

Linköping University Post Print

**Quality of Service and Max-Min Fair Transmit
Beamforming to Multiple Cochannel Multicast
Groups**

Eleftherios Karipidis, Nicholas Sidiropoulos and Zhi-Quan Luo

N.B.: When citing this work, cite the original article.

©2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Eleftherios Karipidis, Nicholas Sidiropoulos and Zhi-Quan Luo, Quality of Service and Max-Min Fair Transmit Beamforming to Multiple Cochannel Multicast Groups, 2008, IEEE Transactions on Signal Processing, (56), 3, 1268-1279.

<http://dx.doi.org/10.1109/TSP.2007.909010>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-66517>

Quality of Service and Max-Min Fair Transmit Beamforming to Multiple Cochannel Multicast Groups

Eleftherios Karipidis, *Student Member, IEEE*, Nicholas D. Sidiropoulos, *Senior Member, IEEE*, and Zhi-Quan Luo, *Fellow, IEEE*

Abstract—The problem of transmit beamforming to multiple cochannel multicast groups is considered, when the channel state is known at the transmitter and from two viewpoints: minimizing total transmission power while guaranteeing a prescribed minimum signal-to-interference-plus-noise ratio (SINR) at each receiver; and a “fair” approach maximizing the overall minimum SINR under a total power budget. The core problem is a multicast generalization of the multiuser downlink beamforming problem; the difference is that each transmitted stream is directed to multiple receivers, each with its own channel. Such generalization is relevant and timely, e.g., in the context of the emerging WiMAX and UMTS-LTE wireless networks. The joint problem also contains single-group multicast beamforming as a special case. The latter (and therefore also the former) is NP-hard. This motivates the pursuit of computationally efficient quasi-optimal solutions. It is shown that Lagrangian relaxation coupled with suitable randomization/cochannel multicast power control yield computationally efficient high-quality approximate solutions. For a significant fraction of problem instances, the solutions generated this way are exactly optimal. Extensive numerical results using both simulated and measured wireless channels are presented to corroborate our main findings.

Index Terms—Broadcasting, convex optimization, downlink beamforming, multicasting, semidefinite relaxation.

I. INTRODUCTION

THE proliferation of streaming media (digital audio, video, IP radio), peer-to-peer services, large-scale software updates, and profiled newscasts over the wireline Internet has

Manuscript received March 17, 2007; revised July 31, 2007. Original version submitted to the IEEE TRANSACTIONS ON SIGNAL PROCESSING on July 4, 2006. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Erik G. Larsson. The work of E. Karipidis was supported in part by the E.U. under FP6 506790 project U-BROAD and by National and Community Funds (75% from E.U.-European Social Fund and 25% from the Greek Ministry of Development - General Secretariat of Research and Technology) under 03ED918 research project, implemented within the framework of the Reinforcement Programme of Human Research Manpower (PENED). The work of N. D. Sidiropoulos was supported in part by the U.S. ARO under ERO Contract N62558-03-C-0012 and by the E.U. under FP6 project NEWCOM. The work of Z.-Q. Luo was supported in part by the National Science Foundation by Grant DMS-0312416. An earlier version of part of this work appears in conference form in the Proceedings of the 1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), Puerto Vallarta, Mexico, December 13–15, 2005, pp. 109–112.

E. Karipidis and N. D. Sidiropoulos are with the Department of Electronic and Computer Engineering, Technical University of Crete, 73100 Chania-Crete, Greece (e-mail: karipidis@telecom.tuc.gr; nikos@telecom.tuc.gr).

Z.-Q. Luo is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: luozq@ece.umn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2007.909010

brought renewed interest in multicast routing protocols. These protocols were originally conceived and have since evolved under the “wireline premise:” the physical network is a graph comprising point-to-point links that do not interfere with each other at the physical layer. Today, multicast routing protocols operate at the network or application layer, using either controlled flooding or minimum spanning tree access.

As wireless networks become ever more ubiquitous, and wireless becomes the choice for not only the “last hop” but also suburban- and metropolitan-area backbones, wireless multicasting solutions are needed to account for and exploit the idiosyncracies of the wireless medium. Wireless is inherently a broadcast medium, where it is possible to reach multiple destinations with a single transmission; different cochannel transmissions are interfering with one another at the intended destination(s); and links are subject to fading and shadowing, in addition to cochannel interference.

The broadcast advantage of wireless has of course been exploited since the early days of radio. The interference problem was dealt with by allocating different frequency bands to the different stations, and transmission was mostly isotropic or focused towards a specific service area.

Today, the situation with wireless networks is much different. First, transmissions need not be “blind.” Many wireless network standards provision the use of transmit antenna arrays. Using baseband beamforming, it is possible to steer energy in the direction(s) of the intended users, whose locations (or, more generally, channels) can often be accurately estimated. Second, the push towards higher capacity and end-user rates necessitates cochannel transmission which exploits the spatial diversity in the user population (*spatial multiplexing*). Third, quality of service (QoS) is an important consideration, especially in wireless backhaul solutions like 802.16e. Finally, due to cochannel interference, wireless multicasting cannot be dealt with in isolation, one group at a time; a joint solution is needed.

The problem of transmit beamforming towards a (single) group of users was first considered in the Ph.D. dissertation of Lopez [2], using the averaged (over all users in the group) received signal-to-noise ratio (SNR) as the design criterion. The solution boils down to a relatively simple eigenvalue problem, but no SNR guarantee is provided this way: some users may get really poor SNR [3]. This is not acceptable in multicasting applications, because it is the worst SNR that determines the *common* information rate. QoS (providing a guaranteed minimum received SNR to every user) and max-min fair (MMF) (maximizing the smallest received SNR) designs

were first proposed in [4] and [3], where it was shown that the core problem is NP-hard, yet high-quality approximate solutions can be obtained using relaxation techniques based on semidefinite programming (SDP). The latter is a class of convex optimization problems which can be solved in polynomial time by powerful interior point methods.

This paper formulates a new and interesting problem: transmit beamforming for multicasting to *multiple* cochannel groups under QoS and MMF criteria. The joint design problem is considered, since designing a transmit beamformer separately for each multicast group can be far from optimal, due to intergroup interference. By simultaneously serving several cochannel groups, the spectral efficiency is much higher than in the single-group case. The extension to multiple groups is nontrivial in the following ways:

- the multigroup QoS problem can be infeasible;
- the QoS and MMF versions are different, unlike the single-group case;
- the approximation step is much more involved: randomization is coupled with multigroup multicast power control, which is of interest in its own right.

We propose two solid and well-motivated (Lagrange dual) algorithms. In addition to semidefinite relaxation ideas, our solutions entail a cochannel multigroup multicast power control component, which can be viewed as a generalization of multiuser power control ideas for the cellular downlink (see, e.g., [5] and references therein). It is important to note that the problem formulation considered here contains as special cases the single-group multicasting (broadcasting) problem [3], as well as the multiuser downlink beamforming problem (see, e.g., [6] and references therein), where each multicast group consists of a single receiver. Our extensive numerical results, including experiments with measured channels, show that in the multigroup case as well, the proposed semidefinite relaxation (SDR)-based algorithms work remarkably well.

Notation: Boldface uppercase letters denote matrices, whereas boldface lowercase letters denote column vectors. The superscripts $(\cdot)^*$, $(\cdot)^T$, and $(\cdot)^H$ denote complex conjugate, transpose, and Hermitian (conjugate) transpose matrix operators, respectively. $\text{tr}(\cdot)$, $\text{rank}(\cdot)$, $|\cdot|$, and $\|\cdot\|_2$ denote the trace, the rank, the absolute value, and the Euclidean norm operators, respectively. By $\mathbf{X} \succeq \mathbf{0}$ we denote that \mathbf{X} is a Hermitian positive-semidefinite matrix. Finally, \mathbf{I}_N and $\mathbf{1}_G$ denote the $N \times N$ identity matrix and the $G \times 1$ all ones vector.

II. DATA MODEL AND PROBLEM STATEMENT

Consider a wireless scenario comprising 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 $i \in \{1, \dots, M\}$. Let there be a total of $1 \leq G \leq M$ multicast groups, $\{\mathcal{G}_1, \dots, \mathcal{G}_G\}$, where \mathcal{G}_k is the index set of the receivers participating in multicast group k , and $k \in \{1, \dots, G\}$. Each receiver listens to a single multicast; thus, $\mathcal{G}_k \cap \mathcal{G}_\ell = \emptyset$, $\ell \neq k$, $\cup_{k=1}^G \mathcal{G}_k = \{1, \dots, M\}$, and, denoting $G_k := |\mathcal{G}_k|$, $\sum_{k=1}^G G_k = M$.




Fig. 1. Cochannel multicast beamforming concept (note that groups need not be spatially clustered).

Let $\mathbf{w}_k^* \in \mathbb{C}^N$ denote the beamforming weight vector applied to the N transmitting antenna elements to generate the spatial channel for transmission to group k (see Fig. 1). Then the signal transmitted by the antenna array is equal to $\sum_{k=1}^G \mathbf{w}_k^* s_k(t)$, where $s_k(t)$ is the temporal information-bearing signal directed to receivers in multicast group k . If each $s_k(t)$ is zero-mean, temporally white with unit variance, and the waveforms $\{s_k(t)\}_{k=1}^G$ are mutually uncorrelated, then the total power radiated by the transmitting antenna array is equal to $\sum_{k=1}^G \|\mathbf{w}_k\|_2^2$.

The joint design of transmit beamformers can then be posed as the problem of minimizing the total radiated power subject to meeting prescribed signal-to-interference-plus-noise ratio (SINR) constraints γ_i at each of the M receivers

$$\begin{aligned} \mathcal{Q} : \\ \min_{\{\mathbf{w}_k \in \mathbb{C}^N\}_{k=1}^G} \sum_{k=1}^G \|\mathbf{w}_k\|_2^2 \\ \text{s.t. : } \frac{|\mathbf{w}_k^H \mathbf{h}_i|^2}{\sum_{\ell \neq k} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq \gamma_i \\ \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\}. \end{aligned}$$

The problem formulation above assumes that the channel vectors $\{\mathbf{h}_i\}_{i=1}^M$ of all intended users and the corresponding noise variances $\{\sigma_i^2\}_{i=1}^M$ are known at the transmitter. Contrary to the single-group case, problem \mathcal{Q} can be infeasible due to interference, if the SINR requirements are too stringent and/or the channels of users listening to different multicasts are highly correlated. Then, in order to render the problem feasible it is necessary to loosen some of the QoS thresholds or deny service to some users in the specific frequency tone/time slot, by means of proper admission control. Beyond feasibility concerns, it is important to note that each beamformer must serve multiple users listening to the same information. When a feasible solution exists, at least one SINR constraint per group will be satisfied with equality at the optimum, whereas the others may be inactive (i.e., oversatisfied); this is contrary to the case of independent information transmission, where all SINR constraints are tight at the optimum.

Problem \mathcal{Q} is a quadratically constrained quadratic programming (QCQP) problem and it can be easily seen that the constraints are *nonconvex*. Furthermore, it contains as a special case the associated broadcasting problem ($G = 1$), which was proven to be NP-hard in [3], thereby implying the following result.

Claim 1: Problem \mathcal{Q} is NP-hard.

This reveals the fundamental difference of the multigroup multicast QoS problem to its other special case: the multiuser downlink beamforming problem ($G = M$), which admits an equivalent *convex* (specifically, second-order cone programming) reformulation [6].¹ Claim 1 motivates (cf. [7]) the pursuit of sensible approximate solutions to the QoS problem \mathcal{Q} .

III. RELAXATION

Towards this end, let us change the optimization variables to $\{\mathbf{X}_k := \mathbf{w}_k \mathbf{w}_k^H\}_{k=1}^G$. Note that $\mathbf{X}_k = \mathbf{w}_k \mathbf{w}_k^H$ for some $\mathbf{w}_k \in \mathbb{C}^N$ if and only if $\mathbf{X}_k \succeq \mathbf{0}$ and $\text{rank}(\mathbf{X}_k) = 1$. Defining $\{\mathbf{Q}_i := \mathbf{h}_i \mathbf{h}_i^H\}_{i=1}^M$ and using that $\text{tr}(\mathbf{A}\mathbf{B}) = \text{tr}(\mathbf{B}\mathbf{A})$ for matrices \mathbf{A} , \mathbf{B} of compatible dimensions, the signal power received at user i by multicast k can be expressed as $|\mathbf{w}_k^H \mathbf{h}_i|^2 = \mathbf{h}_i^H \mathbf{w}_k \mathbf{w}_k^H \mathbf{h}_i = \text{tr}(\mathbf{h}_i^H \mathbf{w}_k \mathbf{w}_k^H \mathbf{h}_i) = \text{tr}(\mathbf{h}_i \mathbf{h}_i^H \mathbf{w}_k \mathbf{w}_k^H) = \text{tr}(\mathbf{Q}_i \mathbf{X}_k)$. Likewise, the power of each beamforming vector can be written as $\|\mathbf{w}_k\|_2^2 = \mathbf{w}_k^H \mathbf{w}_k = \text{tr}(\mathbf{w}_k^H \mathbf{w}_k) = \text{tr}(\mathbf{w}_k \mathbf{w}_k^H) = \text{tr}(\mathbf{X}_k)$. It follows that problem \mathcal{Q} can be *equivalently* reformulated as

$$\begin{aligned} \mathcal{Q} : \\ \min_{\{\mathbf{X}_k \in \mathbb{C}^{N \times N}\}_{k=1}^G} \sum_{k=1}^G \text{tr}(\mathbf{X}_k) \\ \text{s.t. : } \quad \text{tr}(\mathbf{Q}_i \mathbf{X}_k) \geq \gamma_i \sum_{\ell \neq k} \text{tr}(\mathbf{Q}_i \mathbf{X}_\ell) + \gamma_i \sigma_i^2 \\ \quad \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\} \\ \quad \mathbf{X}_k \succeq \mathbf{0} \quad \forall k \in \{1, \dots, G\} \\ \quad \text{rank}(\mathbf{X}_k) = 1 \quad \forall k \in \{1, \dots, G\} \end{aligned}$$

where we have used that the denominator is positive. Note that if the *instantaneous* channel vectors $\{\mathbf{h}_i\}_{i=1}^M$ are unknown, the channel correlation matrices can be used instead as input parameters $\{\mathbf{Q}_i\}_{i=1}^M$. However, in this case the resulting design can only guarantee *average* received SINRs.

¹Note that there exist other special cases of problem \mathcal{Q} that are not NP-hard: e.g., a restriction to Vandermonde channel vectors enables convex reformulation and thereby efficient solution of the problem [8], [9].

Dropping the last G rank-one constraints, which are non-convex, we arrive at the following relaxation of problem \mathcal{Q} :

$$\begin{aligned} \mathcal{Q}_r : \\ \min_{\{\mathbf{X}_k \in \mathbb{C}^{N \times N}\}_{k=1}^G} \sum_{k=1}^G \text{tr}(\mathbf{X}_k) \\ \text{s.t. : } \quad \gamma_i \sum_{\ell \neq k} \text{tr}(\mathbf{Q}_i \mathbf{X}_\ell) - \text{tr}(\mathbf{Q}_i \mathbf{X}_k) + \gamma_i \sigma_i^2 \leq 0 \\ \quad \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\} \\ \quad \mathbf{X}_k \succeq \mathbf{0} \quad \forall k \in \{1, \dots, G\}. \end{aligned}$$

Problem \mathcal{Q}_r consists of a linear objective function, M linear inequality constraints and G positive-semidefinite constraints; hence, it is a standard SDP problem. Modern SDP solvers, such as SeDuMi [10], use interior point methods to efficiently find an optimum solution to the problem, if it is feasible; otherwise, they return a certificate of infeasibility. The SDP problem \mathcal{Q}_r has G matrix variables of size $N \times N$, and M linear constraints. Interior point methods will take $O(\sqrt{GN} \log(1/\epsilon))$ iterations, with each iteration requiring at most $O(G^3 N^6 + MGN^2)$ arithmetic operations [11], where the parameter ϵ represents the solution accuracy at the algorithm's termination. Actual runtime complexity will usually scale far slower with G , N , M than this worst-case bound.

IV. APPROXIMATION

Problem \mathcal{Q} may not admit a feasible solution, but if it does, the aforementioned approach will yield a solution to problem \mathcal{Q}_r . However, due to the relaxation, this solution will not, in general, consist of rank-one matrices. This is because the (convex) feasible set of problem \mathcal{Q}_r is a superset of the (nonconvex) feasible set of problem \mathcal{Q} . In addition, the optimum objective value of problem \mathcal{Q}_r is merely a lower bound on the transmitted power required by the rank-one transmit beamforming scheme. An approximate solution to the original QoS problem \mathcal{Q} can be found using a *randomization* technique (see, e.g., [13], [14]). The idea is to generate candidate sets of beamforming vectors $\{\mathbf{w}_k\}_{k=1}^G$ from the optimum solution matrices $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$ of problem \mathcal{Q}_r and choose the one that can be scaled to satisfy the SINR constraints of problem \mathcal{Q} at the minimum total power cost.

We propose the use of the Gaussian randomization method (see, e.g., [13], [14]) for the generation of the candidate beamformers, motivated by its successful application in related QCQP problems and especially in the single-group multicasting problem [3], [15]. Initially, the eigen-decomposition of each optimal matrix is calculated as $\mathbf{X}_k^{\text{opt}} = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{U}_k^H$. Then, the ℓ th candidate beamformer for multicast k is generated as $\mathbf{w}_k^\ell = \mathbf{U}_k \mathbf{\Sigma}_k^{1/2} \mathbf{v}_\ell$, where $\mathbf{v}_\ell \in \mathbb{C}^N \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$, so that $E[\mathbf{w}_k^\ell (\mathbf{w}_k^\ell)^H] = \mathbf{X}_k^{\text{opt}}$. The main difference relative to the simpler broadcasting case ($G = 1$) considered in [3], is that here we cannot simply "scale up" the candidate beamforming vectors generated during randomization to satisfy the SINR constraints of problem \mathcal{Q} . The reason is that, in contrast to

[3], we herein deal with an interference scenario, and boosting one group's beamforming vector also increases interference to nodes in other groups. Whether it is feasible to satisfy the constraints for a given set of candidate beamforming vectors is also an issue here.

Let $a_{k,i} := |\mathbf{w}_k^H \mathbf{h}_i|^2$ denote the signal power received at receiver i from the stream directed towards users in multicast group k . Let $\beta_k := \|\mathbf{w}_k\|_2^2$, and p_k denote the sought power boost (or reduction) factor for multicast group k . Then the following *multigroup multicast power control* (MMPC) problem emerges in converting candidate beamforming vectors to a candidate solution of problem \mathcal{Q} .

$$\mathcal{S}^{\mathcal{Q}} : \quad \min_{\{p_k \in \mathbb{R}\}_{k=1}^G} \sum_{k=1}^G \beta_k p_k$$

$$\text{s.t. : } \quad \frac{a_{k,i} p_k}{\sum_{\ell \neq k} a_{\ell,i} p_\ell + \sigma_i^2} \geq \gamma_i$$

$$\forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\}$$

$$p_k \geq 0 \quad \forall k \in \{1, \dots, G\}.$$

As in Section III, taking advantage of the fact that the denominator is positive, problem $\mathcal{S}^{\mathcal{Q}}$ can be equivalently reformulated as

$$\mathcal{S}^{\mathcal{Q}} : \quad \min_{\{p_k \in \mathbb{R}\}_{k=1}^G} \sum_{k=1}^G \beta_k p_k$$

$$\text{s.t. : } \quad \gamma_i \sum_{\ell \neq k} a_{\ell,i} p_\ell - a_{k,i} p_k + \gamma_i \sigma_i^2 \leq 0$$

$$\forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\}$$

$$p_k \geq 0 \quad \forall k \in \{1, \dots, G\}.$$

Problem $\mathcal{S}^{\mathcal{Q}}$ is a *linear program* (LP) with G nonnegative variables and M linear inequality constraints. For a feasible instance of the MMPC problem $\mathcal{S}^{\mathcal{Q}}$, interior point methods can generate an ϵ -optimal solution in $O(\sqrt{G} \log(1/\epsilon))$ iterations, each requiring at most $O(G^3 + MG)$ arithmetic operations [11]. Otherwise, they yield an infeasibility certificate. This is a useful property in determining the feasibility of a candidate beamforming configuration. The simplex method could also be used and it will typically be more efficient for small problem sizes.

As noted already, for $G = M$ (independent information transmission to each receiver), problem \mathcal{Q}_r is in fact *equivalent to* (not a relaxation of) problem \mathcal{Q} , see [6]. Likewise, problem $\mathcal{S}^{\mathcal{Q}}$ reduces to the well-known multiuser downlink power control problem, which can be solved using simpler means (see, e.g., [5]): matrix inversion and iterative descent algorithms. In this special case, (in)feasibility can be determined from the spectral radius of a certain ‘‘connectivity’’ matrix. Similar simplifications for the general instance of MMPC are perhaps possible, but nontrivial. In fact, an iterative MMPC algorithm

based on the concept of interference functions was proposed in [12]. However, the power iterations advocated therein are only guaranteed to converge when the problem is feasible. Keeping in mind that the MMPC problem emerges in the context of randomization, it is clear that effective detection of infeasibility is an important issue. Furthermore, even when the problem is feasible, it is not clear whether the power iterations in [12] require smaller overall complexity to find an optimum solution than the available LP routines, which are highly efficient.

The overall algorithm for generating an approximate solution to the original QoS problem \mathcal{Q} can be summarized as follows:

- 1) Relaxation: Solve problem \mathcal{Q}_r using a SDP solver. Denote the solution $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$.
- 2) Randomization/Scaling Loop: For each k , generate a vector using the Gaussian randomization technique. If, for some k , $\text{rank}(\mathbf{X}_k^{\text{opt}}) = 1$, then use the principal component instead. Next, feed the resulting set of candidate beamforming vectors $\{\mathbf{w}_k\}_{k=1}^G$ into the MMPC problem $\mathcal{S}^{\mathcal{Q}}$ and solve it using a LP solver. If the particular instance of problem $\mathcal{S}^{\mathcal{Q}}$ is infeasible or yields a larger objective value than previously checked candidates, discard the proposed set of candidate beamforming vectors; else, record the set of beamforming vectors, the associated scaling factors $\{p_k\}_{k=1}^G$ and the objective value. Repeat for a predetermined number N_{rand} of randomizations.

Assuming that the randomization/scaling loop yields at least one feasible solution, let $\{\mathbf{w}_k, p_k\}_{k=1}^G$ denote the recorded beamvectors and scaling factors. Then, the approximate solution of problem \mathcal{Q} is given by $\{\sqrt{p_k} \mathbf{w}_k\}_{k=1}^G$.

The overall complexity of this solution is that of solving the SDP problem \mathcal{Q}_r once and the LP problem $\mathcal{S}^{\mathcal{Q}}$ N_{rand} times. The choice of N_{rand} is a tradeoff between the extent of sub-optimality of the final solution and the overall complexity of the algorithm. The *quality* of the approximate solution to problem \mathcal{Q} can be measured by the ratio of the minimum objective value of problem $\mathcal{S}^{\mathcal{Q}}$, attained in the randomization/scaling loop, to the lower bound on transmitted power, obtained by the solution of problem \mathcal{Q}_r . The numerical results reported in Section VI show that a few hundred randomizations are adequate, in most scenarios considered, to yield a solution which is at most 3–4 dB away from this lower bound; hence, even less from the (NP-hard to find) optimum. The lower bound obtained by solving problem \mathcal{Q}_r can be further motivated from a duality perspective, as in [3]; that is, the aforementioned relaxation lower bound is in fact the tightest lower bound on the optimum value of problem \mathcal{Q} attainable via Lagrangian duality [16]. This follows from arguments in [17] (see also the single-group case in [3]), due to the fact that \mathcal{Q} is a QCQP problem. For theoretical *a priori* bounds on the extent of the suboptimality of the solution in [3] for the single-group case see [15].

V. JOINT MAX-MIN FAIR BEAMFORMING

In this section, we consider the related problem of maximizing the minimum SINR, received by any of the M intended

users irrespective of the multicast group they belong to, subject to an upper bound P on the total transmission power. Actually, a more general problem is considered, in which each received SINR is scaled by a predetermined (positive real) constant weight factor $1/\gamma_i$, to account for possibly different grades of service. Let $\mathbf{g} := [\gamma_1, \dots, \gamma_M]^T$. The joint (weighted) MMF transmit multicast beamformer design is formulated as

$$\mathcal{F} : \begin{aligned} & \max_{\{\mathbf{w}_k \in \mathbb{C}^N\}_{k=1}^G} \min_{k \in \{1, \dots, G\}} \min_{i \in \mathcal{G}_k} \frac{1}{\gamma_i} \frac{|\mathbf{w}_k^H \mathbf{h}_i|^2}{\sum_{\ell \neq k} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \\ \text{s.t. : } & \sum_{k=1}^G \|\mathbf{w}_k\|_2^2 \leq P. \end{aligned}$$

This problem formulation is important for systems required to comply with a strict upper bound on the total transmitted power, e.g., due to regulation. It is straightforward to see that the inequality power constraint will be met with equality at the optimum. Otherwise, if there is power budget left, one could distribute it evenly, i.e., multiply all beamformers by a constant $c > 1$, thereby increasing the minimum SINR (note that $\sigma_i^2 > 0$), thus contradicting optimality. We may therefore focus on the equality constrained problem and denote this as \mathcal{F} from now on. Introducing an auxiliary (positive real) variable t to lower bound the worst-case scaled SINR, the equality constrained version of problem \mathcal{F} can be *equivalently* rewritten as

$$\mathcal{F} : \begin{aligned} & \max_{\{\mathbf{w}_k \in \mathbb{C}^N\}_{k=1}^G, t \in \mathbb{R}} t \\ \text{s.t. : } & \frac{1}{\gamma_i} \frac{|\mathbf{w}_k^H \mathbf{h}_i|^2}{\sum_{\ell \neq k} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq t \\ & \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\} \\ & \sum_{k=1}^G \|\mathbf{w}_k\|_2^2 = P, \text{ and } t \geq 0. \end{aligned}$$

The design criterion seeks to maximize the worst-case *scaled* SINR, so as to ensure *weighted* fairness among the received multicasts. Obviously, *equal* fairness is a special case that corresponds to the choice of $\mathbf{g} = \mathbf{1}_M$. Contrary to the QoS approach, discussed in Section II, problem \mathcal{F} always admits a feasible solution, apart from the trivial case of zero channel vectors. Denoting as t_0 the optimum value of \mathcal{F} , the optimum beamformers guarantee SINR levels equal to $t_0 \mathbf{g}$. Interpreting the weight factors \mathbf{g} as target SINRs, these are achieved, with total transmitted power P , if and only if $t_0 \geq 1$. In this sense, the MMF problem is more flexible than the QoS one and can be used to determine whether, in a power-constrained system, a specified set of SINR targets \mathbf{g} can be satisfied or not. Moreover, it determines the exact level of under- or over-satisfaction t_0 .

The formulation of problem \mathcal{F} is a generalization of the respective MMF multiuser downlink beamforming problem (see, e.g., [18] and references therein). As for the QoS case, not all SINR inequalities will be tight in general at the optimum. Regarding complexity, problem \mathcal{F} contains as a special case the associated broadcasting problem ($G = 1$), which was proved to be NP-hard in [3]. This immediately implies the following result.

Claim 2: Problem \mathcal{F} is NP-hard.

Claim 2 motivates the pursuit of sensible approximate solutions to the MMF problem \mathcal{F} .²

Before proceeding in proposing an algorithm to find such an approximate solution, let us have a closer look at the connection between problem \mathcal{F} and problem \mathcal{Q} . For a given set of channels and noise powers, \mathcal{F} is parameterized by \mathbf{g} and P . We will use the notation $\mathcal{F}(\mathbf{g}, P)$ to capture this dependence, and, with slight abuse of notation,³ $t = \mathcal{F}(\mathbf{g}, P)$ to denote the associated optimum value (maximum worst-case scaled SINR). Likewise, \mathcal{Q} is parameterized by the vector \mathbf{g} of QoS constraints; we will use the notation $\mathcal{Q}(\mathbf{g})$ to account for this, and $P = \mathcal{Q}(\mathbf{g})$ to denote the associated optimum value (minimum power).

Generalizing the respective results for the two extreme cases of the multigroup multicast beamforming problem (namely, $G = M$ [18] and $G = 1$ [3]), we have the following result:

Claim 3: The QoS problem \mathcal{Q} and the MMF problem \mathcal{F} are related as follows:

$$t = \mathcal{F}(\mathbf{g}, \mathcal{Q}(\mathbf{g})) \quad (1)$$

$$P = \mathcal{Q}(\mathcal{F}(\mathbf{g}, P)\mathbf{g}). \quad (2)$$

Proof: Contradiction can be used to prove (1). Let $\{\mathbf{w}_k^Q\}_{k=1}^G$ and P^Q denote an optimal solution and the associated optimal value to a feasible instance of problem $\mathcal{Q}(t\mathbf{g})$, where $t\mathbf{g}$ are the required SINR targets. Consider the problem instance $\mathcal{F}(\mathbf{g}, P^Q)$. The set $\{\mathbf{w}_k^Q\}_{k=1}^G$ is a feasible solution with objective value t . Assume the existence of another feasible solution $\{\mathbf{w}_k^F\}_{k=1}^G$ with associated optimal value $t^F > t$. Then, it is possible to find a constant $c < 1$ to scale down this solution set, while still fulfilling the SINR constraints of problem $\mathcal{Q}(t\mathbf{g})$. The resulting set $\{c\mathbf{w}_k^F\}_{k=1}^G$ has smaller objective value (total transmitted power) than P^Q , which contradicts optimality of $\{\mathbf{w}_k^Q\}_{k=1}^G$.

A similar procedure can be used to prove (2). Specifically, let $\{\mathbf{w}_k^F\}_{k=1}^G$ and t^F denote an optimal solution and the associated optimal value to a problem instance $\mathcal{F}(\mathbf{g}, P)$. Consider the problem instance $\mathcal{Q}(t^F \mathbf{g})$. The set $\{\mathbf{w}_k^F\}_{k=1}^G$ is a feasible solution with objective value P . Assume the existence of another feasible solution $\{\mathbf{w}_k^Q\}_{k=1}^G$ with associated optimal value $P^Q < P$. This contradicts optimality of $\{\mathbf{w}_k^F\}_{k=1}^G$ for $\mathcal{F}(\mathbf{g}, P)$.

²As for problem \mathcal{Q} , there exist special cases of problem \mathcal{F} that are not NP-hard: e.g., for Vandermonde channel vectors it admits a SDP reformulation [8], [9] and for independent data transmission a generalized eigenvalue problem reformulation [18].

³The meaning will be clear from context.

since the power budget $P - P^Q$ can be distributed evenly to yield an objective value larger than t^F . ■

Another useful property of both formulations is shown in the following claim.

Claim 4: The optimum objective values of the QoS problem $\mathcal{Q}(tg)$ and the MMF problem $\mathcal{F}(\mathbf{g}, P)$ are monotonically non-decreasing in t and P , respectively, for a given \mathbf{g} .

Proof: The feasible set of $\mathcal{Q}(tg)$ is decreasing in t . For $\mathcal{F}(\mathbf{g}, P)$, any additional power can be evenly distributed, thereby increasing all SINRs, provided that all the $\{\sigma_i^2\}_{i=1}^M$ are nonzero. ■

Corollary 1: A solution to $\mathcal{F}(\mathbf{g}, P)$ can be found by iteratively solving $\mathcal{Q}(tg)$ for varying values of t . Claim 3 guarantees optimality of the solution for $P = \mathcal{Q}(tg)$ and Claim 4 enables the use of a simple one-dimensional bisection search for the sought t (see [18] for the special case of multiuser downlink beamforming). Similarly, bisection of $t = \mathcal{F}(\mathbf{g}, P)$ over P can be used to solve $\mathcal{Q}(tg)$.

Corollary 1 suggests a solution to the MMF problem, provided that the QoS problem can be solved *optimally*. However, \mathcal{Q} is NP-hard and we can only find an approximate solution, as proposed in Sections III and IV. Due to this, and keeping in mind that \mathcal{F} is NP-hard (Claim 2), we again pursue a respective sensible approximate solution.

Using the notation introduced in Section III and following similar steps as in the relaxation $\mathcal{Q} \rightarrow \mathcal{Q}_r$, the following relaxation of the original MMF problem \mathcal{F} is obtained by dropping the nonconvex rank-one constraints, associated with the matrices $\{\mathbf{X}_k\}_{k=1}^G$.

$$\begin{aligned}
 \mathcal{F}_r : & \\
 & \max_{\{\mathbf{X}_k \in \mathbb{C}^{N \times N}\}_{k=1}^G, t \in \mathbb{R}} t \\
 \text{s.t. : } & t\gamma_i \left(\sum_{\ell \neq k} \text{tr}(\mathbf{Q}_i \mathbf{X}_\ell) + \sigma_i^2 \right) - \text{tr}(\mathbf{Q}_i \mathbf{X}_k) \leq 0 \\
 & \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\} \\
 & \sum_{k=1}^G \text{tr}(\mathbf{X}_k) = P \\
 & \mathbf{X}_k \succeq \mathbf{0} \quad \forall k \in \{1, \dots, G\}, \text{ and } t \geq 0.
 \end{aligned}$$

At first glance, problem \mathcal{F}_r may seem to be of the same form as \mathcal{Q}_r of Section III, except for the extra linear equality constraint on the total transmission power, the extra nonnegativity constraint, and the different (yet still linear) objective function. However, contrary to \mathcal{Q}_r , \mathcal{F}_r does not admit an equivalent SDP reformulation, because the M inequality constraints on the received SINRs are nonlinear (note that t is a variable).⁴

⁴In the single-group multicast beamforming case, \mathcal{F}_r is a SDP problem [3].

Claim 5: The relaxed QoS problem \mathcal{Q}_r and the relaxed MMF problem \mathcal{F}_r are related as follows:

$$t = \mathcal{F}_r(\mathbf{g}, \mathcal{Q}_r(t\mathbf{g})) \quad (3)$$

$$P = \mathcal{Q}_r(\mathcal{F}_r(\mathbf{g}, P)\mathbf{g}). \quad (4)$$

Proof: Verbatim to the proof of Claim 3, denoting the problem solutions as $\{\mathbf{X}_k\}_{k=1}^G$ instead of $\{\mathbf{w}_k\}_{k=1}^G$. ■

Claim 6: The optimum objective values of the relaxed QoS problem $\mathcal{Q}_r(t\mathbf{g})$ and the relaxed MMF problem $\mathcal{F}_r(\mathbf{g}, P)$ are monotonically nondecreasing in t and P , respectively, for a given \mathbf{g} .

Proof: Verbatim to the proof of Claim 4. ■

Corollary 2: Due to Claims 5 and 6, the relaxed problem $\mathcal{F}_r(\mathbf{g}, P)$ can be solved by a one-dimensional bisection search over t , to attain $P = \mathcal{Q}_r(t\mathbf{g})$.

Specifically, let $\tilde{t} = \mathcal{F}_r(\mathbf{g}, P)$. A feasible solution of $\mathcal{F}_r(\mathbf{g}, P)$ that is at most $\epsilon > 0$ away from \tilde{t} can be generated as follows. Let $[L, U]$ be an interval containing \tilde{t} . Due to the nonnegativity of \tilde{t} , the lower bound is initialized as $L = 0$. Assuming that the total available power is directed towards a single group and using the Cauchy-Schwartz inequality, the upper bound is initialized as $U = \min_{i \in \{1, \dots, M\}} (P \|\mathbf{h}_i\|_2^2 / \gamma_i \sigma_i^2)$.

Given $[L, U]$, the SDP problem $\mathcal{Q}_r(t\mathbf{g})$ is solved at the midpoint $t = (L + U)/2$ of the interval. If it is feasible for the given choice of t and its objective value is lower than P , the solution is stored and $L := t$; otherwise $U := t$. The use of interior point SDP solvers, such as SeDuMi [10], is useful in this context, because they do not only yield an efficient solution to problem $\mathcal{Q}_r(t\mathbf{g})$ when the latter is feasible, but they also provide a certificate of infeasibility otherwise. The aforementioned steps are repeated until $U - L \leq \epsilon$. Since in each iteration the interval is halved, the algorithm requires only $N_{\text{iter}} = \lceil \log_2((U - L)/\epsilon) \rceil$ iterations. In practice, 10–12 iterations are usually enough for typical problem setups. Building on [18], a similar bisection search algorithm was also proposed in [12].

When the algorithm terminates, the resulting matrices $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$ are an ϵ -optimal solution of $\mathcal{F}_r(\mathbf{g}, P)$. The associated optimal value, which is approximately equal to \tilde{t} , is merely an upper bound on the scaled SINR that can be guaranteed to every user, for the specific power budget P . This bound can only be met when all matrices $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$ are rank one, so that their principal components can be chosen as optimum beamforming vectors. However, due to the relaxation, this is generally not the case. As in the QoS approach, postprocessing of the relaxed solution is needed, when the solution matrices $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$ are not all rank one, to yield an approximate solution to the original MMF problem \mathcal{F} . This can be accomplished by a combined randomization/scaling procedure, similar to the one described in Section IV. Specifically, the Gaussian randomization, described in Section IV, may be used in a first step to create candidate sets of beamforming vectors $\{\mathbf{w}_k\}_{k=1}^G$ in the span of the respective transmit covariance matrices

(the optimum solution matrices $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$ of problem \mathcal{F}_r). In a second step, the total available transmission power P is optimally allocated to the candidate beamforming vectors, by means of an appropriate MMPC problem, as explained in the rest of this section. The set $\{\mathbf{w}_k, p_k\}_{k=1}^G$ of beamforming vectors and respective power boost (or back-off) factors, that yields the highest objective value is chosen among all solutions generated this way. The approximate solution to the original MMF problem \mathcal{F} is then equal to $\{\sqrt{p_k}\mathbf{w}_k\}_{k=1}^G$.

Given a candidate set of beamforming vectors $\{\mathbf{w}_k\}_{k=1}^G$, the power budget P can be optimally allocated among them by solving the following MMPC problem:

$$\mathcal{S}^{\mathcal{F}} : \begin{aligned} & \max_{\{p_k \in \mathbb{R}\}_{k=1}^G, t \in \mathbb{R}} t \\ \text{s.t. : } & \frac{1}{\gamma_i \sum_{\ell \neq k} p_\ell \alpha_{\ell,i} + \sigma_i^2} \geq t \\ & \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\} \\ & \sum_{k=1}^G \beta_k p_k = P \\ & p_k \geq 0 \quad \forall k \in \{1, \dots, G\}, \text{ and } t \geq 0. \end{aligned}$$

where the coefficients $\alpha_{k,i}$ and β_k are as introduced in Section IV.

Unlike $\mathcal{S}^{\mathcal{Q}}$, $\mathcal{S}^{\mathcal{F}}$ does not admit an equivalent LP reformulation, because of the M nonlinear inequality constraints (note that t is a variable).

Claim 7: The QoS MMPC problem $\mathcal{S}^{\mathcal{Q}}$ and the MMF MMPC problem $\mathcal{S}^{\mathcal{F}}$ are related as follows:

$$t = \mathcal{S}^{\mathcal{F}}(\mathbf{g}, \mathcal{S}^{\mathcal{Q}}(t\mathbf{g})) \quad (5)$$

$$P = \mathcal{S}^{\mathcal{Q}}(\mathcal{S}^{\mathcal{F}}(\mathbf{g}, P)\mathbf{g}). \quad (6)$$

Proof: Verbatim to the proof of Claim 3, denoting the problem solutions as $\{p_k\}_{k=1}^G$ instead of $\{\mathbf{w}_k\}_{k=1}^G$. ■

Claim 8: The optimum objective values of the QoS MMPC problem $\mathcal{S}^{\mathcal{Q}}(t\mathbf{g})$ and the MMF MMPC problem $\mathcal{S}^{\mathcal{F}}(\mathbf{g}, P)$ are monotonically nondecreasing in t and P , respectively, for a given \mathbf{g} .

Proof: Verbatim to the proof of Claim 4. ■

Corollary 3: Due to Claims 7 and 8, the MMF MMPC problem $\mathcal{S}^{\mathcal{F}}(\mathbf{g}, P)$ can be solved by a one-dimensional bisection search over t , to attain $P = \mathcal{S}^{\mathcal{Q}}(t\mathbf{g})$.

The bisection algorithm, described earlier in this section, can be used again to obtain a solution to problem $\mathcal{S}^{\mathcal{F}}(\mathbf{g}, P)$. The search interval is bounded below by $L = 0$, as before. However, the upper bound may now be further restricted to $U = \tilde{t}$ (the optimal objective value of $\mathcal{F}_r(\mathbf{g}, P)$). The difference is that for each iteration (value of t), the LP problem $\mathcal{S}^{\mathcal{Q}}(t\mathbf{g})$ is solved instead of the SDP problem $\mathcal{Q}_r(t\mathbf{g})$.

The overall complexity of finding an approximate solution to the original MMF problem \mathcal{F} is that of solving N_{iter} times

the SDP problem \mathcal{Q}_r and $N_{\text{rand}}N'_{\text{iter}}$ times the LP problem $\mathcal{S}^{\mathcal{Q}}$, where N_{iter} and N'_{iter} denote the number of bisection iterations required for the solution of \mathcal{F}_r and $\mathcal{S}^{\mathcal{F}}$, respectively, and N_{rand} is the number of Gaussian randomization trials. The quality of the final approximate solution to problem \mathcal{F} can be measured by the ratio of the upper bound obtained by the solution of the relaxed problem \mathcal{F}_r to the maximum value of problem $\mathcal{S}^{\mathcal{F}}$ attained in the randomization/scaling loop.

We remark that $\mathcal{S}^{\mathcal{F}}$ also admits an equivalent reformulation as geometric problem (GP)

$$\mathcal{S}^{\mathcal{F}} : \begin{aligned} & \min_{\{p_k \in \mathbb{R}\}_{k=1}^G, t \in \mathbb{R}} t^{-1} \\ \text{s.t. : } & \sum_{\ell \neq k} \left(\gamma_i \alpha_{\ell,i} \alpha_{k,i}^{-1} \right) t p_\ell p_k^{-1} + \left(\gamma_i \sigma_i^2 \alpha_{k,i}^{-1} \right) t p_k^{-1} \leq 1 \\ & \forall i \in \mathcal{G}_k \quad \forall k, \ell \in \{1, \dots, G\} \\ & \sum_{k=1}^G (\beta_k P^{-1}) p_k \leq 1 \end{aligned}$$

which can be efficiently solved using modern interior point methods, bypassing the need for bisection. We note, however, that the number of bisection steps is rather small in general (at most 12 in all cases that we considered in the reported results of Section VI), so this does not appear to be a big issue.

VI. NUMERICAL RESULTS

A. Monte Carlo Simulation Results

In Sections III and IV, we have derived a two-step polynomial-time algorithm to generate an approximate solution to the joint QoS downlink multicast beamforming problem \mathcal{Q} . The first step of the proposed algorithm consists of a relaxation of the original problem \mathcal{Q} to problem \mathcal{Q}_r . Problem \mathcal{Q} may or may not be feasible; if it is, then so is problem \mathcal{Q}_r . If \mathcal{Q}_r is infeasible, then so is \mathcal{Q} . The converse is generally not true; i.e., if \mathcal{Q}_r is feasible, \mathcal{Q} need not be feasible. In order to establish feasibility of \mathcal{Q} in this case, the randomization/scaling loop should yield at least one feasible solution. This is most often the case, as will be verified in the sequel. If the randomization/scaling loop fails to return at least one feasible solution, then the (in)feasibility of \mathcal{Q} cannot be determined. There is, therefore, a relatively small proportion of problem instances for which (in)feasibility of \mathcal{Q} cannot be decided using the proposed approach. It is evident from the above discussion that feasibility is a key aspect of problem \mathcal{Q} and its proposed solution via problem \mathcal{Q}_r and the randomization/scaling loop. Feasibility depends on a number of factors; namely, the number of transmit antenna elements, the number and the populations of the multicast groups, the channel characteristics, the noise variances, and finally the received SINR constraints.

Beyond feasibility, there are two key issues of interest. The first has to do with cases for which the solution to the relaxed problem \mathcal{Q}_r yields an exact optimum of the original problem

TABLE I
MC SIMULATION RESULTS (RAYLEIGH); 4 Tx ANTENNAS

M	G	SINR	feas.	opt.	feas.	all solutions		appr. solutions	
			\mathcal{Q}_r	\mathcal{Q}_r	appr.	mean	std	mean	std
6	3	6	90	100	-	1	0	-	-
6	3	8	70	100	-	1	0	-	-
6	3	10	45	100	-	1	0	-	-
6	3	12	27	100	-	1	0	-	-
6	3	14	14	100	-	1	0	-	-
6	3	16	7	100	-	1	0	-	-
8	2	6	98	80	98	1.06	0.17	1.29	0.30
8	2	8	91	84	99	1.08	0.38	1.54	0.83
8	2	10	73	83	98	1.19	1.81	2.27	4.54
8	2	12	52	86	99	1.20	2.12	2.55	5.84
8	2	14	32	89	100	1.01	0.06	1.11	0.15
8	2	16	16	90	96	1.04	0.19	1.67	0.44
8	2	18	9	93	100	1.02	0.07	1.22	0.19
8	2	20	3	89	100	1.05	0.16	1.49	0
12	2	6	42	49	79	1.69	1.89	2.82	2.73
12	2	8	10	81	94	1.19	0.51	2.39	0.47
12	2	10	1	100	-	1	0	-	-

\mathcal{Q} . This happens when the solution matrices $\{\mathbf{X}_k^{\text{opt}}\}_{k=1}^G$ turn out all being rank one. Then, the associated principal components solve optimally the original problem \mathcal{Q} , i.e., in such a case \mathcal{Q}_r is not a relaxation after all. It is interesting to find the frequency of occurrence of such an event, whose benefit is twofold: the problem is solved optimally and at a smaller complexity, since the randomization step and the repeated solution of the ensuing MMPC problem $\mathcal{S}^{\mathcal{Q}}$ is avoided. The second issue of interest is the quality of the final approximate solution to problem \mathcal{Q} . A practical figure of merit is the power ratio discussed in Section IV.


We first consider the standard i.i.d. Rayleigh fading model, i.e., the elements of each channel vector are i.i.d. circularly symmetric zero-mean complex Gaussian random variables of variance 1. The results presented in this subsection are obtained by averaging over 300 different channel snapshots, using 300 Gaussian randomization samples in each Monte Carlo run. Tables I and II summarize these results, for number N of transmit antenna elements set to 4 and 8, respectively. The proposed two-step algorithm is tested for a variety of choices for the total number M of single-antenna receivers and the number G of multicast groups, which index the rows in the tables (columns 1 and 2, respectively). The users are considered to be evenly distributed among the multicast groups, i.e., $\{G_k = M/G\}_{k=1}^G$. For each such configuration, the same SINR targets are requested for all users (in the 6–20 dB range, see column 3). The noise variance is set to $\sigma^2 = 1$ for all channels.

The percentage of the 300 Monte Carlo runs for which \mathcal{Q}_r is feasible is shown in column 4. Column 5 reports the percentage

TABLE II
MC SIMULATION RESULTS (RAYLEIGH); 8 Tx ANTENNAS

M	G	SINR	feas.	opt.	feas.	all solutions		appr. solutions	
			\mathcal{Q}_r	\mathcal{Q}_r	appr.	mean	std	mean	std
12	2	6	100	37	95	1.18	0.25	1.30	0.27
12	2	8	100	36	96	1.17	0.24	1.28	0.25
12	2	10	100	35	95	1.17	0.23	1.27	0.24
12	2	12	100	41	96	1.15	0.21	1.26	0.22
12	2	14	100	43	95	1.15	0.22	1.27	0.23
12	2	16	100	45	94	1.13	0.20	1.25	0.21
12	2	18	100	48	96	1.12	0.23	1.25	0.28
12	2	20	100	53	95	1.10	0.18	1.23	0.21
12	3	6	100	79	98	1.04	0.11	1.19	0.17
12	3	8	100	79	98	1.04	0.11	1.19	0.18
12	3	10	99	81	99	1.05	0.14	1.25	0.24
12	3	12	95	85	98	1.04	0.15	1.31	0.29
12	3	14	79	88	99	1.06	0.29	1.52	0.74
12	3	16	52	93	99	1.02	0.11	1.38	0.26
12	3	18	31	94	99	1.03	0.14	1.53	0.37
12	3	20	18	98	100	1.01	0.04	1.29	0
12	4	6	100	93	100	1.01	0.03	1.11	0.08
12	4	8	87	98	100	1.00	0.04	1.24	0.17
12	4	10	42	98	100	1.01	0.06	1.32	0.33
12	4	12	12	97	100	1.01	0.06	1.36	0
12	4	14	3	100	-	1	0	-	-
16	2	6	100	10	93	1.88	1.63	1.99	1.69
16	2	8	100	12	91	2.00	2.27	2.14	2.40
16	2	10	100	15	87	1.88	1.32	2.06	1.38
16	2	12	99	23	88	1.70	1.57	1.94	1.76
16	2	14	95	32	88	1.80	2.30	2.24	2.78
16	2	16	73	46	92	1.71	3.86	2.42	5.40
16	2	18	54	59	93	1.33	1.04	1.91	1.56
16	2	20	33	65	94	1.27	0.82	1.86	1.31
24	2	6	99	0	44	6.79	8.74	6.84	8.76
24	2	8	61	4	30	4.87	6.23	5.53	6.52
24	2	10	12	14	34	3.64	5.20	5.53	6.30

of feasible solutions to problem \mathcal{Q}_r , for which the solution matrices turn out all being (essentially) rank one; defined by the second largest eigenvalue being smaller than 10^{-3} of the sum of all eigenvalues. Column 6 reports the percentage of problem instances for which, once a feasible solution to problem \mathcal{Q}_r is found, the proposed randomization/scaling loop yields at least one feasible solution to the original problem \mathcal{Q} . Columns 7 and 8 hold the mean and the standard deviation of the ratio of the total transmitted power corresponding to the final approximate solution over the lower bound obtained from the SDR solution. This ratio equals 1 when the relaxation is tight, and the reported statistics depend on the frequency (see column 5) of this event.


Fig. 2. Q_r feasibility; 300 MC runs, 4 Tx antennas.

In order to obtain additional insight on the quality of the approximation step, conditional statistics are also reported in columns 9 and 10 after excluding exact optimum solutions from the calculation.

The Q_r feasibility percentages, stored in column 4 of Tables I and II, are also plotted in Figs. 2 and 3, respectively. In all configurations considered, the higher the target SINR, the less likely it is that Q_r is feasible, which is intuitive. Furthermore, Q_r is getting more difficult to solve as the number G of multicast groups increases and/or as more users are added in each group, since in either case interference is higher. Finally, it is seen that, as expected, increasing the number of transmit antenna elements improves service: higher received SINR can be attained by more users in more multicast groups.

The Q_r optimality percentages are also plotted in Figs. 4 and 5, for the case of 4 and 8 transmit antennas, respectively. The most interesting observation is that the optimality percentage increases as the number of users per multicast group decreases; percentages are significant especially when the number of users per group is smaller or equal to the number of transmit antennas. This can be seen in two ways: either by holding the number of groups fixed while decreasing their populations, or by fixing the total number of users and distributing them in more multicast groups. Trying to interpret this fact, note that in both cases the problem is pushed towards the multiuser (independent information) downlink problem, where each user forms a multicast group. The latter is known to be convex, and the associated SDP relaxation has been shown to be tight [6]. In addition, the Q_r optimality percentage also increases with target SINR. It seems as if rank-one solutions are more likely when operating close to the infeasibility boundary.

Regarding the approximation step of the proposed algorithm, we can distinguish two cases. In most of the scenarios considered, the number of users per multicast group was kept smaller or equal to the number of transmit antenna elements, so that a realistic value of the received SINR could be guaranteed, for a significant fraction of the different channel instances. There,

Fig. 3. Q_r feasibility; 300 MC runs, 8 Tx antennas.

the randomization/scaling loop yields a feasible solution with a probability higher than 90% in most cases where Q_r is feasible, as shown in Fig. 6 which illustrates the contents of column 6 of Table II. The approximate solution entails transmission power that is under two times (3 dB from) the possibly unattainable lower bound, on average. The actual numbers for each configuration depend on the number of the Gaussian randomization samples; 300 have proved adequate for most configurations. However, when a relatively low target SINR is to be guaranteed to a number of users per group larger than the number of antennas, the feasibility of the approximation decreases and the power penalty increases. This can be appreciated by looking at the lowest sub-matrices of Tables I and II. Using 1000, instead of 300, Gaussian random samples for these configurations, we have observed a small improvement in the quality of the approximation.

We have repeated the Monte Carlo simulations, under the same setup, in order to validate the performance of the MMF algorithm presented in Section V. The results are very similar to the ones presented so far for the QoS case, and we therefore skip them for brevity. The sole difference is that feasibility is not an issue in the MMF case. Specifically, there is a considerable percentage of problem instances for which the proposed relaxation is tight, so that the optimum solution is found. For all other instances, the proposed algorithm finds a high-quality approximate solution at manageable complexity cost. An interesting observation is that the quality of the approximation for the multi-group case is consistently better than the respective single-group case [3] and that it becomes better as a given number of users is distributed among a larger number of multicast groups; again, moving closer to the multiuser downlink problem.

B. Experiments With Measured Channel Data

The performance of the proposed multicast beamforming algorithms was also tested on measured channel data courtesy of iCORE HCDC Lab, University of Alberta in Edmonton, Canada. Measurements were carried out using a portable 4 ×




Fig. 4. Q_r optimality; 300 MC runs, 4 Tx antennas.




Fig. 5. Q_r optimality; 300 MC runs, 8 Tx antennas.

4 multiple-input multiple-output (MIMO) testbed that operates in the 902–928 MHz (ISM) band. The transmitter (Tx) and the receiver (Rx) were equipped with antenna arrays, each comprising four vertically polarized dipole antennas spaced $\lambda/2$ (≈ 16 cm) apart. The chip rate used for sounding was low enough to safely assume that the channel is not frequency selective. More details on the testbed configuration and the procedure used to estimate the channel gains of the MIMO channel matrix can be found in [19]. Datasets and a detailed description of many measurement campaigns in typical propagation environments are available at the iCORE HCDC Lab website (<http://www.ece.ualberta.ca/~mimo/>). The most pertinent scenario for our purposes is the stationary outdoor one, called Quad and illustrated in Fig. 7. Quad is a 150 by 60 meters lawn surrounded by buildings with heights from approximately 15 to 30 meters. The Tx location was fixed, whereas the Rx was placed in 6 different locations (no measurements are actually




Fig. 6. Approximation feasibility; 300 MC runs, 300 randomizations, 8 Tx antennas.

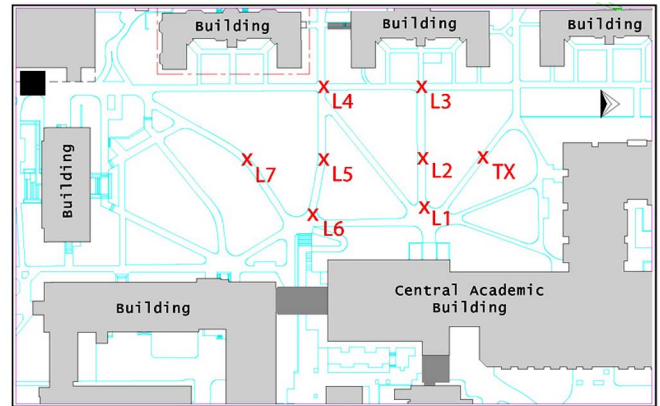


Fig. 7. Sample wireless channel measurement scenario from <http://www.ece.ualberta.ca/~mimo/>.

provided for location 4) as indicated in Fig. 7. For every Rx location, 9 different measurements were taken by shifting the Rx antenna array on a 3×3 square grid with $\lambda/4$ spacing. Each measurement contains about 100 4×4 channel snapshots, recorded 3 per second; thus for each location there are about 900 MIMO channel gain matrices available. We form multicast groups by considering each receive antenna at each location as a separate terminal, and grouping terminals in 1–3 locations. The results reported in Tables III–V for the QoS problem formulation were obtained by averaging over the 900 channel instances. We have tried tens of different configurations and we are only presenting representative results for each scenario considered. All channel gains are normalized before use by the same constant (average amplitude over all channels and all snapshots), in order to facilitate the comparison with the simulated Rayleigh case. Note that this normalization maintains differences in path loss. 300 Gaussian samples are employed in the randomization/scaling loop.

The main findings regarding the performance of our algorithms applied to the measured channel data can be summarized as follows:

TABLE III
MEASURED CHANNEL RESULTS; 2 MULTICAST GROUPS OF 4 USERS EACH

SINR	feas.	opt.	feas.	all solutions		appr. solutions	
	\mathcal{Q}_r	\mathcal{Q}_r	appr.	mean	std	mean	std
Group 1 (4 at L5) & Group 2 (4 at L7)							
6	98	82	98	1.05	0.19	1.30	0.39
8	91	84	88	1.05	0.23	1.33	0.53
10	82	87	98	1.04	0.21	1.40	0.53
12	50	92	99	1.04	0.29	1.61	0.95
14	19	91	98	1.02	0.08	1.23	0.19
16	9	94	100	1.01	0.06	1.16	0.22
18	3	92	96	1.00	0.00	1.01	0

Group 1 (2 at L2 & 2 at L6) & Group 2 (2 at L5 & 2 at L7)

6	100	81	99	1.05	0.18	1.29	0.34
8	100	81	99	1.05	0.18	1.30	0.33
10	96.1	86	99	1.04	0.17	1.32	0.36
12	83	90	99	1.04	0.26	1.43	0.80
14	57	93	99	1.15	2.89	3.73	12.10
16	31	93	99	1.02	0.12	1.36	0.33
18	14	95	99	1.01	0.07	1.31	0.21
20	6	92	100	1.03	0.18	1.44	0.53

Group 1 (1 at L1, 1 at L3 & 2 at L6)

Group 2 (1 at L2, 1 at L5 & 2 at L7)

6	100	72	98	1.12	0.54	1.45	0.98
8	99	75	98	1.09	0.31	1.39	0.55
10	93	80	97	1.18	2.71	2.03	6.44
12	73	87	97	1.05	0.25	1.44	0.63
14	44	89	98	1.07	0.71	1.78	2.19
16	23	93	99	1.03	0.18	1.52	0.59

TABLE IV
MEASURED CHANNEL RESULTS; 2 MULTICAST GROUPS OF 6 USERS EACH

SINR	feas.	opt.	feas.	all solutions		appr. solutions	
	\mathcal{Q}_r	\mathcal{Q}_r	appr.	mean	std	mean	std
Group 1 (3 at L1 & 3 at L3) & Group 2 (3 at L2 & 3 at L6)							
6	100	73	98	1.19	1.50	1.76	2.91
8	90	68	94	1.39	2.45	2.38	4.49
10	61	66	92	1.33	1.06	2.15	1.73
12	18	72	92	1.33	1.15	2.56	2.12
Group 1 (2 at L1, 2 at L2 & 2 at L6)							
Group 2 (2 at L3, 2 at L5 & 2 at L7)							
6	70	24	82	2.25	3.51	2.78	4.06
8	33	41	81	2.02	3.58	3.08	4.90
10	7	48	65	1.18	0.45	1.67	0.69

- For two multicast groups and number of users per group equal to the number of Tx antennas ($N = 4$), the relaxation

TABLE V
MEASURED CHANNEL RESULTS; 3 MULTICAST GROUPS OF 3 OR 4 COLOCATED USERS EACH

SINR	feas.	opt.	feas.	all solutions		appr. solutions	
	\mathcal{Q}_r	\mathcal{Q}_r	appr.	mean	std	mean	std
Group 1 (3 at L1), Group 2 (3 at L2) & Group 3 (3 at L3)							
6	72	98	100	1.01	0.10	1.36	0.66
8	37	99	100	1.00	0.01	1.10	0.07
10	14	100	-	1	0	-	-
Group 1 (4 at L1), Group 2 (4 at L2) & Group 3 (4 at L3)							
6	29	95	99	1.02	0.11	1.36	0.40
8	8	100	-	1	0	-	-

$\mathcal{Q} \rightarrow \mathcal{Q}_r$ is tight very frequently (70%–100%) and the power penalty paid by the approximation step very small. These hold irrespective of the distribution of each group's users in one, two, or even three locations (see Table III).

- For two multicast groups of six users each, evenly distributed in 2 locations, the relaxation $\mathcal{Q} \rightarrow \mathcal{Q}_r$ is tight for more than half of the occasions (see Table IV). There exist channel instances for which SINR up to 14 dB can be guaranteed; such high SINR values are unattainable under the corresponding i.i.d. Rayleigh fading scenario. The quality of approximation is good, even though the number of users per group is larger than the number of transmit antenna elements. When the six users of each group are evenly distributed in three locations, the problem is feasible only up to about 10 dB and the feasibility of the approximation step can drop <80%.
- For three multicast groups (see Table V) of three colocated users each, the relaxation $\mathcal{Q} \rightarrow \mathcal{Q}_r$ is almost always tight (>90%) and feasible up to 10 dB of prescribed SINR. For four users per group, it becomes infeasible for SINR values larger than about 8 dB.

VII. CONCLUSION

The downlink beamforming problem was considered for the general case of multiple cochannel multicast groups, under two design criteria: QoS, in which we seek to minimize the total transmitted power while guaranteeing a prescribed minimum SINR at all receivers; and a fair objective, in which we seek to maximize the minimum received SINR under a total power constraint. Both formulations contain single-group multicast beamforming as a special case, and are therefore NP-hard. Computationally efficient quasi-optimal solutions were proposed by means of SDR and a combined randomization/scaling loop. Extensive numerical results have been presented, using both simulated (i.i.d. Rayleigh) and measured outdoor wireless channel data, showing that the proposed algorithms yield high-quality approximate solutions at a moderate complexity cost. Interestingly, our numerical findings indicate that the solutions generated by our algorithms are often exactly optimal, especially in the case of measured channels. In certain cases, this optimality can be proven beforehand, and alternative convex reformulations of lower complexity can be constructed;

see [8] and [9] for further details. In other cases, a theoretical worst-case bound on approximation accuracy can be derived, and shown to be tight; on this issue, see [15].

REFERENCES

- [1] E. Karipidis, N. D. Sidiropoulos, and Z.-Q. Luo, "Transmit beamforming to multiple co-channel multicast groups," in *Proc. 1st IEEE Int. Workshop on Computat. Adv. Multi-Sensor Adaptive Processing (CAMSAP)*, Puerto Vallarta, Mexico, Dec. 13–15, 2005, pp. 109–112.
- [2] M. J. Lopez, "Multiplexing, scheduling, and multicasting strategies for antenna arrays in wireless networks," Ph.D. dissertation, Dept. Elect. Eng. and Comp. Sci., Mass. Inst. Technol. (MIT), Cambridge, 2002.
- [3] N. D. Sidiropoulos, T. N. Davidson, and Z.-Q. Luo, "Transmit beamforming for physical layer multicasting," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pp. 2239–2251, Jun. 2006.
- [4] N. D. Sidiropoulos and T. N. Davidson, "Broadcasting with channel state information," in *Proc. IEEE Sensor Array and Multichannel (SAM) Signal Process. Workshop*, Sitges, Barcelona, Spain, Jul. 18–21, 2004, vol. 1, pp. 489–493.
- [5] F.-R. Farrokhi, K. J. R. Liu, and L. Tassioulas, "Downlink power control and base station assignment," *IEEE Commun. Lett.*, vol. 1, no. 4, pp. 102–104, Jul. 1997.
- [6] M. Bengtsson and B. Ottersten, "Optimal and suboptimal transmit beamforming," in *Handbook of Antennas in Wireless Communications*, L. C. Godara, Ed. Boca Raton, FL: CRC, 2001, ch. 18.
- [7] M. R. Garey and D. S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-Completeness*. San Francisco, CA: Freeman, 1979.
- [8] E. Karipidis, N. D. Sidiropoulos, and Z.-Q. Luo, "Convex transmit beamforming for downlink multicasting to multiple co-channel groups," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Toulouse, France, May 14–19, 2006, pp. 973–976.
- [9] E. Karipidis, N. D. Sidiropoulos, and Z.-Q. Luo, "Far-field multicast beamforming for uniform linear antenna arrays," *IEEE Trans. Signal Process.*, vol. 55, no. 10, pp. 4916–4927, Oct. 2007.
- [10] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optim. Methods Softw.*, vol. 11–12, pp. 625–653, 1999.
- [11] Y. Ye, *Interior Point Algorithms: Theory and Analysis*. New York: Wiley, 1997.
- [12] Y. Gao and M. Schubert, "Power-allocation for multi-group multicasting with beamforming," in *Proc. IEEE/ITG Workshop on Smart Antennas (WSA)*, Ulm, Germany, Mar. 13–14, 2006.
- [13] S. Zhang and Y. Huang, "Complex quadratic optimization and semidefinite programming," *SIAM J. Optim.*, vol. 16, no. 3, pp. 871–890, Jan. 2006.
- [14] E. Feron, "Nonconvex quadratic programming, semidefinite relaxations and randomization algorithms in information and decision systems," in *System Theory: Modeling Analysis and Control*, T. E. Djaferis and I. Schick, Eds. Norwell, MA: Kluwer, 1999, pp. 255–274.
- [15] Z.-Q. Luo, N. D. Sidiropoulos, P. Tseng, and S. Zhang, "Approximation bounds for quadratic optimization with homogeneous quadratic constraints," *SIAM J. Optim.*, vol. 18, no. 1, pp. 1–28, Feb. 2007.
- [16] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [17] H. Wolkowicz, "Relaxations of Q2P," in *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*, H. Wolkowicz, R. Saigal, and L. Vandenberghe, Eds. Norwell, MA: Kluwer, 2000, ch. 13.4.
- [18] A. Wiesel, Y. C. Eldar, and S. Shamai, "Linear precoding via conic optimization for fixed MIMO receivers," *IEEE Trans. Signal Process.*, vol. 54, no. 1, pp. 161–176, Jan. 2006.
- [19] P. Goud, Jr., R. Hang, D. Truhachev, and C. Schlegel, "A portable MIMO testbed and selected channel measurements," *EURASIP J. Appl. Signal Process.*, vol. 2006, 2006, 10.1155/ASP/2006/51490, Article ID 51490, 10 pp..



Eleftherios Karipidis (S'05) received the Diploma degree in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2001 and the M.Sc. degree in communications engineering from the Technical University of Munich, Germany, in 2003.

He is currently working towards the Ph.D. degree with the Telecommunications Division of the Department of Electronic and Computer Engineering, Technical University of Crete, Chania, Greece. He was an intern from February 2002 to October 2002 with

Siemens ICM, and from December 2002 to November 2003 with the Wireless Solutions Lab of DoCoMo Euro-Labs, both in Munich. His broad research interests are in the area of signal processing for communications, with current emphasis on MIMO VDSL systems, convex optimization, and applications in transmit precoding for wireline and wireless systems.

Mr. Karipidis is a member of the Technical Chamber of Greece.



Nicholas Sidiropoulos (SM'99) received the Diploma degree from the Aristotle University of Thessaloniki, Greece, and the M.S. and Ph.D. degrees from the University of Maryland at College Park (UMCP), in 1988, 1990, and 1992, respectively, all in electrical engineering.

From 1988 to 1992, he was a Fulbright Fellow and a Research Assistant with the Institute for Systems Research (ISR), University of Maryland. From September 1992 to June 1994, he served his military service as a Lecturer in the Hellenic Air

Force Academy. From October 1993 to June 1994, he also was a member of the technical staff, Systems Integration Division, G-Systems Ltd., Athens, Greece. From 1994 to 1995, he was a Postdoctoral Fellow and from 1996 to 1997, he was a Research Scientist with ISR-UMCP, an Assistant Professor with the Department of Electrical Engineering, University of Virginia-Charlottesville from 1997 to 1999, and an Associate Professor with the Department of Electrical and Computer Engineering, University of Minnesota-Minneapolis from 2000 to 2002. Currently, he is a Professor in the Telecommunications Division of the Department of Electronic and Computer Engineering, Technical University of Crete, Chania-Crete, Greece, and Adjunct Professor with the University of Minnesota, Minneapolis. He is an active consultant for industry in the areas of frequency hopping systems and signal processing for xDSL modems. His current research interests are primarily in signal processing for communications, and multiway analysis.

Prof. Sidiropoulos is currently Chair of the Signal Processing for Communications Technical Committee (SPCOM-TC) of the IEEE Signal Processing (SP) Society (2007–2008), where he has served as Member (2000–2005) and Vice-Chair (2005–2006). He is also a member of the Sensor Array and Multichannel processing Technical Committee (SAM-TC) of the IEEE SP Society (2004–2009). He has served as Associate Editor for IEEE TRANSACTIONS ON SIGNAL PROCESSING from 2000 to 2006 and the IEEE SIGNAL PROCESSING LETTERS from 2000 to 2002. He received the NSF/CAREER award (Signal Processing Systems Program) in June 1998, and an IEEE Signal Processing Society Best Paper Award in 2001.



Zhi-Quan Luo (F'07) received the B.Sc. degree in mathematics from Peking University, Peking, China, in 1984 and the Ph.D. degree in operations research from the Operations Research Center and the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, in 1989.

During the academic year of 1984 to 1985, he was with the Nankai Institute of Mathematics, Tianjin, China. In 1989, he joined the Department of Electrical and Computer Engineering, McMaster

University, Hamilton, ON, Canada, where he became a Professor in 1998 and held the Canada Research Chair in Information Processing since 2001. Since April 2003, he has been a Professor with the Department of Electrical and Computer Engineering and holds an endowed ADC Research Chair in Wireless Telecommunications with the Digital Technology Center, University of Minnesota, Minneapolis. His research interests lie in the union of large-scale optimization, information theory and coding, data communications, and signal processing.

Prof. Luo received an IEEE Signal Processing Society Best Paper Award in 2004. He is a member of the Society for Industrial and Applied Mathematics (SIAM) and Mathematical Programming Society (MPS). He is also a member of the Signal Processing for Communications (SPCOM) and Signal Processing Theory and Methods (SPTM) Technical Committees of the IEEE Signal Processing Society. From 2000 to 2004, he served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING and *Mathematics of Computation*. He is presently serving as an Associate Editor for several international journals including *SIAM Journal on Optimization* and *Mathematics of Operations Research*.