

Transmit Beamforming for Physical Layer Multicasting[†]

*Nicholas D. Sidiropoulos*¹, *Timothy N. Davidson*², *Zhi-Quan Luo*³

Abstract

We consider the problem of downlink transmit beamforming for wireless transmission, and down-stream precoding for digital subscriber wireline transmission, in the context of common information broadcasting or multicasting applications wherein Channel State Information (CSI) is available at the transmitter. Unlike the usual “blind” isotropic broadcasting scenario, the availability of CSI allows transmit optimization. We adopt a minimum transmission power criterion, subject to prescribed minimum received Signal-to-Noise Ratios (SNRs) at each of the intended receivers. We also consider a related max-min SNR “fair” problem formulation subject to a transmit power constraint. We prove that both problems are NP-hard; however, suitable reformulation allows the successful application of semidefinite relaxation (SDR) techniques. SDR yields an approximate solution plus a bound on the optimum value of the associated cost/reward. We motivate SDR from a Lagrangian duality perspective, and assess its performance via pertinent simulations for the case of Rayleigh fading wireless channels. We find that SDR typically yields solutions that are within 3–4 dB of the optimum, which is often good enough in practice. In several scenarios, SDR generates exact solutions that meet the associated bound on the optimum value. This is illustrated using measured VDSL channel data, and far-field beamforming for a uniform linear transmit antenna array.

Keywords: Broadcasting, multicasting, minimization of total radiation power, downlink beamforming, VDSL precoding, semidefinite programming, semidefinite relaxation, convex optimization

TSP EDICS: SAM-c (Beamforming) [*Previous EDICS:* 1-ACOM, 3-MCHA, 3-TDSL]

[†] Original manuscript submitted to *IEEE Trans. on Signal Processing*, November 1, 2004. The work of the first author was supported in part by the U.S. ARO under ERO Contract N62558-03-C-0012, the E.U. under FP6 U-BROAD STREP # 506790, and the GSRT. The work of the second author was supported in part by the Natural Sciences and Engineering Research Council of Canada. The work of the third author was supported in part by the U.S. National Science Foundation, Grant No. DMS-0312416. Earlier version of part of this work appears in *Proc. IEEE SAM 2004* workshop.

¹ The (corresponding) author is with the Department of Electronic and Computer Engineering, Technical University of Crete, 73100 Chania - Crete, Greece; Fax: +30-28210-37542, Phone: +30-28210-37227, E-mail: nikos@telecom.tuc.gr

² The author is with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, ON L8S 4K1, Canada; E-mail: davidson@mcmaster.ca

³ The author is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, U.S.A.; E-mail: luozq@ece.umn.edu

I. INTRODUCTION

Consider a transmitter that utilizes an antenna array to broadcast information to multiple radio receivers within a certain service area. The traditional approach to broadcasting is to radiate transmission power isotropically, or with a fixed directional pattern. However, future digital video/audio/data broadcasting and multicasting applications are likely to be based on subscription to services, and hence it is plausible to assume that the transmitter can acquire channel state information (CSI) for all its intended receivers. The goal of this paper is to develop efficient algorithms for the design of broadcasting schemes that exploit this channel information in order to provide better performance than the traditional approaches.

Our design approach is based on providing Quality of Service (QoS) assurance to each of the receivers. Since the received Signal-to-Noise Ratio (SNR) determines the maximum achievable data rate and (essentially) determines the probability of error, it is an effective measure of the QoS. We consider two basic design problems. The first seeks to minimize the total transmission power (and thus leakage to neighboring co-channel transmissions), subject to meeting (potentially different) constraints on the received SNR for each individual intended receiver. The second is a “fair” design problem in which we attempt to maximize the smallest receiver SNR over the intended receivers, subject to a bound on the transmitted power. We will show that both these problems are NP-hard, but we will also show that designs that are close to being optimal can be efficiently obtained by employing Semidefinite Relaxation (SDR) techniques.

Our designs are initially developed for a wireless broadcast scenario in which each user employs a single receive antenna and the channel is modelled as being flat in frequency and quasi-static in time. However, the designs are also appropriate on a per-tone basis for Orthogonal Frequency Division Multiplexing (OFDM) and related multi-carrier systems, and, as we will show, they can be generalized in a straight-forward manner to single-carrier systems transmitting over frequency-selective channels. In addition to wireless systems, applications of the proposed methodology also appear in downstream multicast transmission for multi-carrier and single-carrier Digital Subscriber Line (DSL) systems. In this context, (linear) *precoding* of multiple DSL loops in the same binder that wish to subscribe to a common service (e.g., news feed, video-conference, or movie multicast) can be employed to improve Quality of Service and/or reduce far-end crosstalk (FEXT) interference to other loops in the binder. In scenarios in which the Customer-Premise Equipment (CPE) receivers are not physically co-located (as in residential service), or cannot be coordinated (legacy CPE), multiuser decoding of the downstream transmission is not feasible, while transmit precoding is viable. The most important difference between DSL and the

wireless multicast scenario is that DSL channels are diagonally-dominated. Still, exploiting the crosstalk coupling to reduce FEXT levels to other loops in the binder can provide significant performance gains, especially if (cooperative or competitive) power control is implemented.

It is interesting to note that, as of today, Internet multicasting (using IP's Multicast Backbone — MBone) is performed at the *network layer*, e.g., via packet-level flooding or spanning-tree access of the participant nodes and any intermediate nodes needed to access the participants. To complement that approach, what we advocate herein can be interpreted as judicious *physical layer multicasting*, that is enabled by i) the availability of multiple transmitting elements; ii) exploiting opportunities for joint beamforming/precoding; and iii) the availability of CSI at the transmitting node or one of its proxies. This is a cross-layer optimization approach that exploits information available at the physical layer to reduce relay retransmissions at the network layer, thus providing congestion relief and QoS guarantees. **Organization:** The rest of this paper is structured as follows. Section II contains the statement of the first QoS problem, and a review of the pertinent literature. The basic idea of semidefinite relaxation is explained in Section III. A complete algorithm is developed in Section IV. The case of max-min fair beamforming is treated in Section V (including its relationship to QoS beamforming), while the case of single-carrier transmission over frequency-selective channels is treated in Section VI. Duality theory offers some interesting insights on the structure of the problem at hand; these are presented in Section VII. Extensive simulation results are presented in Section VIII. These include simulations for Rayleigh fading wireless channels, experiments with measured VDSL channels, far-field beamforming for a uniform linear transmit antenna array, and other scenarios meant to unveil different aspects of the problem and the approximate solutions developed herein. Conclusions are drawn in Section IX, where some additional applications of the proposed methodology in problems with a more classical signal processing flavor (e.g., filter design) are also mentioned. Proofs are deferred to the Appendices.

Notation: We use lowercase boldface letters to denote column vectors, and uppercase bold letters to denote matrices. $(\cdot)^T$ denotes transpose, while $(\cdot)^H$ denotes Hermitian (conjugate) transpose. $\text{Re}\{\cdot\}$ extracts the real part of its argument, and $\text{Im}\{\cdot\}$ the imaginary part.

II. DATA MODEL AND PROBLEM STATEMENT

Consider a wireless scenario incorporating a single transmitter with N antenna elements and M receivers each with a single antenna. Let \mathbf{h}_i denote the $N \times 1$ complex vector that models the propagation loss and phase shift of the frequency-flat quasi-static channel from each transmit antenna to the receive antenna of user $i \in \{1, \dots, M\}$, and let \mathbf{w}^H denote the beamforming weight vector applied to the N

transmitting antenna elements. If the signal to be transmitted is zero-mean and white with unit variance, and if the noise¹ at receiver i is zero-mean and white with variance σ_i^2 , then the receiver SNR for the i th user is $|\mathbf{w}^H \mathbf{h}_i|^2 / \sigma_i^2$. Therefore, the design of the beamformer that minimizes the transmit power, subject to (possibly different) constraints on the received SNR of each user can be written as

$$\begin{array}{l} \mathcal{Q}: \\ \min_{\mathbf{w} \in \mathbb{C}^N} \|\mathbf{w}\|_2^2 \\ \text{subject to: } |\mathbf{w}^H \mathbf{h}_i|^2 \geq c_i, \quad i \in \{1, \dots, M\}. \end{array}$$

Remark 1: One could think of imposing the stricter constraints $\mathbf{w}^H \mathbf{h}_i = \sqrt{c_i}$, $\forall i$, in order to avoid the need for single-tap equalization at the receivers. However, we are interested in the practically important case of $M > N$, wherein the stricter constraints generically yield an over-determined system of equations, and thus an infeasible problem. On the other hand, it is easy to see that problem \mathcal{Q} is always feasible, provided of course that none of the channel vectors is identically zero.

Problem \mathcal{Q} is formulated under the assumption that the design centre (usually the transmitter) has knowledge of the channel vector \mathbf{h}_i (and the noise variance σ_i^2) for each user. This can be accomplished in a straight-forward manner in fixed wireless systems and Time-Division-Duplex (TDD) systems. In other systems it can be accomplished through the use of beacon signals, periodically transmitted from the broadcasting station (and typically embedded in the transmission). The receiving radios can then feed back their CSI through a feedback channel. For the purposes of this paper, we will assume that the design centre has perfect knowledge of the channel vectors, but extensions to cases of imperfect knowledge are under development.

Problem \mathcal{Q} is a quadratically constrained quadratic programming (QCQP) problem, but unfortunately the constraints are not convex². Non-convexity *per se* does not mean that the problem is hard to solve; however, we have the following claim, whose proof can be found in Appendix I:

Claim 1: The QoS problem \mathcal{Q} is NP-hard.

The implication of Claim 1 is that if an algorithm could solve an arbitrary instance of problem \mathcal{Q} in polynomial time, it would then be possible to solve a whole class of computationally very difficult problems in polynomial time [4]. The current scientific consensus indicates that this is unlikely.

¹which may include unmodelled interference

²This is easy to see for $N = 1$, in which case each constraint requires that the magnitude of w be greater than a constant.

A. Review of Pertinent Prior Art

The above problem is reminiscent of some closely-related problems. For $M = 1$, the optimum \mathbf{w} is a matched filter. When the channel vectors \mathbf{h}_i span a ball or ellipsoid about a “nominal” channel vector³, the problem can be transformed *exactly* into a second-order cone program, and hence can be efficiently solved [12]. Unfortunately, this transformation cannot be employed in the case of finitely-many channel vectors (intended receivers).

Another closely-related work is that in [1] (and references therein), which considers the problem of multiuser transmit beamforming for the cellular downlink. The key difference between [1] and our formulation is that the authors of [1] consider the transmission of independent information to each of the downlink users, whereas we focus on (common information) broadcast. The mathematical problems are not equivalent. A fundamental difference is that our problem is NP-hard, whereas the formulation in [1] can be efficiently solved. To further appreciate the difference intuitively, we point out that in the generic case of our formulation most of the SNR constraints will be inactive at the optimum (i.e., most of the constraints will be over-satisfied). Consider, for example, the case of two closely-located receivers with different SNR requirements: one of the two associated constraints will be over-satisfied at the optimum. On the other hand, it is proven in [1] that in the formulation of [1] the constraints are always met with equality at the optimum. The important common denominator of our work and [1] is the use of semidefinite programming tools.

III. RELAXATION

Towards solving our problem, we first re-cast it as follows:

$$\begin{aligned} & \min_{\mathbf{w}} \text{trace}(\mathbf{w}\mathbf{w}^H) \\ & \text{subject to: } \text{trace}(\mathbf{w}\mathbf{w}^H \mathbf{Q}_i) \geq c_i, \quad i \in \{1, \dots, M\}, \end{aligned}$$

where we have used the fact that $\mathbf{h}_i^H \mathbf{w}\mathbf{w}^H \mathbf{h}_i = \text{trace}(\mathbf{h}_i^H \mathbf{w}\mathbf{w}^H \mathbf{h}_i) = \text{trace}(\mathbf{w}\mathbf{w}^H \mathbf{h}_i \mathbf{h}_i^H)$, and we have defined $\mathbf{Q}_i := \mathbf{h}_i \mathbf{h}_i^H$. Now consider the following reformulation of the problem:

$$\begin{aligned} & \min_{\mathbf{X} \in \mathbb{C}^{N \times N}} \text{trace}(\mathbf{X}) \\ & \text{subject to: } \text{trace}(\mathbf{X} \mathbf{Q}_i) \geq c_i, \quad i \in \{1, \dots, M\}, \\ & \mathbf{X} \succeq \mathbf{0}, \\ & \text{rank}(\mathbf{X}) = 1, \end{aligned}$$

³This implies a continuum of intended receivers.

where now \mathbf{X} is an $N \times N$ complex matrix, and the inequality $\mathbf{X} \succeq \mathbf{0}$ means that the matrix \mathbf{X} is symmetric positive semidefinite. Note that, in the above *equivalent* formulation of our problem, the cost function is linear in \mathbf{X} ; the trace constraints are linear inequalities in \mathbf{X} , and the set of symmetric positive semidefinite matrices is convex; however the rank constraint on \mathbf{X} is not convex⁴. The important observation is that the above problem is in a form suitable for semidefinite relaxation (SDR) (see, e.g., [8]); that is, dropping the rank-one constraint, one obtains the relaxed problem

$$\begin{aligned} & \min_{\mathbf{X} \in \mathbb{C}^{N \times N}} \quad \text{trace}(\mathbf{X}) \\ & \text{subject to: } \text{trace}(\mathbf{X}\mathbf{Q}_i) \geq c_i, \quad i \in \{1, \dots, M\}, \quad \text{and } \mathbf{X} \succeq \mathbf{0}, \end{aligned}$$

which is a semidefinite programming problem (SDP), albeit not yet in standard form. In order to put it in standard form, we add M “slack” variables $s_i \in \mathbb{R}$, $i \in \{1, \dots, M\}$, one for each trace constraint. In this way, we obtain the program

$$\begin{aligned} \mathcal{Q}_r: & \\ & \min_{\mathbf{X} \in \mathbb{C}^{N \times N}, s_i \in \mathbb{R}} \quad \text{vec}(\mathbf{I}_N)^T \text{vec}(\mathbf{X}) \\ & \text{s.t.: } \text{vec}(\mathbf{Q}_i^T)^T \text{vec}(\mathbf{X}) - s_i = c_i, \quad i \in \{1, \dots, M\} \\ & \quad s_i \geq 0, \quad i \in \{1, \dots, M\}, \quad \text{and } \mathbf{X} \succeq \mathbf{0} \end{aligned}$$

which is now expressed in a standard form used by SDP solvers, such as SeDuMi [10]. Here, \mathbf{I}_N is the identity matrix of size $N \times N$.

SDP problems can be efficiently solved using interior point methods, at a complexity cost that is at most $O((M + N^2)^{3.5})$, and is usually much less. SeDuMi [10] is a MATLAB implementation of modern interior point methods for SDP that is particularly efficient for up to moderate-sized problems, as is the case in our context. Typical run times for realistic choices of N and M are about 1/10 sec, on a typical PC.

IV. ALGORITHM

Due to the relaxation, the matrix \mathbf{X}_{opt} obtained by solving the SDP in Problem \mathcal{Q}_r will not be rank-one in general. If it is, then its principal component will be the optimal solution to the original problem. If not, then $\text{trace}(\mathbf{X}_{\text{opt}})$ is a lower bound on the power needed to satisfy the constraints. This comes from the fact that we have removed one of the original problem’s constraints. Researchers in optimization have recently

⁴The sum of two rank-one matrices has generic rank two.

developed ways of generating good solutions to the original problem, \mathcal{Q} , from \mathbf{X}_{opt} , [8], [13], [11]. This process is based on *randomization*: using \mathbf{X}_{opt} to generate a set of candidate weight vectors, $\{\mathbf{w}_\ell\}$, from which the “best” solution will be selected. We consider three methods for generating the \mathbf{w}_ℓ ’s, which have been designed so that their computational cost is negligible compared to that of computing \mathbf{X}_{opt} . (For consistency, the principal component is also included in the set of candidates.) In the first method (**randA**), we calculate the eigen-decomposition of $\mathbf{X}_{\text{opt}} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^H$ and choose \mathbf{w}_ℓ such that $\mathbf{w}_\ell = \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{e}_\ell$, where \mathbf{e}_ℓ is uniformly distributed on the unit sphere. This ensures that $\mathbf{w}_\ell^H\mathbf{w}_\ell = \text{trace}(\mathbf{X}_{\text{opt}})$, irrespective of the particular realization of \mathbf{e}_ℓ . In the second method (**randB**), inspired by Tseng [11], we choose \mathbf{w}_ℓ such that $[\mathbf{w}_\ell]_i = \sqrt{[\mathbf{X}_{\text{opt}}]_{ii}} e^{j\theta_{\ell,i}}$, where the $\theta_{\ell,i}$ are independent and uniformly distributed on $[0, 2\pi)$. This randomization ensures that $|[\mathbf{w}_\ell]_i|^2 = [\mathbf{X}_{\text{opt}}]_{ii}$. The third method (**randC**), motivated by successful applications in related QCQP problems [7], uses $\mathbf{w}_\ell = \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{v}_\ell$, where \mathbf{v}_ℓ is a vector of zero-mean, unit-variance complex circularly symmetric uncorrelated Gaussian random variables. This ensures that $E[\mathbf{w}_\ell\mathbf{w}_\ell^H] = \mathbf{X}_{\text{opt}}$ [7].

For both **randA** and **randB**, $\|\mathbf{w}_\ell\|_2^2 = \text{trace}(\mathbf{X}_{\text{opt}})$, and hence when $\text{rank}(\mathbf{X}_{\text{opt}}) > 1$, at least one of the constraints $|\mathbf{w}_\ell^H\mathbf{h}_i|^2 \geq c_i$ will be violated.⁵ However, a feasible weight vector can be found by simply scaling \mathbf{w}_ℓ so that all the constraints are satisfied. Under **randC**, $\|\mathbf{w}_\ell\|_2^2$ depends on the particular realization of \mathbf{v}_ℓ , but again the resulting \mathbf{w}_ℓ can be scaled to the minimum length necessary to satisfy the constraints. The “best” of these randomly generated weight vectors is the one that requires the smallest scaling. For convenience, we have summarized the algorithm in Table I, which includes a simple MATLAB interface to SeDuMi [10] for the solution of the semidefinite relaxation, \mathcal{Q}_r . We point out that we have not yet been able to obtain theoretical *a priori* bounds on the extent of the sub-optimality of solutions generated in this way, but our simulation results are quite encouraging.

V. MAX-MIN FAIR BEAMFORMING

We now consider the related problem of maximizing the minimum received SNR over all receivers, subject to a bound on the transmitted power. That is,

$$\begin{array}{l} \mathcal{F}: \\ \max_{\mathbf{w} \in \mathbb{C}^N} \min_i \{ |\mathbf{w}^H \mathbf{h}_i|^2 / \sigma_i^2 \}_{i=1}^M \\ \text{subject to: } \|\mathbf{w}\|_2^2 \leq P \end{array}$$

⁵Recall that because of the relaxation, $\text{trace}(\mathbf{X}_{\text{opt}})$ is a lower bound on the energy of the optimal weight vector for the original problem.

It is easy to see that the constraint in Problem \mathcal{F} should be met with equality at an optimum, for otherwise \mathbf{w} could be scaled up, thereby improving the objective and contradicting optimality. Thus we can focus on the equality-constrained problem. With a scaling of the optimization variable $\mathbf{w} = \sqrt{P}\tilde{\mathbf{w}}$, and scaling of the channel vectors $\tilde{\mathbf{h}}_i = \mathbf{h}_i/\sigma_i$, the equality-constrained problem can be written as

$$\begin{aligned} \max_{\tilde{\mathbf{w}}} \min_i \left\{ P |\tilde{\mathbf{w}}^H \tilde{\mathbf{h}}_i|^2 \right\}_{i=1}^M &= P \max_{\tilde{\mathbf{w}}} \min_i \left\{ |\tilde{\mathbf{w}}^H \tilde{\mathbf{h}}_i|^2 \right\}_{i=1}^M \\ \text{subject to: } \|\tilde{\mathbf{w}}\|_2^2 &= 1. \end{aligned}$$

It is clear that P is immaterial with respect to optimization; the solution scales up with \sqrt{P} , while the optimum value scales up with P . We can therefore restrict our attention to the problem (dropping the tildes for brevity):

$$\begin{aligned} \max_{\mathbf{w}} \min_i \left\{ |\mathbf{w}^H \mathbf{h}_i|^2 \right\}_{i=1}^M \\ \text{subject to: } \|\mathbf{w}\|_2^2 &= 1. \end{aligned}$$

Some discussion is due at this point on the relationship between the two problem formulations: the original QoS formulation that seeks to minimize the total transmit power subject to prescribed lower bounds, c_i , on the received signal powers; and the max-min ‘‘fair’’ one that aims to maximize the received signal power of the weakest user⁶ subject to an overall transmit power constraint. Suppose that all c_i ’s are equal to c , and the QoS formulation yields a beamformer \mathbf{w}_q and associated minimum transmit power P_q . Then we can scale the solution of the max-min fair beamformer to power P_q , and this scaled max-min fair solution, denoted \mathbf{w}_f , will be optimum for

$$\begin{aligned} \max_{\mathbf{w}} \min_i \left\{ |\mathbf{w}^H \mathbf{h}_i|^2 \right\}_{i=1}^M \\ \text{subject to: } \|\mathbf{w}\|_2^2 &= P_q. \end{aligned}$$

As a result, since \mathbf{w}_q already attains $|\mathbf{w}_q^H \mathbf{h}_i|^2 \geq c, \forall i$, it follows that $|\mathbf{w}_f^H \mathbf{h}_i|^2 \geq c, \forall i$. Hence \mathbf{w}_f also satisfies the constraints of the QoS formulation, and at the same power as \mathbf{w}_q . It follows that \mathbf{w}_f is equivalent to \mathbf{w}_q . This shows that

Claim 2: The QoS problem formulation and the max-min fair problem formulation are equivalent up to scaling in the case that all the c_i ’s are equal.

⁶Note that this weakest user is not known beforehand; instead, it is determined by the optimization.

When the c_i 's are different, however, the two problem formulations generally yield different beamformers. This can be intuitively appreciated by noting that the max-min fair formulation aims to maximize the minimum received signal power, without regard to the *distribution* of signal powers which are above the attained minimum. In particular, to conserve power, the max-min fair beamformer will tend to equalize the received powers, if possible. On the other hand, the QoS formulation explicitly guarantees the desired received power level at each node.

Claim 2 implies an indirect way of solving the max-min fair problem:

Corollary 1: One way to solve the max-min fair problem is to solve the QoS problem with $c_i = 1$, $\forall i \in \{1, \dots, M\}$, then scale the resulting solution to the desired power P .

Interestingly, even though QoS and max-min fair *optimal solutions* are not equivalent up to scaling when the c_i 's are not equal, the two problem formulations are still equivalent in a broader sense. To see this, note that we can simply divide $\sqrt{c_i}$ into \mathbf{h}_i in the QoS formulation, thus obtaining an equivalent problem where all the constraints are equal. This can be solved as a max-min fair problem, and the solution can be scaled to satisfy the constraints at minimum power. This establishes the following.

Claim 3: The QoS and max-min-fair problems are equivalent, in the sense that i) each instance of one is equivalent, up to scaling, to some instance of the other; and ii) an exact algorithm that solves one can also be used to solve the other.

From the above, and Claim 1, it follows that

Claim 4: The max-min fair problem \mathcal{F} is NP-hard.

If the QoS problem could be solved *exactly*, there would have been no need for a separate algorithm for the max-min fair problem. However, we can only solve the QoS problem approximately (c.f., randomization post-processing of the generally higher-rank solution). Due to this, it is of interest to develop a customized SDR algorithm directly for the max-min fair problem formulation. Again using $\mathbf{h}_i^H \mathbf{w} \mathbf{w}^H \mathbf{h}_i = \text{trace}(\mathbf{w} \mathbf{w}^H \mathbf{h}_i \mathbf{h}_i^H)$, and $\mathbf{Q}_i := \mathbf{h}_i \mathbf{h}_i^H$, we re-cast the max-min fair problem as follows

$$\begin{aligned} & \max_{\mathbf{X} \in \mathbb{C}^{N \times N}} \min_{i=1, \dots, M} \text{trace}(\mathbf{X} \mathbf{Q}_i) \\ & \text{subject to: } \text{trace}(\mathbf{X}) = P, \quad \mathbf{X} \succeq \mathbf{0}, \\ & \text{rank}(\mathbf{X}) = 1. \end{aligned}$$

Dropping the rank constraint, we obtain the relaxation

$$\begin{aligned} & \max_{\mathbf{X} \in \mathbb{C}^{N \times N}} \min_{i=1, \dots, M} \text{trace}(\mathbf{X}\mathbf{Q}_i) \\ & \text{subject to: } \text{trace}(\mathbf{X}) = P, \quad \mathbf{X} \succeq \mathbf{0}. \end{aligned}$$

Introducing an additional variable, t , this relaxation can be equivalently written as

$$\begin{aligned} & \max_{\mathbf{X} \in \mathbb{C}^{N \times N}, t \in \mathbb{R}} t \\ & \text{subject to: } \text{trace}(\mathbf{X}\mathbf{Q}_i) \geq t, \quad \forall i \in \{1, \dots, M\}, \\ & \text{trace}(\mathbf{X}) = P, \quad \mathbf{X} \succeq \mathbf{0}. \end{aligned}$$

Further introducing M non-negative real slack variables, one for each inequality constraint, we convert the problem to an equivalent one involving only equality, non-negativity, and positive-semidefinite constraints

\mathcal{F}_r :

$$\begin{aligned} & \min_{\substack{t, s_i \in \mathbb{R} \\ \mathbf{X} \in \mathbb{C}^{N \times N}}} -t \\ & \text{subject to: } -t - s_i + \text{vec}(\mathbf{Q}_i^T)^T \text{vec}(\mathbf{X}) = 0, \quad \forall i, \\ & \text{vec}(\mathbf{I}_N)^T \text{vec}(\mathbf{X}) = P, \\ & \mathbf{X} \succeq \mathbf{0}, \quad s_i \geq 0, \quad \forall i, \quad t \geq 0. \end{aligned}$$

This problem is formatted for direct solution via SeDuMi [10]. Table II provides a suitable MATLAB interface for solving this relaxation. Post-processing of the solution of the relaxed problem to approximate the solution of the original max-min-fair problem can be accomplished using randA, randB, randC; but the selection criterion is different, see Table II.

In closing this section, we would like to point out connections between Problems \mathcal{F} and \mathcal{F}_r and the problem of maximizing the common mutual information of the (non-degraded) Gaussian broadcast channel in which the transmitter has N antennas and each of the M (non-cooperative) receivers has a single antenna. If \mathbf{X} denotes the covariance of the transmitted signal, then the maximum achievable common information rate (in the sense of Shannon) can be written as (see, e.g., [6] and references therein)

$$C := \max_{\substack{\mathbf{X} \succeq \mathbf{0}, \\ \text{trace}(\mathbf{X}) \leq P}} \min_i \left\{ \log(1 + \tilde{\mathbf{h}}_i^H \mathbf{X} \tilde{\mathbf{h}}_i) \right\}_{i=1}^M,$$

where, as earlier, $\tilde{\mathbf{h}}_i = \mathbf{h}_i / \sigma_i$. Alternatively, we can rewrite this max-min problem as

$$\max t$$

$$\text{subject to } \mathbf{X} \succeq \mathbf{0}, \text{ trace}(\mathbf{X}) \leq P, \log(1 + \tilde{\mathbf{h}}_i^H \mathbf{X} \tilde{\mathbf{h}}_i) \geq t, \quad \forall i.$$

By the monotonicity of the “log” function, the above problem is further equivalent to

$$\max \tilde{t}$$

$$\text{subject to } \mathbf{X} \succeq \mathbf{0}, \text{ trace}(\mathbf{X}) \leq P, \tilde{\mathbf{h}}_i^H \mathbf{X} \tilde{\mathbf{h}}_i \geq \tilde{t}, \quad \forall i,$$

in the sense that they yield the same optimal transmit covariance matrix \mathbf{X} . Since $\tilde{\mathbf{Q}}_i = \tilde{\mathbf{h}}_i \tilde{\mathbf{h}}_i^H$ has rank 1, the latter problem is identical to Problem \mathcal{F}_r if the noise powers $\{\sigma_i\}$ are the same for all users. In other words, in this case the semidefinite relaxation of Problem \mathcal{F} actually yields a transmit covariance matrix that achieves the maximum common information rate C . In a similar manner, we can argue that the rank-1 transmit covariance matrix obtained from Problem \mathcal{F} achieves the maximum common information rate under the restriction that beamforming is employed. However, the latter rate can be significantly lower than C for a large number of users [6]. Nonetheless, from a practical perspective, beamforming is attractive because it is simple to implement,⁷ requiring only a single standard AWGN channel encoder at the transmitter. In contrast, achieving the maximum common information rate C in general requires higher-rank transmit covariance matrix \mathbf{X} . In that case, a weighted sum of multiple independent signals is transmitted from each antenna, with each independent signal requiring a separate AWGN channel encoder, therefore resulting in significant implementation complexity. Hence, the beamforming strategy considered in this paper trades off a potential reduction in the maximum common information rate for implementation simplicity.

VI. THE CASE OF FREQUENCY-SELECTIVE MULTIPATH

Although we have focused our attention so far on frequency-flat fading channels, the situation is quite similar for frequency-selective (intersymbol-interference) channels. Let $\mathbf{h}_i^{(\ell)}$ denote the ℓ -th $N \times 1$ vector tap of the baseband-equivalent discrete-time impulse response of the multipath channel between the transmitter antenna array and the (single) receive antenna of receiver i . Assume that delay spread is limited⁸ to L non-zero vector channel taps. Define the channel matrix for the i -th receiver as

$$\mathbf{H}_i := \left[\mathbf{h}_i^{(0)}, \dots, \mathbf{h}_i^{(L-1)} \right].$$

⁷a properly weighted common temporal signal is transmitted from each antenna

⁸or, essentially limited; the remaining taps can be treated as interference.

Beamforming the transmit array with a fixed (time-invariant) \mathbf{w}^H yields a scalar equivalent channel from the viewpoint of the i -th receiver, whose scalar taps are given by

$$\left[\bar{h}_i^{(0)}, \dots, \bar{h}_i^{(L-1)}\right]^T = \left[\mathbf{w}^H \mathbf{h}_i^{(0)}, \dots, \mathbf{w}^H \mathbf{h}_i^{(L-1)}\right]^T,$$

or, in vector form,

$$\bar{\mathbf{h}}_i^T = \mathbf{w}^H \mathbf{H}_i.$$

Now, if a Viterbi equalizer is used for sequence estimation at the receiver, then the parameter that determines performance is [3]:

$$\|\bar{\mathbf{h}}_i\|_2^2 = \mathbf{w}^H \mathbf{H}_i \mathbf{H}_i^H \mathbf{w} = \text{trace}(\mathbf{w} \mathbf{w}^H \mathbf{Q}_i),$$

where now $\mathbf{Q}_i := \mathbf{H}_i \mathbf{H}_i^H$ and is generally of higher rank than before, but otherwise things remain conceptually the same. In particular, the relaxations \mathcal{Q}_r , \mathcal{F}_r and the algorithms in Tables I, II can be employed as they were in the frequency-flat case – the only difference is in the definition of the \mathbf{Q}_i matrices.

VII. INSIGHTS AFFORDED VIA DUALITY

Let us return to our original problem, \mathcal{Q} :

$$\begin{aligned} & \min_{\mathbf{w} \in \mathbb{C}^N} \|\mathbf{w}\|_2^2 \\ & \text{subject to: } |\mathbf{w}^H \mathbf{h}_i|^2 \geq c_i, \quad i \in \{1, \dots, M\}. \end{aligned}$$

We will now gain some insight into the quality of the solution generated by the semidefinite relaxation of \mathcal{Q} using bounds obtained from duality. For convenience, we first convert the problem to real-valued form; this yields a $2N \times 1$ vector of real variables, $\mathbf{x} := \left[\text{Re}\{\mathbf{w}\}^T \text{Im}\{\mathbf{w}\}^T\right]^T$, and the \mathbf{Q}_i 's are now $2N \times 2N$ symmetric matrices of rank 2: $\mathbf{Q}_i := \mathbf{g}_i \mathbf{g}_i^T + \bar{\mathbf{g}}_i \bar{\mathbf{g}}_i^T$, where $\mathbf{g}_i := \left[\text{Re}\{\mathbf{h}_i\}^T \text{Im}\{\mathbf{h}_i\}^T\right]^T$, and $\bar{\mathbf{g}}_i := \left[\text{Im}\{\mathbf{h}_i\}^T - \text{Re}\{\mathbf{h}_i\}^T\right]^T$. Problem \mathcal{Q} can then be re-written as

\mathcal{P} :

$$\begin{aligned} & \min_{\mathbf{x} \in \mathbb{R}^{2N \times 2N}} \mathbf{x}^T \mathbf{x} \\ & \text{subject to: } \mathbf{x}^T \mathbf{Q}_i \mathbf{x} \geq c_i, \quad i \in \{1, \dots, M\}. \end{aligned}$$

The Lagrangian of Problem \mathcal{P} is [2]

$$\begin{aligned}\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) &= \mathbf{x}^T \mathbf{x} + \sum_{i=1}^M \lambda_i (c_i - \mathbf{x}^T \mathbf{Q}_i \mathbf{x}) \\ &= \mathbf{x}^T (\mathbf{I} - \sum_{i=1}^M \lambda_i \mathbf{Q}_i) \mathbf{x} + \sum_{i=1}^M \lambda_i c_i\end{aligned}$$

and the dual problem is

$$\max_{\boldsymbol{\lambda} \succeq \mathbf{0}} \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}),$$

where $\boldsymbol{\lambda} \succeq \mathbf{0}$ denotes $\lambda_i \geq 0$. If the symmetric matrix $(\mathbf{I} - \sum_{i=1}^M \lambda_i \mathbf{Q}_i)$ has a negative eigenvalue, then it is easy to see that the quadratic term in $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ is unbounded from below (e.g., choose \mathbf{x} proportional to the corresponding eigenvector). If, on the other hand, all eigenvalues are greater than or equal to zero, then the said matrix is positive semidefinite and the minimum over \mathbf{x} is attained, e.g., at $\mathbf{x} = \mathbf{0}$. This yields the following equivalent of the dual problem:

$$\begin{aligned}\max_{\lambda_i \in \mathbb{R}} \quad & \sum_{i=1}^M \lambda_i c_i \\ \text{subject to: } & \mathbf{I} - \sum_{i=1}^M \lambda_i \mathbf{Q}_i \succeq \mathbf{0}, \\ & \lambda_i \geq 0, \quad i = 1, \dots, M,\end{aligned}$$

which is a semidefinite program.

The dual problem is interesting, because the maximum of the dual problem is a lower bound on the minimum of the original (primal) problem [2]. The dual problem is convex by virtue of its definition, however the particular dual studied above is special in the sense that optimization over \mathbf{x} for a given $\boldsymbol{\lambda}$ can be carried out analytically, and the residual $\boldsymbol{\lambda}$ -optimization problem is an SDP. This means that we can solve the dual problem and thus obtain the tightest bound obtainable via duality. This duality-derived bound can be compared to the SDR bound we used earlier. Let $\mathcal{D}(\cdot)$ denote the dual of a given optimization problem, and let $\mathcal{R}(\mathcal{P})$ denote the semidefinite relaxation of \mathcal{P} , obtained by dropping the associated rank-one constraint. Furthermore, let $\beta(\cdot)$ denote the optimal value of a given optimization problem. The following result shows that the duality and SDR bounds are in fact equal for the problem at hand:

Claim 5: $\mathcal{D}(\mathcal{D}(\mathcal{P})) = \mathcal{R}(\mathcal{P})$ and $\beta(\mathcal{R}(\mathcal{P})) = \beta(\mathcal{D}(\mathcal{P}))$.

More specifically, Claim 5 states that the dual of the dual of \mathcal{P} is the SDR of \mathcal{P} , and that the optimal objective value of the SDR of \mathcal{P} is the same as the optimal objective value of the dual of \mathcal{P} . Hence, SDR yields the same lower bound on the optimal value of \mathcal{P} as that obtained from duality, and the associated gap between this bound and the optimal value is equal to the duality gap. This result is apparently known among researchers in the convex optimization/semidefinite relaxation community, although we have been unable to pinpoint a published proof. In [7], the claim is proven for the general inhomogeneous QCQP problem, which entails extra obstructions in the proof. We have derived an independent short proof for our specific (homogeneous) problem, which we include in Appendix II for completeness.

Claim 5 along with Claim 3 directly yield the following corollary for the max-min fair problem, \mathcal{F} :

Corollary 2: $\mathcal{D}(\mathcal{D}(\mathcal{F})) = \mathcal{R}(\mathcal{F})$ and $\beta(\mathcal{R}(\mathcal{F})) = \beta(\mathcal{D}(\mathcal{F}))$.

VIII. SIMULATION RESULTS

An appropriate figure of merit for the performance of the proposed algorithm for the QoS beamforming problem, \mathcal{Q} , would be the ratio of the minimum transmitted power achieved by the proposed algorithm and $\beta(\mathcal{Q})$, the transmitted power achieved by the (true) optimal solution. Unfortunately, Problem \mathcal{Q} is NP-hard, and thus $\beta(\mathcal{Q})$ can be difficult to compute. However, we can replace $\beta(\mathcal{Q})$ in the figure of merit by the lower bound obtained from the SDR; i.e., $\beta(\mathcal{Q}) \geq \beta(\mathcal{Q}_r) = \text{trace}(\mathbf{X}_{\text{opt}})$. If we let $\{\mathbf{w}_\ell\}$ denote the sequence of candidate weight vectors generated via randomization, and $\check{\mathbf{w}}_\ell$ denote the minimally scaled version of \mathbf{w}_ℓ that satisfies the constraints of Problem \mathcal{Q} , then a meaningful and easily computable figure of merit is $(\min_\ell \|\check{\mathbf{w}}_\ell\|_2^2) / \text{trace}(\mathbf{X}_{\text{opt}})$. We will call this ratio the upper bound on the power boost required to satisfy the constraints. If our algorithm achieves a power boost of ρ , then the transmitted power is guaranteed to be within a factor ρ of that of the optimal solution, $\beta(\mathcal{Q})$, and will often be closer.

A. Rayleigh fading wireless channels

We consider the standard i.i.d. Rayleigh fading model described in the caption of Table III. That table summarizes the results obtained using the direct QoS relaxation algorithm in Table I with all three randomization options (randA, randB, and randC) employed in parallel, for a fixed number of 1000 randomization samples each. Table IV summarizes results for the same scenario, except that $30NM$ randomization samples are drawn for each randomization strategy — thus the number of randomizations grows linearly in the problem size. Note that, in many cases, our solutions are within 3–4 dB from the (generally optimistic) lower bound on transmit power provided by SDR, and thus are guaranteed

to be at most 3–4 dB away from optimal; this is often good enough from an engineering perspective. Comparing the corresponding entries in Tables III and IV, it is evident that switching from 1000 to $30NM$ randomizations per channel realization only yields a minor performance improvement in the cases considered.

Table V summarizes our simulation results for max-min fair beamforming, using the direct algorithm in Table II. Table V presents averages for the upper bound on minimum SNR (the optimum attained in Problem \mathcal{F}_r), the SDR-attained minimum SNR (after randomization), and the minimum SNR for the case of no beamforming. For the latter, we have used $\mathbf{w} = \frac{1}{\sqrt{N}}\mathbf{1}_{N \times 1}$, which fixes transmit power to 1. Under the i.i.d. Rayleigh fading assumption, this is *equivalent* to selecting an arbitrary transmit antenna, allocating the entire power budget to it, and shutting off all others. To see this, note that the sum channel $\frac{1}{\sqrt{N}}\mathbf{1}_{N \times 1}^T \mathbf{h}_i$ viewed by any particular receiver i will still be Rayleigh, of the same variance as the elements of \mathbf{h}_i . For this reason, we can view the beamforming vector $\mathbf{w} = \frac{1}{\sqrt{N}}\mathbf{1}_{N \times 1}$ as corresponding to no beamforming at all. All three randomization options (randA, randB, and randC) were employed in parallel, for $30NM$ samples each. It is satisfying to note that the SDR solution attains a significant fraction of the (possibly unattainable) upper bound. Furthermore, the SDR technique provides substantially better performance than the approach that does not employ any beamforming.

We observe from Tables III–V, that as N and/or M increase, the quality of the approximate solution drifts away from the respective relaxation/duality bound. This could be due to a variety of factors, or combination thereof. First, the relaxation bound may become more optimistic at higher N and/or M — remember that it is only a bound, not necessarily a tight bound. If this is true, then the apparent degradation may in fact be much milder in reality. Second, the number of randomizations required to attain a quasi-optimal solution may increase faster than linearly in the product NM . Third, the approximation quality of the method *per se* may degrade as the problem size grows. In a related, but distinct, problem the quality of the SDR approximation degrades logarithmically in the problem size [9].

B. Far-field beamforming for a uniform linear transmit antenna array

In several scenarios the solutions generated by the SDR technique are essentially optimal. This is illustrated in Figure 1, which shows the optimized transmit beampattern for a particular far-field multicasting scenario using a Uniform Linear antenna Array (ULA); the details of the simulation setup are included in the figure caption for ease of reference.

C. Measured VDSL channels

In this section, we test the performance of our algorithms using measured VDSL channel data collected by France Telecom R&D as part of the EU-FP6 U-BROAD project # 506790.

Gigabit VDSL technology for very short twisted copper loops (in the order of 100-500m) is currently under development in the context of fiber to the basement (FTTB) or fiber to the curb/cabinet (FTTC) hybrid access solutions. Multiple-input multiple-output (MIMO) transmission modalities are an important component of gigabit VDSL. These so-called *vectoring* techniques rely on transmit precoding and/or multiuser detection to provide reliable communication at very high transmission rates [5]. Transmit precoding is particularly appealing when the targeted receivers are not physically co-located, or when legacy equipment is being used at the receive site. In both cases, multiuser detection is not feasible. In this context, media streaming (e.g., news-feed, pay-per-view, or video-conferencing) may involve multiple recipients in the same binder.

Let N denote the number of loops subscribing to a given multicast. With multicarrier transmission, each tone can be viewed as a flat-fading MIMO channel with N inputs and N outputs, plus noise and alien interference. The diagonal of the channel matrix consists of samples of the N direct (insertion loss - IL) channel frequency responses, while off-diagonal elements are drawn from the corresponding far-end crosstalk (FEXT) channel frequency responses. Due to the non-coherent combining of the self-FEXT coupling coefficients, the useful signal power received at each output terminal is reduced, even when all inputs are fed with the same information-bearing signal. That is, the equivalent channel tap at frequency f is $h_e(f) = h_{IL}(f) + \sum_{n_{FEXT}=1}^{N-1} h_{n_{FEXT}}(f)$, where $h_{IL}(f)$ denotes the direct (insertion loss) channel, and $h_{n_{FEXT}}(f)$ denotes a generic FEXT interference channel.

Conceptually, the scenario is very similar to the wireless scenario considered earlier, but with two key differences: now $N = M$, and the channel matrix $\mathbf{H} := [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_M]$ is diagonally-dominated, because FEXT coupling is much weaker than insertion loss. The question then is whether transmit precoding can provide a meaningful benefit relative to simply ignoring FEXT altogether.

We use IL and far-end FEXT measured data for S88 cable comprising 14 quads, i.e., 28 loops. The length of the cable is 300 meters. For each channel, a log-frequency sweeping scheme was used to measure the I/Q components of the frequency response from 10 kHz – 30 MHz, yielding 801 complex samples per channel. Cubic spline complex interpolation was used to convert these samples to a linear frequency scale.

We consider 17 $N \times N$ channel matrices, with $N = 14$, in the frequency range 21.5 to 30 MHz. Insertion

loss drops between -40 and -45 dB in this range of frequencies, while FEXT coupling is between -77 and -82 dB in the mean, with over 10 dB standard deviation and significant variation across frequency as well. For each channel matrix, we apply our max-min fair beamforming algorithm with $P = 1$. Figure 2 shows the resulting plots of minimum received signal power, the associated relaxation/duality bound, and the minimum received signal power when no precoding is used. We observe that SDR can almost double the minimum received signal power relative to no precoding, and it often attains zero gap relative to the relaxation/duality bound. For shorter loops (e.g., 100m) the situation is even more in favor of SDR, because then FEXT resembles near-end crosstalk (NEXT) and is relatively more pronounced.

D. Further observations

1) *Comparison of the two relaxations:* We have shown theoretically that the two problem formulations (QoS, \mathcal{Q} , and max-min-fair, \mathcal{F}) are algorithmically equivalent, i.e., had we had an optimal algorithm that provides the exact solution to one, it could have also been used to obtain the exact solution to the other. What we have instead is two generally approximate algorithms, obtained by direct relaxation of the respective problems. The link between the two formulations can still be exploited. For example, we may obtain an approximate solution to the max-min-fair problem by first running the QoS algorithm in Table I with all the c_i 's equal to 1, then scaling the resulting solution to the desired power level P . Of course, we can also use the direct relaxation of the max-min-fair problem in Table II. Due to approximation, there is no *a priori* reason to expect that the two solutions will be identical, even in the mean.

In order to address this issue, we have compared the two strategies by means of Monte-Carlo simulation. We chose $N = 4$, $M = 16$, $P = 1$, and ran both algorithms for 300 i.i.d. Rayleigh fading channels. All three randomizations (randA, randB, randC) were employed in parallel, for $30NM$ randomization samples each. For each channel, we recorded the percent gap (100 times the gap over the relaxation bound) of each algorithm. Figure 3 shows a portion of the results, along with the mean percent gap attained by each algorithm (averaged over all 300 channels). By 'direct' we refer to the algorithm in Table II, whereas by 'indirect' to the algorithm in Table I with all c_i equal to 1, followed by scaling.

We observe that the mean percent gaps of the two algorithms are virtually identical, and in fact most of the respective percent gaps are very close on a sample-by-sample basis. However, there are instances wherein each algorithm is significantly better than the other (over 10 % difference in the gap). Two pronounced cases are highlighted by arrows in Figure 3. We conclude that, while both approaches are equally effective on average, it pays to use both, if possible, in certain cases.

2) *On the dependence of gap statistics on channel statistics:* We have seen that, for i.i.d. circular Gaussian (Rayleigh) channel matrices, the gap between our relaxation - randomization approximate solutions and the relaxation / duality bound might not be insignificant. We have also seen cases wherein the gap is very small, c.f., the far-field uniform linear transmit antenna array example, and a good proportion of the VDSL channels tested earlier.

It is evident that the gap statistics depend on the channel statistics. Interestingly, the gap statistics are far more favorable for *real* (as opposed to complex circular) i.i.d. Gaussian channels. This is illustrated in Figure 4, using the QoS algorithm in Table I for $N = 4$, $M = 8$, $c_i = 1$, $\forall i$, and 300 real i.i.d. Gaussian channels. All three randomizations (randA, randB, randC) are employed in parallel, for $30NM$ randomization samples each. For each channel, we recorded the percent gap (100 times the gap over the relaxation bound) of the algorithm in Table I. Observe that for about 95% of the channels the percent gap is down to numerical accuracy in this case. Contrast this situation with Figure 5, which shows the respective results for complex circular Gaussian channel matrices — the difference is remarkable.

There are other cases where we have observed that the relaxation approach operates close to zero gap. One somewhat contrived case is when the real and imaginary parts of the channel coefficients are non-negative. This is illustrated in Figure 6, where it is worth noting that the scaling of the y-axis is 10^{-8} . In this case, the gap hovers around numerical accuracy, without exhibiting any bad runs at all for the 300 channel matrices considered.

In conclusion, the complex circular Gaussian channel case appears to be the least favorable of the scenarios considered.

IX. CONCLUSIONS

We have taken a new look at the broadcasting/multicasting problem when channel state information is available at the transmitter. We have proposed two pertinent problem formulations: minimizing transmitted power under multiple minimum received power constraints, and maximizing the minimum received power subject to a bound on the transmitted power. We have shown that both formulations are NP-hard optimization problems, however their solution can often be well-approximated using semidefinite relaxation tools. We have explored the relationship between the two formulations, and also insights afforded by Lagrangian duality theory. In view of i) our extensive numerical experiments with simulated and measured data, verifying that semidefinite relaxation consistently yields good performance, ii) proof that the basic problem is NP-hard, thus approximation is unavoidable, and iii) corroborating motivation provided by duality theory, we conclude that the approximate solutions provided herein offer useful

designs across a broad range of applications.

It would be useful to analyze the duality gap for the problem at hand, for this would yield *a priori* bounds on the degree of suboptimality introduced by relaxation, as opposed to the *a posteriori* bound that we now have by virtue of Claim 5. Our numerical results indicate that the degree of suboptimality is often acceptable in our intended applications. In an effort to understand the apparent success of the SDR approach (e.g., in the case where the channel vectors have nonnegative real and imaginary parts), one can consider the following simple linearly constrained convex quadratic program (QP) *restriction* of the QoS problem:

$$\begin{array}{l}
 \mathcal{Q}_s: \\
 \min_{\mathbf{w} \in \mathbb{C}^N} \|\mathbf{w}\|_2^2 \\
 \text{subject to: } \operatorname{Re}\{\mathbf{h}_i^H \mathbf{w}\} \geq 1, \text{ for all } i.
 \end{array}$$

Notice that the feasible region of this problem is a subset of that of the original non-convex (and NP-hard) QoS formulation, \mathcal{Q} . Thus, $P^* \leq \bar{P}$, where P^* and \bar{P} denote the minimum beamforming power obtained from optimal solutions of \mathcal{Q} and \mathcal{Q}_s , respectively. We have recently shown that the gap between P^* and \bar{P} is never more than $1/\cos^2(\alpha/2)$, where α is the maximum phase spread across the different users measured at each transmit antenna, and is assumed to be less than π . Notice that the two cases where channel vectors i) are real and nonnegative; ii) have nonnegative real and imaginary parts, correspond to $\alpha = 0$ and $\alpha \leq \pi/2$. Thus, \mathcal{Q}_s provides an exact solution in the first case and a factor of 2 approximation in the second case. These results indicate that problem \mathcal{Q} is well approximated by \mathcal{Q}_s if the phase spread α is small.

There are many other interesting extensions to the algorithms developed herein: e.g., robustness issues, and multiple co-channel multicasting groups. These are subjects of on-going work, and will be reported elsewhere. Furthermore, aside from transmit beamforming/precoding, there are also more traditional signal processing applications of the proposed methodology. One is linear filter design; in particular, the design of a linear “batch” filter that responds to certain prescribed frequencies in its input and attenuates all other frequencies. In this setting, the \mathbf{h}_i vectors will be Vandermonde, with generators $e^{j\omega_i}$, $\omega_i \in [-\pi, \pi)$. One may easily envision scenarios wherein such a problem formulation can be appropriate: radio-astronomy applications, frequency-diversity combining, and frequency-hopping communications. The context can be further generalized: design a linear filter that responds to prescribed but otherwise arbitrary signals in its

input, while attenuating all else.

Acknowledgment: We would like to thank Dr. Meryem Ouzzif and Dr. Rabah Tarafi of France Telecom R&D, who conducted the VDSL channel measurements under the auspices of the U-BROAD project.

APPENDIX I PROOF OF CLAIM 1

Before dealing with Claim 1 directly, we first consider the following restriction of the QoS problem \mathcal{Q} : the case when all \mathbf{h}_i are real, and optimization is over \mathbb{R}^N . We will show that

\mathcal{S} :

$$\begin{aligned} & \min_{\mathbf{x} \in \mathbb{R}^N} \mathbf{x}^T \mathbf{x} \\ & \text{subject to: } |\mathbf{x}^T \mathbf{h}_m| \geq 1, \quad m \in \{1, \dots, M\} \end{aligned}$$

contains

\mathcal{A} :

$$\begin{aligned} & \min_{y_n \in \mathbb{R}} y_1^2 + \dots + y_N^2 + \left(\sum_{n=1}^N a_n y_n \right)^2 \\ & \text{subject to: } |y_n|^2 \geq 1, \quad n \in \{1, \dots, N\} \end{aligned}$$

as a special case; and that Problem \mathcal{A} is at least as hard as the

Partition Problem II: Given integers a_1, \dots, a_N , do there exist binary variables $\{x_n\}_{n=1}^N \in \{+1, -1\}^N$, such that $\sum_{n=1}^N a_n x_n = 0$?

which is known to be NP-complete [4].

It is easy to check that the optimal value of Problem \mathcal{A} is equal to N if and only if the answer to Problem II is affirmative. Thus, solving Problem \mathcal{A} is at least as hard as solving Problem II.

To show that Problem \mathcal{S} contains Problem \mathcal{A} (i.e., an arbitrary instance of Problem \mathcal{A} can be posed as a special instance of Problem \mathcal{S}), note that $|y_n|^2 \geq 1$ can be written as $|\mathbf{y}^T \mathbf{e}_n| \geq 1$, where $\mathbf{y} := [y_1, \dots, y_N]^T$ and \mathbf{e}_n contains 1 in the n -th position and zeros elsewhere. Furthermore,

$$y_1^2 + \dots + y_N^2 + \left(\sum_{n=1}^N a_n y_n \right)^2 = \mathbf{y}^T (\mathbf{I} + \mathbf{a}\mathbf{a}^T) \mathbf{y} = \mathbf{y}^T \mathbf{Q} \mathbf{y},$$

where $\mathbf{a} := [a_1, \dots, a_N]^T$, and $\mathbf{Q} := \mathbf{I} + \mathbf{a}\mathbf{a}^T$. The matrix \mathbf{Q} is positive. Let $\mathbf{Q} = \mathbf{S}^T \mathbf{S}$, and $\mathbf{x} := \mathbf{S} \mathbf{y}$. Then $\mathbf{y}^T \mathbf{Q} \mathbf{y} = \mathbf{x}^T \mathbf{x}$, $\mathbf{y} = \mathbf{S}^{-1} \mathbf{x}$, and $|\mathbf{y}^T \mathbf{e}_n| \geq 1$ can be written as $|\mathbf{x}^T \mathbf{S}^{-T} \mathbf{e}_n| \geq 1$, or, with $\mathbf{h}_n := \mathbf{S}^{-T} \mathbf{e}_n$, as $|\mathbf{x}^T \mathbf{h}_n| \geq 1$. This shows that an arbitrary instance of Problem \mathcal{A} can be transformed to a special instance

of Problem \mathcal{S} (with $M = N$). Thus, \mathcal{S} is at least as hard as \mathcal{A} , which is at least as hard as the partition problem. \square

Proof of Claim 1: The QoS problem \mathcal{Q} is NP-hard: Consider the problem

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{C}^N} \quad & \mathbf{w}^H \mathbf{w} \\ \text{subject to:} \quad & |\mathbf{w}^H \mathbf{h}_i| \geq 1, \quad i = 1, \dots, M. \end{aligned} \quad (1)$$

Define the $N \times M$ matrix $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_M]$, and the $M \times 1$ vector \mathbf{z} , with $\mathbf{z}^H := \mathbf{w}^H \mathbf{H}$. Consider the case that $M \geq N$, and \mathbf{H} is full row-rank (N). Then $\mathbf{w}^H = \mathbf{z}^H \mathbf{H}^\dagger$, where $\mathbf{H}^\dagger = \mathbf{H}^H (\mathbf{H} \mathbf{H}^H)^{-1}$ denotes the right pseudo-inverse of \mathbf{H} , and the problem in (1) is equivalent to

$$\begin{aligned} \min_{\mathbf{z} \in \mathbb{C}^M} \quad & \mathbf{z}^H \mathbf{Q} \mathbf{z} \\ \text{subject to:} \quad & |z_k| \geq 1, \quad k = 1, \dots, M, \end{aligned} \quad (2)$$

where $\mathbf{Q} := \mathbf{H}^\dagger (\mathbf{H}^\dagger)^H \succeq \mathbf{0}$, a positive semidefinite matrix of rank $N \leq M$; and z_k denotes the k -th element of the vector \mathbf{z} . We will show that problem (2) is NP-hard in general. To this end, we consider a reduction from the NP-complete partition problem [4]; i.e., given $a_1 > 0, a_2 > 0, \dots, a_P > 0$, decide whether or not a subset, say I , of $\{1, \dots, P\}$ exists, such that

$$\sum_{k \in I} a_k = \frac{1}{2} \sum_{k=1}^P a_k. \quad (3)$$

Let $M = 2P + 1$ and let the complex-valued decision vector be

$$\mathbf{z} = [z_0, z_1, \dots, z_P, z_{P+1}, \dots, z_{2P}]^T \in \mathbb{C}^M.$$

Let us denote

$$\begin{aligned} \mathbf{a} &:= \begin{bmatrix} a_1 & a_2 & \dots & a_P \end{bmatrix}^T, \\ \mathbf{A} &:= \begin{bmatrix} -\mathbf{1}_P & \mathbf{I}_P & \mathbf{I}_P \\ -\frac{1}{2} \mathbf{1}_P^T \mathbf{a} & \mathbf{a}^T & \mathbf{0}_P^T \end{bmatrix}, \\ \mathbf{Q} &:= \mathbf{A}^T \mathbf{A} + \mathbf{I}_M, \end{aligned}$$

where $\mathbf{1}_P$ denotes the length- P vector of ones, and $\mathbf{0}_P$ is the length- P vector of zeros.

Next we show that a partition I satisfying (3) exists if and only if the optimization problem (2) has a minimum value of M . In other words, the existence of I is equivalent to the fact that there is $\mathbf{z} \in \mathbb{C}^M$ such that $\mathbf{z}^H \mathbf{Q} \mathbf{z} = M$ and $|z_k| \geq 1$, for all k . Since

$$\mathbf{z}^H \mathbf{Q} \mathbf{z} = \|\mathbf{A} \mathbf{z}\|_2^2 + \sum_{k=0}^{2P} |z_k|^2 \geq 2P + 1 = M, \quad \text{for } |z_k| \geq 1 \quad \forall k,$$

it follows that

$$\mathbf{z}^H \mathbf{Q} \mathbf{z} = M, \quad |z_k| \geq 1 \quad \text{for all } k$$

is equivalent to

$$\mathbf{A} \mathbf{z} = \mathbf{0}, \quad |z_k| = 1 \quad \text{for all } k.$$

The latter gives rise to a set of linear equations:

$$-z_0 + z_k + z_{P+k} = 0, \quad k = 1, \dots, P, \quad (4)$$

$$-\frac{1}{2} \left(\sum_{k=1}^P a_k \right) z_0 + \sum_{k=1}^P a_k z_k = 0. \quad (5)$$

The z_k 's are all constrained to be on the unit circle; thus let $z_k/z_0 = e^{i\theta_k}$ for $k = 1, \dots, 2P$. Using (4) we have

$$\cos \theta_k + \cos \theta_{P+k} = 1, \quad (6)$$

$$\sin \theta_k + \sin \theta_{P+k} = 0, \quad (7)$$

where $k = 1, \dots, P$. These two equations imply that $\theta_k \in \{-\pi/3, \pi/3\}$ for all k . This in particular means that $\cos \theta_k = \cos \theta_{P+k} = 1/2$ for $k = 1, \dots, P$, implying that

$$\operatorname{Re} \left\{ -\frac{1}{2} \left(\sum_{k=1}^P a_k \right) + \sum_{k=1}^P a_k z_k / z_0 \right\} = 0.$$

Therefore, (5) is satisfied if and only if

$$\begin{aligned} \operatorname{Im} \left\{ -\frac{1}{2} \left(\sum_{k=1}^P a_k \right) + \sum_{k=1}^P a_k z_k / z_0 \right\} &= \\ \operatorname{Im} \left\{ \sum_{k=1}^P a_k z_k / z_0 \right\} &= \sum_{k=1}^P a_k \sin \theta_k = 0, \end{aligned}$$

with $\theta_k \in \{-\pi/3, \pi/3\}$ for all k , and thus $\sin \theta_k \in \{\frac{\sqrt{3}}{2}, -\frac{\sqrt{3}}{2}\}$, which is equivalent to the existence of a partition I of $\{a_1, \dots, a_P\}$ such that (3) holds. In fact, we can simply take $I = \{k \mid \theta_k = \pi/3\}$. \square

APPENDIX II PROOF OF CLAIM 5

The dual problem $\mathcal{D}(\mathcal{P})$ involves two Linear Matrix Inequality (LMI) constraints. (Note that the positivity constraint on the individual λ_i 's can be written as $\operatorname{diag}(\boldsymbol{\lambda}) \succeq \mathbf{0}$, where $\operatorname{diag}(\cdot)$ constructs a

diagonal matrix with the elements of its argument on the diagonal.) We can combine the two LMIs into a single block-diagonal LMI:

$$\begin{bmatrix} \sum_{i=1}^M \lambda_i \mathbf{Q}_i - \mathbf{I}_{2N} & \mathbf{0} \\ \mathbf{0} & -\text{diag}(\boldsymbol{\lambda}) \end{bmatrix} \preceq \mathbf{0}.$$

Upon defining matrices

$$\mathbf{F}_i := \begin{bmatrix} \mathbf{Q}_i & \mathbf{0} \\ \mathbf{0} & -\mathbf{e}_i \mathbf{e}_i^T \end{bmatrix}_{(2N+M) \times (2N+M)} \quad \text{for } i \in \{1, \dots, M\},$$

where the i -th element of the vector \mathbf{e}_i is equal to 1 while the remaining elements are 0, and

$$\mathbf{G} := \begin{bmatrix} -\mathbf{I}_{2N} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}_{(2N+M) \times (2N+M)}$$

the block-diagonal LMI can be written as

$$\sum_{i=1}^M \lambda_i \mathbf{F}_i + \mathbf{G} \preceq \mathbf{0}.$$

Therefore, $\mathcal{D}(\mathcal{P})$ is

$$\begin{aligned} & \max_{\boldsymbol{\lambda} \in \mathbb{R}^M} \mathbf{c}^T \boldsymbol{\lambda} \\ & \text{subject to: } \sum_{i=1}^M \lambda_i \mathbf{F}_i + \mathbf{G} \preceq \mathbf{0}. \end{aligned}$$

The dual of this latter problem (i.e., $\mathcal{D}(\mathcal{D}(\mathcal{P}))$) is [2, pp. 265-266]:

$$\begin{aligned} & \min_{\mathbf{Z}} -\text{trace}(\mathbf{G}\mathbf{Z}) \\ & \text{subject to: } \text{trace}(\mathbf{F}_i \mathbf{Z}) = c_i, \quad i \in \{1, \dots, M\}, \\ & \mathbf{Z} \succeq \mathbf{0}, \end{aligned}$$

where the matrix variable $\mathbf{Z} \in \mathbb{R}^{(2N+M) \times (2N+M)}$. Partition \mathbf{Z} as follows:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{X}_{2N \times 2N} & \mathbf{X}_{\text{off}}^T \\ \mathbf{X}_{\text{off}} & \mathbf{Y}_{M \times M} \end{bmatrix}$$

and observe that

$$\begin{aligned} -\text{trace}(\mathbf{G}\mathbf{Z}) &= \text{trace} \left(\begin{bmatrix} \mathbf{X}_{2N \times 2N} & \mathbf{X}_{\text{off}}^T \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \right) \\ &= \text{trace}(\mathbf{X}_{2N \times 2N}), \end{aligned}$$

and that

$$\begin{aligned} \text{trace}(\mathbf{F}_i \mathbf{Z}) &= \text{trace} \left(\begin{bmatrix} \mathbf{Q}_i \mathbf{X}_{2N \times 2N} & \mathbf{Q}_i \mathbf{X}_{\text{off}}^T \\ -\mathbf{e}_i \mathbf{e}_i^T \mathbf{X}_{\text{off}} & -\mathbf{e}_i \mathbf{e}_i^T \mathbf{Y}_{M \times M} \end{bmatrix} \right) \\ &= \text{trace}(\mathbf{Q}_i \mathbf{X}_{2N \times 2N}) - \mathbf{Y}_{M \times M}(i, i). \end{aligned}$$

It is therefore clear that \mathbf{X}_{off} and the off-diagonal elements of $\mathbf{Y}_{M \times M}$ play no role in terms of the cost function and the equality constraints. We have the following claim.

Claim 6: In the context of $\mathcal{D}(\mathcal{D}(\mathcal{P}))$, we may assume, without loss of generality,

$$\mathbf{X}_{\text{off}} = \mathbf{0}, \quad \mathbf{Y}_{M \times M} = \text{diag}([\mu_1, \dots, \mu_M]),$$

and enforce non-negativity of the μ_i 's to ensure that the block matrix is positive semidefinite.

Proof of Claim 6: Let \mathbf{Z}^o be a solution to $\mathcal{D}(\mathcal{D}(\mathcal{P}))$. Since $\mathbf{Z}^o \succeq \mathbf{0}$, it can be factorized as $\mathbf{Z}^o = \mathbf{C}\mathbf{C}^T$, i.e., a sum of symmetric rank-1 outer products. It follows that the $\mathbf{X}_{2N \times 2N}^o$ part of \mathbf{Z}^o is a sum of symmetric rank-1 outer products, and thus $\mathbf{X}_{2N \times 2N}^o \succeq \mathbf{0}$. $\mathbf{Y}_{M \times M}^o$ should also be a sum of symmetric rank-1 outer products, and thus the elements along its diagonal should all be non-negative (sums of non-negative quantities). Starting from \mathbf{Z}^o and setting $\mathbf{X}_{\text{off}}^o = \mathbf{0}$ and all off-diagonal elements of $\mathbf{Y}_{M \times M}^o$ to zero, we can therefore produce another positive semidefinite matrix, $\bar{\mathbf{Z}}^o$, which also satisfies the equality constraints and entails the same cost. \square

Denoting $\mathbf{X} := \mathbf{X}_{2N \times 2N}$ for brevity, Claim 6 results in the following reformulation of $\mathcal{D}(\mathcal{D}(\mathcal{P}))$:

$$\begin{aligned} & \min_{\substack{\mathbf{X} \in \mathbb{R}^{2N \times 2N} \\ \mu_i \in \mathbb{R}}} \text{trace}(\mathbf{X}) \\ & \text{subject to: } \text{trace}(\mathbf{Q}_i \mathbf{X}) - \mu_i = c_i, \quad i \in \{1, \dots, M\}, \\ & \mu_i \geq 0, \quad i \in \{1, \dots, M\}, \quad \mathbf{X} \succeq \mathbf{0}. \end{aligned}$$

The first and second constraints are equivalent to $\text{trace}(\mathbf{Q}_i \mathbf{X}) \geq c_i$, $i \in \{1, \dots, M\}$, and thus $\mathcal{D}(\mathcal{D}(\mathcal{P}))$ can be written as

$$\begin{aligned} & \min_{\mathbf{X} \in \mathbb{R}^{2N \times 2N}} \text{trace}(\mathbf{X}) \\ & \text{subject to: } \text{trace}(\mathbf{Q}_i \mathbf{X}) \geq c_i, \quad i \in \{1, \dots, M\}, \quad \mathbf{X} \succeq \mathbf{0}, \end{aligned}$$

which is exactly $\mathcal{R}(\mathcal{P})$; thus we have shown that $\mathcal{D}(\mathcal{D}(\mathcal{P})) = \mathcal{R}(\mathcal{P})$. Now, since $\mathcal{D}(\mathcal{P})$ is convex (and Slater's condition is satisfied [2]), there is zero duality gap between $\mathcal{D}(\mathcal{P})$ and $\mathcal{D}(\mathcal{D}(\mathcal{P}))$; hence $\beta(\mathcal{D}(\mathcal{D}(\mathcal{P}))) = \beta(\mathcal{D}(\mathcal{P}))$. \square

REFERENCES

- [1] M. Bengtsson and B. Ottersten, "Optimal and suboptimal transmit beamforming", ch. 18 in *Handbook of Antennas in Wireless Communications*, L. C. Godara, Ed., CRC Press, Aug. 2001.
- [2] S. Boyd, and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004 (see also <http://www.stanford.edu/~boyd/cvxbook.html>).
- [3] G.D. Forney, "Maximum Likelihood Sequence Estimation of Digital Sequences in the Presence of Intersymbol Interference", *IEEE Trans. Informat. Theory*, vol. 18, no. 3, pp. 363–378, May 1972.
- [4] M.R. Garey, and D.S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [5] G. Ginis and J.M. Cioffi, "Vectored transmission for digital subscriber line systems," *IEEE J. Select Areas Commun.*, vol. 20, no. 5, pp. 1085–1104, June 2002.
- [6] N. Jindal, and A. Goldsmith, "Optimal Power Allocation for Parallel Gaussian Broadcast Channels with Independent and Common Information", in *Proc. Int. Symp. Informat. Theory*, Chicago, IL, June 2004. (See also http://www.ece.umn.edu/users/nihar/common_info_isit.pdf.)
- [7] Z.-Q. Luo, Lecture 13 in *Lecture Notes for EE8950: Engineering Optimization*, University of Minnesota, Minneapolis, Spring 2004. Available upon request to luozq@ece.umn.edu.
- [8] W.-K. Ma, T.N. Davidson, K.M. Wong, Z.-Q. Luo, P.-C. Ching, "Quasi-ML multiuser detection using semi-definite relaxation with application to synchronous CDMA", *IEEE Trans. Signal Processing*, vol. 50, no. 4, pp. 912–922, Apr. 2002.
- [9] A. Nemirovski, C. Roos and T. Terlaky, "On maximization of quadratic form over intersection of ellipsoids with common center", *Math. Program., Ser. A*, vol. 86, pp. 463–473, 1999.
- [10] J.F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones", *Optimization Methods and Software*, vol. 11-12, pp. 625–653, 1999. See also <http://fewcal.kub.nl/sturm/software/sedumi.html>
- [11] P. Tseng, "Further results on approximating nonconvex quadratic optimization by semidefinite programming relaxation", *SIAM J. Optimization*, vol. 14, no. 1, pp. 268–283, July 2003.
- [12] S.A. Vorobyov, A.B. Gershman, Z.-Q. Luo, "Robust Adaptive Beamforming Using Worst-Case Performance Optimization: A Solution to the Signal Mismatch Problem," *IEEE Trans. on Signal Processing*, Vol. 51, No. 2, pp. 313–324, Feb. 2003.
- [13] S. Zhang, "Quadratic maximization and semidefinite relaxation", in *Math. Program., Ser. A*, vol. 87, pp. 453–465, 2000.

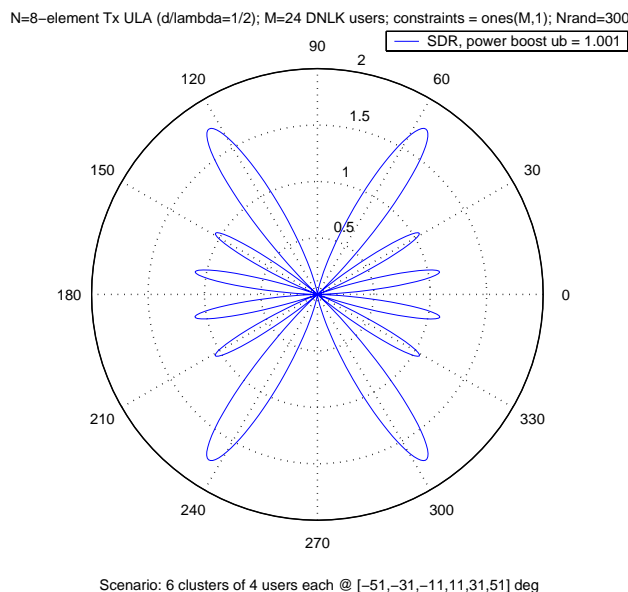


Fig. 1. Broadcast beamforming example using algorithm in Table I. $N = 8$ -element transmit ULA ($d/\lambda = 1/2$); $M = 24$ downlink users, in 6 clusters of 4 users each. Clusters centered at $[-51, -31, -11, 11, 31, 51]^\circ$ with extent $\pm 2^\circ$. Symmetric lobes appear due to the inherent ULA ambiguity. All received SNR constraints set to 1. randA , # post-SDR randomizations = 300. In this case, the solution is guaranteed to be within 0.1% of the optimum.




Fig. 2. Transmit precoding for VDSL multicasting.




Fig. 3. Comparison of direct and indirect solutions to the max-min fair problem.




Fig. 4. Percent gap outcomes for 300 real Gaussian channel realizations.




Fig. 5. Percent gap outcomes for 300 complex circular Gaussian (Rayleigh) channel realizations.




Fig. 6. Percent gap outcomes for 300 channel realizations with positive real and imaginary parts (uniformly distributed between 0 and 1). Note that the scaling of the y-axis is 10^{-8} .

TABLE I
BROADCAST QOS BEAMFORMING VIA SDR: ALGORITHM

- Solve the relaxed problem:

A simple MATLAB interface for SeDuMi is as follows:

```
% Input Data:
% H: N by M, holding the channel vectors;
% constraints: M by 1, receive power constraints
% Output Data:
% Xopt: the solution to the SDR
vecQs = [];
for i=1:M,
    Qi = H(:,i)*H(:,i)';
    vecQs = [vecQs, vec(Qi.')];
end;
A=[-eye(M), vecQs.'];
b=constraints;
c=[zeros(M,1); vec(eye(N))];
K.l=M; K.s=N; K.scomplex=1;
[x_opt,y_opt,info]=sedumi(A,b,c,K);
Xopt=mat(x_opt(M+1:end));
```

- Randomization:

Use randA, and/or randB, randC to generate the candidates, w_e11.
For each w_e11, find the most violated constraint.
Scale w_e11 so that that constraint is satisfied with equality.
Pick the w_e11 with the smallest norm.

TABLE II
BROADCAST MAX-MIN BEAMFORMING VIA SDR: ALGORITHM

• Solve the relaxed problem:
A suitable MATLAB interface for SeDuMi is as follows:

```
% Input Data:
% H: N by M, holding the scaled channel vectors;
% P: scalar, the total transmit power constraint
% Output Data:
% Xopt: the solution to the SDR
% t_opt: the minimum objective value of the SDR
vecQs = [];
for i=1:M,
    Qi = H(:,i)*H(:,i)';
    vecQs = [vecQs vec(Qi.')];
end;
A1=[-ones(M,1); 0];
A2=[-eye(M); zeros(1,M)];
A3=[vecQs.'; vec(eye(N)).'];
A=[A1 A2 A3];
b=[zeros(M,1); P];
c = [-1; zeros(M+N*N,1)];
K.l=M+1; K.s=N; K.scomplex=1;
[x_opt,y_opt,info]=sedumi(A,b,c,K);
Xopt=mat(x_opt(M+2:end));
t_opt = x_opt(1)
```

• Randomization:
Use randA, and/or randB, randC to generate the candidates w_{ell} .
Scale each w_{ell} to norm P.
Pick the one that yields the largest $\min(\text{abs}(w_{ell}' * H))$.

TABLE III

MC SIMULATION RESULTS FOR QoS BEAMFORMING: MEAN AND STANDARD DEVIATION OF UPPER BOUND ON POWER BOOST. EACH ELEMENT OF \mathbf{h}_i IS I.I.D. WITH A CIRCULARLY SYMMETRIC COMPLEX GAUSSIAN (RAYLEIGH) DISTRIBUTION OF VARIANCE 1. ALL THREE RANDOMIZATION TECHNIQUES (randA,randB,randC) ARE USED IN PARALLEL, FOR 1000 RANDOMIZATIONS EACH. ALL RECEIVED POWER CONSTRAINTS ARE FIXED TO $c_i = 1$.

| N/M | mean | std |
|-------|------|------|
| 4/8 | 1.12 | 0.16 |
| 4/16 | 1.47 | 0.30 |
| 8/16 | 1.82 | 0.37 |
| 8/32 | 2.79 | 0.47 |

TABLE IV

MC SIMULATION RESULTS FOR QoS BEAMFORMING: MEAN AND STANDARD DEVIATION OF UPPER BOUND ON POWER BOOST. HERE, THE NUMBER OF POST-SDR RANDOMIZATIONS = $30NM$. THE REMAINING PARAMETERS ARE AS IN

TABLE III.

| N/M | mean | std |
|-------|------|------|
| 4/8 | 1.12 | 0.16 |
| 4/16 | 1.44 | 0.29 |
| 8/16 | 1.76 | 0.34 |
| 8/32 | 2.49 | 0.38 |

TABLE V

MC SIMULATION RESULTS FOR MAX-MIN FAIR BEAMFORMING: AVERAGES FOR THE UPPER BOUND ON $\min_i \text{SNR}_i$, THE RELAXATION-ATTAINED $\min_i \text{SNR}_i$, AND THE $\min_i \text{SNR}_i$ FOR THE CASE OF NO BEAMFORMING. THE RESULTS ARE AVERAGED OVER 1000 MC RUNS. FOR EACH MC RUN, THE ELEMENTS OF \mathbf{h}_i ARE INDEPENDENTLY RE-DRAWN FROM A CIRCULARLY SYMMETRIC COMPLEX GAUSSIAN DISTRIBUTION OF VARIANCE 1, AND THE NOISE VARIANCE OF EACH RECEIVER IS 1. ALL THREE RANDOMIZATION TECHNIQUES (**randA**,**randB**,**randC**) ARE USED IN PARALLEL, FOR $30NM$ RANDOMIZATIONS EACH. $P = 1$.

| N/M | upper bound | SDR | no BMF |
|-------|-------------|------|--------|
| 4/8 | 1.05 | 0.94 | 0.12 |
| 4/16 | 0.73 | 0.51 | 0.06 |
| 8/16 | 1.43 | 0.86 | 0.06 |
| 8/32 | 1.07 | 0.45 | 0.03 |