

A Reinforcement Learning Approach for the Multichannel Rendezvous Problem

Jen-Hung Wang, Ping-En Lu, Cheng-Shang Chang, and Duan-Shin Lee
Institute of Communications Engineering
National Tsing Hua University
Hsinchu 30013, Taiwan, R.O.C.

Email: f0960778676@gmail.com; j94223@gmail.com; cschang@ee.nthu.edu.tw; lds@cs.nthu.edu.tw

Abstract—In this paper, we consider the multichannel rendezvous problem in cognitive radio networks (CRNs) where the probability that two users hopping on the same channel have a successful rendezvous is a function of channel states. The channel states are modelled by two-state Markov chains that have a good state and a bad state. These channel states are not observable by the users. For such a multichannel rendezvous problem, we are interested in finding the optimal policy to minimize the expected time-to-rendezvous (ETTR) among the class of *dynamic blind rendezvous policies*, i.e., at the t^{th} time slot each user selects channel i independently with probability $p_i(t)$, $i = 1, 2, \dots, N$. By formulating such a multichannel rendezvous problem as an adversarial bandit problem, we propose using a reinforcement learning approach to learn the channel selection probabilities $p_i(t)$, $i = 1, 2, \dots, N$. Our experimental results show that the reinforcement learning approach is very effective and yields comparable ETTRs when comparing to various approximation policies in the literature.

Index Terms—reinforcement learning, multichannel rendezvous

I. INTRODUCTION

The multichannel rendezvous problem that asks two secondary users (SU) to find a common available channel (not used by primary users (PU)) is one of the fundamental problems in cognitive radio networks (CRNs) (see e.g., the book [1] and references therein). In view of possible jamming attacks [2], the multichannel rendezvous problem is commonly solved by having each SU hopping on its available channels over time. When both SUs hop on a common available channel at the same time, it is assumed that a successful rendezvous occurs. For such a rendezvous problem, the objective is to minimize the time-to-rendezvous (TTR), i.e., the first time that the two SUs have a successful rendezvous. In the literature, there are various *deterministic* channel hopping (CH) sequences that can guarantee finite maximum time-to-rendezvous (MTTR) under various assumptions for CRNs, e.g., QCH [3], DRSEQ [4], Modular Clock [5], JS [6], DRDS [7], FRCH [8], ARCH [9], CBH [10], and Two-prime Modular Clock [11]. As pointed out in [12], there is one practical factor that is not considered in these CH sequences, i.e., the channel states. Due to channel fading and interferences from other SUs, two SUs might not have a successful rendezvous even when they both hop on a common available channel at the

same time. As such, it might be more practical to focus on the expected time-to-rendezvous (ETTR), instead of the MTTR.

In [12], the authors considered a random channel state model, in which each channel has several random states and the probability that two SUs hopping on a common channel have a successful rendezvous is a function of the channel state. Specifically, for a CRN with N channels, the states of the N channels are characterized by a stochastic process $\{\mathbf{X}(t) = (X_1(t), X_2(t), \dots, X_N(t)), t \geq 0\}$, where $X_i(t)$, $i = 1, 2, \dots, N$, is the random variable that represents the state of channel i at time t . The channel states are assumed to be not observable by a SU. When two SUs hopping on a channel in state x , they will rendezvous with probability $r(x)$. The authors in [12] considered the class of *blind rendezvous policies* in which each user selects channel i independently with probability p_i , $i = 1, 2, \dots, N$, in every time slot. They showed that there does not exist a channel selection policy (in terms of the channel selection probabilities, p_i , $i = 1, 2, \dots, N$) that is universally optimal for any time-varying channel state model. For a fast time-varying channel model, the optimal policy is the single channel policy that only selects one particular channel. On the other hand, for a slow time-varying channel model, SUs should avoid selecting a single channel as that channel could be in a bad state for a long period of time.

Even though the channel states are not observable, one question is whether they can be implicitly learned (from either failed attempts or successful rendezvous) so as to speed up the rendezvous process in the future. To address such a question, we adopt a reinforcement learning approach to learn the channel selection probabilities of a SU. Reinforcement learning (see, e.g., the book [13] and the recent survey [14]) is a field of machine learning that addresses the problems of how to behave in an environment by performing certain actions and observing the reward from those actions. In these problems, the fixed limited resources must be allocated to maximize their expected gain. The reward of choice is only known at the time of allocation and may become better understood as time passes. Our problem is then to treat the channel selection probabilities as the fixed limited resources and learn how to allocate the channel selection probabilities to minimize ETTR. Specifically, our approach is to consider the multichannel ren-

dezdvous problem as the multi-armed bandit problem, and each successful rendezvous on a channel renders a reward for that channel. We then use the adversarial bandit algorithm, **Exp3**, in [15] to learn the channel selection probabilities. When the N channels are independent and identically distributed (i.i.d.), our numerical experiments show that **Exp3** yields comparable ETTRs to various approximation algorithms proposed in [12]. On the other hand, when channels are not i.i.d., **Exp3** is capable of learning the “best” channel. To the best of our knowledge, it seems that our paper is the first to study the multichannel rendezvous problem by a reinforcement learning approach.

II. SYSTEM MODEL

A. The multichannel rendezvous problem

In this paper, we consider a cognitive radio network (CRN) with N channels (with $N \geq 2$), indexed from 1 to N , in the discrete-time setting where time is slotted and indexed from $t = 0, 1, 2, \dots$. We assume that there are two states for each channel, state 0 for the *bad* state and state 1 for the *good* state. Denote by $r(x)$ the rendezvous probability when a channel in state x , $x = 0$ or 1. Then when two users hop on a channel in state x at the same time, these two users will rendezvous with probability $r(x)$, and this is independent of everything else. Since state 0 is the bad state and state 1 is the good state, we assume that

$$r(0) \leq r(1).$$

The states of the N channels are characterized by the stochastic process $\{\mathbf{X}(t) = (X_1(t), X_2(t), \dots, X_N(t)), t \geq 0\}$, where $X_i(t)$, $i = 1, 2, \dots, N$, is the random variable that represents the state of channel i at time t . The exact state of a channel at any time is not observable by a user. As discussed in [12], the reason for that is because it is in general difficult for a user to know the congestion level of a channel (the number of users in a channel).

We consider the class of *dynamic blind rendezvous policies*, i.e., at the t^{th} time slot each user selects channel i with probability $p_i(t)$, $i = 1, 2, \dots, N$. Such a channel selection is independent of everything else. Suppose that the channel state of the i^{th} channel at time t is x_i , $i = 1, 2, \dots, N$. Then under the dynamic blind rendezvous policy, the probability that these two users will have a successful rendezvous at time t on channel i is simply $(p_i(t))^2 \cdot r(x_i)$. This is because the two users have to hop on channel i at time t and the rendezvous is successful on channel i with probability $r(x_i)$. As such, the two users will have a successful rendezvous at time t is $\sum_{i=1}^N (p_i(t))^2 \cdot r(x_i)$. The objective is to learn a dynamic blind rendezvous policy (and the corresponding channel selection probabilities) that minimizes the expected time-to-rendezvous (ETTR).

B. A Markov channel model with two states

For the model of channel states, we consider the Markov chain with two states in [12]. We assume that the states of these N channels are *independent*. The probability that the i^{th}

channel is in the good (resp. bad) state is ρ_i (resp. $1 - \rho_i$) for some $0 \leq \rho_i \leq 1$. As such, we have the following stationary joint distribution for the channel states

$$\begin{aligned} P(X_1(t) = x_1, X_2(t) = x_2, \dots, X_N(t) = x_N) \\ = \prod_{i=1}^N \rho_i^{x_i} (1 - \rho_i)^{1-x_i}, \end{aligned} \quad (1)$$

where x_i (with the value being 0 or 1) is the state of channel i . For the i^{th} channel, its channel state is characterized by a Markov chain with the transition probabilities:

$$P(X_i(t+1) = 1 | X_i(t) = 1) = p_{1,1}^{(i)}, \quad (2)$$

$$P(X_i(t+1) = 0 | X_i(t) = 1) = 1 - p_{1,1}^{(i)}, \quad (3)$$

$$P(X_i(t+1) = 0 | X_i(t) = 0) = p_{0,0}^{(i)}, \quad (4)$$

$$P(X_i(t+1) = 1 | X_i(t) = 0) = 1 - p_{0,0}^{(i)}, \quad (5)$$

where $0 < p_{1,1}^{(i)}, p_{0,0}^{(i)} < 1$. Clearly, we have

$$\rho_i = P(X_i(t) = 1) = \frac{1 - p_{0,0}^{(i)}}{(1 - p_{1,1}^{(i)}) + (1 - p_{0,0}^{(i)})}.$$

Note that

$$\text{Var}[X_i(t+1)] = \text{Var}[X_i(t)] = \rho_i(1 - \rho_i)$$

and thus the correlation coefficient between $X_i(t+1)$ and $X_i(t)$, denoted by ω_i , is

$$\begin{aligned} \frac{E[X_i(t+1)X_i(t)] - E[X_i(t+1)]E[X_i(t)]}{\sqrt{\text{Var}[X_i(t+1)]\text{Var}[X_i(t)]}} \\ = p_{1,1}^{(i)} + p_{0,0}^{(i)} - 1. \end{aligned} \quad (6)$$

We say that the Markov chain $\{X_i(t), t \geq 0\}$ is positively correlated if $\omega_i \geq 0$. In this paper, we only consider positively correlated two-state Markov chains and we note the transition probabilities of the i^{th} Markov chain can be characterized by the two parameters ρ_i and ω_i . It is shown in [12] that the ETTR of a blind rendezvous policy is bounded below when $\omega_i = 0$ for all i and it is bounded above when $\omega_i = 1$ for all i . The argument used there can also be extended to show that the ETTR of a blind rendezvous policy is in fact increasing in ω_i when $\omega_i \geq 0$. Based on such structural results, various approximation algorithms for choosing the channel selection probabilities p_1, p_2, \dots, p_N of a (fixed) blind rendezvous policy were proposed in [12]. These policies include

- (i) Single selection policy: $p_1 = 1$ and $p_i = 0$ for $i = 2, \dots, N$.
- (ii) Uniform selection policy: $p_i = 1/N$ for $i = 1, 2, \dots, N$.
- (iii) $(1 + \epsilon)$ -approximation policy: $p_i = \frac{\sqrt{u_i}}{\sum_{j=1}^N \sqrt{u_j}}$, where $u_1 = 1 - (N - 1)\delta$, $u_i = \delta$, $i = 2, \dots, N$, and $\delta = (\frac{\epsilon}{3(N-1)})^2$.
- (iv) Harmonic selection policy: $p_i = c/i$, $i = 1, 2, \dots, N$, where c is the normalization constant so that the sum of p_i 's is 1.

- (v) Square selection policy: $p_i = c/i^2$, $i = 1, 2, \dots, N$, where c is the normalization constant so that the sum of p_i 's is 1.
- (vi) Sqrt selection policy: $p_i = c/i^{1/2}$, $i = 1, 2, \dots, N$, where c is the normalization constant so that the sum of p_i 's is 1.

These 6 blind rendezvous policies will serve as the benchmarks for the comparison with our reinforcement learning approach. In particular, it was shown in [12] that the $(1 + \epsilon)$ -approximation policy achieves an asymptotic $(1 + \epsilon)$ -approximation ratio in the setting where either $r(0) \rightarrow r(1)$ or $r(0) \rightarrow 0$ and $\omega_i = 1$ for all i . In our experiments, we set $\epsilon = 0.2$.

III. REINFORCEMENT LEARNING

In this section, we adopt a reinforcement learning approach to learn the channel selection probabilities so as to minimize the ETTR. It is assumed that each user cannot observe the channel states of the (hidden) Markov chain. This is similar to the multi-armed bandit problem where a gambler does not know the success probability of a slot machine. For this, we formulate the multichannel rendezvous problem as an adversarial bandit problem [16] in which there are N possible actions, indexed from $1, 2, \dots, N$, in each time slot. The i^{th} action corresponds to the selection of the i^{th} channel. When two SUs rendezvous, one unit of reward is given to both users. Otherwise, there is no reward for the two SUs. For such an adversarial bandit problem, a famous algorithm to choose actions is the **Exp3** algorithm [16] (which stands for "Exponential-weight algorithm for Exploration and Exploitation"). In Algorithm 1, we show the detailed steps of the **Exp3** algorithm for the multichannel rendezvous problem.

ALGORITHM 1: **Exp3** for the multichannel rendezvous problem

Input: A real parameter $\gamma \in (0, 1]$ and a time horizon T

Output: The channel selection probabilities $p_i(t)$, $i = 1, 2, \dots, N$ and $t = 1, 2, \dots, T$.

Initialization: Set $w_i(1) = 1$ for $i = 1, \dots, N$.

For each $t = 1, 2, \dots, T$

1: Set $p_i(t) = (1 - \gamma) \frac{w_i(t)}{\sum_{j=1}^N w_j(t)} + \frac{\gamma}{N}$, $i = 1, \dots, N$.

2: Select a channel i_t randomly accordingly to the probabilities $p_1(t), \dots, p_N(t)$.

3: Receive reward $z_{i_t} = 1$ if there is a successful rendezvous, and $z_{i_t} = 0$ otherwise.

4: For $j = 1, \dots, N$, set

$$\hat{z}_j(t) = \begin{cases} z_j(t)/p_j(t) & \text{if } j = i_t, \\ 0 & \text{otherwise.} \end{cases}$$

5: Set $w_i(t+1) = w_i(t) \exp(\gamma \hat{z}_i(t)/N)$.

To see the intuition of Algorithm 1, we note that there are two terms in the channel selection probabilities $p_i(t)$'s in Step

1. These two terms represent two fundamental concepts of reinforcement learning, exploitation, and exploration. The first term in $p_i(t)$ is the "exploitation" term that makes the "best" decision given the current information. The second terms in $p_i(t)$ is the "exploration" term that allows us to gather more information that might lead to better decisions. These two concepts are rather intuitive for the channel selection problem. The exploitation term leads to a "good" channel. On the other hand, as the channel might change its state in the next time slot, the exploration term allows us to find another good channel. The parameter $w_i(t)$ is the weight of the channel i at time t and they are set to be 1 at time 1. When a successful rendezvous occurs, both SUs receive one unit of reward. We do not give a penalty to a channel selected by an SU that does not lead to a successful rendezvous. Therefore, two SUs have the same weights for all t and thus the same channel selection probabilities $p_i(t)$, $i = 1, 2, \dots, N$ for all t . The weight update rules in Steps 4 and 5 are the softmax update [17] that increases the weight of a channel with a successful rendezvous.

The reward in Step 3 of the original **Exp3** algorithm in [16] is assigned by an adversary. As there are two SUs in the multichannel rendezvous problem, SU 2 can be viewed as the adversary of SU 1. Intuitively, one might think such an adversarial viewpoint might be used for deriving an upper bound on the expected weak regret (defined as the difference between the maximum accumulated reward and the accumulated reward from the **Exp3** algorithm) like Theorem 3.1 of [16]. However, as the channel selection probabilities of these two SUs are coupled through Algorithm 1, the rewards of these two SUs are not independent of each other and the analysis in Theorem 3.1 of [16] cannot be directly applied.

Another insight of Algorithm 1 is to view it as a stochastic game [18]. If $r(0) = r(1) = 1$, then it is clear that the single selection policy that selects channel 1 all the time is optimal as both users rendezvous in every time slot. Through the process of exploration and exploitation, one expects that Algorithm 1 converges to the channel selection probabilities with $p_1 = (1 - \gamma) + \gamma/N$, $p_i = \gamma/N$, $i = 2, \dots, N$. This is exactly the $(1 + \epsilon)$ -approximation policy (for some ϵ that is a function of γ) that achieves the $(1 + \epsilon)$ -approximation ratio. Such an intuitive observation will be further verified in our experiments in Section IV.

IV. EXPERIMENTAL RESULTS

In this section, we report our experimental results. For Algorithm 1, we set $\gamma = 0.02$. If γ is set to be very large, then the probability distribution will be similar to the uniform selection policy. On the other hand, if γ is very small, then the update of $w_i(t)$ is very small and that leads to a very slow convergence of the algorithm.

In our first experiment, we consider a system of 16 independent two-state Markov channels with the same parameters, i.e., $N = 16$. The rendezvous probability at state 0 (resp. state 1) is $r(0) = 0.001$ (resp. $r(1) = 1$). There are 9 parameter settings for the two-state Markov channel model, the steady

state probability $\rho = 0.1, 0.5$ and 0.9 , and the correlation coefficient $\omega = 0.1, 0.5$, and 0.9 .

For all the 9 settings, we find that the probability distributions learned by the algorithm when it converges are the same in every simulation. They all converge to the probability distribution $[0.98125, 0.00125, 0.00125, \dots]$ (after sorting in the descending order of the channel selection probabilities). This is similar to the $(1 + \epsilon)$ -approximation policy in [12], which the $\epsilon = 0.05732$. In Figure 1, we plot the channel selection probability $p_i(t)$ with $\rho = 0.1, 0.5, 0.9$ and $\omega = 0.5$. Each curve (marked with various colors) in this figure corresponds to the channel selection probability of a channel with respect to time.

To see whether the **Exp3** algorithm converges to a good blind rendezvous policy, we measure the ETTR for the blind rendezvous policy with the channel selection probabilities with $p_1 = 0.98125$ and $p_i = 0.00125, i = 2, \dots, 16$ and compare that with the 6 blind rendezvous policies described in Subsection II-B. The ETTRs are obtained by averaging over 1000 independent runs for these blind rendezvous policies. In Table I, we show the comparison results for the 9 settings. For the fast time-varying channel model, i.e., $\omega \approx 0$, the optimal policy is the single channel policy [12]. For the slow time-varying channel model, i.e., $\omega \approx 1$, the $(1 + \epsilon)$ -approximation policy has the asymptotic approximation ratio $1 + \epsilon$ [12]. From these numerical results, The ETTRs of **Exp3** are comparable to the best blind rendezvous policy among the 6 blind rendezvous policies described in Subsection II-B for all the 9 settings.

Policy \ ω	0.1	0.5	0.9
Single	11.097	18.325	81.849
Uniform	156.968	156.007	159.818
Harmonic	74.290	79.734	100.212
$1 + \epsilon$	12.041	19.865	92.220
Square	23.572	29.714	81.369
Sqrt	134.378	134.256	144.121
Exp3	11.480	17.594	87.198

(a) $\rho = 0.1$

Policy \ ω	0.1	0.5	0.9
Single	2.089	2.884	10.724
Uniform	32.060	33.599	32.591
Harmonic	14.958	14.619	17.665
$1 + \epsilon$	2.449	3.459	11.565
Square	4.485	5.471	10.603
Sqrt	25.062	26.952	27.184
Exp3	2.282	2.957	10.616

(b) $\rho = 0.5$

Policy \ ω	0.1	0.5	0.9
Single	1.130	1.228	2.256
Uniform	17.994	17.477	17.515
Harmonic	7.894	7.727	8.271
$1 + \epsilon$	1.280	1.368	2.150
Square	2.735	2.661	3.280
Sqrt	15.173	14.748	13.678
Exp3	1.148	1.265	2.249

(c) $\rho = 0.9$

Table I: Comparisons of the ETTRs of various blind rendezvous policies.

To show the effectiveness of Algorithm 1, we measure the ETTRs for the channel selection probabilities $p_i(t), i = 1, 2, \dots, 16$ learned at time t (by averaging over 1000 independent runs). In Figure 2, we plot the ETTRs as a function of t . As shown in this figure, all the ETTR curves are decreasing in time. This shows that Algorithm 1 is indeed learning better blind rendezvous policies with respect to time. One notable difference in these 9 settings is the convergence time of the algorithm. When ρ is small, the probability that a channel is in a good state is also small. Hence, it is difficult to receive a reward for each channel. As such, it is more difficult to learn when ρ is small and that leads to a longer convergence time. Moreover, we note that the fluctuation of ETTRs is much larger when $\omega = 0.9$ (the yellow curves). The intuition behind this is that the channel state changes slowly when ω is large. But when a channel changes its state, it will take some time for the algorithm to learn such a change of states.

Even though the channel states are not directly observable, they can be implicitly learned. This is an additional advantage of Algorithm 1. Instead of assuming that all the channels are identically distributed in [12], we consider the setting with 10 channels that have different values of being in a good state. Specifically, We set $\rho_1 = 0, \rho_2 = 0.1, \dots, \rho_{10} = 0.9$. In Figure 3, we show the channel selection probability $p_i(t)$ with $\omega = 0.1, 0.5, 0.9$, respectively. As shown in this figure, Algorithm 1 learns that channel 10 is the best channel in the long run and sticks to that channel with the probability 0.982. The channel selection probabilities for all the other channels are 0.002.

V. CONCLUSION

In this paper, we proposed a reinforcement learning approach for the multichannel rendezvous problem. When the channel states are not observable, we showed that the **Exp3** algorithm is very effective and yields comparable ETTRs when comparing to various approximation policies in the literature. One future work is to extend the reinforcement learning approach to the setting where the channel states are either observable or partially observable. In that setting, we need to develop effective Q -learning algorithms.

REFERENCES

- [1] Z. Gu, Y. Wang, Q.-S. Hua, and F. C. M. Lau, *Rendezvous in Distributed Systems: Theory, Algorithms and Applications*. Springer, 2017.
- [2] M. J. Abdel-Rahman, H. Rahbari, and M. Krunz, "Multicast rendezvous in fast-varying DSA network," *IEEE Transactions on Mobile Computing*, vol. 14, no. 7, pp. 1449–1462, 2015.
- [3] K. Bian, J.-M. Park, and R. Chane, "A quorum-based framework for establishing control channels in dynamic spectrum access networks," *ACM MobiCom'09*, 2009.
- [4] D. Yang, J. Shin, and C. Kim, "Deterministic rendezvous scheme in multichannel access networks," *Electronics Letters*, vol. 46, no. 20, pp. 1402–1404, 2010.
- [5] N. C. Theis, R. W. Thomas, and L. A. DaSilva, "Rendezvous for cognitive radios," *IEEE Transactions on Mobile Computing*, vol. 10, no. 2, pp. 216–227, 2011.
- [6] Z. Lin, H. Liu, X. Chu, and Y.-W. Leung, "Jump-stay based channel-hopping algorithm with guaranteed rendezvous for cognitive radio networks," in *Proc. IEEE INFOCOM 2011*.

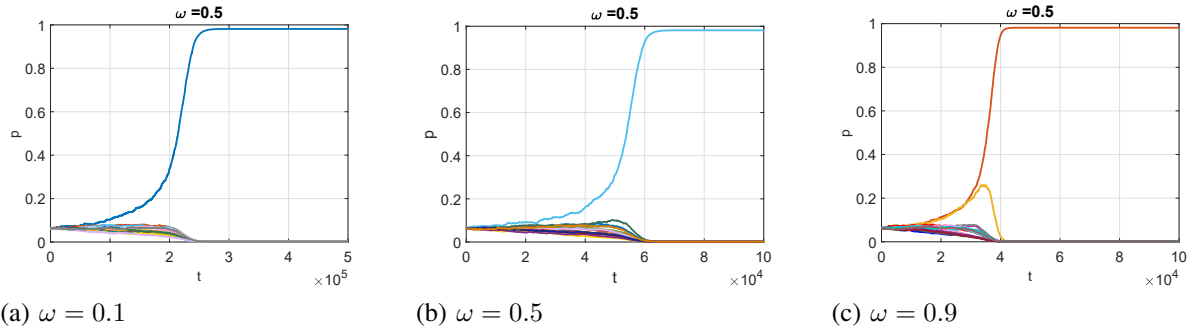


Figure 1: The channel selection probability $p_i(t)$ with $\rho = 0.1, 0.5, 0.9$ and $\omega = 0.5$.

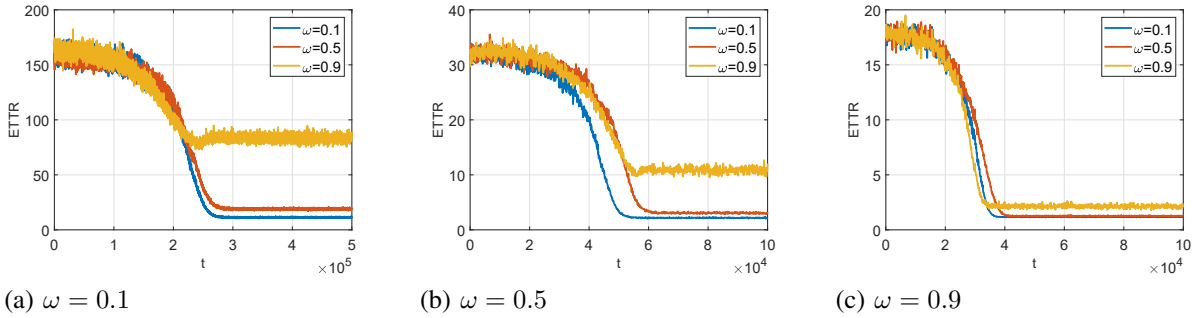


Figure 2: The ETTRs (as a function of t) with $\rho = 0.1, 0.5, 0.9$, respectively.

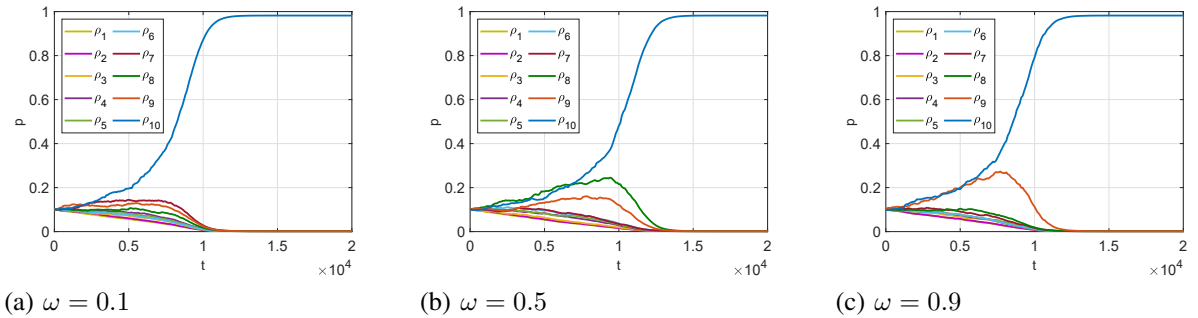


Figure 3: The channel selection probability $p_i(t)$ with $\omega = 0.1, 0.5, 0.9$, respectively.

- [7] Z. Gu, Q.-S. Hua, Y. Wang, and F. C. M. Lau, "Nearly optimal asynchronous blind rendezvous algorithm for cognitive radio networks," in *Proc. IEEE SECON*, 2013.
- [8] G.-Y. Chang and J.-F. Huang, "A fast rendezvous channel-hopping algorithm for cognitive radio networks," *IEEE Communications Letters*, vol. 17, no. 7, pp. 1475–1478, 2013.
- [9] G.-Y. Chang, W.-H. Teng, H.-Y. Chen, and J.-P. Sheu, "Novel channel-hopping schemes for cognitive radio networks," *IEEE Transactions on Mobile Computing*, vol. 13, pp. 407–421, Feb. 2014.
- [10] Z. Gu, Q.-S. Hua, and W. Dai, "Fully distributed algorithm for blind rendezvous in cognitive radio networks," in *Proc. ACM MobiHoc*, pp. 155–164, 2014.
- [11] C.-S. Chang, C.-Y. Chen, D.-S. Lee, and W. Liao, "Efficient encoding of user IDs for nearly optimal expected time-to-rendezvous in heterogeneous cognitive radio networks," *IEEE/ACM Transactions on Networking*, vol. 25, no. 6, pp. 3323–3337, 2017.
- [12] C.-S. Chang, D.-S. Lee, Y.-L. Lin, and J.-H. Wang, "ETTR Bounds and Approximation Solutions of Blind Rendezvous Policies in Cognitive Radio Networks with Random Channel States," *arXiv e-prints*, p. arXiv:1906.10424, Jun 2019.
- [13] R. S. Sutton and A. G. Barto, *Reinforcement learning: an introduction*. Massachusetts: The MIT Press, 2012.
- [14] N. C. Luong, D. T. Hoang, S. Gong, D. Niyato, P. Wang, Y. C. Liang, and D. I. Kim, "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Communications Surveys & Tutorials*, 2019.
- [15] J. C. Gittins, K. D. Glazebrook, R. Weber, and R. Weber, *Multi-armed bandit allocation indices*. Wiley Online Library, 1989, vol. 25.
- [16] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2002.
- [17] S. Gold and A. Rangarajan, "Softmax to softassign: Neural network algorithms for combinatorial optimization," *Journal of Artificial Neural Networks*, vol. 2, no. 4, pp. 381–399, 1996.
- [18] L. S. Shapley, "Stochastic games," *Proceedings of the National Academy of Sciences*, vol. 39, no. 10, pp.1095–1100, 1953.