users 4-6 have the same delay sensitivity fixed at $\delta_4 = \delta_5 = \delta_6 = 0.99$, while users 1-3 have the same delay sensitivity which is changing. The stationary policy is infeasible when $N = 6$. Again, the proposed policy achieves a significant performance improvement over the round-robin policy (up to 100% when $\delta_1 \in [0.92, 0.98]$).

VI. Conclusion

In this paper, we studied nonstationary spectrum sharing policies among users with heterogeneous delay-sensitivity. The proposed policy significantly outperforms existing stationary policies and simple nonstationary policies (i.e. the round-robin TDMA policy). Although the nonstationary policy is difficult to analyze and design, we propose a low-complexity algorithm, which can be run by each user in a distributed manner, to construct the optimal nonstationary policy.

Appendix I
Proof of Theorem 1

A. Preliminaries on The Theory of Self-generating Sets

The proof builds on the theory of self-generating sets [16], which can be considered as the multi-user extension of Bellman equations in dynamic programming. We can decompose each user i's discounted average payoff into the current payoff and the continuation payoff as follows:

$$\begin{aligned} R_i(\pi) &= (1-\delta_i)\sum_{t=0}^{\infty}\delta_i^t\cdot(\mathbf{1}_{i=\pi(t)}\cdot\bar{v}_i) \\ &= (1-\delta_i)\cdot \underbrace{\mathbf{1}_{i=\pi(0)}\cdot\bar{v}_i}_{\text{current payoff at } t=0} \\ &+ \delta_i\cdot\underbrace{\left[(1-\delta_i)\sum_{t=1}^{\infty}\delta_i^{t-1}\cdot(\mathbf{1}_{i=\pi(t)}\cdot\bar{v}_i)\right]}_{\text{continuation payoff starting from } t=1}. \end{aligned}$$

We can see that the continuation payoff starting from $t=1$ is the discounted average payoff as if the system starts from $t=1$. In general, we can define user i's continuation payoff starting from t, written as $r_i(t)$, as follows:

$$v_i(t) = (1-\delta_i)\sum_{\tau=t}^{\infty}\delta_i^{\tau-t}\cdot(\mathbf{1}_{i=\pi(\tau)}\cdot\bar{v}_i).$$

Then the decomposition at time t can be simply written as $v_i(t) = (1-\delta_i)\cdot(\mathbf{1}_{i=\pi(t)}\cdot\bar{v}_i)+\delta_i\cdot v_i(t+1)$. Note that here we allow $\pi(t)=0$, which means that no user transmits in time t.

Based on the above decomposition, the theory of self-generating sets [16] characterizes the set of feasible payoffs. Write the vector of payoffs as $\boldsymbol{v}=\{v_1,\ldots,v_N\}$. Then the self-generating set is defined as follows.

Definition 1 (Self-generating Set [16]): A set of payoff vectors $\mathcal{V}$ is a self-generating set under the discount factors $\boldsymbol{\delta}$, if for any payoff vector $\boldsymbol{v}\in\mathcal{V}$, there exists a $i^*\in\mathcal{N}\cup\{0\}$ and a continuation payoff $\boldsymbol{v}'\in\mathcal{V}$ such that

$$v_i = (1-\delta_i)\cdot(\mathbf{1}_{i=i^*}\bar{v}_i)+\delta_i\cdot v_i'. \tag{8}$$

The key property of a self-generating set is that when we decompose the current payoff vector $\boldsymbol{v}$ by letting user $i^*\in\mathcal{N}$ or no user ($i^*=0$) utilize the resource, the continuation payoff vector $\boldsymbol{v}'$ should also lie in the set $\mathcal{V}$. Since any payoff vector in the set can be decomposed by a continuation payoff in the set, all the payoff vectors in the set are achievable [16].

B. Characterization of Achievable Payoff Profiles

Now we focus on the set of payoff vectors of the following form:

$$\mathcal{V}=\left\{\boldsymbol{v}:\sum_{i\in\mathcal{N}}\frac{v_i}{c_i\cdot\bar{v}_i}\le 1,\frac{v_i}{\bar{v}_i}\ge\mu_i,\forall i\right\} \tag{9}$$

with $c_i>0$ and $\mu_i\ge 0$, and derive the conditions on $\{c_i\}$ and $\{\mu_i\}$ under which $\mathcal{V}$ is a self-generating set.

For a given vector $\boldsymbol{v}\in\mathcal{V}$, if we let user $i\in\mathcal{N}$ to use the resource, the continuation payoff vector $\boldsymbol{v}'$ is

$$v_i'=\frac{v_i}{\delta_i}-\frac{1-\delta_i}{\delta_i}\cdot\bar{v}_i,\ \text{and}\ v_j'=\frac{v_j}{\delta_j},\ \forall j\ne i. \tag{10}$$

To satisfy $\boldsymbol{v}'\in\mathcal{V}$, the following conditions must be fulfilled:

$$\frac{v_i}{\bar{v}_i}\ge 1-\delta_i(1-\mu_i)\ \text{and}\ \sum_{j\in\mathcal{N}}\frac{v_j/\bar{v}_j}{c_j\delta_j}-\frac{1-\delta_i}{c_i\delta_i}\le 1. \tag{11}$$

Similarly, if we let no user to use the resource, the continuation payoff vector $\boldsymbol{v}'$ is

$$v_i'=\frac{v_i}{\delta_i},\ \forall i\in\mathcal{N}. \tag{12}$$

To satisfy $\boldsymbol{v}'\in\mathcal{V}$, the following conditions must be fulfilled:

$$\sum_{j\in\mathcal{N}}\frac{v_j/\bar{v}_j}{c_j\delta_j}\le 1. \tag{13}$$

Hence, for $\mathcal{V}$ to be a self-generating set, we need to ensure that for any $\boldsymbol{v}\in\mathcal{V}$, there exists a $i\in\mathcal{N}$ such that (11) is fulfilled, or if such a i does not exist, (13) is fulfilled. We can divide $\mathcal{V}$ into two subsets $\mathcal{V}_1$ and $\mathcal{V}_2$, where $\mathcal{V}_1$ contains the payoff profiles for which there may exist a $i\in\mathcal{N}$ such that $\frac{v_i}{\bar{v}_i}\ge 1-\delta_i(1-\mu_i)$, and $\mathcal{V}_2$ contains the payoff profiles for which $\frac{v_i}{\bar{v}_i}<1-\delta_i(1-\mu_i)$ for all $i\in\mathcal{N}$. Formally, we have

$$\mathcal{V}_1\triangleq\{\boldsymbol{v}\in\mathcal{V}:\exists i\in\mathcal{N}\ \text{such that}\ \frac{v_i}{\bar{v}_i}\ge 1-\delta_i(1-\mu_i)\}, \tag{14}$$

and

$$\mathcal{V}_2\triangleq\{\boldsymbol{v}\in\mathcal{V}:\frac{v_i}{\bar{v}_i}<1-\delta_i(1-\mu_i),\forall i\in\mathcal{N}\}. \tag{15}$$

In other words, for any payoff profile $v\in\mathcal{V}_2$, we need to decompose it by letting no user transmit.

Now we find conditions on $\{c_i\},\{\mu_i\}$ under which $\mathcal{V}$ is a self-generating set. We look at the payoff profiles in $\mathcal{V}_1$ and those in $\mathcal{V}_2$ separately. First, we look at the payoff profiles in $\mathcal{V}_2$. For any $\boldsymbol{v}\in\mathcal{V}_2$, (13) must be fulfilled. Since $\sum_{j\in\mathcal{N}}\frac{v_j/\bar{v}_j}{c_j\delta_j}$ is increasing in v_j, we must have

$$\sum_{j\in\mathcal{N}}\frac{1-\delta_j(1-\mu_j)}{c_j\delta_j}\le 1. \tag{16}$$

Then we look at the payoff profiles in $\mathcal{V}_1$. We define the following optimization problem:

$$\kappa(\{c_i\},\{\mu_i\})=\max_{\boldsymbol{v}\in\mathcal{V}_1}\min_{i\in\mathcal{I}(\boldsymbol{v})}\sum_{j\in\mathcal{N}}\frac{v_j/\bar{v}_j}{c_j\delta_j}-\frac{1-\delta_i}{c_i\delta_i}, \tag{17}$$

where $\mathcal{I}(\boldsymbol{v}) \triangleq \{i \in \mathcal{N} : \frac{v_i}{\bar{v}_i} \geq 1-\delta_i(1-\mu_i)\}$. We can see that if $\kappa(\{c_i\},\{\mu_i\}) \leq 1$, then any $\boldsymbol{v} \in \mathcal{V}_1$ can be decomposed properly.

Instead of solving the complicated optimization problem (17) directly, we simplify it by focusing on the $\{c_i\}$ that satisfy:

$$\frac{1-\delta_i}{c_i\delta_i} = \beta,\ \forall i \in \mathcal{N}, \tag{18}$$

where $\beta > 0$ is a constant. In this way, given $\boldsymbol{v}$, the objective function in (17) is independent of the choice of i, as long as $i \in \mathcal{I}(\boldsymbol{v})$. Hence, the optimization problem can be simplified into

$$\bar{\kappa}(\beta,\{\mu_i\}) = \quad \max_{\boldsymbol{v}\in\mathcal{V}_1} \quad \beta \cdot \sum_{j\in\mathcal{N}} \frac{v_j/\bar{v}_j}{1-\delta_j} - \beta. \tag{19}$$

Now, instead of trying to fulfill $\kappa(\{c_i\},\{\mu_i\}) \leq 1$, we need to find conditions under which $\bar{\kappa}(\beta,\{\mu_i\}) \leq 1$. First, the objective function is increasing in v_j. As a result, at the solution $\boldsymbol{v}^*$ to the above problem, the constraint $\sum_j \frac{v_j^*/\bar{v}_j}{c_j} \leq 1$ must hold as an equality. In other words, we must have $\sum_j \frac{v_j^*/\bar{v}_j}{c_j} = 1$. Moreover, the solution $\boldsymbol{v}^*$ to the above linear programming must have $v_{i^*}^*$ as large as possible, where $i^* = \arg\max_i \frac{c_i}{1-\delta_i}$. Since $\frac{c_i}{1-\delta_i} = \frac{\beta}{\delta_i}$ and δ_i is increasing with i, we have $i^* = 1$. In summary, to maximize $\sum_{j\in\mathcal{N}} \frac{v_j/\bar{v}_j}{1-\delta_j}$, we need to set $v_1/\bar{v}_1$ as large as possible. Hence, the solution to the above optimization problem is

$$v_1^* = c_1\bar{v}_1 \cdot \left(1-\sum_{j=2}^N \frac{\mu_j}{c_j}\right),\ v_j^* = \mu_j\bar{v}_j, j=2,\ldots,N,$$

for which we make another restriction

$$c_1\left(1-\sum_{j=2}^N \mu_j/c_j\right) \geq 1-\delta_1(1-\mu_1) \tag{20}$$

such that the solution $\boldsymbol{v}^*$ is in $\mathcal{V}_1$.

Then $\bar{\kappa}(\beta,\{\mu_i\}) \leq 1$ is equivalent to

$$\beta \cdot \left[\sum_{j=2}^N \left(1-\frac{\delta_j}{\delta_N}\right)\frac{\mu_j}{1-\delta_j}\right] - \beta \leq 1 - \frac{1}{\delta_1}. \tag{21}$$

The condition (16) to satisfy for the payoff profiles in $\mathcal{V}_2$ can be rewritten as

$$\beta \cdot \left(N + \sum_{j=1}^N \frac{\mu_j}{1-\delta_j}\delta_j\right) \leq 1. \tag{22}$$

In summary, we need to find $\{c_i\}$, $\{\mu_i\}$, and β such that the conditions in (18), (20), (21), and (22) are fulfilled. First, we want to have $\mathcal{V}$ as large as possible such that the set of achievable payoff profiles is large. Hence, we set $\mu_1 = \cdots = \mu_N = 0$. Note that this will also make the inequalities in (20), (21), and (22) easier to hold. Given $\mu_1 = \cdots = \mu_N = 0$, (21) reduces to $\beta \geq \frac{1}{\delta_1} - 1$, which gives a lower bound for β. We should let $\beta = \frac{1}{\delta_1} - 1$ such that $\{c_i\}$ can be large. Then we have $c_1 = 1$ and $c_i = \frac{1/\delta_1 - 1}{1/\delta_i - 1}$. We can check that (20) is satisfied. Finally, to fulfill (22), we must have $\beta = \frac{1}{\delta_1} - 1 \leq \frac{1}{N}$, which leads to $\delta_1 \geq \frac{N}{N+1}$. Although we can achieve any payoff profiles in $\mathcal{V}$, we only care about the Pareto boundary of $\mathcal{V}$, which is $\mathcal{B}(\boldsymbol{\delta})$ in Theorem 1.

Appendix II
Proof of Theorem 2

In Appendix I, we have characterized the self-generating set $\mathcal{V}$. In the algorithm in Table III, we start with the target payoff $\boldsymbol{v}^\star$ as the discounted average payoff at time 0, and decompose it into a current payoff and a continuation payoff. The decomposition tells us who should transmit in time 0. Then we decompose the continuation payoff and determine who should transmit to play in time 1. By performing the decomposition in every time slot, we can determine the complete transmission schedule.

Specifically, suppose that the continuation payoff at time t is $\boldsymbol{v}(t)$. Then if $\boldsymbol{v}(t) \in \mathcal{V}_2$, then no user transmit. Otherwise, since $\boldsymbol{v}(t) \in \mathcal{V}_1$, we can let any user i, who satisfies $\frac{v_i}{\bar{v}_i} \geq 1 - \delta_i$, to transmit. As proved in Appendix I, such a decomposition is always feasible.

References

[1] X. Kang, R. Zhang, Y.-C. Liang, and H. K. Garg, "Optimal power allocation strategies for fading cognitive radio channels with primary user outage constraint," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 2, pp. 374-383, 2011.

[2] C. W. Tan and S. H. Low, "Spectrum management in multiuser cognitive wireless networks: Optimality and algorithm," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 2, pp. 421-430, 2011.

[3] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 5, pp. 1074-1084, May 2006.

[4] N. Gatsis, A. G. Marques, G. B. Giannakis, "Power control for cooperative dynamic spectrum access networks with diverse QoS constraints," *IEEE Trans. Commun.*, vol. 58, no. 3, pp. 933-944, 2010.

[5] J. Huang, R. A. Berry, and M. L. Honig, "Auction-based spectrum sharing," *Mobile Network and Applications*, vol. 11, pp. 405-418, 2006.

[6] S. Sorooshyari, C. W. Tan, M. Chiang, "Power control for cognitive radio networks: Axioms, algorithms, and analysis," *IEEE/ACM Trans. Netw.*, vol. 20, no. 3, pp. 878–891, 2012.

[7] S. Sharma and D. Teneketzis, "An externalities-based decentralized optimal power allocation algorithm for wireless networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 6, pp. 1819–1831, 2009.

[8] S. Sharma and D. Teneketzis, "A game-theoretic approach to decentralized optimal power allocation for cellular networks," *Telecommunication Systems*, pp. 1–16, 2010.

[9] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," *IEEE J. Sel. Areas Commun.*, vol. 25, pp. 517–528, 2007.

[10] Y. Wu, B. Wang, K. J. R. Liu, and T. C. Clancy, "Repeated open spectrum sharing game with cheat-proof strategies," *IEEE Trans. Wireless Commun.*, vol. 8, no. 4, pp. 1922-1933, 2009.

[11] M. van der Schaar and F. Fu, "Spectrum access games and strategic learning in cognitive radio networks for delay-critical applications," *Proc. of IEEE, Special issue on Cognitive Radio*, vol. 97, no. 4, pp. 720–740, Apr. 2009.

[12] M. Le Treust and S. Lasaulce, "A repeated game formulation of energy-efficient decentralized power control," *IEEE Trans. on Wireless Commun.*, vol. 9, no. 9, pp. 2860–2869, 2010.

[13] Y. Xiao, J. Park, and M. van der Schaar, "Repeated games with intervention: Theory and applications in communications," *IEEE Trans. Commun.*, vol. 60, no. 10, pp. 3123–3132, 2012.

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

[15] F. Fu and M. van der Schaar, "A systematic framework for dynamically optimizing multi-user video transmission," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 3, pp. 308–320, Apr. 2010.

[16] D. Abreu, D. Pearce, and E. Stacchetti, "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, vol. 58, no. 5, pp. 1041–1063, 1990.