

Robust and Reliable Stochastic Resource Allocation via Tail Waterfilling

Gokberk Yaylali and Dionysis Kalogerias
Department of EE – Yale University

Abstract—Stochastic allocation of resources in the context of wireless systems ultimately demands reactive decision making for meaningfully optimizing network-wide random utilities, while respecting certain resource constraints. Standard ergodic-optimal policies are however susceptible to the statistical variability of fading, often leading to systems which are severely unreliable and spectrally wasteful. On the flip side, minimax/outage-optimal policies are too pessimistic and often hard to determine. We propose a new *risk-aware* formulation of the resource allocation problem for standard multi-user point-to-point power-constrained communication with no cross-interference, by employing the Conditional Value-at-Risk (CV@R) as a measure of fading risk. A remarkable feature of this approach is that it is a convex generalization of the ergodic setting while inducing robustness and reliability in a fully tunable way, thus bridging the gap between the (naive) ergodic and (conservative) minimax approaches. We provide a closed-form expression for the CV@R-optimal policy given primal/dual variables, extending the classical stochastic waterfilling policy. We then develop a primal-dual *tail-waterfilling* scheme to recursively learn a globally optimal risk-aware policy. The effectiveness of the approach is verified via detailed simulations.

Index Terms—Resource Allocation, Waterfilling, Conditional Value-at-Risk (CV@R), Risk-Aware Optimization.

I. INTRODUCTION

We revisit the classical problem of allocating resources in point-to-point communication networks operating over realizations of random fading channels $\mathbf{h} \in \mathcal{H} \subseteq \mathbb{R}^n$. Resources such as transmission power and/or channel access are allotted among users to meaningfully optimize certain network-wide random utilities. Traditionally, such resources are allocated either *deterministically* by essentially disregarding the statistical variability of fading as an integral characteristic of the system, including minimax formulations [1], [2], or *stochastically* in an ergodic sense by considering performance averages [2]–[6], i.e., expectations of random network objectives in an attempt to optimize performance in the “long-term”.

However, being optimal in expectation, ergodic stochastic resource allocation lacks the ability to effectively quantify relatively infrequent though statistically significant fading events causing performance drops, e.g., deep(er) fades. Indeed, the statistical dispersion of a communication medium with a fatter-tailed distribution is quite likely to result in rather undesirable channel realizations, leading to potentially major service losses. This happens because expectations of random services do not capture such risky tail events. In other words, ergodic-optimal resource allocation policies are *risk-neutral*.

In fact, it is well-known that optimal ergodic policies are often channel-opportunistic [3], and prone to sporadic channel realizations that negatively affect performance, leading to unreliable systems suffering from substantial spectrum underutilization. On the other extent, minimax-type (sometimes called “robust”) resource policies aim for maximally reliable system performance [1], [2]. Still, such policies are known for being overcautious and for achieving conservative system performance, on top of the often unreasonable difficulty of the resulting optimization problems.

While approaches based on outage probability optimization try to bridge the two extremes [7], they exhibit counterfactual issues: Outage probability targets required for performing allocation of resources might not even be feasible to begin with, and even if they are, they might not result in operationally meaningful performance. Quantile-based resource allocation, such as outage rate/capacity optimization, aims for alleviating those issues, however the resulting problems still suffer from other limitations, mainly related to interpretability and lack of favorable structure (e.g., convexity).

In this paper, we introduce a *risk-aware* formulation of the resource allocation problem –such ideas/approaches have recently started getting traction [8]–[10]– for standard multi-user point-to-point resource-constrained communication with no cross-interference, by capitalizing on the *Conditional Value-at-Risk (CV@R)* [11] as a measure of fading risk. CV@R, deeply rooted in mathematical finance, is a *coherent risk measure* [12], trades naturally between the (naive) ergodic and (conservative) minimax settings, and allows formulating the proposed risk-aware problem as a convex, well-structured extension of its ergodic counterpart, liberated from counterfactual issues and inducing robustness and reliability in a fully tunable way.

After obtaining a closed-form expression for the optimal Lagrangian-relaxed CV@R policy, we propose the *tail waterfilling algorithm*, a primal-dual scheme to learn a globally optimal risk-aware policy in a recursive fashion. Indeed, tail waterfilling continuously extends classical stochastic waterfilling [13] to the risk-aware universe. We present detailed numerical simulations, empirically corroborating the effectiveness of our approach for two standard utilities, namely, weighted sumrate and proportional fairness.

II. SYSTEM MODEL

We consider a n -terminal parallel point-to-point communication channel model; some examples of relevant networking scenarios may be visualized as in Fig. 1. We assume perfect




Fig. 1: Examples of multi-user one-to-one communication channels. Left: Multiplexed (e.g., time or frequency) star uplink or downlink model. Right: Classical parallel channel model.

channel state information (CSI) at transmission time (mainly for simplicity), which is leveraged to allocate resources via a policy $\mathbf{p}(\mathbf{h})$, where \mathbf{h} is the channel fading vector. The rate for terminal or user $i = 1, \dots, n$ in the network is

$$r_i(p_i(h_i), h_i) \triangleq \log \left(1 + \frac{h_i p_i(h_i)}{\sigma_i^2} \right), \quad (1)$$

where $\sigma_i^2 > 0$ is the noise variance of the corresponding link. Under this setting, optimal resource allocation in an ergodic sense may be achieved by solving the convex problem [3]

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{X}, \mathbf{p} \succeq \mathbf{0}}{\text{maximize}} && f_0(\mathbf{x}) \\ & \text{subject to} && \mathbf{x} \preceq \mathbb{E}[\mathbf{r}(\mathbf{p}(\mathbf{h}), \mathbf{h})], \\ & && \|\mathbb{E}[\mathbf{p}(\mathbf{h})]\|_1 \leq P_0 \end{aligned} \quad (2)$$

where f_0 is a given concave utility, \mathbf{x} is the mean-ergodic rate vector, \mathcal{X} is a convex set, \mathbf{r} is the instantaneous rate vector, and P_0 is a total mean power budget. Problem (2) is very well-studied; in fact, a globally optimal solution may be obtained via the well-known (stochastic) *waterfilling algorithm* [13], which is the same as (stochastic) dual descent [3]. However, as mentioned in Section I, it is also known that optimal policies obtained by solving (2) are channel-opportunistic, unavoidably leading to systems which are severely unreliable and spectrally wasteful. This is because of the *risk-neutral* quantification of channel uncertainty in (2), which discards information about the higher-order or *tail* behavior of the rate as a function of random fading. On the other hand, minimax-optimal policies or policies minimizing outage probabilities are either overly pessimistic [7], or result in problems that are difficult to handle, or suffer from counterfactual issues.

To effectively address those shortcomings, we take a fundamentally distinct approach to stochastic resource allocation by replacing the expectation of rates in (2) by a (vector) *risk measure* [12], specifically the CV@R [11], defined for an integrable random *cost* z as

$$\text{CV@R}^\alpha[z] \triangleq \inf_{t \in \mathbb{R}} t + \frac{1}{\alpha} \mathbb{E}[(z - t)_+], \quad (3)$$

where $\alpha \in (0, 1]$ is the corresponding *confidence level*. CV@R is a strict and tractable generalization of expectation, because

$$\begin{aligned} \text{CV@R}^1[z] &= \mathbb{E}[z] \leq \text{CV@R}^\alpha[z], \quad \forall \alpha \in (0, 1] \text{ and} \\ \text{CV@R}^0[z] &\triangleq \lim_{\alpha \downarrow 0} \text{CV@R}^\alpha[z] = \text{ess sup } z. \end{aligned} \quad (4)$$




Fig. 2: $R_1 = -\text{CV@R}^\alpha[-z]$ (blue) and $R_2 = \text{CV@R}^\alpha[z]$ (red) calculated at level $\alpha = 0.6$, for a Rayleigh distributed z with scale $\sigma = 1$.

Intuitively, CV@R measures *expected losses restricted to the upper tail* of z of probability equal to α ; see Fig. 2. Therefore, it provides an interpretable and tunable tradeoff bridging risk-neutrality and minimax robustness.

To make CV@R suitable for maximizing rewards –cf. (2)– rather than minimizing losses, it is sufficient to reflect it as

$$-\text{CV@R}^\alpha[-z] = \sup_{t \in \mathbb{R}} t - \frac{1}{\alpha} \mathbb{E}[(t - z)_+], \quad (5)$$

now measuring *expected rewards restricted to the lower tail* of z of probability equal to α ; again, see Fig. 2 for a comparison. Using (5), we may formulate our proposed risk-aware resource allocation problem as

$$\begin{aligned} P^* &= \underset{\mathbf{x} \in \mathcal{X}, \mathbf{p} \succeq \mathbf{0}}{\text{maximize}} && f_0(\mathbf{x}) \\ & \text{subject to} && \mathbf{x} \preceq -\text{CV@R}^\alpha[-\mathbf{r}(\mathbf{p}(\mathbf{h}), \mathbf{h})], \\ & && \|\mathbb{E}[\mathbf{p}(\mathbf{h})]\|_1 \leq P_0 \end{aligned} \quad (6)$$

where \mathbf{x} is now interpreted as a *risk-ergodic rate* vector, and the vector operator $\text{CV@R}^\alpha[\cdot]$ with a confidence level vector α evaluates the risk of the corresponding rate vector in an elementwise manner. Note that we have tacitly not enforced risk-aware behavior on the policy itself (resource constraint), as this is operationally unnecessary. It is a standard exercise to show that problem (6) can be equivalently expressed as

$$\begin{aligned} P^* &= \underset{\mathbf{x} \in \mathcal{X}, \mathbf{p} \succeq \mathbf{0}, \mathbf{t}}{\text{maximize}} && f_0(\mathbf{x}) \\ & \text{subject to} && \mathbf{x} \preceq \mathbf{t} - \frac{1}{\alpha} \odot \mathbb{E}[(\mathbf{t} - \mathbf{r}(\mathbf{p}(\mathbf{h}), \mathbf{h}))_+], \\ & && \|\mathbb{E}[\mathbf{p}(\mathbf{h})]\|_1 \leq P_0 \end{aligned} \quad (7)$$

where “ \odot ” stands for elementwise multiplication, while $(\cdot)_+$ and division with a vector are similarly overloaded.

Observe that problem (6) is infinite-dimensional; in general such problems are challenging to tackle. Nonetheless, CV@R is a convex and monotone (in fact, coherent) risk measure [12], which preserves the convexity of both (6) and (7). Therefore, under the effect of some appropriate constraint qualification, such as Slater’s condition –assumed hereafter–, problems (6) and (7) exhibit no duality gap (in fact, strong duality). This fact suggests that we handle (7) in the dual domain, i.e., within the framework of Lagrangian duality.

III. LAGRANGIAN DUALITY

The Lagrangian of problem (7) is defined as

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \mathbf{p}, \mathbf{t}, \Lambda, \mu) & \triangleq f_0(\mathbf{x}) + \mu (P_0 - \|\mathbb{E}[\mathbf{p}(\mathbf{h})]\|_1) \\ & + \Lambda^T \left[\mathbf{t} - \frac{1}{\alpha} \odot \mathbb{E}[(\mathbf{t} - \mathbf{r}(\mathbf{p}(\mathbf{h}), \mathbf{h}))_+] - \mathbf{x} \right], \end{aligned} \quad (8)$$

where $\Lambda \succeq 0$ and $\mu \geq 0$ are the dual variables associated with the explicit constraints of (7). Accordingly, the *dual function* is defined as the maximization of the Lagrangian function over the primal variable triplet $(\mathbf{x}, \mathbf{p}, \mathbf{t})$, i.e.,

$$D(\Lambda, \mu) = \sup_{\mathbf{x} \in \mathcal{X}, \mathbf{p} \succeq \mathbf{0}, \mathbf{t}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \mathbf{t}, \Lambda, \mu). \quad (9)$$

Subsequently, the *dual problem* is the minimization of the dual function over the dual variable pair (Λ, μ) , i.e.,

$$\begin{aligned} D^* & = \inf_{(\Lambda, \mu) \succeq \mathbf{0}} D(\Lambda, \mu), \\ & = \inf_{(\Lambda, \mu) \succeq \mathbf{0}} \sup_{\mathbf{x} \in \mathcal{X}, \mathbf{p} \succeq \mathbf{0}, \mathbf{t}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \mathbf{t}, \Lambda, \mu). \end{aligned} \quad (10)$$

As mentioned previously, problem (7) exhibits strong duality (under Slater's condition), which means that $P^* = D^*$ and, what is more, optimal dual variables are guaranteed to exist. We also observe that even though problem (7) is infinite-dimensional, its dual (10) is finite-dimensional, which is a very useful fact if we are able to tackle the maximization involved in the dual function (9) –in particular over \mathbf{p} – adequately.

Leveraging strong duality, we hereafter focus on devising an efficient primal-dual algorithm for solving the minimax problem (10), hopefully providing an optimal solution to the constrained convex risk-aware problem (7), as well [3].

IV. RISK-AWARE RESOURCE ALLOCATION: THE TAIL WATERFILLING ALGORITHM

The dual problem can be separated into several subproblems with respect to the primal variables. In particular, (10) can be equivalently expressed as

$$\begin{aligned} \inf_{(\Lambda, \mu) \succeq \mathbf{0}} \left\{ \mu P_0 + \sup_{\mathbf{x} \in \mathcal{X}} f_0(\mathbf{x}) - \Lambda^T \mathbf{x} + \sup_{\mathbf{t}} \sum_{i=1}^n \lambda_i t_i \right. \\ \left. + \mathbb{E} \left[\sup_{p_i \geq 0} -\mu p_i - \frac{\lambda_i}{\alpha_i} (t_i - r_i(p_i(h_i), h_i))_+ \right] \right\}, \end{aligned} \quad (11)$$

where interchanging the sup over \mathbf{p} with expectation (integration) is justified in light of the *interchangeability principle*; see, e.g., [12, Theorem 7.92]. This fact allows us to derive an optimal policy in closed-form.

A. Optimal Resource Policy

The particular policy subproblem for each user i is

$$\sup_{p_i \geq 0} -\mu p_i - \frac{\lambda_i}{\alpha_i} \left(t_i - \log \left(1 + \frac{h_i p_i}{\sigma_i^2} \right) \right)_+. \quad (12)$$

The next result provides an optimal solution of subproblem (12), determining the behavior of the optimal risk-aware policy $\mathbf{p}^*(\mathbf{h})$, parameterized by primal/dual variables $(\mathbf{t}, \Lambda, \mu)$.




Fig. 3: Optimal resource policies for risk-aware (RA, solid) and risk-neutral (RN, dashed) settings in a 3-user network, with $\mu = 0.07$, $\alpha_i = 0.53$ and $\lambda_i = 0.33$ for all i , $t_1 = 2.9$, $t_2 = 2.15$ and $t_3 = 2.45$, respectively.

Theorem 1 (Optimal Resource Policy): An optimal solution to the i -th policy subproblem (12) may be expressed as

$$p_i^*(h_i, \cdot) = \begin{cases} \left[\frac{\lambda_i}{\mu \alpha_i} - \frac{\sigma_i^2}{h_i} \right]_+, & \text{if } \frac{\lambda_i}{\mu \alpha_i e^{t_i}} - \frac{\sigma_i^2}{h_i} < 0 \\ \frac{\sigma_i^2 (e^{t_i} - 1)}{h_i}, & \text{if } \frac{\lambda_i}{\mu \alpha_i e^{t_i}} - \frac{\sigma_i^2}{h_i} \geq 0 \end{cases}, \quad (13)$$

whenever μ and λ_i are not simultaneously zero; otherwise, it is optimal to choose $p_i^*(h_i) = 0$.

The risk-aware policy presented in Theorem 1 is a genuine extension of the classical risk-neutral waterfilling power policy [3], [13], which for user i may be expressed as

$$p_i^N(h_i, \cdot) = \left[\frac{\lambda_i}{\mu} - \frac{\sigma_i^2}{h_i} \right]_+. \quad (14)$$

This reduction is obtained by setting $\alpha_i = 1$, and then sending the CV@R variable t_i to $+\infty$ in (13) (this can be explained by the construction of the CV@R; see [12, Chapter 6]).

Presence of a confidence level $\alpha_i \in (0, 1)$ (and the related rate target t_i) makes the policy *less opportunistic* as compared with its risk-neutral counterpart. *First*, power is allocated more aggressively to more faded channels (smaller values of h_i) in the $[\cdot]_+$ -related part of the policy (cf. classical waterfilling), as α_i decreases and the corresponding branch of (13) is active (also depending on t_i). Further, if α_i (resp. t_i) is small enough to make the lower branch of (13) active, constant relative to α but now proportional to t_i power equal to $\sigma_i^2 (e^{t_i} - 1) / h_i$ is allocated. We observe a striking similarity of this term to outage-optimal policies; see, e.g., [7]. This might point to intricate information-theoretic properties of the CV@R approach, which can be the subject of future investigation.

Second, when the $[\cdot]_+$ -related part of the policy is nonpositive, then the policy indeed becomes opportunistic. However, given relevant (or optimized) values of t_i , the opportunism of the policy is effectively *restricted to the lower tail* (occurring with probability equal to α_i) of the random rate $r_i(p_i(h_i), h_i)$.

These remarks may also be readily observed in Fig. 3. Having explicitly determined \mathbf{p}^* , we may now derive primal –in (\mathbf{x}, \mathbf{t}) – and dual –in (Λ, μ) – iterations, purposed to *recursively learn* a globally optimal solution to (7).

B. CV@R Target / Risk-Ergodic Rate Updates–Primal Ascent

The remaining subproblems of (11) relative to primal variables (\mathbf{x}, \mathbf{t}) can be solved separately. Regarding maximization over \mathbf{t} , and with $p_i^* = p_i^*(h_i, t_i, \cdot)$, we end up with the problem

$$\sup_{t_i} \mathbb{E} \left[\lambda_i t_i - \mu p_i^* - \frac{\lambda_i}{\alpha_i} \left(t_i - \log \left(1 + \frac{h_i p_i^*}{\sigma_i^2} \right) \right)_+ \right], \quad (15)$$

for each i , where we explicitly denote the dependence of p_i^* on t_i . It can be easily seen that the function inside the expectation of (15) is jointly concave relative to (t_i, p_i) . Therefore, it is also concave in t_i under partial maximization over $p_i \geq 0$.

By [12, Theorem 7.52], it then follows that every subgradient of the latter is a stochastic subgradient of the objective of (15). Due to initial joint concavity in (t_i, p_i) , it can be shown that such a stochastic subgradient may be selected as

$$g_i(h_i, t_i, \cdot) = \lambda_i - \frac{\lambda_i}{\alpha_i} H \left[t_i - \log \left(1 + \frac{h_i p_i^*(h_i, t_i, \cdot)}{\sigma_i^2} \right) \right], \quad (16)$$

where $H(\cdot)$ is the step (Heaviside) multifunction. Then, provided an iteration index $n \in \mathbb{N}$ and processes $\{h_i^n\}$ and $\{t_i^n\}$ (and implicitly meant $\{\lambda_i^n\}$ and $\{\mu^n\}$; see below), we may formulate a stochastic subgradient ascent scheme for t_i as

$$t_i^n = t_i^{n-1} + \varepsilon_t g_i(h_i^n, t_i^{n-1}, \cdot), \quad n \geq 1, \quad (17)$$

with stepsize $\varepsilon_t > 0$, and starting from some initial value t_i^0 .

Maximization over \mathbf{x} , on the other hand, depends on the dual variables Λ and concave utility f_0 . Hereafter, we assume that an optimal solution as a function of Λ –or $\Lambda^n, n \geq 0$ –

$$\mathbf{x}^*(\Lambda) \in \arg \max_{\mathbf{x} \in \mathcal{X}} f_0(\mathbf{x}) - \Lambda^T \mathbf{x} \quad (18)$$

exists, and f_0 is such that $\mathbf{x}^*(\Lambda^n)$ is available (e.g., in closed form). Standard derivations and variable eliminations for popular utility functions, namely, sumrates and proportional fairness are provided later on, for completeness.

C. Dual Variable Updates

Lastly, we can formulate stochastic (quasi-)subgradient descent updates for dual variables (Λ, μ) . This is done along the lines of [3], by exploiting the corresponding constraint gaps. Note that the dual function D is convex and separable in (Λ, μ) . For the power constraint multiplier μ , we have

$$\mu^n = \left[\mu^{n-1} - \varepsilon_\mu \left(P_0 - \sum_{i=1}^n p_i^*(h_i^n, \mu^{n-1}, \cdot) \right) \right]_+, \quad (19)$$

with stepsize $\varepsilon_\mu > 0$, starting from μ^0 . Similarly, for the rate constraint vector of multipliers Λ , we get, for each i ,

$$\lambda_i^n = \left[\lambda_i^{n-1} - \varepsilon_\Lambda \left(-x_i^*(\Lambda^{n-1}) + t_i^{n-1} - \frac{1}{\alpha_i} \left(t_i^{n-1} - \log \left(1 + \frac{h_i^n p_i^*(h_i^n, \lambda_i^{n-1}, \cdot)}{\sigma_i^2} \right) \right)_+ \right) \right]_+, \quad (20)$$

with stepsize $\varepsilon_\Lambda > 0$, and starting from λ_i^0 .

The complete description of the proposed primal-dual algorithm, which we suggestively call *tail waterfilling*, is presented in Algorithm 1.

Algorithm 1 Tail Waterfilling

Choose initial values $\mathbf{t}^0, \mathbf{p}^0, \mathbf{x}^0, \mu^0, \Lambda^0$.

for $n = 1$ **to** Process **End do**

Observe \mathbf{h}^n .

Primal Variables

→ Set $p_i^*(\cdot)$ using (13), for all i .

→ Update t_i^n using (17) and (16), for all i .

→ Obtain $\mathbf{x}^*(\Lambda^{n-1})$ from (18).

Dual Variables

→ Update μ^n using (19).

→ Update λ_i^n using (20), for all i .

end for

D. Common Utilities

Sumrate: If $f_0(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$, $\mathbf{w} \succeq \mathbf{0}$, $\mathbf{x} \in \mathbb{R}^n$, the subproblem relative to \mathbf{x} becomes $\sup_{\mathbf{x}} (\mathbf{w} - \Lambda)^T \mathbf{x}$, which is unbounded for any \mathbf{w} and Λ , except for the optimal dual choice $\Lambda = \mathbf{w}$. Of course, this step eliminates both variables \mathbf{x} and Λ .

Proportional Fairness: In this case, we can choose $f_0(\mathbf{x}) = \sum_{i=1}^n \log(x_i)$, $\mathbf{x} \in \mathbb{R}^n$, and the subproblem in \mathbf{x} becomes

$$\sup_{\mathbf{x}} \sum_{i=1}^n \log(x_i) - \Lambda^T \mathbf{x} = \sup_{\mathbf{x}} \sum_{i=1}^n \log(x_i) - \lambda_i x_i, \quad (21)$$

which gives $x_i^* = 1/\lambda_i$ for each i .

V. PERFORMANCE EVALUATION

Let us now verify and discuss the efficacy of the proposed tail waterfilling algorithm, summarized in Algorithm 1. We consider a basic 3-terminal network consisting of Rayleigh fading point-to-point links with different noise variances. We then apply the tail waterfilling scheme for two utilities, namely, sumrate and proportional fairness. For the sumrate utility, the weights \mathbf{w} are selected to average the individual risk-ergodic services x_i per terminal, i.e., $w_i = 1/3, \forall i$. For both utilities, the stepsizes are set as $\varepsilon_t = 10^{-3}$ and $\varepsilon_\mu = 10^{-4}$, respectively. For proportional fairness, we set $\varepsilon_\Lambda = 10^{-4}$.

As shown in the histograms of Figs. 4 (top) and 6 (top), decreasing the (common) CVaR level α restricts the achievable rates in program (6), while constraining their volatility. This induces system robustness, since the system sustains a consistent and reliable level of performance, incurring infrequent rate drops in the long-term. Especially for proportional fairness, we observe that, *simultaneously* with being robust, the rates are more fairly distributed among users with different noise levels.

Equivalent remarks are in order regarding Figs. 4 (bottom) and 6 (bottom), which show the user outage probabilities, i.e., the cumulative distribution function $P_{\text{out}}(r_o) = P\{r \leq r_o\}$, as another intuitive measure to evaluate system robustness. In fact, we observe that for lower values of α , the system exhibits *lower and sharper* probabilities of outage –optimally tuned in the CV@R sense– at *always attainable* channel rates. On the other hand, less risk-aware settings corresponding to higher values for α result in much higher variability in the corresponding optimal channel rates.

The latter observations are also evident from Figs. 5 and 7, which highlight the vast difference in rate variability between

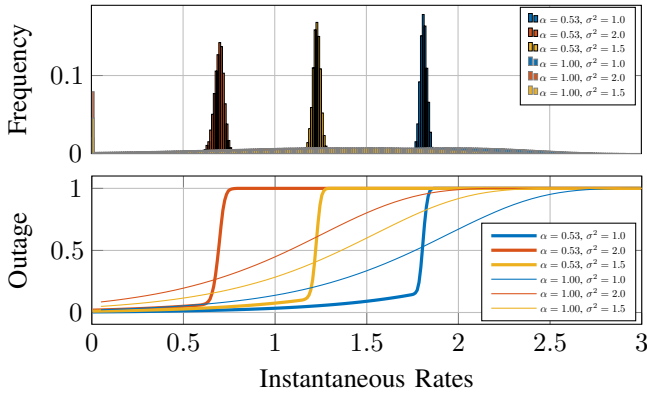


Fig. 4: Results for a network with 3 users and a sum rate utility, $P_0 = 10$. Top: Rate histograms. Bottom: Outage probabilities.




Fig. 5: Achieved rates for the 3-user network with a sum rate utility, $P_0 = 10$. Top: $\alpha = 0.53$ (risk-aware). Bottom: $\alpha = 1$ (risk-neutral).

the risk-aware and risk-neutral policies, through their evolution in time (channel use). We see that the optimal CV@R policy exhibits quasi-invariant communication rate trends, keeping the rates at certain reliability levels. Further, the proportional fairness utility achieves more evenly distributed rates among users, as shown in Fig. 7 (top). In other words, the combination of risk-awareness and proportional fairness achieves system performance that is *both* user-fair, *and* aware of fading risk.

VI. CONCLUSION

We proposed a new risk-aware reformulation of the classical resource allocation problem for point-to-point networks. Utilizing the CV@R as a measure of risk generalizing expectations, we developed the tail waterfilling algorithm, which extends classical stochastic waterfilling in an interpretable and tunable way, and induces network robustness and reliability rigorously and tractably. The effectiveness of tail waterfilling was corroborated via detailed numerical simulations.

REFERENCES

- [1] S. Parsaeefard and A. R. Sharafat, "Robust Distributed Power Control in Cognitive Radio Networks," *IEEE Trans. on Mobile Computing*, vol. 12, no. 4, pp. 609–620, Apr. 2013.
- [2] N. Mokari, S. Parsaeefard, P. Azmi, H. Saedi, and E. Hossain, "Robust Ergodic Uplink Resource Allocation in Underlay OFDMA Cognitive Radio Networks," *IEEE Trans. on Mobile Computing*, vol. 15, no. 2, pp. 419–431, Feb. 2016.




Fig. 6: Results for a network with 3 users and a proportional fairness utility, $P_0 = 10$. Top: Rate histograms. Bottom: Outage probabilities.




Fig. 7: Achieved rates for the 3-user network with a proportional fairness utility, $P_0 = 10$. Top: $\alpha = 0.51$ (risk-aware). Bottom: $\alpha = 1$ (risk-neutral).

- [3] A. Ribeiro, "Optimal resource allocation in wireless communication and networking," *EURASIP Journal on Wireless Commun. and Networking*, vol. 2012, no. 1, Aug. 2012.
- [4] R. AliHemmati, M. Dong, B. Liang, G. Boudreau, and S. H. Seyedmehdi, "Multi-Channel Resource Allocation Toward Ergodic Rate Maximization for Underlay Device-to-Device Communications," *IEEE Trans. on Wireless Communications*, vol. 17, no. 2, pp. 1011–1025, Feb. 2018.
- [5] D. S. Kalogieras, M. Eisen, G. J. Pappas, and A. Ribeiro, "Model-Free Learning of Optimal Ergodic Policies in Wireless Systems," *IEEE Trans. on Signal Processing*, vol. 68, pp. 6272–6286, 2020.
- [6] H. Hashmi and D. S. Kalogieras, "Model-Free Learning of Optimal Deterministic Resource Allocations in Wireless Systems via Action-Space Exploration," in *2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, Oct. 2021.
- [7] L. Li, N. Jindal, and A. Goldsmith, "Outage capacities and optimal power allocation for fading multiple-access channels," *IEEE Trans. on Information Theory*, vol. 51, no. 4, pp. 1326–1347, 2005.
- [8] T. K. Vu, M. Bennis, M. Debbah, M. Latva-aho, and C. S. Hong, "Ultra-Reliable Communication in 5G mmWave Networks: A Risk-Sensitive Approach," *IEEE Commun. Letters*, vol. 22, no. 4, pp. 708–711, 2018.
- [9] Y. Li, D. Guo, Y. Zhao, X. Cao, and H. Chen, "Efficient Risk-Averse Request Allocation for Multi-Access Edge Computing," *IEEE Communications Letters*, vol. 25, no. 2, pp. 533–537, 2021.
- [10] M. Bennis, M. Debbah, and H. V. Poor, "Ultrareliable and Low-Latency Wireless Communication: Tail, Risk, and Scale," *Proceedings of the IEEE*, vol. 106, no. 10, pp. 1834–1853, 2018.
- [11] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.
- [12] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*, 2nd ed. Society for Industrial and Applied Mathematics, Jan. 2014.
- [13] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, Apr. 2005.