

Bi-Criteria Radio Spectrum Sharing With Subspace-Based Pareto Tracing

Zachary J. Grey¹, Susanna Mosleh¹, Jacob D. Rezac¹, Yao Ma¹, *Senior Member, IEEE*,
Jason B. Coder¹, and Andrew M. Dienstfrey

Abstract—Radio spectrum is a scarce resource. To meet demands, new wireless technologies must operate in shared spectrum over unlicensed bands (coexist). We consider coexistence of Long-Term Evolution (LTE) License-Assisted Access (LAA) with incumbent Wi-Fi systems. Our scenario consists of multiple LAA and Wi-Fi links sharing an unlicensed band; we aim to simultaneously optimize performance of both coexistence systems. To do this, we present a technique to continuously estimate the Pareto frontier of parameter sets (traces) which approximately maximize all convex combinations of network throughputs over network parameters. We use a dimensionality reduction approach known as active subspaces to determine that this near-optimal parameter set is primarily composed of two physically relevant parameters. A choice of two-dimensional subspace enables visualizations augmenting explainability and the reduced-dimension convex problem results in approximations which dominate random grid search.

Index Terms—LTE, LAA, Wi-Fi, wireless coexistence, active subspace, Pareto trace.

I. INTRODUCTION

AS WIRELESS communications evolve and proliferate into our daily lives, the demand for radio spectrum grows dramatically. To accommodate this growth, wireless device protocols are beginning to transition from a predominantly-licensed spectrum to a shared approach in which use of the unlicensed spectrum bands is increasing

Manuscript received September 10, 2021; revised January 28, 2022; accepted March 6, 2022. Date of publication March 22, 2022; date of current version May 18, 2022. This work was funded by individual departments within the National Institute of Standards and Technology. This work is U.S. Government work and not protected by U.S. copyright. An earlier version of this work was presented at the IEEE International Conference on Communications (ICC 2021) [1] [DOI: 10.1109/ICC42927.2021.9500533]. The associate editor coordinating the review of this article and approving it for publication was C. Jiang. (*Corresponding author: Zachary J. Grey.*)

Zachary J. Grey and Andrew M. Dienstfrey are with the Information Technology Laboratory, Applied and Computational Mathematics Division, National Institute of Standards and Technology, Boulder, CO 80305 USA (e-mail: zachary.grey@nist.gov).

Susanna Mosleh is with the Communications Technology Laboratory, Spectrum Technology and Research Division, National Institute of Standards and Technology, Boulder, CO 80305 USA, and also with the Department of Physics, University of Colorado Boulder, Boulder, CO 80309 USA.

Jacob D. Rezac is with the Communications Technology Laboratory, RF Technology Division, National Institute of Standards and Technology, Boulder, CO 80305 USA.

Yao Ma and Jason B. Coder are with the Communications Technology Laboratory, Spectrum Technology and Research Division, National Institute of Standards and Technology, Boulder, CO 80305 USA.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TCOMM.2022.3161516>.

Digital Object Identifier 10.1109/TCOMM.2022.3161516

rapidly. The main bottleneck of this approach, however, is balancing new network paradigms with incumbent networks, such as Wi-Fi.

Previously, unlicensed bands were dominated by Wi-Fi traffic and, occasionally, used by commercial cellular carriers for offloading data that would otherwise have been communicated via Long-Term Evolution (LTE) in the licensed spectrum. In order to address spectrum scarcity in new operating paradigms, mobile network operators are choosing to operate in unlicensed bands (such as LAA) in addition to data offloading [2]. Even though operating LAA in unlicensed bands improves spectral-usage efficiency, it could have a significant influence on Wi-Fi operation and thereby create a number of challenges for spectrum sharing. Understanding and addressing these challenges calls for a deep dive into the operations and parameter selection of both networks in the medium access control (MAC) and physical (PHY) layers.

There have been many investigations of fairness in spectrum sharing among LAA and Wi-Fi networks [3]–[5]. Critically, these works do not consider optimizing key performance indicators (KPIs). In contrast to [3]–[5], the authors in [6] and [7] maximize LAA throughput and total network sum rate, respectively, over contention window sizes of both networks while guaranteeing the Wi-Fi throughput satisfies a threshold. Ignoring the constraint on Wi-Fi throughput, the authors in [8] maximize the overall network throughput over the same variables—contention window size of both networks. These studies [6]–[8], however, optimize only a single MAC layer parameter and do not consider optimizing over a set of MAC and PHY layer parameters. A multi-criteria optimization problem was formulated in [9] to satisfy the quality of service requirements of LAA eNodeBs by investigating the trade-off between the co-channel interference in the licensed band and the Wi-Fi collision probability in the unlicensed band. The line of work in [9] is further expanded in [10]. Considering both PHY and MAC layer parameters, [10] maximizes the weighted sum rate of an LAA network subject to Wi-Fi throughput constraint with respect to the fraction of time that LAA is active. However, Wi-Fi throughput was not simultaneously optimized in [9] nor [10]. For a scenario with both LTE base stations and Wi-Fi APs connected to a central controller, [11] and [12] maximized the overall throughput subject to constraints that Wi-Fi throughput does not degrade significantly. Their results were compared to the case that the unlicensed spectrum was not shared with LTE devices such that Wi-Fi and LTE-U operated on separate and

non-overlapping channels. The issue of efficient coexistence of LTE operators in the unlicensed bands through an optimization framework is considered in [13] using stochastic geometry. However, this study considers a scenario where an operator would like to optimize performance but has no constraints on the performance of other operators (cost). Moreover, the overall network throughput maximization is not discussed in [13].

In its most general form, the spectrum sharing problem can be modeled as a bi-criteria optimization problem where the KPIs of all operators, in a heterogeneous network that coexists on the unlicensed band, are maximized simultaneously. In this context, we explore optimal trade-offs between Wi-Fi and LAA throughputs that are simultaneously maximized over an aggregate of PHY and MAC layer parameters. This motivates a bi-criteria optimization formalism in which the input space has high-dimensionality [1]. The model we investigated in [1], and in this paper as well, uses 17 MAC and PHY layer variables to characterize the LAA and Wi-Fi coexistence performance.

Previous experience suggests that not all of these variables are equally important in determining the quality of network KPIs. To address this difficulty, we use a mathematical formalism known as *active subspaces* to determine parameter combinations that change KPI values the most on average—and those that do not [14]. The sets of parameter combinations defined by the active subspaces inform KPI approximations and visualizations over a low-dimension subspace—simplifying and regularizing the bi-criteria optimization. Here, we expand our investigation in [1] and study the impact of these parameters via analytical models and simulation results. This approach explores the system behavior and provides deep insights into related spectrum sharing and communication systems—LAA and Wi-Fi coexistence is merely an example in this work.

In this more detailed theoretical exposition of the initial work in [1], we take the model exploration further by exploring and comparing against additional traces through parameter space summarized in Section IV-B. We also incorporate an improved space filling sampling scheme over the reduced dimension subspace to improve the quadratic surrogates informing the resulting traces in Section IV-C. Moreover, we compare eigenspaces of two alternative approaches to the subspace dimension reduction in Section V to emphasize degeneracy in the fit of convex quadratics over the full dimensional parameter space. Lastly, we emphasize improvements in the stability of the proposed quadratic trace formulated over the reduced dimension subspace versus a full dimensional quadratic trace in Section V.

We incorporate active subspace dimension reduction into the bi-criteria optimization framework to analyze, interpret, and explain the shared spectrum coexistence problem. The set of maximizing arguments quantify the inherent trade-off between LAA and Wi-Fi throughputs. The dimension reduction supplements a trade-off analysis of network throughputs by computing a Pareto trace. The Pareto trace provides a *continuous* approximation of Pareto optimal (non-dominated) points in a common domain of a bi-criteria problem [15], [16]—resulting in a parameter manifold consisting of near-best trade-offs

between differing throughputs. Facilitated by the dimension reduction, our work provides a continuous description of this parameter manifold that quantifies high-quality performance of both networks.

This work differs from those previously mentioned by facilitating *new interpretations and explanations* of results. The main contributions of this paper can be summarized as follows:

- We incorporate *active subspace dimension reduction* into a *bi-criteria optimization* problem for radio spectrum sharing which *simplifies and regularizes a bi-criteria optimization*.
- We provide *ridge approximations* (defined in Section IV) of the KPIs in a coexistence scenario over an unlicensed band.
- We determine two parameter combinations that are most important in changing the quality of network KPIs versus those which are not. Picking two sets of parameter combinations lead to *visualizations over a low-dimension parameter subspace* (augmenting explainability of the approximations).
- We calculate *convex quadratic ridge approximations* of network throughputs that inform a *continuous quadratic trace* describing the trade-off between near-optimal network throughput combinations. This offers a *continuous description of a parameter manifold* that quantifies high quality performance of both networks as opposed to a set of point-wise approximations of optimal values.
- We demonstrate that *the resulting quadratic trace defined in terms of the low-dimensional mixed active subspace is more stable* than the comparable quadratic trace defined over the original (high-dimensional) variables. In this sense, the subspace-based dimension reduction facilitates a regularization of the Pareto trace.

The paper is organized as follows: Section II describes the system model and presents the problem formulation. Section III formalizes solutions to the problem statement and introduces the concept of a Pareto trace. Section IV introduces active subspaces and offers technical considerations for quantifying spectrum sharing as a continuous Pareto trace of near Pareto optimal parameters. Simulation results are shown and discussed in Section V. Finally, we conclude with an overview of the results and remarks about future work.

Notation: Standard math-font letters are used to denote scalars. Boldface capital and boldface lower-case letters denote matrices and vectors, respectively. All vectors, e.g., $\mathbf{u}, \mathbf{v} \in \mathbb{R}^K$, are assumed to be tall (column) vectors with K entries and all expressions correspond to standard matrix vector multiplication with the transpose of matrix \mathbf{A} is denoted by \mathbf{A}^\top . For example, matrix \mathbf{Q} , vector \mathbf{a} , and scalar c represent a matrix-vector format for general quadratic polynomials as $f(\boldsymbol{\theta}) = \boldsymbol{\theta}^\top \mathbf{Q} \boldsymbol{\theta} + \mathbf{a}^\top \boldsymbol{\theta} + c$. The operator \odot represents the Hadamard product and the operator $(\cdot)_k$ represents the k -th index of a vector—i.e., $(\mathbf{u} \odot \mathbf{v})_k = (\mathbf{u})_k (\mathbf{v})_k$. Given a vector of probabilities \mathbf{p} , complementary probabilities are represented by the operator \cdot^c such that $\mathbf{p}^c = \mathbf{1} - \mathbf{p}$ where $\mathbf{1}$ is an appropriately sized vector of ones. The gradient ∇ and

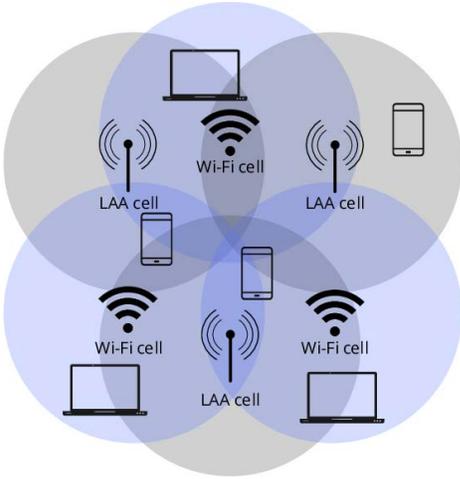


Fig. 1. System model consists of multiple LAA and Wi-Fi links sharing an unlicensed band.

Hessian ∇^2 are taken with respect to model parameters θ , if not otherwise decorated by a label. Finally, \mathcal{L} refers to the set LAA variables while \mathcal{W} refers to the set of Wi-Fi variables. When necessary, \mathcal{K} is used as a placeholder to relate operators or quantities applied analogously to either set. Parameters which are varied or examined for the purposes of model exploration, optimization, and transformation are denoted using the Greek alphabet.

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a downlink coexistence scenario where two mobile network operators (MNOs) operate over the same shared unlicensed industrial, scientific, and medical (ISM) radio band. We are primarily focused on the operation of cellular base stations in the unlicensed bands. However, LTE base stations may have permission to utilize a licensed band as well. We assume the MNOs use time sharing to simultaneously operate in the unlicensed band and we aim to analyze competing trade-offs in throughputs of the Wi-Fi and LAA systems. A network throughput is a function of both MAC and PHY layer parameters. In this section, we introduce the parameters defining the network topology, MAC layer protocols, the PHY layer, and briefly discuss the relation of these variables to network throughput.

A. Network Topology

We consider a coexistence scenario in which the LAA network consists of L Evolved Node B (eNodeBs) indexed by the set \mathcal{L} , $i \in \mathcal{L} = \{1, 2, \dots, L\}$, while the Wi-Fi network is composed of W access points (APs) indexed by the set \mathcal{W} , $j \in \mathcal{W} = \{1, 2, \dots, W\}$, as depicted in Figure 1. Note that our proposed subspace-based Pareto tracing approach could be applied to many types of communication systems (or alternative applications), but LTE is used here as an example. The eNodeBs and APs are randomly distributed over a rectangular area while LAA user equipment (UEs) and Wi-Fi clients/stations (STAs) are, respectively, distributed

around each eNodeB and AP independently and uniformly. Each transmission node serves one single antenna UE/STA. We assume (i) both Wi-Fi and LAA are in the saturated traffic condition, i.e., at least one packet is waiting to be sent, (ii) there are neither hidden nodes nor false alarm/miss detection problems in the network,¹ (iii) the channel knowledge is ideal, i.e., perfect channel sensing is assumed among the links, so the only source of unsuccessful transmission (packet failure) is collision, (iv) a successful transmission happens if only one link transmits at a time, i.e., exclusive channel access (ECA) model [17] is considered which is consistent with the IEEE 802.11 channel access protocols, and (v) each link is subject to Rayleigh fading and Log-normal shadowing.

B. MAC Layer Protocols

While Wi-Fi systems count on a contention-based medium access with a random back off process, called Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) [18], [19], LTE-LAA networks use a listen before talk (LBT) channel access mechanism as a means of sharing the band with Wi-Fi. Among different LAA-LBT schemes, the Category 4 (Cat 4) LBT—which is based on the same Wi-Fi CSMA/CA scheme and is well-suited in a coexistence scenario [2]—is considered in this paper. The medium access key feature in both Wi-Fi and LAA involves the station accessing the medium to sense the channel by performing clear channel assessment prior to transmitting. The station only transmits if the medium is determined to be idle. Otherwise, the transmitting station refrains from transmitting data until it senses the channel is available. Although LAA and Wi-Fi technologies follow similar channel access procedures, they utilize different carrier sense schemes, different channel sensing threshold levels, and different channel contention parameters, leading to different unlicensed channel access probabilities and thus, different throughputs.

Conforming with the analytical model in [20], [21], the probability of either network transmitting a packet in a randomly chosen time slot can be expressed as

$$p_k = \frac{2(1 - 2c_k)}{(1 - 2c_k)(1 + \omega_k) + c_k\omega_k(1 - (2c_k)^{\mu_k})}, \quad (1)$$

where k is representative of an index from either \mathcal{L} or \mathcal{W} , c_k denotes the probability of collision experienced by the k -th transmission node, and ω_k and μ_k indicate the minimum contention window size and the maximum back-off stage, respectively, of the k -th transmission node on the unlicensed channel.

To simplify notation, we aggregate the stationary transmission probability model (1) into entries of a vector \mathbf{p} . With this notation, the Wi-Fi stationary transmission probability of AP $j \in \mathcal{W}$ is considered $(\mathbf{p})_j : (c_j, \omega_j, \mu_j) \mapsto p_j(c_j; \omega_j, \mu_j)$ such that $\mathbf{p}_{\mathcal{W}} \triangleq \mathbf{p}(c_{\mathcal{W}}; \omega_{\mathcal{W}}, \mu_{\mathcal{W}}) \in \mathbb{R}^W$ represent the set of all Wi-Fi probabilities for all APs. The Wi-Fi probabilities depend

¹We assume perfect spectrum sensing in both systems. The impact of imperfect sensing is beyond the scope of this paper and investigating the effect of sensing errors is an important topic for future work.

explicitly on $\mathbf{c}_{\mathcal{W}} \in \mathbb{R}^W$ and parameters $\boldsymbol{\omega}_{\mathcal{W}}, \boldsymbol{\mu}_{\mathcal{W}} \in \mathbb{R}^W$. Similarly, for the LAA eNodeB's, we assign $\mathbf{p}_{\mathcal{L}} \triangleq \mathbf{p}(\mathbf{c}_{\mathcal{L}}; \boldsymbol{\omega}_{\mathcal{L}}, \boldsymbol{\mu}_{\mathcal{L}})$ such that $\mathbf{p}_{\mathcal{L}}, \mathbf{c}_{\mathcal{L}}, \boldsymbol{\omega}_{\mathcal{L}}, \boldsymbol{\mu}_{\mathcal{L}} \in \mathbb{R}^L$. Note, for brevity, we are dropping the explicit dependencies, $p_k(c_k; \omega_k, \mu_k)$, and supplement with an index-set subscript, $\mathbf{p}_{\mathcal{K}}$, to indicate the length of the vector-valued map as the cardinality of \mathcal{K} in addition to the network association when \mathcal{K} is either \mathcal{W} or \mathcal{L} . Moreover, these explicit dependencies are coupled with a set of complementary probabilities, $\mathbf{p}_{\mathcal{K}}^{\mathbb{C}} = \mathbf{1}_{\mathcal{K}} - \mathbf{p}_{\mathcal{K}}$, resulting in a set of equations which must be made consistent.

We write the LAA overall probability of transmission as $1 - \prod_{i \in \mathcal{L}} (\mathbf{p}_{\mathcal{L}}^{\mathbb{C}})_i$ and similarly for Wi-Fi $1 - \prod_{j \in \mathcal{W}} (\mathbf{p}_{\mathcal{W}}^{\mathbb{C}})_j$. In a complementary fashion, the collision probability of the transmitting Wi-Fi AP $j \in \mathcal{W}$ and LAA eNodeB $i \in \mathcal{L}$ on a shared unlicensed band are expressed as entries

$$(\mathbf{c}_{\mathcal{W}})_j = 1 - \frac{1}{(\mathbf{p}_{\mathcal{W}}^{\mathbb{C}})_j} \prod_{k \in \mathcal{W}} (\mathbf{p}_{\mathcal{W}}^{\mathbb{C}})_k \prod_{i \in \mathcal{L}} (\mathbf{p}_{\mathcal{L}}^{\mathbb{C}})_i, \quad (2)$$

and

$$(\mathbf{c}_{\mathcal{L}})_i = 1 - \frac{1}{(\mathbf{p}_{\mathcal{L}}^{\mathbb{C}})_i} \prod_{k \in \mathcal{L}} (\mathbf{p}_{\mathcal{L}}^{\mathbb{C}})_k \prod_{j \in \mathcal{W}} (\mathbf{p}_{\mathcal{W}}^{\mathbb{C}})_j. \quad (3)$$

respectively, stored in the vectors $\mathbf{c}_{\mathcal{W}} \in \mathbb{R}^W$ and $\mathbf{c}_{\mathcal{L}} \in \mathbb{R}^L$. Notice, (2) and (3) are expressions dependent on both $\mathbf{p}_{\mathcal{L}}$ and $\mathbf{p}_{\mathcal{W}}$ given parameters $\boldsymbol{\omega}_{\mathcal{W}}, \boldsymbol{\mu}_{\mathcal{W}}, \boldsymbol{\omega}_{\mathcal{L}}, \boldsymbol{\mu}_{\mathcal{L}}$. Consequently, (2) and (3) represent a coupling of probabilities subject to $\mathbf{0} \leq \mathbf{p}_{\mathcal{K}}, \mathbf{c}_{\mathcal{K}} \leq \mathbf{1}$.

The probability of a successful transmission for the Wi-Fi AP j (resp. LAA eNodeB i) on the unlicensed band is the j -th entry of $\mathbf{p}_{\mathcal{W}} \odot \mathbf{c}_{\mathcal{W}}^{\mathbb{C}}$ (resp. i -th entry of $\mathbf{p}_{\mathcal{L}} \odot \mathbf{c}_{\mathcal{L}}^{\mathbb{C}}$). Additionally, the average duration to support one successful transmission in the unlicensed band can be calculated as $\mathbf{p}_{\mathcal{T}}^{\top} \mathbf{t}$, as in [22]. The variable $\mathbf{p}_{\mathcal{T}}$ contains probabilities of an LAA and a Wi-Fi successful transmissions, and a collision among the Wi-Fi transmissions, among the LAA transmissions, and between the Wi-Fi and the LAA transmissions. Lastly, entries of \mathbf{t} are composed of the time that the unlicensed channel is being occupied by: an LAA successful transmission, a Wi-Fi successful transmission, a collision among the Wi-Fi transmissions, a collision among the LAA transmissions, and a collision between the Wi-Fi and the LAA transmissions. The vector $\mathbf{p}_{\mathcal{T}}$ is computed based on $\mathbf{p}_{\mathcal{W}}, \mathbf{c}_{\mathcal{W}}, \mathbf{p}_{\mathcal{L}}, \mathbf{c}_{\mathcal{L}}$ —as in [22]—and the vector \mathbf{t} is dictated by an access mechanism which in this paper the basic access mechanism is considered, as in [20], [23].

C. Physical Layer Parameters and Data Rates

To calculate network throughput, we also need to introduce data rates parametrized by physical layer parameters. The achievable physical data rate of the \mathcal{L} or \mathcal{W} operators is a function of link signal-to-noise ratio (SNR) that depends on changes with the link distances and propagation model; we refer to [22] for a detailed discussion of these relationships. The link distances and propagation model depend on parameters identified in Table I. The table summarizes entries of $\boldsymbol{\theta}$ representing the set of parameters that can be varied to study the model behavior along with entries of \mathbf{x} as the remaining

fixed scenario parameters. We write the parametrized SNR as $\text{SNR}(\boldsymbol{\theta}; \mathbf{x})$ while the data rate is expressed as $\log_2(1 + \text{SNR}_{\mathcal{K}}(\boldsymbol{\theta}; \mathbf{x}))$.

D. Parametrized Model

We parametrize MAC layer parameters by assuming common minimum contention window sizes and maximum back-off stages for each network. In other words, $\boldsymbol{\omega}_{\mathcal{W}} = \theta_1 \mathbf{1}_{\mathcal{W}}$, $\boldsymbol{\omega}_{\mathcal{L}} = \theta_2 \mathbf{1}_{\mathcal{L}}$, $\boldsymbol{\mu}_{\mathcal{W}} = \theta_3 \mathbf{1}_{\mathcal{W}}$, and $\boldsymbol{\mu}_{\mathcal{L}} = \theta_4 \mathbf{1}_{\mathcal{L}}$ for parameters $\theta_k \in \mathbb{R}$. Note that the channel access of each network depends upon these $(\boldsymbol{\theta})_1$ through $(\boldsymbol{\theta})_4$ parameters. It is conceivable that we may consider each independently for a total of $2(L + W)$ parameters as entries in a vector $\boldsymbol{\theta}$. However, we opt for a simplification to four total parameters—two common parameters per network. These parametrizations result in subsequent dependencies as $\mathbf{p}_{\mathcal{K}}(\boldsymbol{\theta})$ and $\mathbf{c}_{\mathcal{K}}(\boldsymbol{\theta})$ for either network \mathcal{K} which now depend on variations in $\boldsymbol{\theta}$. Thus, in general, we consider $\boldsymbol{\mu}_{\mathcal{K}}(\boldsymbol{\theta})$ and $\boldsymbol{\omega}_{\mathcal{K}}(\boldsymbol{\theta})$ where \mathcal{K} is either \mathcal{W} or \mathcal{L} and $\boldsymbol{\theta} \in \mathbb{R}^D$ is a vector of all parameters—the first four defined as MAC parameters and the remaining parameters described below and in Table I.

We also append the physical parameters to the MAC parameter vector by reassigning $\boldsymbol{\theta}$ as a vector representing the full set of parameters that can be varied to study the model behavior. These additional PHY parameters are summarized with appropriate bounds and description as the remaining entries of $\boldsymbol{\theta}$ in Table I. Any remaining physical parameters are held fixed and constitute a *scenario* for a particular model evaluation. We aggregate these remaining scenario parameters into a vector \mathbf{x} , summarized in Table I.

The LAA and Wi-Fi throughputs, indicated respectively by $f_{\mathcal{L}}$ and $f_{\mathcal{W}}$, are functions of D total MAC and PHY layer parameters in a vector $\boldsymbol{\theta}$ conditioned on fixed values in a vector \mathbf{x} ,

$$f_{\mathcal{L}} : \mathbb{R}^D \times \{\mathbf{x}\} \rightarrow \mathbb{R} : (\boldsymbol{\theta}, \mathbf{x}) \mapsto f_{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{x}), \quad (4)$$

and

$$f_{\mathcal{W}} : \mathbb{R}^D \times \{\mathbf{x}\} \rightarrow \mathbb{R} : (\boldsymbol{\theta}, \mathbf{x}) \mapsto f_{\mathcal{W}}(\boldsymbol{\theta}; \mathbf{x}). \quad (5)$$

Given a set of consistent probabilities for parameters $\boldsymbol{\theta}$, the throughputs are computed as

$$f_{\mathcal{K}}(\boldsymbol{\theta}; \mathbf{x}) = T_{\mathcal{K}} \log_2(1 + \text{SNR}_{\mathcal{K}}(\boldsymbol{\theta}; \mathbf{x})) \frac{\mathbf{p}_{\mathcal{K}}^{\top}(\boldsymbol{\theta}) \mathbf{c}_{\mathcal{K}}^{\mathbb{C}}(\boldsymbol{\theta})}{T^{\text{avg}}(\boldsymbol{\theta})} \quad (6)$$

for arbitrary \mathcal{K} representing either \mathcal{L} or \mathcal{W} where $T_{\mathcal{K}}$ indicates the (Wi-Fi or LAA) operator payload duration and $T^{\text{avg}}(\boldsymbol{\theta}) = \mathbf{p}_{\mathcal{T}}^{\top}(\boldsymbol{\theta}) \mathbf{t}$ denotes the average time duration to achieve a successful transmission in the network \mathcal{K} [22]. This numerical study and choice of model considers $D = 17$. However, parameters and models will be further generalized as part of on-going research efforts.

E. Problem Formulation

The problem of interest is to maximize a convex combination of network throughputs for the fixed scenario \mathbf{x} over the MAC and PHY parameters $\boldsymbol{\theta}$ in a bi-criteria optimization.

A *Pareto front* is defined by the following optimization problem:

$$\underset{\boldsymbol{\theta} \in \mathcal{D} \subset \mathbb{R}^D}{\text{maximize}} \quad \tau f_{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{x}) + (1 - \tau) f_{\mathcal{W}}(\boldsymbol{\theta}; \mathbf{x}), \quad (7)$$

for all $\tau \in [0, 1]$ where \mathcal{D} is the parameter domain defined by the ranges in Table I. We refer to $\tau = 0$ and $\tau = 1$ as *left* and *right* solutions, respectively. The goal is to quantify a smooth trajectory $\boldsymbol{\theta}(\tau; \mathbf{x})$ through MAC and PHY parameter space, or *trace* [16], such that the convex combination of throughputs is maximized over a map $\boldsymbol{\theta} : [0, 1] \times \{\mathbf{x}\} \rightarrow \mathcal{D}$. Note that the objective function in (7) is a combination of both throughputs, thus approximation requires global information such as channel state information. We consider a cooperative scenario, as in [24]–[27], in which the information of both networks can be exchangeable—i.e., the throughput information of both systems can be obtained by each network. In Section III we formalize this notion of a trace.

III. PARETO TRACING

We refer to (7) as a maximization of the convex total objective or *scalarization* [16]. In this case, we have a single degree of freedom to manipulate the scalarization parametrized by $\tau \in [0, 1]$ such that $\varphi_{\tau}(\boldsymbol{\theta}; \mathbf{x}) = (1 - \tau) f_{\mathcal{W}}(\boldsymbol{\theta}; \mathbf{x}) + \tau f_{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{x})$. In order to satisfy the necessary conditions for a (locally) Pareto optimal solution, we must determine $\boldsymbol{\theta} \in \mathbb{R}^D$ critical for φ_{τ} . This condition, also known as *stationarity*, requires $\nabla \varphi_{\tau}(\boldsymbol{\theta}; \mathbf{x}) = (1 - \tau) \nabla f_{\mathcal{W}}(\boldsymbol{\theta}; \mathbf{x}) + \tau \nabla f_{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{x}) = \mathbf{0}$. Moreover, $\boldsymbol{\theta}$ is a locally unique Pareto optimal solution if $\nabla^2 \varphi_{\tau}(\boldsymbol{\theta}; \mathbf{x}) = (1 - \tau) \nabla^2 f_{\mathcal{W}}(\boldsymbol{\theta}; \mathbf{x}) + \tau \nabla^2 f_{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{x})$ is symmetric negative definite [15], [16].

In order to find a continuous (in τ) solution to (7) we first examine the aforementioned necessary conditions. Provided the set of all Pareto optimal solutions to (7) is convex, we can continuously parametrize this set as $\tau \mapsto \boldsymbol{\theta}(\tau)$ for all $\tau \in [0, 1]$ (dropping scenario parameters \mathbf{x} for brevity). As an analogy, consider τ as pseudo-time describing an analogous trajectory through parameter space, i.e., launching from one maximizing argument to another. This interpretation is facilitated by the following Proposition:

Proposition 1 ([1], [16]): Given full rank $\nabla^2 \varphi_{\tau}(\boldsymbol{\theta}(\tau)) \in \mathbb{R}^{D \times D}$, the one-dimensional immersed submanifold parametrized by $\boldsymbol{\theta}(\tau) \in \mathbb{R}^D$ for all $\tau \in [0, 1]$ is necessarily Pareto optimal such that

$$\nabla^2 \varphi_{\tau}(\boldsymbol{\theta}(\tau)) \dot{\boldsymbol{\theta}}(\tau) = \nabla f_{\mathcal{W}}(\boldsymbol{\theta}(\tau)) - \nabla f_{\mathcal{L}}(\boldsymbol{\theta}(\tau)).$$

Proof: Differentiating the stationarity condition by composing in pseudo-time, $\nabla \varphi_{\tau} \circ \boldsymbol{\theta}(\tau) = \mathbf{0}$, results in

$$\begin{aligned} \mathbf{0} &= \frac{d}{d\tau} (\nabla \varphi_{\tau} \circ \boldsymbol{\theta}(\tau)) \\ &= \frac{d}{d\tau} (1 - \tau) (\nabla f_{\mathcal{W}} \circ \boldsymbol{\theta}(\tau)) + \frac{d}{d\tau} \tau (\nabla f_{\mathcal{L}} \circ \boldsymbol{\theta}(\tau)) \\ &= -\nabla f_{\mathcal{W}} \circ \boldsymbol{\theta}(\tau) + (1 - \tau) (\nabla_{\boldsymbol{\theta}}^2 f_{\mathcal{W}} \circ \boldsymbol{\theta}(\tau)) \dot{\boldsymbol{\theta}}(\tau) \\ &\quad + \nabla f_{\mathcal{L}} \circ \boldsymbol{\theta}(\tau) + \tau (\nabla_{\boldsymbol{\theta}}^2 f_{\mathcal{L}} \circ \boldsymbol{\theta}(\tau)) \dot{\boldsymbol{\theta}}(\tau) \\ &= \nabla^2 \varphi_{\tau}(\boldsymbol{\theta}(\tau)) \dot{\boldsymbol{\theta}}(\tau) - (\nabla f_{\mathcal{W}} \circ \boldsymbol{\theta}(\tau) - \nabla f_{\mathcal{L}} \circ \boldsymbol{\theta}(\tau)). \end{aligned}$$

Then, the flowout along necessary Pareto optimal solutions constitutes an immersed submanifold of \mathbb{R}^D nowhere tangent

to the integral curve generated by the system of differential equations (See [28], Thm. 9.20)—i.e., assuming $\nabla^2 \varphi_{\tau}(\boldsymbol{\theta}(\tau))$ is full rank, $\nabla^2 \varphi_{\tau}(\boldsymbol{\theta}(\tau))^{-1} (\nabla f_{\mathcal{W}}(\boldsymbol{\theta}(\tau)) - \nabla f_{\mathcal{L}}(\boldsymbol{\theta}(\tau)))$ is the infinitesimal generator of a submanifold in \mathbb{R}^D of locally Pareto optimal solutions contained in the flowout. \square

We note that the system of equations in Prop. 1, proposed in [16], constitutes a set of *necessary* conditions for optimality. The utility of Prop. 1 offers an interpretation that the solution set (if it exists) constitutes elements of a submanifold in \mathbb{R}^D . This formalism establishes a theoretical foundation for the use of manifold learning or splines over sets of points that are approximately Pareto optimal. Our future research efforts are motivated by drawing comparisons and complementary analysis with the aforementioned methods—Prop. 1 serves as the theoretical motivation and interpretation for such efforts.

Suppose $f_{\mathcal{W}}$ and $f_{\mathcal{L}}$ are well approximated by convex quadratics as *surrogates*,

$$-f_{\mathcal{L}}(\boldsymbol{\theta}; \mathbf{x}) \approx \boldsymbol{\theta}^T \mathbf{Q}_{\mathcal{L}} \boldsymbol{\theta} + \mathbf{a}_{\mathcal{L}}^T \boldsymbol{\theta} + c_{\mathcal{L}} \quad (8)$$

and

$$-f_{\mathcal{W}}(\boldsymbol{\theta}; \mathbf{x}) \approx \boldsymbol{\theta}^T \mathbf{Q}_{\mathcal{W}} \boldsymbol{\theta} + \mathbf{a}_{\mathcal{W}}^T \boldsymbol{\theta} + c_{\mathcal{W}}, \quad (9)$$

such that $\mathbf{Q}_{\mathcal{L}}, \mathbf{Q}_{\mathcal{W}} \in S_{++}^D$, $\mathbf{a}_{\mathcal{L}}, \mathbf{a}_{\mathcal{W}} \in \mathbb{R}^D$, and $c_{\mathcal{L}}, c_{\mathcal{W}} \in \mathbb{R}$ where S_{++}^D denotes the collection of D -by- D symmetric positive definite matrices—note the change in sign convention. In this case, the Pareto trace defined in Prop. 1 can be solved in closed form,

$$\boldsymbol{\theta}(\tau; \mathbf{x}) = \frac{1}{2} [\tau \mathbf{Q}_{\mathcal{L}} + (1 - \tau) \mathbf{Q}_{\mathcal{W}}]^{-1} [(\tau - 1) \mathbf{a}_{\mathcal{W}} - \tau \mathbf{a}_{\mathcal{L}}], \quad (10)$$

referred to in this context as a *quadratic trace*. The challenge in our context given the possibility of degeneracy in quadratic Hessians associated with $f_{\mathcal{W}}$ and $f_{\mathcal{L}}$ is: *how can we assess conditions informing (10) or regularize the solve to guarantee these conditions?* We offer an approach that assess these conditions by regularization through *subspace-based dimension reduction* to inform convex quadratic surrogates as approximations satisfying Prop. 1.

Naively, for arbitrary dimension D and number of samples N , we can also pose a convex optimization problem over the cone of positive semi-definite matrices,

$$\underset{\mathbf{Q} \in S_+^D, \mathbf{a} \in \mathbb{R}^D, c \in \mathbb{R}}{\text{minimize}} \quad \frac{1}{2} \sum_{n=1}^N \left(c + \boldsymbol{\theta}_n^T \mathbf{a} + \boldsymbol{\theta}_n^T \mathbf{Q} \boldsymbol{\theta}_n - f_n \right)^2. \quad (11)$$

Approximations resulting from problem (11) allow us to compute the quadratic trace (10) as a global data-driven surrogate given $\{(\boldsymbol{\theta}_n, f_n)\}$ of paired parameters and function responses where f_n represents the throughput response—modeled or measured—for either network; e.g., $f_n = f_{\mathcal{W}}(\boldsymbol{\theta}_n; \mathbf{x})$ or $f_n = f_{\mathcal{L}}(\boldsymbol{\theta}_n; \mathbf{x})$. Solutions to problem (11) are described as a *fit* $\mathcal{CVX}(\{(\boldsymbol{\theta}_n, f_n)\})$. Following an approximation to (11), the global quadratic Pareto trace is predicated on the rank of the minimizing arguments \mathbf{Q} —in particular, the rank of the convex combination of matrices given by independent solutions to (11) using data for either

TABLE I
MAC AND PHY PARAMETERS INFLUENCING THROUGHPUTS

Parameters	Description	Bounds	Nominal
$(\theta)_1$	Wi-Fi min contention window size	(8, 1024)	516
$(\theta)_2$	LAA min contention window size	(8, 1024)	516
$(\theta)_3$	Wi-Fi max back-off stage	(0, 8)	4
$(\theta)_4$	LAA max back-off stage	(0, 8)	4
$(\theta)_5$	Distance between transmitters	(10 m, 20 m)	15 m
$(\theta)_6$	Minimum distance between transmitters and receivers	(10m, 35 m)	22.5 m
$(\theta)_7$	Height of each LAA eNodeB and Wi-Fi AP	(3 m, 6 m)	4.5 m
$(\theta)_8$	Height of each LAA UEs and Wi-Fi STAs	(1 m, 1.5 m)	1.25 m
$(\theta)_9$	Standard deviation of shadow fading	(8.03, 8.29)	8.16
$(\theta)_{10}$	k_{LoS}^*	(45.12, 46.38)	45.75
$(\theta)_{11}$	k_{NLoS}^*	(34.70, 46.38)	40.54
$(\theta)_{12}$	α_{LoS}^*	(17.3, 21.5)	19.4
$(\theta)_{13}$	α_{NLoS}^*	(31.9, 38.3)	35.1
$(\theta)_{14}$	Antenna gain at each transmitter	(0 dBi, 5 dBi)	2.5 dBi
$(\theta)_{15}$	Noise figure at each receiver	(5 dB, 9 dB)	7 dB
$(\theta)_{16}$	Transmit power at each LAA eNodeB and Wi-Fi AP	(18 dBm, 23 dBm)	20.5 dBm
$(\theta)_{17}$	Carrier channel bandwidth	(10 MHz, 20 MHz)	15 MHz
$(\mathbf{x})_1$	Number of LAA eNodeBs (L)	–	6
$(\mathbf{x})_2$	Number of Wi-Fi APs (W)	–	6
$(\mathbf{x})_3$	Number of LAA UEs	–	6
$(\mathbf{x})_4$	Number of Wi-Fi STAs	–	6
$(\mathbf{x})_5$	Number of unlicensed channels	–	1
$(\mathbf{x})_8$	Scenario width	–	120 m
$(\mathbf{x})_9$	Scenario height	–	80 m

Note: parameter ranges are established by 3GPP TS 36.213 V15.6.0 and 3GPP TR. 36.889 v13.0.0. MAC parameters are highlighted in blue. Scenario parameters are highlighted in gray. PHY parameters are not highlighted.

*The path-loss for both line-of-sight (LoS) and non-LoS scenarios can be computed as $k + \alpha \log_{10}(d)$ in dB, where d is the distance in meters between the transmitter and the receiver.

throughput, i.e., $\{\mathbf{Q}_{\mathcal{L}}, \mathbf{a}_{\mathcal{L}}, c_{\mathcal{L}}\} = \mathcal{CVX}\{(\theta_n, f_{\mathcal{L}}(\theta_n; \mathbf{x}))\}$ and $\{\mathbf{Q}_{\mathcal{W}}, \mathbf{a}_{\mathcal{W}}, c_{\mathcal{W}}\} = \mathcal{CVX}\{(\theta_n, f_{\mathcal{W}}(\theta_n; \mathbf{x}))\}$. Note that in (11) we relax the condition of strict positive definiteness to account for potentially degenerate quadratics—which can act as alternative models suggesting low-dimensional approximations [14], [29], [30]—ensuring the solution space of (11) is closed [15]. Efficient solutions (fits) to problem (11) can be computed with “CVX,” an open-source package for defining and solving convex programs [31], [32].

Evidently, from equation (10) we see that the stability of the quadratic trace depends on the condition number of $\tau \mathbf{Q}_{\mathcal{L}} + (1 - \tau) \mathbf{Q}_{\mathcal{W}}$. In other words, subsequent fits $\mathcal{CVX}\{(\theta_n, f_{\mathcal{K}}(\theta_n; \mathbf{x}))\}$ dependent on various data sets as input may result in perturbations to the coefficient matrices $\mathbf{Q}_{\mathcal{K}}$ for either throughput.² Consequently, the otherwise unknown (and potentially variable) conditioning of $\mathbf{Q}_{\mathcal{K}}$ may result in instabilities of the quadratic trace subject to small perturbations—despite a reliable constrained optimization (11). We desire solutions that are *stable* (better conditioned) against variations in the quadratic surrogate when

²Notice also that the fits $\mathcal{CVX}\{(\theta_n, f_{\mathcal{K}}(\theta_n; \mathbf{x}))\}$ depend on the fixed scenario \mathbf{x} and consequently any subsequent reference to these approximations should reflect this dependency.

presented with new or perturbed data. In this work, we study numerical experiments in section V emphasizing a significant improvement in stability after dimension reduction. Formalizing improvements in stability induced by dimension reduction is a topic for future research.

IV. ACTIVE SUBSPACES

Following the development in [14], we introduce a dimension-reduction method to approximate the functions in (7). This dimension-reduction technique draws from active subspace analysis to identify linear subspaces of parameters that lead to significant changes in a function. To describe the details of the active subspaces approach, we introduce a scalar-valued function $f_{\mathcal{K}} : \mathcal{D} \subset \mathbb{R}^D \rightarrow \mathbb{R}$ defined on a compact domain \mathcal{D} . Again, where \mathcal{K} denotes either throughput functions $f_{\mathcal{L}}$ or $f_{\mathcal{W}}$ and dropping scenario parameters for brevity.

The main results of the section rely on an eigendecomposition of the symmetric positive semi-definite matrix $\mathbf{C}_{\mathcal{K}} \in \mathbb{R}^{D \times D}$ defined as

$$\mathbf{C}_{\mathcal{K}} = \int_{\mathcal{D}} \nabla f_{\mathcal{K}}(\theta) \nabla f_{\mathcal{K}}^{\top}(\theta) d\theta. \quad (12)$$

In the throughput maximization application discussed in this article, the compact domain \mathcal{D} is a D -dimensional rectangle constructed from the Cartesian product of lower and upper bounds, $(\theta_\ell)_i \leq \theta_i \leq (\theta_u)_i$ for all $i = 1, \dots, D$. We also assume a uniform measure for integration over \mathcal{D} . Additional discussion is available in [1].

If $\text{rank}(\mathbf{C}_\mathcal{K}) = r < D$, its eigendecomposition $\mathbf{C}_\mathcal{K} = \mathbf{W}_\mathcal{K} \mathbf{\Lambda}_\mathcal{K} \mathbf{W}_\mathcal{K}^\top$ with orthogonal $\mathbf{W}_\mathcal{K}$ satisfies $\mathbf{\Lambda}_\mathcal{K} = \text{diag}(\lambda_1, \dots, \lambda_D)$ with $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_r > \lambda_{r+1} = \dots = \lambda_D = 0$. This defines two sets of important $\mathbf{W}_{r,\mathcal{K}} = [\mathbf{w}_1 \dots \mathbf{w}_r]_\mathcal{K} \in \mathbb{R}^{D \times r}$ and unimportant $\mathbf{W}_{r,\mathcal{K}}^\perp = [\mathbf{w}_{r+1} \dots \mathbf{w}_D]_\mathcal{K} \in \mathbb{R}^{D \times (D-r)}$ directions over the domain. The column span of $\mathbf{W}_{r,\mathcal{K}}$ and $\mathbf{W}_{r,\mathcal{K}}^\perp$ constitute the *active* and *inactive* subspaces, respectively. Note that (12) depends on a single scalar-valued response and, therefore, the active subspaces potentially differ for the separate throughputs in (4) and (5).

What do we mean by *important directions*? Using the eigenvectors $\mathbf{w}_j \in \mathbb{R}^D$, we can simplify $\mathbf{w}_j^\top \mathbf{C}_\mathcal{K} \mathbf{w}_j$ to obtain an expression for the eigenvalues,

$$\lambda_j = \int_{\mathcal{D}} (\mathbf{w}_j^\top \nabla f_\mathcal{K}(\boldsymbol{\theta}))^2 d\boldsymbol{\theta}. \quad (13)$$

In other words, the j -th eigenvalue can be interpreted as the mean squared directional derivative of $f_\mathcal{K}$ in the direction of $\mathbf{w}_j \in \mathbb{R}^D$ denoted $\mathbb{E}[df_\mathcal{K}^2[\mathbf{w}_j]]$ [1], [14]. Thus, the ordering of the eigenpairs $\{(\lambda_j, \mathbf{w}_j)\}_{j=1}^D$ indicate directions \mathbf{w}_j over which the function $f_\mathcal{K}$ changes more, on average, up to the $r+1, \dots, D$ directions that *do not change the function at all* [14]. In fact, either throughput response from (4) or (5) is referred to as a *ridge function* over $\boldsymbol{\theta}$'s if and only if the directional derivatives are zero over all $\mathbf{w} \in \text{Null}(\mathbf{W}_{r,\mathcal{K}}^\top)$. We formalize this interpretation in the following:

Proposition 2 ([14]): Given \mathbf{w}_j for $j = r+1, \dots, D$ as the trailing eigenvectors of $\mathbf{C}_\mathcal{K}$, the paired eigenvalues $\lambda_j = 0$ if and only if $f_\mathcal{K}$ does not change over $\text{span}\{\mathbf{w}_{r+1}, \dots, \mathbf{w}_D\}$.

Proof: An alternative presentation of the result is offered for completeness. (\implies) Take a linear combination of any two eigenvectors \mathbf{w}_i and \mathbf{w}_j with identically zero eigenvalues, λ_i and λ_j , for any $i, j \in \{r+1, \dots, D\}$. Then, for arbitrary $a, b \in \mathbb{R}$, $\boldsymbol{\theta} \in \mathcal{D}$, and differentiable function $f_\mathcal{K} : \mathcal{D} \rightarrow \mathbb{R}$,

$$\begin{aligned} |\mathbb{E}[df_\mathcal{K}^2[a\mathbf{w}_i + b\mathbf{w}_j]]| &\stackrel{(i)}{=} |\mathbb{E}[(adf_\mathcal{K}[\mathbf{w}_i] + bdf_\mathcal{K}[\mathbf{w}_j])^2]| \\ &\stackrel{(ii)}{=} |2ab| |\mathbb{E}[df_\mathcal{K}[\mathbf{w}_i] df_\mathcal{K}[\mathbf{w}_j]]| \\ &\stackrel{(iii)}{\leq} |2ab| (\mathbb{E}[df_\mathcal{K}^2[\mathbf{w}_i]] \mathbb{E}[df_\mathcal{K}^2[\mathbf{w}_j]])^{1/2} \\ &\stackrel{(iv)}{=} |2ab| \sqrt{\lambda_i \lambda_j} \\ &\stackrel{(v)}{=} 0. \end{aligned}$$

The differentiability of $f_\mathcal{K}$ implies continuity so that $\mathbb{E}[df_\mathcal{K}^2[a\mathbf{w}_i + b\mathbf{w}_j]] = 0 \implies df_\mathcal{K}(\boldsymbol{\theta})[a\mathbf{w}_i + b\mathbf{w}_j] = 0$ for all $\boldsymbol{\theta}$. Equality (i) follows from the linearity of the differential over directions. Equality (ii) is a simplification of the expanded quadratic that follows by assumption—i.e., $\mathbb{E}[df_\mathcal{K}^2[\mathbf{w}_i]] = \mathbb{E}[df_\mathcal{K}^2[\mathbf{w}_j]] = \lambda_i = \lambda_j = 0$ for $\mathbf{w}_i, \mathbf{w}_j \in \text{Null}(\mathbf{C}_\mathcal{K})$. Inequality (iii) is Cauchy-Schwarz for a Hilbert space of square integrable (measurable) functions. Equality (iv) follows from

expression (13) and equality (v) follows from assumption. Finally, by recursively assigning $\mathbf{w}_i \leftarrow a\mathbf{w}_i + b\mathbf{w}_j$ and taking the next eigenvector from the set to be \mathbf{w}_j , we can repeat the above for all remaining eigenvectors. The converse (\impliedby) follows directly from (13) and linearity of the differential. \square

Naturally, if the trailing eigenvalues are merely small as opposed to zero, then the function changes much less over the inactive directions with smaller directional derivatives. This lends itself to a framework for reduced-dimension approximation of the function such that we only approximate changes in the function over the first r active directions and take the approximation to be constant over the trailing $D-r$ inactive directions [14]. Such an approximation to $f_\mathcal{K}$ is called a *ridge approximation* by a function $h_\mathcal{K}$ referred to as the *ridge profile* [33], i.e.,

$$f_\mathcal{K}(\boldsymbol{\theta}) \approx h_\mathcal{K}(\mathbf{W}_{r,\mathcal{K}}^\top \boldsymbol{\theta}). \quad (14)$$

In the event that the trailing eigenvalues of $\mathbf{C}_\mathcal{K}$ are zero, then the approximation is exact for a particular $h_\mathcal{K}$ [14].

In either case, approximation or an exact ridge profile, the possibility of reducing dimension by projection to fewer, $r < D$, *active coordinates* $\boldsymbol{\gamma} = \mathbf{W}_{r,\mathcal{K}}^\top \boldsymbol{\theta} \in \mathbb{R}^r$ can enable higher-order polynomial approximations for a given data set of coordinate-output pairs and an ability to *visualize* the approximation. For example, we can visualize the approximation by projection to the active coordinates when r is chosen to be 1 or 2 based on the decay and gaps in the eigenvalues. These subsequent visualizations are referred to as *shadow plots* [29] or graphs $\{(\mathbf{W}_{r,\mathcal{K}}^\top \boldsymbol{\theta}_n, f_\mathcal{K}(\boldsymbol{\theta}_n))\}_{n=1}^N$ for N samples $\{\boldsymbol{\theta}_n\}$ drawn uniformly. A strong decay leading to a small sum of trailing eigenvalues implies an improved approximation over relatively few important directions while larger gaps in eigenvalues imply an improved approximation to the low-dimensional subspace [14]. Identifying if this structure exists depends on the decay and gaps in eigenvalues. We can subsequently exploit any reduced dimensional visualization and approximation to simplify our problem (7). However, we must reconcile that our problem of interest involves two separate computations of throughput, $f_\mathcal{L}$ and $f_\mathcal{W}$. These considerations are addressed in subsection IV-A. Computational considerations (parameter scalings and selection of r) are discussed in [1]. We present algorithm 1 from [1] again for completeness. Additionally, convergence guarantees and error bounds on the subspace approximations resulting from Algorithm 1 are discussed at length in [14] (Thm. 3.13).

A. Subspace Mixing

Independently approximating active subspaces for the objectives $f_\mathcal{L}$ and $f_\mathcal{W}$ generally results in different subspaces of the shared parameter domain. The next challenge is to define a common subspace that, while not active for each objective, is nevertheless sufficient to capture variability of both KPI simultaneously. Assume that we can reduce important parameter combinations to a common dimension r of potentially distinct subspaces. These subspaces are spanned by the column spaces of $\mathbf{W}_{r,\mathcal{L}}$ and $\mathbf{W}_{r,\mathcal{W}}$ chosen as the first r eigenvectors resulting from separate approximations of (12)

Algorithm 1 Monte Carlo Approximation of Throughput Active Subspaces Using Forward Differences

Require: Forward maps $f_{\mathcal{L}}$ and $f_{\mathcal{W}}$, small coordinate perturbation $\delta \geq 0$, fixed scenario parameters \mathbf{x} , parameter bounds $\theta_{\ell}, \theta_u \in \mathbb{R}^D$, and the number of Monte Carlo samples N .

- 1: Generate N i.i.d. uniform samples, $\{\tilde{\theta}_n\}_{n=1}^N \sim U_D[-1, 1]$.
- 2: Compute \mathbf{M} , \mathbf{M}^{-1} , and \mathbf{b} according to a uniform distribution of log-scale parameters given θ_{ℓ} and θ_u .
- 3: **for** $n = 1$ to N **do**
- 4: Transform the uniform log-scale sample to the original scale $\theta_n = \exp(\mathbf{M}^{-1}(\tilde{\theta}_n - \mathbf{b}))$ where the exponential is taken component-wise.
- 5: Evaluate forward maps $(f_{\mathcal{L}})_n = f_{\mathcal{L}}(\theta_n; \mathbf{x})$ and $(f_{\mathcal{W}})_n = f_{\mathcal{W}}(\theta_n; \mathbf{x})$.
- 6: **for** $i = 1$ to D **do**
- 7: Transform the i -th coordinate perturbation to the original input scale $\theta_{\delta} = \exp(\mathbf{M}^{-1}(\tilde{\theta}_n + \delta e_i - \mathbf{b}))$, where e_i is the i -th column of the D -by- D identity matrix.
- 8: Approximate the i -th entry of the gradient (i -th partial) at the n -th sample as

$$(\tilde{\nabla} f_{\mathcal{L}})_{n,i} = \frac{f_{\mathcal{L}}(\theta_{\delta}; \mathbf{x}) - (f_{\mathcal{L}})_n}{\delta},$$

similarly for $(\tilde{\nabla} f_{\mathcal{W}})_{n,i}$.

- 9: **end for**
- 10: **end for**
- 11: Compute the average of the outer product of approximate gradients as

$$\tilde{\mathcal{C}}_{\mathcal{L}} = \frac{1}{N} \sum_{n=1}^N (\tilde{\nabla} f_{\mathcal{L}})_{n,:} \otimes (\tilde{\nabla} f_{\mathcal{L}})_{n,:},$$

similarly for $\tilde{\mathcal{C}}_{\mathcal{W}}$, where the tensor (outer) product is taken over the i -th index replaced by:

- 12: Compute the eigenvalue decompositions

$$\tilde{\mathcal{C}}_{\mathcal{L}} = \tilde{\mathbf{W}}_{\mathcal{L}} \tilde{\Lambda}_{\mathcal{L}} \tilde{\mathbf{W}}_{\mathcal{L}}^{\top} \quad \text{and} \quad \tilde{\mathcal{C}}_{\mathcal{W}} = \tilde{\mathbf{W}}_{\mathcal{W}} \tilde{\Lambda}_{\mathcal{W}} \tilde{\mathbf{W}}_{\mathcal{W}}^{\top}$$

ordered by decreasing eigenvalues.

- 13: Observe the eigenvalue decay and associated gaps to inform a reasonable choice of r .
 - 14: **return** The first r columns of $\tilde{\mathbf{W}}_{\mathcal{L}}$ and $\tilde{\mathbf{W}}_{\mathcal{W}}$, denoted $\tilde{\mathbf{W}}_{r,\mathcal{L}}$ and $\tilde{\mathbf{W}}_{r,\mathcal{W}}$.
-

for LAA and Wi-Fi throughputs, respectively. The challenge is to appropriately “mix” the subspaces so we may formulate a solution to (7) over a common dimension reduction.

One method to find an appropriate subspace mix is to take the union of both subspaces. However, if $r \geq 2$ and $\text{Range}(\mathbf{W}_{r,\mathcal{L}}) \cap \text{Range}(\mathbf{W}_{r,\mathcal{W}}) = \{\mathbf{0}\}$ then the combined subspace dimension is inflated. This inflation betrays the goals of dimension reduction and, furthermore, hinders prospects for visualization. We use interpolation between the two subspaces to overcome these difficulties and retain the common reduction to an r -dimensional subspace. The space of all r -dimensional subspaces in \mathbb{R}^D is the $r(D - r)$ -dimension Grassmann manifold (Grassmannian) denoted $\text{Gr}(r, D)$ [34]. Utilizing the analytic form of a geodesic over the Grassmannian [34], we can smoothly interpolate from one subspace to another—an interpolation that is, in general, non-linear. This is particularly useful because the *distance* [34] between any two subspaces along such a path, $[\mathbf{U}_r] : \mathbb{R} \rightarrow \text{Gr}(r, D) : s \mapsto [\mathbf{U}_r(s)]$ for all $s \in$

$[0, 1]$, minimizes the distance between the two subspaces $\text{Range}(\mathbf{W}_{r,\mathcal{L}}), \text{Range}(\mathbf{W}_{r,\mathcal{W}}) \in \text{Gr}(r, D)$ defining the geodesic. That is, the geodesic $[\mathbf{U}_r(s)]$ minimizes the distance between $\text{Range}(\mathbf{W}_{r,\mathcal{L}})$ and $\text{Range}(\mathbf{W}_{r,\mathcal{W}})$ while still constituting an r -dimensional subspace in \mathbb{R}^D .

In an effort to improve the ridge approximations while retaining the ability to visualize the response and trace of the convex quadratic polynomial, we fix $r = 2$ and mix the subspaces according to a quadratic approximation with corresponding coefficients of determination $R_{\mathcal{L}}^2$ and $R_{\mathcal{W}}^2$. These are computed over the Grassmannian geodesic using representative subspace coordinates and throughput (coordinate-output) pairs, i.e.,

$$R_{\mathcal{K}}^2(s) = 1 - \frac{\sum_{n=1}^N (f_{\mathcal{K}}(\theta_n) - h_{\mathcal{K}}(\mathbf{U}_r^{\top}(s)\theta_n))^2}{\sum_{n=1}^N (f_{\mathcal{K}}(\theta_n) - 1/N \sum_{n=1}^N f_{\mathcal{K}}(\theta_n))^2} \quad (15)$$

for either network represented by \mathcal{K} and quadratic ridge profiles $h_{\mathcal{K}}$. Moreover, the dimension reduction down to $r = 2$ is anticipated to help regularize fits $h_{\mathcal{K}}$ and provide more stable approximations resulting from a quadratic trace. We select a criterion to mix subspaces achieving a balanced approximation when $r = 2$. This offers the subproblem,

$$\text{maximize}_{s \in [0,1]} \min\{R_{\mathcal{L}}^2(s), R_{\mathcal{W}}^2(s)\}, \quad (16)$$

where the separate throughput coefficients of determination are parametrized over a consistent set of subspace coordinates $\gamma_n = \mathbf{U}_r^{\top}(s)\theta_n$ for all $n = 1, \dots, N$ and quadratic ridge profiles $h_{\mathcal{K}}$.

B. Tracing Ridge Profiles

After approximating $\mathbf{W}_{r,\mathcal{L}}$ and $\mathbf{W}_{r,\mathcal{W}}$ we must make an informed decision to take the union of subspaces or compute a new subspace $\text{Range}(\mathbf{U}_r)$ against some criteria parametrized over the Grassmannian geodesic. Then we may restate the original problem with a common dimension reduction, $\gamma = \mathbf{U}_r^{\top}\theta$, utilizing updated approximations over $r < D$ mixed active coordinates,

$$\text{maximize}_{\gamma \in \mathcal{Y}} \tau h_{\mathcal{L}}(\gamma; \mathbf{x}) + (1 - \tau) h_{\mathcal{W}}(\gamma; \mathbf{x}), \quad (17)$$

for all $\tau \in [0, 1]$. Although somewhat intuitive, we formalize stationarity of (17) as a Corollary (again, dropping the scenario variables for clarity):

Corollary 1: If $h_{\mathcal{L}}$ and $h_{\mathcal{W}}$ are ridge functions over a mutual r -subspace $\text{Range}(\mathbf{U}_r)$, $\mathbf{U}_r \in \mathbb{R}^{D \times r}$ such that $\mathbf{U}_r^{\top}\mathbf{U}_r = \mathbf{I}_r$, then the necessary conditions of Prop. 1 map to subspace coordinates.

Proof: By assumption, $f_{\mathcal{K}}(\theta) = h(\mathbf{U}_r^{\top}\theta)$ for either objective \mathcal{K} . Consequently, the gradient of $f_{\mathcal{K}}$ with respect to parameters θ is spanned by the subspace $\text{Range}(\mathbf{U}_r)$, $\nabla_{\theta} f_{\mathcal{K}} = \mathbf{U}_r \nabla_{\gamma} h_{\mathcal{K}} \implies \mathbf{U}_r^{\top} \nabla_{\theta} f_{\mathcal{K}} = \nabla_{\gamma} h_{\mathcal{K}}$. Similarly, the Hessian becomes $\nabla_{\theta}^2 f_{\mathcal{K}} = \mathbf{U}_r \nabla_{\gamma}^2 h_{\mathcal{K}} \mathbf{U}_r^{\top}$. Rewriting the necessary conditions per Prop. 1 such that

$$\nabla_{\theta}^2 \varphi_t(\theta(\tau)) = \mathbf{U}_r ((1 - \tau) \nabla_{\gamma}^2 h_{\mathcal{W}}(\tau) + \tau \nabla_{\gamma}^2 h_{\mathcal{L}}(\tau)) \mathbf{U}_r^{\top},$$

and assigning $\nabla_{\gamma}^2 \tilde{\varphi}(\tau) = (1 - \tau) \nabla_{\gamma}^2 h_{\mathcal{W}}(\tau) + \tau \nabla_{\gamma}^2 h_{\mathcal{L}}(\tau)$ implies $\mathbf{U}_r \nabla_{\gamma}^2 \tilde{\varphi}(\tau) \mathbf{U}_r^{\top} \dot{\theta}(\tau) = \nabla_{\theta} f_{\mathcal{W}}(\tau) - \nabla_{\theta} f_{\mathcal{L}}(\tau)$. For

active coordinates $\gamma(\tau) = \mathbf{U}_r^\top \boldsymbol{\theta}(\tau)$, the necessary conditions simplify given $\mathbf{U}_r^\top \nabla_{\boldsymbol{\theta}} f_{\mathcal{K}} = \nabla_{\gamma} h_{\mathcal{K}}$ such that $\nabla_{\gamma}^2 \tilde{\varphi}(\tau) \dot{\gamma}(\tau) = \nabla_{\gamma} h_{\mathcal{W}}(\tau) - \nabla_{\gamma} h_{\mathcal{L}}(\tau)$ and according to pseudo-time derivative that commutes with matrix multiplication $\dot{\gamma}(\tau) = \mathbf{U}_r^\top \dot{\boldsymbol{\theta}}(\tau)$. \square Once again, this optimization problem involves a closed and bounded feasible domain of parameter values $\mathcal{Y} = \{\gamma \in \mathbb{R}^r : \gamma = \mathbf{U}_r^\top \boldsymbol{\theta}, \boldsymbol{\theta} \in \mathcal{D}\}$ that remains convex for convex \mathcal{D} and a new subspace $\text{Range}(\mathbf{U}_r)$ —for hypercube \mathcal{D} , \mathcal{Y} is referred to as a *zonotope* [14], [35]. The utility of the dimension reduction is the ability to formulate a continuous trace of the Pareto front [16]—involving the inverse of a convex combination of Hessians—in fewer dimensions. In the context of a low-dimensional quadratic trace, the ridge approximations are $h_{\mathcal{K}}(\gamma; \mathbf{x}) = \gamma^\top \mathbf{Q}_{\mathcal{K}} \gamma + \mathbf{a}_{\mathcal{K}}^\top \gamma + c_{\mathcal{K}}$ given by $\mathcal{CVX}(\{(U_r^\top \boldsymbol{\theta}_n, f_{\mathcal{K}}(\boldsymbol{\theta}_n; \mathbf{x}))\})$ for corresponding network \mathcal{K} . Consequently, the resulting active coordinate trace is given by replacing $\boldsymbol{\theta}(\tau; \mathbf{x})$ with the parametrization $\gamma(\tau; \mathbf{x})$ in (10) for all $\gamma(\tau; \mathbf{x}) \in \mathcal{Y}$. This is supplemented by visualization in the case $r = 1$ or $r = 2$ providing *empirical evidence of convexity* and the ability to *visualize the resulting trace*—a powerful tool for facilitating *exploration, explainability, and comparison*.

However, given the transformation to active coordinates, we are now afforded the flexibility of selecting from an infinite number of inactive coordinates orthogonal to the quadratic trace of (17). Precisely, restricting approximations to an r -dimensional subspace defines a submanifold of approximately Pareto optimal solutions $\mathcal{M} \subseteq \mathbb{R}^D$ given as the product manifold $\mathcal{M} = \gamma([0, 1]; \mathbf{x}) \times \mathcal{Z}_{\gamma}$ where $\gamma([0, 1]; \mathbf{x})$ is the image of the quadratic trace over active coordinates and $\mathcal{Z}_{\gamma} = \{\zeta \in \mathbb{R}^{D-r} : \mathbf{U}_r \gamma + \mathbf{U}_r^\perp \zeta \in \mathcal{D}\}$. This offers a parametrization over a $(D - r + 1)$ -submanifold of approximately Pareto optimal solutions,

$$\boldsymbol{\theta}(\tau, \zeta; \mathbf{x}) = \mathbf{U}_r \gamma(\tau; \mathbf{x}) + \mathbf{U}_r^\perp \zeta, \quad (18)$$

for all $\zeta \in \mathcal{Z}_{\gamma(\tau; \mathbf{x})}$. Provided we remain off the boundary of \mathcal{M} such that we are only interested in a trace over the interior of \mathcal{D} , we can formulate a *geodesic* over \mathcal{M} of near Pareto optimal solutions,

$$\boldsymbol{\theta}(\tau; \mathbf{x}) = \mathbf{U}_r \gamma(\tau; \mathbf{x}) + \mathbf{U}_r^\perp \zeta(\tau), \quad (19)$$

where $\zeta(\tau)$ is a straight-line segment over the flat portion of the submanifold \mathcal{M} . It is conceivable to parametrize this line segment as the convex combination $\zeta(\tau) = \tau \zeta_1 + (1 - \tau) \zeta_0$ where $\zeta_0 = \underset{\zeta \in \mathcal{Z}_{\gamma(0; \mathbf{x})}}{\text{argmax}} (f_{\mathcal{W}}(\mathbf{U}_r \gamma(0; \mathbf{x}) + \mathbf{U}_r^\perp \zeta; \mathbf{x}))$

and $\zeta_1 = \underset{\zeta \in \mathcal{Z}_{\gamma(1; \mathbf{x})}}{\text{argmax}} (f_{\mathcal{L}}(\mathbf{U}_r \gamma(1; \mathbf{x}) + \mathbf{U}_r^\perp \zeta; \mathbf{x}))$ are the left ($\tau = 0$) and right ($\tau = 1$) inactive solutions, respectively. However, the subsequent inactive geodesic is expected to change throughputs significantly less—resulting in marginal improvements corresponding to inaccuracies induced by a choice of convex quadratic ridge profile.

Lastly, we write the image of throughputs over specific choices of the map $\boldsymbol{\theta}(\tau; \mathbf{x})$ as planar curves through Pareto space (fronts) $\mathcal{P} = \{\mathbf{f} \in \mathbb{R}^2 : \mathbf{f} = (f_{\mathcal{W}}, f_{\mathcal{L}}) \circ \boldsymbol{\theta}, \boldsymbol{\theta} \in \mathcal{D}\}$ containing the true (unknown) Pareto front. In this study, we consider three choices for attempting to approximate the

Pareto front: (i) the *geodesic front*

$$\{\mathbf{f}(\tau; \mathbf{x}) \in \mathbb{R}^2 : \mathbf{f}(\tau; \mathbf{x}) = (f_{\mathcal{W}}, f_{\mathcal{L}}) \circ \boldsymbol{\theta}(\tau; \mathbf{x})\} \quad (20)$$

per (19), (ii) the *linear front*

$$\{\mathbf{f}(\tau; \mathbf{x}) \in \mathbb{R}^2 : \mathbf{f}(\tau; \mathbf{x}) = (f_{\mathcal{W}}, f_{\mathcal{L}}) \circ \boldsymbol{\ell}(\tau; \mathbf{x})\} \quad (21)$$

where $\boldsymbol{\ell}(\tau; \mathbf{x}) = (1 - \tau) \boldsymbol{\theta}_0 + \tau \boldsymbol{\theta}_1$ defined by approximated left solution $\boldsymbol{\theta}_0$ (resp. right solution $\boldsymbol{\theta}_1$) to (7) using a state-of-the-art interior point solver defining $\boldsymbol{\ell}(\tau; \mathbf{x})$ as a geodesic over Euclidean Pareto manifold, and (iii) the *conditional front*

$$\left\{ \mathbf{f}(\tau; \mathbf{x}) \in \mathbb{R}^2 : \mathbf{f}(\tau; \mathbf{x}) = \int_{\mathcal{Z}_{\gamma(\tau; \mathbf{x})}} (f_{\mathcal{W}}, f_{\mathcal{L}}) \circ \boldsymbol{\theta}(t, \zeta; \mathbf{x}) dV(\zeta | \gamma(\tau; \mathbf{x})) \right\} \quad (22)$$

for conditional integral measure over inactive coordinates, $dV(\zeta | \gamma(\tau; \mathbf{x}))$. Additionally, the Pareto front approximations summarized in (17)-(19) can be contrasted with a “brute force” procedure which approximates (5) using a state-of-the-art interior point solver over discrete values of $\tau \in [0, 1]$. The resulting approximations and visualizations are compared in Section V.

C. Stretch Sampling

Projection of randomly sampled points $\{\boldsymbol{\theta}_n\}$ to subspace coordinates suffers from a *concentration of distances* phenomenon. This is demonstrated empirically, for our case, in Fig. 2. Generally, for otherwise unknown (random) subspace, we anticipate—with high probability—that distances between our random samples will uniformly scale by a factor of $\sqrt{r/D}$ for a random projection to subspace coordinates [36], [37]. Hence, given an otherwise unknown set of subspaces to be mixed resulting from Algorithm 1, we desire a procedure to systematically improve samples that *may concentrate* over the reduced dimension subspace.

To inform a supplementary space-filling design that produces additional samples for $\mathcal{CVX}(\{(U_r^\top \boldsymbol{\theta}_n, f_{\mathcal{K}}(\boldsymbol{\theta}_n; \mathbf{x}))\})$ such that $r \leq 2$, we implement a heuristic procedure referred to as *stretch sampling*. In brief, we systematically sample active coordinates supplemented by random draws over corresponding inactive coordinates, \mathcal{Z}_{γ} . Given data $\{\boldsymbol{\theta}_n\}$ and a basis for the mixed subspace \mathbf{U}_r , we consider the boundary of the convex hull of all projected samples $\{\mathbf{U}_r^\top \boldsymbol{\theta}_n\}$ and zonotope boundary $\partial \mathcal{Y}$. Utilizing linear interpolation of an ordered (clockwise or counterclockwise) set of vertices from either boundary, we take a chosen number of points to uniformly discretize—e.g., compute 25 points per boundary via piecewise linear interpolation of the ordered set of vertices over separate boundaries. The uniformly sampled boundary points define a Delaunay triangulation, and we return corresponding Voronoi centers not contained in the interior of the projected-data convex hull (exterior Voronoi centers) as new samples that are *stretched* beyond the extent of the projected data. These new samples—aggregated with the original data—fill out the remainder of the zonotope while inactive coordinates are drawn randomly from \mathcal{Z}_{γ} [14] at each new sample. The result produces improved sampling to

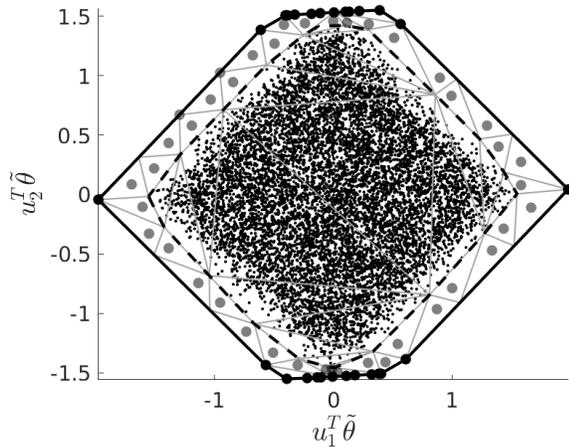


Fig. 2. Stretch sampling a two-dimensional zonotope. The boundary of the zonotope \mathcal{Y} (solid black curve) is depicted along with the boundary of the convex hull of data (dashed-black curve) and $N = 10,000$ projected random samples (black dots). Axes correspond to independent active coordinates. The resulting triangulation (gray mesh) is informed by 25 uniformly sampled points per boundary and exterior Voronoi centers (gray dots) constitute the new stretch samples used to improve the ridge approximations.

mitigate concerns of extrapolating surrogates over the reduced dimension subspace. A visualization of the new samples using 25 points per boundary is shown in Fig. 2.

V. SIMULATION AND RESULTS

We demonstrate the ideas proposed in Section IV on the LAA-Wi-Fi coexistence scenario described in Section II to maximize both throughputs simultaneously per (7) and subsequent simplification (17). We apply active subspaces to simplify the bi-criteria optimization problem (7) by focusing on a reduced set of mixed MAC and PHY layer parameter combinations informing a manifold of near Pareto optimal solutions (18). In particular, we contrast the geodesic trace (19) with a simple linear interpolation of left and right approximations to (7) computed with a state-of-the-art interior-point solver over the full dimensional space—informing (20) and (21), respectively. We also visualize a conditional approximation of the Pareto front by taking the conditional average of throughput responses over inactive coordinates subordinate to \mathcal{Z}_γ using (18)—computed as a Monte Carlo approximation of (22). The subsequent visualization of the Monte Carlo approximation of (22) emphasizes the expected range of throughput values captured, on average, over the near-optimal Pareto manifold (18). Table I lists the scenario parameters and parameter bounds (feasible region) used by the model for all computations and optimizations.

The baseline for all comparisons is a random grid search with solutions defined by sorting random evaluations (Pareto sort) to identify non-dominated designs in the random set [38]. These random samples are readily obtained from the random function evaluations utilized in Algorithm 1. The sorted random grid search is also supplemented by a “brute force” approximation of successive solutions to (7) with state-of-the-art interior-point solver over a uniform discretization of $\tau \in [0, 1]$ —referred to as the “common approach” in [39].

The numerical experiment utilizes $N = 10,000$ samples resulting in $N(m+1) = 180,000$ total function evaluations to compute forward differences with $\delta = 10^{-6}$. The eigenvalues resulting from Algorithm 1 indicate a steady decay devoid of dramatically different gaps in eigenvalues (see Fig. 3). This is contrasted with the resulting decay of eigenvalues associated with the principal Hessian directions [30] and quadratic active subspace approximations [29], both defined with respect to full-dimensional fits of throughputs (11), $\mathcal{CVX}(\{(\theta_n, f_n)\})$. Interestingly, the dimension reductions informed by the convex quadratic fits admit stronger decays in trailing eigenvalues (in both cases). This indicates a bias resulting from the choice of convex quadratic model over full-dimensional parameter space in contrast to the unbiased estimates of finite differences in Algorithm 1 [14].

Despite the steady decay in eigenvalues resulting from Algorithm 1, the separate active subspaces for corresponding throughputs informed relatively accurate degree-2 to degree-5 polynomial approximations—computed utilizing sets of coordinate-output pairs $\{(\tilde{\mathbf{W}}_{r,\mathcal{L}}^\top \tilde{\theta}_n, (f_{\mathcal{L}})_n)\}_{n=1}^N$ and $\{(\tilde{\mathbf{W}}_{r,\mathcal{W}}^\top \tilde{\theta}_n, (f_{\mathcal{W}})_n)\}_{n=1}^N$ —with varying coefficients of determination between 0.83 – 0.98 for both throughputs when $r = 1$ or $r = 2$.

The univariate subproblem in (16) can be visualized and, in this experiment, achieved a unique maximizing argument $s^* \in [0, 1]$ resulting in a mixed subspace with orthonormal basis given by two columns in a matrix $\mathbf{U}_r(s^*) = [\mathbf{u}_1 \ \mathbf{u}_2]$ over which separate quadratic ridge functions obtained roughly equal accuracy as approximations to their respective throughputs. The coefficients of determination varied monotonically and intersected over the Grassmannian parametrization. Consequently, the subproblem (16) results in an approximately equal criteria for the accuracy of the quadratic ridge profiles $h_{\mathcal{W}}$ and $h_{\mathcal{L}}$, i.e., $R_{\mathcal{L}}^2(s^*) \approx R_{\mathcal{W}}^2(s^*) \approx 0.86$.

Fig 4 depicts the condition numbers of convex combinations of quadratic Hessians resulting from a full-dimensional fit (11), i.e., $r = D$, and the chosen ($r = 2$)-dimensional subspace over the quadratic trace. Examining Fig. 4, we observe that the condition number over the two-dimensional subspace is more than two orders of magnitude lower than the full dimensional fit—achieving near optimal conditioning in the middle of the trace. We also note that at least some portion of the 2-dimensional trace passes through the domain while the full-dimensional trace does not. Consequently, the implicit regularization over the mixed subspace informed by subproblem (16) gives a more stable (and feasible) quadratic trace than simply solving (11) in full-dimensional space. The convex quadratic ridge approximations, Pareto trace approximations, non-dominated designs from the set of $N = 10,000$ random parameters, along with projected random samples and mixed subspace zonotope are shown in Fig. 5. The Pareto front approximation resulting from various traces are shown with the non-dominated designs in Fig. 6.

Observing Fig. 5, we depict traces and approximations over the two-dimensional subspace along with corresponding convex quadratic approximations and the data constituting $N = 10,000$ projected samples augmented by stretch sampling. The

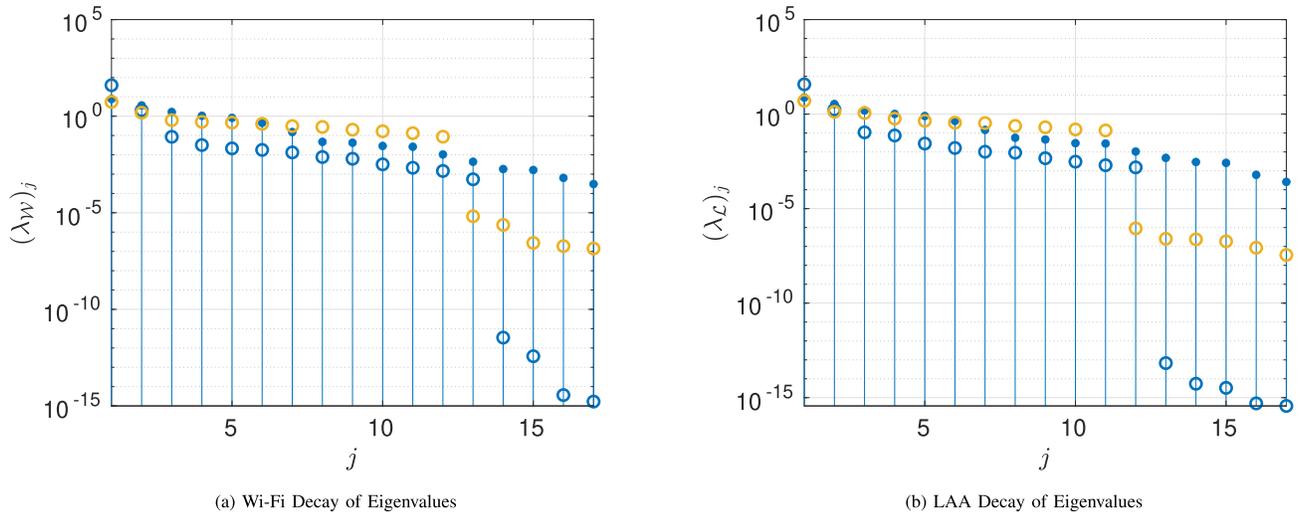


Fig. 3. Comparison of the eigenvalue decay for Wi-Fi and LAA throughputs. The eigenvalues resulting from Algorithm 1 are shown as the blue stem plot (solid blue dots) corresponding to $N = 10,000$ Monte Carlo samples. This is contrasted with eigenvalues of the convex quadratic Hessian (yellow circles) computed using the $N = 10,000$ unperturbed random function evaluations defining the convex problem (11) over the full-dimensional parameter space. Additionally, the eigenvalues of the active subspaces resulting from the full-dimensional convex quadratic fit as a surrogate (blue circles) is contrasted to the unbiased finite difference approach (solid blue dots).

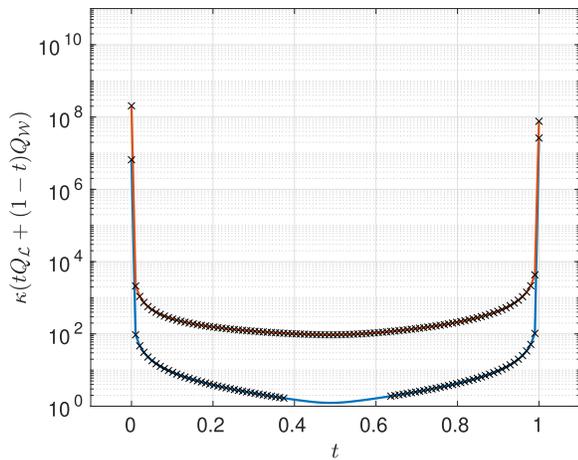


Fig. 4. Condition number of convex combinations of quadratic Hessians. The convex combination resulting from (11) over the full dimensional space (orange) is contrasted with the significantly improved conditioning over the mixed two-dimensional subspace (blue). Results are depicted at 100 points along the corresponding quadratic traces (10). Black crosses indicate points along the trace that do not pass through the domain.

red curve corresponds to the quadratic trace (19) while the blue curve corresponds to the linear trace projected to the subspace for the purposes of visualization and comparison. As a ground truth, we depict the collection of projected non-dominated designs (black circles) resulting from sorting the full data set—effectively this may be viewed as a Pareto optimality solution obtained via random grid search. The non-dominated designs are determined from the full set of $N = 10,000$ random samples sorted according to [38]. Note, it is not clear through this visualization that the non-dominated designs constitute elements of an alternative continuous approximation of the Pareto front—perhaps represented by an alternative low-dimensional manifold informed by a machine learning procedure.

In Fig. 6, we depict the corresponding throughput evaluations from conditional inactive samples over the quadratic trace as red dots along with associated Monte Carlo approximation of the conditional Pareto front (22). In other words, the solid-red line (overlapping the cloud of red dots) connects conditional averages of throughputs over 25 inactive samples along a corresponding discretization of 15 active coordinates over the subspace-based quadratic trace in Fig. 5. The visualization emphasizes that the throughputs change significantly less over the inactive coordinates in contrast to the range of values observed over the trace. Contrasting the red dots (conditional samples) and solid-red curve (approximated conditional front) to the colored scatter (averaged throughput response) paired with all random evaluations emphasizes that the near-Pareto optimal manifold (18) satisfies, on average, an averaged (summed) throughput, (i.e., $\tau = 0.5$) which is approximately greater than or equal to seven. The conditional front (22) moves approximately through the non-dominated designs of a random grid search. Hence, we have supplemented with a near-optimal (predominately flat) Pareto manifold (18) which is regularized as a solution over a low-dimensional subspace, and nearly captures the non-dominated designs from a random grid search, on average.

However, there are infinite θ in the original parameter space that correspond to points along the quadratic trace with a subset depicted in Fig. 5 (solid-red curve and dots)—i.e., infinitely many $D - r$ inactive coordinate values that may change throughputs albeit significantly less (roughly an additional 15%–20%) than the two mixed active coordinates, $\gamma_1 = \mathbf{u}_1^T \tilde{\theta}$ and $\gamma_2 = \mathbf{u}_2^T \tilde{\theta}$. To reconcile the choice of infinitely many inactive coordinates, we consider a discretization of 15 points along the geodesic trace (19) (curvature and discretization depicted in Fig. 5) and associated Pareto front approximation (20) (red-dashed line in Fig. 6). This is contrasted with τ discretized uniformly to approximate

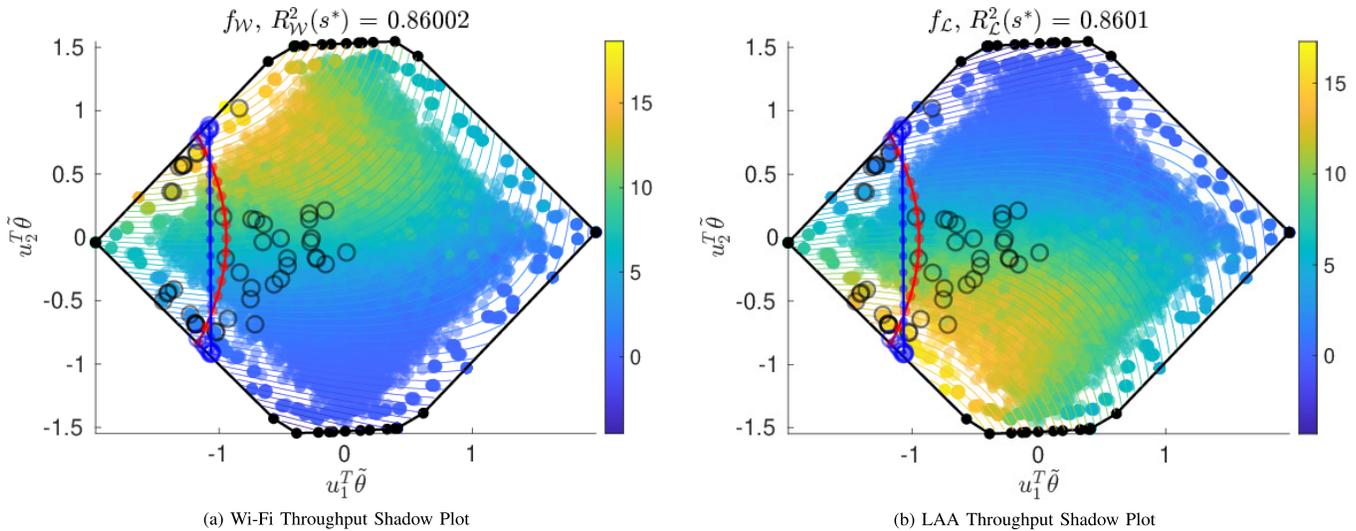


Fig. 5. Pareto trace of convex quadratic ridge profiles. The quadratic Pareto trace (red curve and dots) is overlaid on a shadow plot over the mixed coordinates (colored scatter) with the projected bounds and vertices (zonotope) of the domain (black dots and lines), \mathcal{Y} . The quadratic approximations (colored contours) are contrasted against the true function evaluations represented by the colors of the scatter. Also depicted is the projection of the non-dominated domain values from the $N = 10,000$ random samples (black circles). The active coordinate trace (red dots and curve) begins at the upper left-most boundary with near maximum quadratic Wi-Fi throughput and we move (smoothly) along a trajectory to the lower left-most boundary obtaining near maximum quadratic LAA throughput—maintaining an approximately best trade-off over the entire curve restricted to \mathcal{Y} . This is contrasted with a linear interpolation of full dimensional left ($\tau = 0$) and right ($\tau = 1$) approximations to (7) (blue dots and curve) and 15 successive approximations to (7) with uniform discretization of $\tau \in [0, 1]$ (blue circles) as the “brute force” solution.

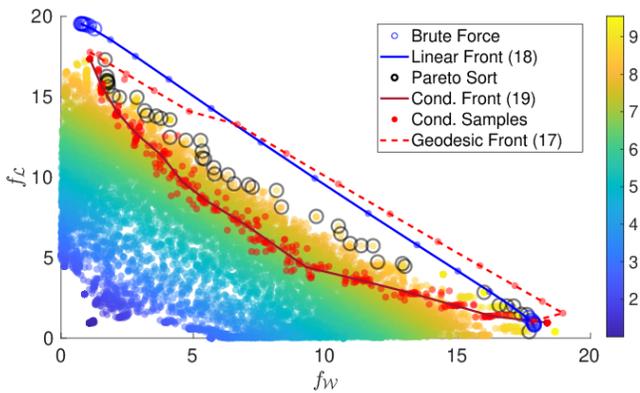


Fig. 6. Approximation of the Pareto front resulting from the quadratic trace. The conditional Pareto front (solid-red curve) is contrasted with the non-dominated random throughput values (black circles) and scatter of $N = 10,000$ random responses colored according to the averaged throughput ($\tau = 0.5$) response. The solid-red curve is a Monte Carlo approximation of the conditional Pareto front (22) with corresponding evaluations (overlaid red dots) used to compute conditional means. The dashed-red curve and coincident dots represent a discretization over geodesic (19) between left ($\tau = 0$) and right ($\tau = 1$) inactive maximizing arguments on the manifold of approximately Pareto optimal solutions (18). Also shown is the linear Pareto front (21) through the full parameter space (blue curve and coincident dots) and a set of 15 successive maximizations (7) over a uniform discretization of $\tau \in [0, 1]$ (blue circles) as the “brute force” solution.

15 successive optimizations solved in the full dimensional space (7) (blue circles) and the linear trace (21) (blue line and dots). There is reasonable agreement in the solutions produced by all three approaches. However, the geodesic trace (19) produces marginally better solutions from strict Wi-Fi optimization until the interesting intermediate design region (over $\tau \in [0, 0.5]$). The naive approach of aggregating successive optimizations (7) over uniform discretization of τ struggles to identify any intermediate combinations of

throughputs with solutions clustering towards one maximum or the other—a recognized issue in bi-criteria problems [39] remedied by our alternative parameterizations (20) and (21). Interestingly, and unexpectedly, the linear submanifold and subsequent front (21) perform comparatively well. This may suggest that the curved portion of the Pareto optimal manifold only affords minor improvements and an alternative Pareto manifold of near-optimal solutions could be built from a tubular neighborhood of the line segment interpolating left and right approximations to (7).

We additionally facilitate a comparison between multiple techniques of Pareto trace estimation in Fig. 6. In particular, the subspace-based quadratic trace (red-dashed curve) obtains larger pairs of throughput values than the random grid search (black circles) and fills out the front better than the brute force approach (blue circles). Indeed, the continuous traces from (20)-(21) all outperform the 10,000 non-dominated random values computed per [37] while the conditional front (22) serves as an approximation of these non-dominated (Pareto sort) random values. Additionally, the “brute force” state-of-the-art interior point solutions cluster on either edge of the Pareto front. The red-dashed curve does better than the majority of the linearly interpolated left and right interior-point solves (linear front, blue curve) although these curves are similar.

There is some bias in the approximation of the conditional front (solid-red curve) in Fig. 6 that is not a least-squares curve of non-dominated throughput values (black circles) potentially due in part to the quadratic ridge approximations or regularization by virtue of simplifying over a subspace. However, this issue is reconciled by constructing the geodesic trace (19) (with corresponding Pareto front as the dashed-red curve in Fig. 6) that dominates the random grid search

(black circles in Fig. 6) and the majority of the linear trace (blue curve in Fig. 6). We expect further refinements to these approximations (higher-order polynomials), increased subspace dimension, or numerical integration of the ODE in Prop. 1 with simplifications per Cor. 1 will further improve the near-optimal Pareto manifold and subsequent traces.

VI. CONCLUSION & FUTURE WORK

We have proposed a technique to simultaneously optimize the performance of two MNOs sharing limited unlicensed spectrum resources. An exploratory analysis utilizing an example of LAA coexistence with Wi-Fi network identified a common subspace-based dimension reduction of a basic model of network behavior. This enabled visualizations and low-dimensional approximations that led to a continuous approximation of the Pareto frontier for the bi-criteria problem of maximizing all convex combinations of network throughputs over MAC and PHY parameters. Such a result simplifies and regularizes the search for parameters that enable high quality performance of both networks, particularly compared to approaches that do not operate on a reduced parameter space. Analysis of the LAA-Wi-Fi example revealed an explainable and interpretable solution to an otherwise challenging problem—*devoid of any known convexity until subsequent exploration*.

Future work will incorporate alternative low-dimensional approximations including both cases of Grassmannian mixing and subspace unions to improve the trace. Additionally, these considerations for alternative subspace approximations will investigate the possibility of a dynamic subspace for application to dynamic environments interested in optimal control and further improvement in the Pareto trace approximation. We will also study alternative methods of subspace and non-linear dimension reduction to accelerate reinforcement learning over the various near-optimal Pareto manifolds. Future approaches will enable spectrum sharing for a variety of wireless communications models over unlicensed bands by simplifying the design of wireless network operation and architecture—ultimately quantifying model parameter combinations giving near-optimal KPI trade-offs.

REFERENCES

- [1] Z. J. Grey, S. Mosleh, J. D. Rezac, Y. Ma, J. B. Coder, and A. M. Dienstfrey, "Optimizing unlicensed band spectrum sharing with subspace-based Pareto tracing," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2020, pp. 1–7.
- [2] *Study on Licensed-Assisted Access to Unlicensed Spectrum (Release 13)*, DOCUMENT3GPP TR. 36.889 v13.0.0 ed., Jun. 2015.
- [3] H. He, H. Shan, A. Huang, L. Cai, and T. Quek, "Proportional fairness-based resource allocation for LTE-U coexisting with Wi-Fi," *IEEE Access*, vol. 5, pp. 4720–4731, 2016.
- [4] C. Cano, D. Leith, A. Garcia-Saavedra, and P. Serrano, "Fair coexistence of scheduled and random access wireless networks: Unlicensed LTE/Wi-Fi," *IEEE/ACM Trans. Netw.*, vol. 25, no. 6, pp. 3267–3281, Dec. 2017.
- [5] M. Mehrmouh, S. Roy, V. Sathya, and M. Ghosh, "On the fairness of Wi-Fi and LTE-LAA coexistence," *IEEE Trans. Cogn. Commun. Netw.*, vol. 4, no. 4, pp. 735–748, Dec. 2018.
- [6] Y. Gao, B. Chen, X. Chu, and J. Zhang, "Resource allocation in LTE-LAA and WiFi coexistence: A joint contention window optimization scheme," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2017, pp. 735–748.
- [7] Y. Gao, "LTE-LAA and WiFi in 5G NR unlicensed: Fairness, optimization and win-win solution," in *Proc. IEEE SmartWorld, Ubiquitous Intell. Comput., Adv. Trusted Comput., Scalable Comput. Commun., Cloud Big Data Comput., Internet People Smart City Innov. (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/IOP/SCI)*, Aug. 2019, pp. 1638–1643.
- [8] M. Han, S. Khairy, L. X. Cai, Y. Cheng, and R. Zhang, "Reinforcement learning for efficient and fair coexistence between LTE-LAA and Wi-Fi," *IEEE Trans. Veh. Technol.*, vol. 69, no. 8, pp. 8764–8776, Aug. 2020.
- [9] R. Yin, G. Yu, A. Maaref, and G. Y. Li, "A framework for co-channel interference and collision probability tradeoff in LTE licensed-assisted access networks," *IEEE Trans. Wireless Commun.*, vol. 15, no. 9, pp. 6078–6090, Sep. 2016.
- [10] S. Mosleh, Y. Ma, J. B. Coder, E. Perrins, and L. Liu, "Enhancing LAA co-existence using MIMO under imperfect sensing," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, Dec. 2019, pp. 1–7.
- [11] A. D. Shoaib, M. Derakhshani, and T. Le-Ngoc, "Efficient LTE/Wi-Fi coexistence in unlicensed spectrum using virtual network entity: Optimization and performance analysis," *IEEE Trans. Commun.*, vol. 66, no. 6, pp. 2617–2629, Jun. 2018.
- [12] W. Wu, Q. Yang, R. Liu, T. Q. S. Quek, and K. S. Kwak, "Online spectrum partitioning for LTE-U and WLAN coexistence in unlicensed spectrum," *IEEE Trans. Commun.*, vol. 68, no. 1, pp. 506–520, Jan. 2020.
- [13] A. K. Ajami, H. Artail, and K. G. Shin, "An optimization method for multi-operator and mixed-traffic LTE deployments in the unlicensed bands," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2018, pp. 1–6.
- [14] P. G. Constantine, *Active Subspaces: Emerging Ideas in Dimension Reduction for Parameter Studies*. Philadelphia, PA, USA: SIAM, Mar. 2015.
- [15] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [16] M. Bolten, O. Tanil Doganay, H. Gottschalk, and K. Klamroth, "Tracing locally Pareto optimal points by numerical integration," 2020, *arXiv:2004.10820*.
- [17] Y. Ma, S. Mosleh, and J. Coder, "A unified analytical approach to multi-cell LBT-based spectrum sharing systems," in *Proc. IEEE 91st Veh. Technol. Conf. (VTC-Spring)*, May 2020, pp. 1–6.
- [18] A. Babaei, J. Andreoli-Fang, Y. Pang, and B. Hamzeh, "On the impact of LTE-U on Wi-Fi performance," *Int. J. Wireless Inf. Netw.*, vol. 22, no. 4, pp. 336–344, Dec. 2015.
- [19] E. Perahia and R. Stacey, *Next Generation Wireless LANs: 802.11n and 802.11ac*. Cambridge, U.K.: Cambridge Univ. Press, Jun. 2013.
- [20] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, Mar. 2000.
- [21] Y. Song, K. W. Sung, and Y. Han, "Coexistence of Wi-Fi and cellular with listen-before-talk in unlicensed spectrum," *IEEE Commun. Lett.*, vol. 20, no. 1, pp. 161–164, Jan. 2016.
- [22] S. Mosleh, Y. Ma, J. D. Rezac, and J. B. Coder, "Dynamic spectrum access with reinforcement learning for unlicensed access in 5G and beyond," in *Proc. IEEE 91st Veh. Technol. Conf. (VTC-Spring)*, May 2020, pp. 1–7.
- [23] Y. Ma, D. G. Kuester, J. Coder, and W. F. Young, "Slot-jamming effect and mitigation between LTE-LAA and WLAN systems with heterogeneous slot durations," *IEEE Trans. Commun.*, vol. 67, no. 6, pp. 4407–4422, Jun. 2019.
- [24] W. Dong *et al.*, "IDEAL: Incentivized dynamic cellular offloading via auctions," in *Proc. IEEE INFOCOM*, Apr. 2013, pp. 755–763.
- [25] L. Gao, G. Iosifidis, J. Huang, L. Tassiulas, and D. Li, "Bargaining-based mobile data offloading," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1114–1125, Jun. 2014.
- [26] G. Iosifidis, L. Gao, J. Huang, and L. Tassiulas, "A double-auction mechanism for mobile data-offloading markets," *IEEE/ACM Trans. Netw.*, vol. 23, no. 5, pp. 1634–1647, Oct. 2015.
- [27] H. Yu, G. Iosifidis, J. Huang, and L. Tassiulas, "Auction-based competition between LTE unlicensed and Wi-Fi," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 1, pp. 79–90, Jan. 2017.
- [28] J. M. Lee, *An Introduction to Smooth Manifolds*, 2nd ed. New York, NY, USA: Springer, 2003.
- [29] Z. J. Grey and P. G. Constantine, "Active subspaces of airfoil shape parameterizations," *AIAA J.*, vol. 56, no. 5, pp. 2003–2017, May 2018.
- [30] K.-C. Li, "On principal Hessian directions for data visualization and dimension reduction: Another application of Stein's lemma," *J. Amer. Stat. Assoc.*, vol. 87, no. 420, pp. 1025–1039, Dec. 1992.

- [31] M. Grant and S. Boyd. (Mar. 2014). *CVX: MATLAB Software for Disciplined Convex Programming, Version 2.1*. [Online]. Available: <http://cvxr.com/cvx>
- [32] M. C. Grant and S. P. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control* (Lecture Notes in Control and Information Sciences), V. Blondel, S. Boyd, and H. Kimura, Eds. London, U.K.: Springer-Verlag, 2008, pp. 95–110. [Online]. Available: http://stanford.edu/~boyd/graph_dcp.html
- [33] O. Zahm, P. G. Constantine, C. Prieur, and Y. M. Marzouk, "Gradient-based dimension reduction of multivariate vector-valued functions," *SIAM J. Sci. Comput.*, vol. 42, no. 1, pp. A534–A558, Jan. 2020.
- [34] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM J. Matrix Anal. Appl.*, vol. 20, no. 2, pp. 303–353, Apr. 1999.
- [35] K. Fukuda, "From the zonotope construction to the Minkowski addition of convex polytopes," *J. Symbolic Comput.*, vol. 38, no. 4, pp. 1261–1272, Oct. 2004.
- [36] A. Eftekhari and M. Babaie, "Two-dimensional random projection," *Signal Process.*, vol. 91, no. 7, pp. 1589–1603, 2011. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016516841100003X>
- [37] R. G. Baraniuk and M. B. Wakin, "Random projections of smooth manifolds," *Found. Comput. Math.*, vol. 9, no. 1, pp. 51–77, 2009.
- [38] H. T. Kung, F. Luccio, and F. P. Preparata, "On finding the maxima of a set of vectors," *J. ACM*, vol. 22, no. 4, pp. 469–476, 1975.
- [39] I. Das and J. E. Dennis, "A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems," *Struct. Optim.*, vol. 14, no. 1, pp. 63–69, Aug. 1997.



Zachary J. Grey received the B.S. degree in aerospace engineering from Embry-Riddle Aeronautical University, Prescott, AZ, USA, in 2010, the M.S. degree in aero/astro engineering from Purdue University, West Lafayette, IN, USA, in 2014, and the Ph.D. degree in aerospace engineering from the University of Colorado Boulder in 2019.

As a full time (2010–2015) and a part-time (2015–2019) employee, he worked at Rolls-Royce Corporation, where he is designing gas-turbine engine flow-path components and modeling uncertainty. From 2019 to 2021, he completed the National Academies of Science, Engineering, and Medicine NRC Post-Doctoral Research Associateship at the National Institute of Standards and Technology (NIST), Applied and Computational Mathematics Division (ACMD—Mathematical Analysis and Modeling Group). Following his post-doctoral work, he joined NIST ACMD as a Staff at the NIST Boulder Campus. His research interests include theoretical considerations of differential geometry for a variety of applications involving model dimension reduction, uncertainty quantification, and manifold learning.



Susanna Mosleh received the Ph.D. degree in electrical engineering from The University of Kansas, Lawrence, KS, USA, in 2019.

Currently, she is a Professional Research Experience Program Research Associate with the Shared Spectrum Metrology Group, Communications Technology Laboratory, National Institute of Standards and Technology, Boulder, CO, USA, developing new methods for dynamic spectrum sharing and wireless coexistence networks. Her general research interests include emerging technologies for 5G cellular networks and the application of machine learning/deep learning in wireless communications.



Jacob D. Rezac received the B.S. and M.S. degrees in applied mathematics from the Colorado School of Mines, Golden, CO, USA, in 2011 and 2012, respectively, and the Ph.D. degree in applied mathematics from the University of Delaware, Newark, DE, USA, in 2017.

From 2017 to 2019, he was a Post-Doctoral Researcher with the University of Colorado at Boulder, Boulder, CO, USA, and the Communications Technology Laboratory, National Institute of Standards and Technology (NIST), Boulder. He has been a Mathematical Statistician with NIST since 2019. His research interests include theoretical and computational aspects of inverse problems as applied to problems in communications and the physical sciences



Yao Ma (Senior Member, IEEE) received the Ph.D. degree in electrical engineering from the National University of Singapore, in 2000.

From April 2000 to July 2015, he took various government, academic, and industrial positions, including a member of Technical Staff at the Centre for Wireless Communications, Singapore; a Post-Doctoral Fellow at the University of Toronto; an Assistant Professor at Iowa State University; a Research Assistant Professor at Wright State University; an Electronics Engineer at the U.S. Air Force Research Laboratory; and a Senior Computer Scientist at Infocortex Inc. Since July 2015, he has been with the Communication Technology Laboratory, National Institute of Standards and Technology, Boulder, Colorado. He has been a Project Leader of wireless coexistence since 2018. His work covers general area of wireless communication and networks, with recent focus on spectrum sharing of 4G, 5G-and-beyond and other networks, and testbed programming and experiments. He has been an Associate Editor for the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY since 2004 and a Past Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS.



Jason B. Coder currently leads the Shared Spectrum Metrology Group, National Institute of Standards and Technology's Communications Technology Laboratory. During his tenure at NIST, he has worked in fundamental EM measurements, electromagnetic compatibility, and antenna measurements. During his career, he has authored more than 70 publications, including peer-reviewed journal articles, refereed conference proceedings, invited papers and presentations, and NIST technical notes. His current research focuses on developing new measurement methods for spectrum sharing, wireless coexistence, and interference.

Mr. Coder currently serves as the Chair for the IEEE/ANSI C63.27 working group on wireless coexistence and the C63 Subcommittee on Spectrum Etiquette.



Andrew M. Dienstfrey received the Ph.D. degree in applied mathematics from the Courant Institute of Mathematical Sciences, New York University. He joined the Mathematical and Computational Sciences Division, NIST/Boulder, in 2000. His research interests include applications of mathematics to computational physics and numerical analysis.