

Ultra-Reliable Communication in 5G mmWave Networks: A Risk-Sensitive Approach

Trung Kien Vu, *Student Member, IEEE*, Mehdi Bennis, *Senior Member, IEEE*,
M  rouane Debbah, *Fellow, IEEE*, Matti Latva-aho, *Senior Member, IEEE*,
and Choong Seon Hong, *Senior Member, IEEE*

Abstract—In this letter, we investigate the problem of providing gigabit wireless access with reliable communication in 5G millimeter-Wave (mmWave) massive multiple-input multiple-output (MIMO) networks. In contrast to the classical network design based on average metrics, we propose a distributed *risk-sensitive* reinforcement learning-based framework to jointly optimize the beamwidth and transmit power, while taking into account the sensitivity of mmWave links due to blockage. Numerical results show that our proposed algorithm achieves more than 9 Gbps of user throughput with a guaranteed probability of 90%, whereas the baselines guarantee less than 7.5 Gbps. More importantly, there exists a rate-reliability-network density tradeoff, in which as the user density increases from 16 to 96 per km², the fraction of users that achieve 4 Gbps are reduced by 11.61% and 39.11% in the proposed and the baseline models, respectively.

Index Terms—URLLC, URC, reliable communication, mmWave communications, risk-sensitive learning, reinforcement learning.

I. INTRODUCTION

To enable gigabit wireless access with reliable communication, a number of candidate solutions are currently investigated for 5G: 1) higher frequency spectrum, e.g., millimeter wave (mmWave); 2) advanced spectral-efficient techniques, e.g., massive multiple-input multiple-output (MIMO); and 3) ultra-dense small cells [1]. This work explores the above techniques to enhance the wireless access [1]–[3]. Massive MIMO yields remarkable properties such as high signal-to-interference-plus-noise ratio due to large antenna gains, and extreme spatial multiplexing gain [3], [4]. Specially, mmWave frequency bands offer huge bandwidth [5], while it allows for packing a massive antennas for highly directional beamforming [5]. A unique peculiarity of mmWave is that mmWave links are very sensitive to blockage, which gives rise to unstable connectivity and unreliable communication [5]. To overcome such challenge,

Manuscript received Dec 29, 2017; revised Jan 20, 2018; accepted Jan 31, 2018. Date of publication Feb 06, 2018; This work was supported in part by Tekes, Nokia, Huawei, MediaTek, Keysight, Bittium, Kyynel, in part by the Academy of Finland via the grant 307492 and the CARMA grants 294128 and 289611. The associate editor coordinating the review of this letter and approving it for publication was A. S. Cacciapuoti. (Corresponding author: Trung Kien Vu.)

T. K. Vu and M. Latva-aho are with the Centre for Wireless Communications, University of Oulu, Oulu 90014, Finland (e-mail: trungkien.vu@oulu.fi; matti.latva-aho@oulu.fi).

M. Bennis is with the Centre for Wireless Communications, University of Oulu, 90014 Oulu, Finland, and also with the Department of Computer Science and Engineering, Kyung Hee University, Yongin 446-701, South Korea (e-mail: mehdi.bennis@oulu.fi).

M. Debbah is with the Large Networks and System Group (LANEAS), CentraleSup  lec, Universit   Paris-Saclay, 91192 Gif-sur-Yvette, France and also with the Mathematical and Algorithmic Sciences Laboratory, Huawei France R&D, 92100 Paris, France (e-mail: merouane.debbah@huawei.com).

C. S. Hong is with the Department of Computer Science and Engineering, Kyung Hee University, Yongin 446-701, South Korea (email: cshong@khu.ac.kr).

Citation information: 10.1109/LCOMM.2018.2802902, IEEE Communications Letters

we leverage principles of risk-sensitive reinforcement learning (RSL) and exploit the multiple antennas diversity and higher bandwidth to optimize transmission to achieve gigabit data rates, while considering the sensitivity of mmWave links to provide ultra-reliable communication (URC). The prime motivation behind using RSL stems from the fact that the *risk-sensitive* utility function to be optimized is a function of not only the average but also the variance [6], and thus it captures the tail of rate distribution to enable URC. While our proposed algorithm is fully distributed, which does not require full network observation, and thus the cost of channel estimation and signaling synchronization is reduced. Via numerical experiments, we showcase the inherently key trade-offs between (i) reliability/data rates and network density, and (ii) availability and network density.

Related work: In [7], [8] authors provided the principles of ultra-reliable and low latency communication (URLLC) and described some techniques to support URLLC. Recently, the problem of low latency communication [9] and URLLC [10], [11] for 5G mmWave network was studied to evaluate the performance under the impact of traffic dispersion and network densification. Moreover, a reinforcement learning (RL) approach to power control and rate adaptation was studied in [12]. All these works focus on maximizing the time average of network throughput or minimizing the mean delay without providing any guarantees for higher order moments (e.g., variance, skewness, kurtosis, etc.). In this work, we depart from the classical average-based system design and instead take into account higher order moments in the utility function to formulate a RSL framework through which every small cell optimizes its transmission while taking into account signal fluctuations.

II. SYSTEM MODEL

Let us consider a mmWave downlink (DL) transmission of a small cell network consisting of a set \mathcal{B} of B small cells (SCs), and a set \mathcal{K} of K user equipments (UEs) equipped with N_k antennas. We assume that each SC is equipped with a large number of N_b antennas to exploit massive MIMO gain and adopt a hybrid beamforming architecture [13], and assume that $N_b \gg N_k \geq 1$. Without loss of generality, one UE per one SC is considered¹. The data traffic is generated from SC to UE via mmWave communication. A co-channel time-division duplexing protocol is considered, in which the DL channel can be obtained via the uplink training phase.

Each SC adopts the hybrid beamforming architecture, which enjoys both analog and digital beamforming techniques [13]. Let $g_{bk}^{(tx)}$ and $g_{bk}^{(rx)}$ denote the analog transmitter and receiver beamforming gains at the SC b and UE k , respectively. In addition, we use $\omega_{bk}^{(tx)}$ and $\omega_{bk}^{(rx)}$ to represent

¹For the multiple UEs case, addition channel estimation and user scheduling need to be considered, one example was studied in [3].

the angles deviating from the strongest path between the SC b and UE k . Also, let $\theta_{bk}^{(tx)}$ and $\theta_{bk}^{(rx)}$ denote the beamwidth at the SC and UE, respectively. We denote $\boldsymbol{\theta}$ as a vector of the transmitter beamwidth of all SCs. We adopt the widely used antenna radiation pattern model [13] to determine the analog beamforming gain as

$$g_{bk}(\omega_{bk}, \theta_{bk}) = \begin{cases} \frac{2\pi - (2\pi - \theta_{bk})\eta}{\theta_{bk}}, & \text{if } |\omega_{bk}| \leq \frac{\theta_{bk}}{2}, \\ \eta, & \text{otherwise,} \end{cases} \quad (1)$$

where $0 < \eta \ll 1$ is the side lobe gain.

Let $\mathbf{h}_{bk} \in \mathbb{C}^{N_b \times N_k}$ denote the channel state from the SC b to UE k . We assume a time-varying channel state described by a Markov chain and there are $T \in \mathbb{Z}^+$ states, i.e., for each $\mathbf{h}_{bk}(t), t = \{1, \dots, T\}$. Considering imperfect channel state information (CSI), the estimated channel state between the SC b and UE k is modeled as [10]

$$\hat{\mathbf{h}}_{bk} = \sqrt{N_b \times N_k} \Theta_{bk}^{1/2} (\sqrt{1 - \tau_k^2} \mathbf{w}_{bk} + \tau_k \hat{\mathbf{w}}_{bk}),$$

where $\Theta_{bk} \in \mathbb{C}^{N_b \times N_b}$ is the spatial channel correlation matrix that accounts for path loss and shadow fading. Here, $\mathbf{w}_{bk} \in \mathbb{C}^{N_b \times N_k}$ is the small-scale fading channel matrix, modeled as a random matrix with zero mean and variance of $\frac{1}{N_b \times N_k}$. Here $\tau_k \in [0, 1]$ reflects the estimation accuracy for UE k , if $\tau_k = 0$, we assume that perfect channel state information. $\hat{\mathbf{w}}_{bk} \in \mathbb{C}^{N_b \times N_k}$ is the estimated noise vector, also modeled as a random matrix with zero mean and variance of $\frac{1}{N_b \times N_k}$. We denote $\mathbf{h} = \{\mathbf{h}_{bk} | \forall b \in \mathcal{B}, \forall k \in \mathcal{K}\}$ as the network state.

By applying a linear precoding scheme $\mathbf{f}_{bk}(\hat{\mathbf{h}}_{bk})$ [13], i.e., $\mathbf{f}_{bk}(\hat{\mathbf{h}}_{bk}) = \hat{\mathbf{h}}_{bk}$ for the conjugate precoding, the achievable rate² of UE k from SC b can be calculated as

$$r_b(t) = W \log_2 \left(1 + \frac{p_b g_{bk}^{(tx)} g_{bk}^{(rx)} |\mathbf{h}_{bk}^\dagger \mathbf{f}_{bk}|^2}{\sum_{b' \neq b} p_{b'} g_{b'k}^{(tx)} g_{b'k}^{(rx)} |\mathbf{h}_{b'k}^\dagger \mathbf{f}_{b'k}|^2 + \eta_{bk}} \right),$$

where p_b and $p_{b'}$ are the transmit powers of SC b and SC b' , respectively. In addition, W denotes the system bandwidth of the mmWave frequency band. The thermal noise of user k served by SC b is η_{bk} . Here, we denote P_b^{\max} as the maximum transmit power of SC b and $\mathbf{p} = (p_b | \forall b \in \mathcal{B}, 0 \leq p_b \leq P_b^{\max})$ as the transmit power vector.

III. PROBLEM FORMULATION

We model a decentralized optimization problem and harness tools from RSL to solve, whereby SCs autonomously respond to the network states based on the historical data. Let us consider a joint optimization of transmitter beamwidth³ $\boldsymbol{\theta}$ and transmit power allocation \mathbf{p} . We denote $\mathbf{z}(t) = (\boldsymbol{\theta}(t), \mathbf{p}(t))$, which takes values in $\mathcal{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_B\}$, where $\mathbf{z}_b = (\theta_b, p_b)$. Assume that each SC b selects its beamwidth and transmit power drawn from a given probability distribution $\boldsymbol{\pi}_b = (\pi_b^1, \dots, \pi_b^m, \dots, \pi_b^{Z_b})$ in which Z_b is the cardinality of the set of all combinations (θ_b, p_b) , i.e., $\sum_{m=1}^{Z_b} \pi_b^m = 1$. For each $m = \{1, \dots, Z_b\}$ and $\mathbf{z}_b^m = (\theta_b^m, p_b^m)$ the mixed-strategy probability is defined as

$$\pi_b^m(t) = \Pr(\mathbf{z}_b(t) = \mathbf{z}_b^m | \mathbf{z}_b(0:t-1), \boldsymbol{\pi}_b(0:t-1)). \quad (2)$$

²Note that we omit the beam search/track time, since it can be done in a short time as compared with transmission time [14]. We assume that each BS sends a single stream to its users via the main beams.

³As studied in [13], for $\eta \leq \frac{1}{2}$, the problem of selecting beamwidth for the transmitter and receiver can be done by adjusting the transmitter beamwidth with a fixed receiver beamwidth.

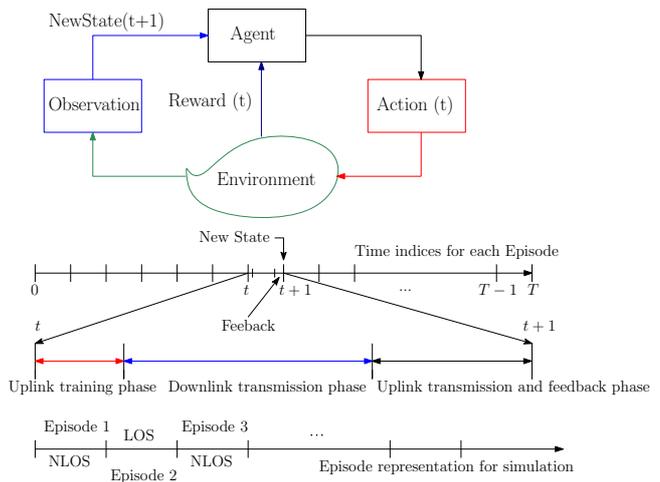


Fig. 1. A reinforcement learning model.

We denote $\boldsymbol{\pi} = \{\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_b, \dots, \boldsymbol{\pi}_B\} \in \Pi$, in which Π is the set of all possible probability mass functions (PMF). Let $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_B)$ denote the instantaneous rates, in which $\mathbf{r}_b = (r_b(0), \dots, r_b(T))$. Let \mathcal{R} denote the rate region, which is defined as the convex hull of the rates [15], i.e., $\mathbf{r} \in \mathcal{R}$. Inspired by the RSL [6], we consider the following utility function, given by

$$\bar{u}_b = \frac{1}{\mu_b} \log \mathbb{E}_{\mathbf{h}, \boldsymbol{\pi}} \left[\exp(\mu_b \sum_{t=0}^T r_b(t)) \right], \quad (3)$$

where the parameter $\mu_b < 0$ denotes the desired risk-sensitivity, which will *penalize the variability* [6] and the operator \mathbb{E} denotes the expectation operation.

Remark 1: The Taylor expansion of the utility function given in (3) yields

$$\bar{u}_b \triangleq \mathbb{E}_{\mathbf{h}, \boldsymbol{\pi}} \left[\sum_{t=0}^T r_b(t) \right] + \frac{\mu_b}{2} \text{Var}_{\mathbf{h}, \boldsymbol{\pi}} \left[\sum_{t=0}^T r_b(t) \right] + \mathcal{O}(\mu_b^2).$$

Remark 1 basically shows that the utility function (3) considers both mean and variance terms (Var) of the mmWave links. We formulate the following distributed optimization problem for every SC as:

$$\max_{\boldsymbol{\pi}_b} \quad \frac{1}{\mu_b} \log \mathbb{E}_{\mathbf{h}, \boldsymbol{\pi}_b} \left[\exp(\mu_b \sum_{t=0}^T r_b(t)) \right] \quad (4a)$$

$$\text{subject to} \quad \mathbf{r}_b \in \mathcal{R}, \boldsymbol{\pi}_b \in \Pi, p_b \leq P_b^{\max}. \quad (4b)$$

It is challenging to solve (4) if each SC does not have full network observation. This work does not assume an explicit knowledge of the state transition probabilities. Here, we leverage principles of RL to optimize the transmit beam in a totally decentralized manner [6], [12], [16].

IV. PROPOSED ALGORITHM

In Fig. 1 each SC acts as an agent which selects an action to maximize a long-term reward based on user feedback and probability distribution for each action. The action is defined as the selection of \mathbf{z}_b , while the long-term utility in (4) is the reward, and the environment here contains the network state. To this end, we build the probability distribution for every action and provide a RL procedure to solve (4). We denote $u_b^m = u_b^m(\mathbf{z}_b^m, \mathbf{z}_{-b})$ as a utility function of SC b when selecting \mathbf{z}_b^m . Here, \mathbf{z}_{-b} denotes the composite variable of other agents' actions excluding SC b . From (3),

the utility $u_b(t)$ of SC b at time slot t , i.e., $\bar{u}_b = \sum_{t=0}^T u_b(t)$, is rewritten as

$$u_b(t) = \frac{1}{\mu_b} \log \left(\sum_{m=1}^{Z_b} \pi_b^m \exp(\mu_b r_b^m(\mathbf{z}_b^m(t), \mathbf{z}_{-b})) \right), \quad (5)$$

where $r_b^m(\mathbf{z}_b^m(t), \mathbf{z}_{-b})$ is the instantaneous rate of SC b when choosing $\mathbf{z}_b^m(t) = (\theta_b^m(t), p_b^m(t))$ with probability $\pi_b^m(t)$.

Remark 2: For a small μ_b , (3) is approximated via the Taylor approximation⁴ of \bar{r}_b around $\mu_b \rightarrow 0$ as

$$\bar{u}_b = \frac{1}{\mu_b} \mathbb{E} \left[\sum_{t=0}^T (\exp(\mu_b r_b(t)) - 1) \right], \quad (6)$$

$$= \frac{1}{(T+1)} \sum_{t=0}^T \frac{\exp(\mu_b r_b(t)) - 1}{\mu_b}, \quad (7)$$

where (7) is obtained by expanding the time average of (6). Each SC determines (θ_b^m, p_b^m) from \mathcal{Z}_b based on the probability distribution from the previous stage $t-1$, i.e.,

$$\pi_b(t-1) = (\pi_b^1(t-1), \dots, \pi_b^{Z_b}(t-1)). \quad (8)$$

We introduce the Boltzmann-Gibbs distribution to capture the exploitation and exploration, $\beta_b(\mathbf{u}_b(t))$, given by

$$\beta_b^m(\mathbf{u}_b(t)) = \operatorname{argmax}_{\pi_b \in \Pi} \sum_{m \in \mathcal{Z}_b} [\pi_b^m u_b^m(t) - \kappa_b \pi_b^m \ln(\pi_b^m)], \quad (9)$$

where $\mathbf{u}_b(t) = (u_b^1(t), \dots, u_b^{Z_b}(t))$ is the utility vector of SC b for $\mathbf{z}_b \in \mathcal{Z}_b$, and the trade-off factor κ_b is used to balance between exploration and exploitation. If κ_b is small, the SC selects \mathbf{z}_b with highest payoff. For $\kappa_b \rightarrow \infty$ all decisions have equal chance.

For a given $\mathbf{u}_b(t)$ and κ_b , we solve (9) to find the probability distribution, by adopting the notion of logit equilibrium [16], we have

$$\beta_b^m(\mathbf{u}_b(t)) = \frac{\exp\left(\frac{1}{\kappa_b} [u_b^m]^+\right)}{\sum_{m' \in \mathcal{Z}_b} \exp\left(\frac{1}{\kappa_b} [u_b^{m'}]^+\right)}, \quad (10)$$

where $[x]^+ \equiv \max[x, 0]$. Finally, we propose two coupled RL processes that run in parallel and allow SCs to decide their optimal strategies at each time instant t as follows [16].

Risk-Sensitive Learning procedure: We denote $\hat{u}_b(t)$ as the estimate utility of SC b , in which the estimate utility and probability mass function are updated for each action $m \in \mathcal{Z}_b$ as follows:

$$\begin{cases} \hat{u}_b^m(t) = \hat{u}_b^m(t-1) + \\ \quad \zeta_b(t) \mathbb{I}_{\{\mathbf{z}_b(t) = \mathbf{z}_b^m\}} \times (u_b(t-1) - \hat{u}_b^m(t-1)), \\ \pi_b^m(t) = \pi_b^m(t-1) + \iota_b(t) (\beta_b^m(\mathbf{u}_b(t)) - \pi_b^m(t-1)), \end{cases}$$

where $\zeta_b(t)$ and $\iota_b(t)$ are the learning rates which satisfy the following conditions (due to space limits please see [16] for convergence proof):

$$\begin{cases} \lim_{T \rightarrow \infty} \sum_{t=0}^T \zeta_b(t) = +\infty, \lim_{T \rightarrow \infty} \sum_{t=0}^T \iota_b(t) = +\infty. \\ \lim_{T \rightarrow \infty} \sum_{t=0}^T \zeta_b^2(t) = +\infty, \lim_{T \rightarrow \infty} \sum_{t=0}^T \iota_b^2(t) = +\infty. \\ \lim_{t \rightarrow \infty} \frac{\iota_b(t)}{\zeta_b(t)} = 0. \end{cases}$$

Finally, each SC determines \mathbf{z}_b^m as per (8).

⁴For a small $x > 0$, the Taylor approximation of $\log(x)$ is $x - 1$.

V. NUMERICAL RESULTS

A dense SCs are randomly deployed in a 0.5×0.5 km² area and we assume one UE per each SC and a fixed user association. We assume that each SC adjusts its beamwidth with a step of 0.02 radian from the range $[\theta^{\min}, \theta^{\max}]$, where $\theta^{\min} = 0.2$ radian and $\theta^{\max} = 0.4$ radian denote the minimum and maximum beamwidths of each SC, respectively. The transmit power level set of each SC is $\{21, 23, 25\}$ dBm and the SC antenna gain is 5 dBi. The number of transmit antennas N_b and receive antennas N_k at the SC and UE are set to 64 and 4, respectively. The blockage is modeled as a distance-dependent probability state where the channel is either line-of-sight (LOS) or non-LOS for urban environments at 28 GHz and the system bandwidth is 1 GHz [17]. Numerical results are obtained via Monte-Carlo simulations over 50 different random topologies. The risk-sensitive parameter is set to $\mu_b = -2$. For the learning algorithm, the trade-off factor κ_b is set to 5, while the learning rates $\zeta_b(t)$ and $\iota_b(t)$ are set to $\frac{1}{(t+1)^{0.55}}$ and $\frac{1}{(t+1)^{0.6}}$, respectively [16]. Furthermore, we compare our proposed **RSL** scheme with the following baselines:

- **Classical Learning (CSL)** refers to the RL framework in which the utility function only considers the mean value of mmWave links [16].
- **Baseline 1 (BL1)** refers to [13] optimizing the beamwidth with maximum transmit power.

In Fig. 3, we plot the complementary cumulative distribution function (tail distribution - CCDF) of user throughput (UT) at 28 GHz when the number of SCs is 24 per km². The CCDF curves reflect the reliable probability (in both linear and logarithmic scales), defined as the probability that the UT is higher than a target rate r_0 Gbps, i.e., $\Pr(\text{UT} \geq r_0)$. We also study the impact of imperfect CSI with $\tau_k = 0.3$ and feedback with noise from UEs. We observe that the performance of our proposed RSL framework is reduced under these impacts. We next compare our proposed RSL method with other baselines with perfect CSI and user feedback. It is observed that the **RSL** scheme achieves better reliability, $\Pr(\text{UT} \geq 10$ Gbps), of more than 85%, whereas the baselines **CSL** and **BL1** obtain less than 75% and 65%, respectively. However, at very low rate (less than 2 Gbps) or very high rate (10.65 – 11 Gbps) captured by the cross-point, the **RSL** obtains a lower probability as compared to the baselines. In other words, our proposed solution provides a UT which is more concentrated around its median in order to provide uniformly great service for all users. For instance, the UT distribution of our proposed algorithm has a small variance of 0.4846, while the **CSL** has a higher variance of 2.6893.

A. Impact of network density

Fig. 2 reports the impact of network density on the reliability, which is defined as the fraction of UEs who achieve a given target rate r_0 , i.e., $\frac{K_{r>r_0}}{K}$. Here, the number of SCs is varying from 16 to 128 per km². For given target rates of 2, 3, and 4 Gbps, our proposed algorithm guarantees higher *reliability* as compared to the baselines. Moreover, the higher the target rate, the bigger the performance gap between our proposed algorithm and the baselines. A linear increase in network density decreases reliability, for example, when the density increases from 16 to 96, the fraction

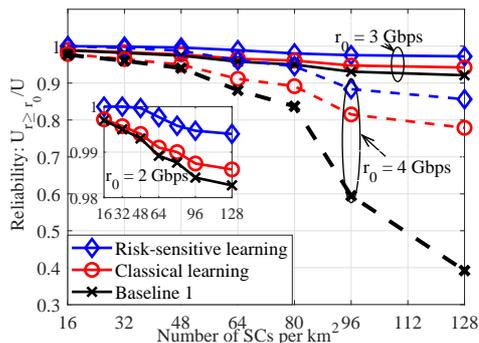


Fig. 2. Reliability versus network density.

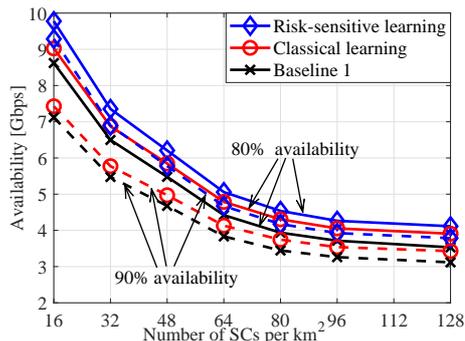


Fig. 4. Availability versus network density.

of users that achieve 4 Gbps of the **RSL**, **CSL**, and **BL1** are reduced by 11.61%, 16.72%, and 39.11%, respectively. This highlights a key tradeoff between reliability and network density.

In Fig. 4 we show the impact of network density on the availability, which defines how much rate is obtained for a target probability. We plot the 80% and 90% probabilities in which the system achieves a rate of at least r Gbps. For a given target probability of 90%, our proposed algorithm guarantees more than 9 Gbps of UT, whereas the baselines guarantee less than 7.5 Gbps of UT for $B = 16$, while if we lower the target probability to 80%, the achievable rate is increased by 5%. This gives rise to a tradeoff between the reliability and the data rate. In addition, for a given probability, the achievable rate r is reduced with the increase in network density. For instance, when the network density increases from 16 to 80, the achievable rate is reduced by 50%. This highlights the tradeoff between availability and network density.

We numerically observe that $T = 4000$ is long enough for agents to learn and enjoy the optimal solution. We assume that the channel condition is changed after every $T = 4000$. Our proposed algorithm converges faster than the classical learning baseline as shown in Fig. 5. By harnessing the notion of risk-averse, the agents try to find the best strategy subject to the variations of the mmWave rates.

VI. CONCLUSIONS

In this letter, we studied the problem of providing multi-gigabit wireless access with reliable communication by optimizing the transmit beam and considering the link sensitivity in 5G mmWave networks. A distributed risk-sensitive RL based approach was proposed taking into account both mean and variance values of the mmWave

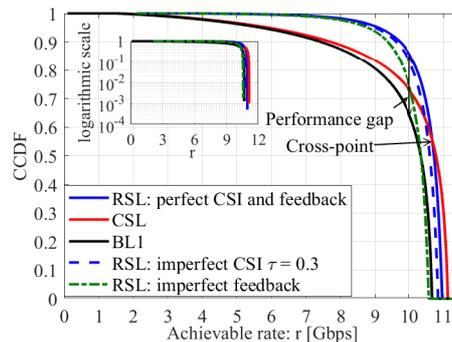
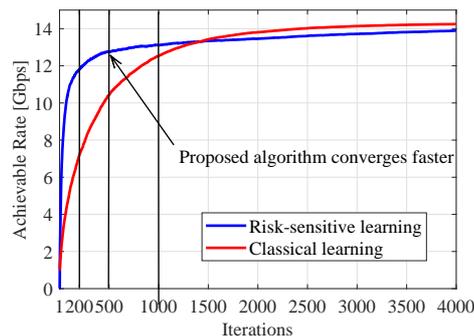
Fig. 3. Tail distribution of the achievable rate, $B = 24$.

Fig. 5. Convergence of the proposed RSL and classical RL.

links. Numerical results show that our proposed approach provides better services for all users. For instance, our proposed approach achieves a $\Pr(UT \geq 10\text{Gbps})$ is higher than 85%, whereas the baselines obtain less than 75% and 65% with 24 small cells.

REFERENCES

- [1] J. G. Andrews et al., "What Will 5G Be?" *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 6, pp. 1065–1082, June 2014.
- [2] A. Anpalagan, M. Bennis, and R. Vannithamby, *Design and Deployment of Small Cell Networks*. Cambridge University Press, 2015.
- [3] T. K. Vu et al., "Joint load balancing and interference mitigation in 5G heterogeneous networks," *IEEE Transactions on Wireless Communications*, vol. 16, no. 9, pp. 6032–6046, Sep. 2017.
- [4] Y. Wu, R. Schober, D. W. K. Ng, C. Xiao, and G. Caire, "Secure massive MIMO transmission with an active eavesdropper," *IEEE Transactions on Information Theory*, vol. 62, no. 7, pp. 3880–3900, 2016.
- [5] T. S. Rappaport et al., "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.
- [6] O. Mihatsch and R. Neuneier, "Risk-sensitive reinforcement learning," *Machine learning*, vol. 49, no. 2-3, pp. 267–290, 2002.
- [7] P. Popovski et al., "Wireless access for ultra-reliable low-latency communication (urllc): Principles and building blocks," *submitted to IEEE Network*, 2017.
- [8] M. Bennis, M. Debbah, and H. V. Poor, "Ultra-Reliable and Low-Latency Wireless Communication: Tail, Risk and Scale," *submitted to Proceedings of the IEEE*, 2018.
- [9] G. Yang, M. Xiao, and H. V. Poor, "Low-latency millimeter-wave communications: Traffic dispersion or network densification?" *submitted to IEEE Transactions on Communication*, 2017.
- [10] T. K. Vu et al., "Ultra-reliable and low latency communication in mmwave-enabled massive MIMO networks," *IEEE Communications Letters*, vol. 21, no. 9, pp. 2041–2044, Sep. 2017.
- [11] —, "Path Selection and Rate Allocation in Self-Backhauled mmWave Networks," in *Proc. IEEE Int. Conf. on Wireless Communications and Networking Conference (WCNC)*, Barcelona, Spain, 2018, pp. 1–6.
- [12] E. Ghadimi, F. D. Calabrese, G. Peters, and P. Soldati, "A reinforcement learning approach to power control and rate adaptation in cellular networks," in *2017 IEEE International Conference on Communications*, Paris, France, 2017, pp. 1–7.

- [13] J. Liu and E. S. Bentley, "Hybrid-Beamforming-Based Millimeter-Wave Cellular Network Optimization," in *Proc. 15th IEEE Int. Sym. on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Paris, France, 2017, pp. 1–8.
- [14] J. Palacios et al., "Tracking mm-Wave Channel Dynamics: Fast Beam Training Strategies under Mobility," in *Proc. 36th Annual IEEE Int. Conf. on Computer Communications (INFOCOM)*, Atlanta, GA, USA, 2017, pp. 1–9.
- [15] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [16] M. Bennis, S. M. Perlaza, P. Blasco, Z. Han, and H. V. Poor, "Self-organization in small cell networks: A reinforcement learning approach," *IEEE Transactions on Wireless Communications*, vol. 12, no. 7, pp. 3202–3212, 2013.
- [17] T. Bai, V. Desai, and R. W. Heath, "Millimeter wave cellular channel models for system evaluation," in *2014 IEEE Int. Conf. on Computing, Networking and Communications (ICNC)*, Honolulu, HI, USA, 2014, pp. 178–182.