

# Efficient Batch and Adaptive Approximation Algorithms for Joint Multicast Beamforming and Admission Control

Evaggelia Matskani, Nicholas D. Sidiropoulos, *Fellow, IEEE*, Zhi-Quan Luo, *Fellow, IEEE*, and Leandros Tassioulas, *Fellow, IEEE*

**Abstract**—Wireless multicasting is becoming increasingly important for efficient distribution of streaming media and location-aware services to mobile and hand-held devices, network management, and software updates over cellular (UMTS-LTE) and indoor/outdoor wireless networks (e.g., 802.11/16). Multicast beamforming was recently proposed as a means of exploiting the broadcast nature of the wireless medium to boost spectral efficiency and meet Quality of Service (QoS) requirements. Infeasibility is a key issue in this context, due to power or mutual interference limitations. We therefore consider the joint multicast beamforming and admission control problem for one or more co-channel multicast groups, with the objective of maximizing the number of subscribers served and minimizing the power required to serve them. The problem is NP-hard even for an isolated multicast group and no admission control; but drawing upon our earlier work for the multiuser SDMA downlink, we develop an efficient approximation algorithm that yields good solutions at affordable worst-case complexity. For the special case of an isolated multicast, Lozano proposed a particularly simple adaptive algorithm for implementation in UMTS-LTE. We identify strengths and drawbacks of Lozano's algorithm, and propose two simple but worthwhile improvements. All algorithms are carefully tested on publicly available indoor/outdoor measured channel data.

**Index Terms**—Admission control, beamforming, convex approximation, E-MBMS, multicasting, NP-hard, UMTS-LTE.

Manuscript received June 27, 2008; accepted April 20, 2009. First published June 30, 2009; current version published November 18, 2009. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Subhrakanti Dey. Conference versions of parts of this work appear in the *Proceedings of the IEEE CAMSAP*, St. Thomas, U.S. Virgin Islands, December 12–14, 2007, and the *Proceedings of the IEEE ICASSP*, Las Vegas, NV, April 15–20, 2008. The work of E. Matskani and N. D. Sidiropoulos is supported in part by ARL/ERO contract N62558-06-0340, and EC/FP6 project COOPCOM and WIP. The work of Z.-Q. Luo is supported in part by the U.S. NSF Grants DMS-0312416 and DMS-0610037. The work of L. Tassioulas is supported in part by the ARO under Grant W911NF-04-1-0306, and EC/FP6 project WIP.

E. Matskani and N. D. Sidiropoulos are with the Department of Electronic and Computer Engineering, Technical University of Crete, 73100 Chania, Crete, Greece (e-mail: 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@umn.edu).

L. Tassioulas is with the Department of Computer Engineering and Telecommunications, University of Thessaly, 38221 Volos, Greece (e-mail: leandros@uth.gr).

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.2009.2026621

## I. INTRODUCTION

WIRELESS multicasting is gaining ground as an enabling technology for mass content distribution (Internet TV, streaming media, pay-per-view, network management, software updates) over wireless networks. Multicasting lies in between two widely used information dissemination modalities: broadcasting, where common information is addressed to all nodes in a network, and parallel (orthogonal or co-channel) unicast transmissions. The middle ground between the two is important for existing and emerging applications.

Multicasting has been traditionally viewed as a network-layer issue, addressed by multicast routing. This viewpoint is natural for wired networks, but wireless is different due to the so-called *broadcast advantage* (a node's transmission may reach multiple receivers) and its flip-side: co-channel interference.

Unlike traditional broadcast radio, wireless networks nowadays incorporate feedback mechanisms that provide varying grades of channel state information at the transmitter (CSI-T). The availability of CSI-T coupled with the proliferation of antenna arrays open the door for multicast beamforming at the transmitter. The idea is to shape the transmit beampattern in a way that steers power in the directions of multicast subscribers while minimizing leakage in other directions. It is also possible to design multiple beampatterns to transmit simultaneously to more than one multicast groups, provided that the different groups are spatially well-separated. Multicast beamforming under Quality of Service (QoS) constraints, measured via the received signal-to-noise ratio (SNR), was treated in [17], where it was shown that the problem is NP-hard, yet amenable to convex approximation tools. The case of multiple co-channel multicast groups was treated in [6]. A key problem that arises in this context is potential infeasibility, due to power or inter-group crosstalk limitations. If the QoS constraints are rigid, one is led to consider the problem of joint transmit beamforming and admission control. This has been considered in [13] for the multiuser space-division multiple access (SDMA) downlink scenario, comprising interfering co-channel unicast transmissions.

In this paper, we consider the joint co-channel multicast beamforming and admission control problem from the viewpoint of maximizing the number of subscribers served at or above a prescribed signal-to-interference-plus-noise ratio (SINR), under an overall power constraint. We pay special attention to the case of a single multicast group, which is

important in view of ongoing standardization activity [9], [16]. The Evolved Multimedia Broadcast/Multicast Service (E-MBMS) in the context of 3GPP<sup>1</sup>/UMTS-LTE<sup>2</sup> specifically provisions point-to-multipoint physical layer multicasting [16]. Motivated by [17], Lozano [9] proposed a particularly simple alternating gradient iteration for the case of a single multicast group. This being an NP-hard problem, the results reported in [9] are intriguing given the simplicity of Lozano’s algorithm.

Our specific contributions in this paper can be summarized as follows.

- We generalize the convex approximation framework developed in [13] for the multiuser SDMA downlink to the co-channel multicast context. This yields an efficient approximation algorithm that is directly applicable in a far broader range of problems: from single-group multicast to the SDMA downlink and everything in between.
- We take a closer look at Lozano’s alternating gradient iteration, and identify strengths and weaknesses. We show, by means of simple but instructive examples, that it is sensitive with respect to initialization and can exhibit limit cycle behavior.
- We propose two simple improvements of Lozano’s iteration that mitigate these drawbacks and significantly boost overall performance. The resulting algorithm is an ideal candidate for practical implementation in next-generation cellular systems.
- We present a comprehensive suite of carefully designed numerical experiments using publicly available measured channel data, for both indoor and outdoor scenarios.

## II. PROBLEM FORMULATION

Consider a base station or access point equipped with  $N$  transmit antennas and a population of  $K$  subscribers, each with a single receive antenna. Let  $\mathbf{h}_i$  be the  $N \times 1$  complex channel vector from the transmit antenna array to receiver  $i$ ,  $i \in \{1, \dots, K\}$ . We begin by assuming instantaneous CSI-T, but our formulation and algorithms can be readily adapted to work with long-term CSI-T, as will be explained in the sequel.

Consider  $1 \leq G \leq K$  multicast groups,  $\{\mathcal{G}_1, \dots, \mathcal{G}_G\}$ , where  $\mathcal{G}_m$  contains the indexes of receivers that wish to subscribe to multicast  $m$ . Since the transmissions are co-channel, we may assume without loss of generality that  $\mathcal{G}_m \cap \mathcal{G}_l = \emptyset$ ,  $l \neq m$ ,  $\cup_m \mathcal{G}_m = \{1, \dots, K\}$ , and with  $G_m := |\mathcal{G}_m|$ ,  $\sum_{m=1}^G G_m = K$ . Let  $\mathbf{w}_m^H$  be the weight vector applied to the  $N$  transmitting elements to beamform towards group  $m$ ,  $m \in \{1, \dots, G\}$ , where  $(\cdot)^H$  denotes Hermitian transpose. Streaming media applications demand a minimum instantaneous<sup>3</sup> SINR. With this

in mind, the multicast beamformer design problem can be cast as follows:

$$\begin{aligned} \min_{\{\mathbf{w}_m \in \mathbb{C}^N\}_{m=1}^G} & \sum_{m=1}^G \|\mathbf{w}_m\|_2^2 & (1) \\ \text{subject to:} & \sum_{m=1}^G \|\mathbf{w}_m\|_2^2 \leq P, & (2) \\ & \frac{|\mathbf{w}_m^H \mathbf{h}_i|^2}{\sum_{\ell \neq m} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i, & \\ & \forall i \in \mathcal{G}_m, \forall m \in \{1, \dots, G\} & (3) \end{aligned}$$

where  $\sigma_i^2$  is the additive noise power at receiver  $i$  and  $c_i$  stands for the associated minimum SINR requirement. The objective function reflects the desire to pick the minimum power solution when the problem is feasible, while the explicit sum power constraint accounts for regulatory and equipment limitations.<sup>4</sup>

The above problem is NP-hard, for it contains the case of a single multicast group ( $G = 1$ ), which is already NP-hard as shown in [17]. Still, it has been shown [6], [17] that it is possible to compute high-quality approximate solutions via *convex (semidefinite) approximation*. The idea is to approximate the non-convex and NP-hard problem using a suitable convex problem, then use the solution of the convex problem to guide the search for a good feasible solution of the original NP-hard problem. The second step in [6] and [17] is based on *Gaussian randomization*, which can provide provably good approximate solutions in the sense that the distance to the optimum can be analytically bounded [10]. The case of  $G = K$  corresponds to the multiuser space-division multiple-access (SDMA) downlink [4] whose convexity was shown in [1].

(In)Feasibility is a key issue with the above formulation. Infeasibility may arise for a number of reasons: due to proximity of channel vectors of users interested in different multicast streams, scattering of users belonging to a given multicast group, degrees of freedom (few transmit antennas relative to the number of multicast groups), spatial group interleaving, and power limitations. The power constraint alone may limit coverage even for a single multicast group. While in most cases infeasibility can be detected using the tools developed in [6], what one does next is far less obvious. If the SINR constraints are infeasible, then some form of admission control is needed; but this should ideally be considered together with the beamformer design problem, for the two are obviously coupled.

Towards this end, it makes sense to consider maximizing the number of subscribers that can be served at their desired SINR, and then minimizing the power required to serve those selected in the first step. This approach can be mathematically formulated in two stages:

$$S_o = \operatorname{argmax}_{S \subseteq \{1, \dots, K\}, \{\mathbf{w}_m \in \mathbb{C}^N\}_{m=1}^G} |S| \quad (4)$$

$$\text{subject to: } \sum_{m \mid \mathcal{G}_m \cap S \neq \emptyset} \|\mathbf{w}_m\|_2^2 \leq P, \quad (5)$$

$$\frac{|\mathbf{w}_m^H \mathbf{h}_i|^2}{\sum_{\ell \mid \mathcal{G}_\ell \cap S \neq \emptyset, \ell \neq m} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i, \forall i \in \mathcal{G}_m \cap S, \forall m \quad (6)$$

<sup>4</sup>Per-antenna power amplifier constraints can be easily incorporated in our approach; we skip them for brevity.

<sup>1</sup>Third Generation Partnership Project.

<sup>2</sup>Universal Mobile Telecommunications System—Long Term Evolution.

<sup>3</sup>Lower-priority services such as software updates can operate under a long-term average SINR constraint, which can be guaranteed given only long-term CSI-T, as will be discussed in the sequel.

where  $|S|$  denotes the cardinality of  $S$ . Note that if  $\mathcal{G}_m \cap S = \emptyset$ , then no constraints are imposed for the given  $m$ . Given  $S_o$ , we then wish to

$$\min_{\{\mathbf{w}_m \in \mathbb{C}^N\}_m} \sum_{m \mid \mathcal{G}_m \cap S_o \neq \emptyset} \|\mathbf{w}_m\|_2^2 \quad (7)$$

$$\text{s.t. } \frac{|\mathbf{w}_m^H \mathbf{h}_i|^2}{\sum_{\ell \mid \mathcal{G}_\ell \cap S_o \neq \emptyset, \ell \neq m} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i, \forall i \in \mathcal{G}_m \cap S_o, \forall m. \quad (8)$$

The second optimization stage in (7)–(8) is feasible (provided  $S_o$  is a solution of (4)–(6)), but remains NP-hard, as per [6], [17]. This is different from the multiuser SDMA downlink case, considered in [13]. In addition, we have the following result.

*Claim 1:* The problem in (4)–(6) is NP-hard.

*Proof:* Consider the decidability version of problem (4)–(6) for  $G = 1$ ,  $S = \{1, \dots, K\}$  (i.e.,  $|S| = K$ , the maximal value), and  $c_i = c$ ,  $\forall i$ : does there exist a vector  $\mathbf{w}$  such that

$$\|\mathbf{w}\|_2^2 \leq P$$

and for which

$$\frac{|\mathbf{w}^H \mathbf{h}_i|^2}{\sigma_i^2} \geq c, \forall i \in S?.$$

This problem is the decidability version of the following problem:

$$\max \min_{i \in S} \frac{|\mathbf{w}^H \mathbf{h}_i|^2}{\sigma_i^2}$$

subject to:  $\|\mathbf{w}\|_2^2 \leq P$

which has been shown to be NP-hard in [17]. ■

The two-stage formulation in (4)–(6) and (7)–(8) is not particularly convenient, for a number of reasons. Nested optimization problems are awkward to work with; perhaps more importantly, the nested formulation does not suggest a way to approach the problem from a convex approximation perspective. Towards this end, we will use a technique originally developed in [13] for the multiuser SDMA downlink—which is a special case of our present formulation for  $G = K$ .

Introduce binary admission control (“slack”) variables  $s_i$ , one for each receiver  $i \in \{1, \dots, K\}$ . When  $s_i = -1$  ( $s_i = +1$ ) receiver  $i$  is scheduled (rejected, respectively). Consider the following optimization problem:

$$\min_{\{\mathbf{w}_m \in \mathbb{C}^N, \{s_i \in \{-1, +1\}\}_{i \in \mathcal{G}_m}\}_{m=1}^G} \mathcal{J} \left( \{\mathbf{w}_m, \{s_i\}_{i \in \mathcal{G}_m}\}_{m=1}^G \right)$$

$$:= \epsilon \sum_{m=1}^G \|\mathbf{w}_m\|_2^2 + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (s_i + 1)^2 \quad (9)$$

$$\text{subject to: } \sum_{m=1}^G \|\mathbf{w}_m\|_2^2 \leq P \quad (10)$$

$$\frac{|\mathbf{w}_m^H \mathbf{h}_i|^2 + \delta^{-1}(s_i + 1)^2}{\sum_{\ell \neq m} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i, \forall i \in \mathcal{G}_m, \forall m \in \{1, \dots, G\}. \quad (11)$$

*Claim 2:* With

$$\delta \leq \min_{m \in \{1, \dots, G\}} \min_{i \in \mathcal{G}_m} \frac{4c_i^{-1}}{P \max_{n \in \{1, \dots, K\}} \|\mathbf{h}_n\|_2^2 + \sigma_i^2}$$

and  $\epsilon < 1/(P/4 + 1)$ , the problem in (9)–(11) is always feasible, and solution of (9)–(11) is *equivalent* to first solving (4)–(6) and then solving (7)–(8). If there are multiple solutions of (4)–(6), i.e., if the maximal subset of users that can be served is not unique, then (9)–(11) will automatically pick a maximal subset requiring minimal total power.

Proof is deferred to the Appendix.

By virtue of Claim 2, solution of (9)–(11) could be used to obtain a solution of (4)–(6), which is NP-hard per Claim 1. It follows that:

*Claim 3:* The problem in (9)–(11) is NP-hard.

This means that if we insist on polynomial complexity in the worst case, we have to give up optimality (unless  $P = NP$ ); that is, we can only hope for an approximate solution, rather than an optimal one. A key benefit of the single-stage formulation in (9)–(11) is that it is naturally amenable to convex approximation, as explained next.

### III. A SEMIDEFINITE RELAXATION APPROACH

Define  $\mathbf{H}_i := \mathbf{h}_i \mathbf{h}_i^H$ , and introduce rank-one positive semidefinite matrix variables  $\mathbf{W}_m := \mathbf{w}_m \mathbf{w}_m^H$ , and  $\mathbf{S}_i := \mathbf{s}_i \mathbf{s}_i^T$ , where  $\mathbf{s}_i := [s_i \ 1]^T$ . Then (e.g., cf. [11] and [13]) the optimization problem in (9)–(11) can be transformed to the following *equivalent* form:

$$\min_{\{\mathbf{W}_m \in \mathbb{C}^{N \times N}, \{\mathbf{S}_i \in \mathbb{R}^{2 \times 2}\}_{i \in \mathcal{G}_m}\}_{m=1}^G} \phi \left( \{\mathbf{W}_m, \{\mathbf{S}_i\}_{i \in \mathcal{G}_m}\}_{m=1}^G \right)$$

$$:= \epsilon \sum_{m=1}^G \text{Tr}(\mathbf{W}_m) + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} \text{Tr}(\mathbf{1}_{2 \times 2} \mathbf{S}_i) \quad (12)$$

$$\text{subject to: } \sum_{m=1}^G \text{Tr}(\mathbf{W}_m) \leq P, \quad (13)$$

$$\frac{\text{Tr}(\mathbf{H}_i \mathbf{W}_m) + \delta^{-1} \text{Tr}(\mathbf{1}_{2 \times 2} \mathbf{S}_i)}{\sum_{\ell \neq m} \text{Tr}(\mathbf{H}_i \mathbf{W}_\ell) + \sigma_i^2} \geq c_i, \forall i \in \mathcal{G}_m, \forall m, \quad (14)$$

$$\mathbf{W}_m \geq 0, \text{rank}(\mathbf{W}_m) = 1, \forall m, \quad (15)$$

$$\mathbf{S}_i \geq 0, \text{rank}(\mathbf{S}_i) = 1, \mathbf{S}_i(1, 1) = \mathbf{S}_i(2, 2) = 1, \forall i \in \mathcal{G}_m, \forall m. \quad (16)$$

The above is a quadratically constrained quadratic programming (QCQP) problem, in which only the rank-one constraints are non-convex. Wolkowicz [19] has shown that simply dropping the rank-one constraints in this case yields the Lagrange bi-dual problem, which is the strongest convex relaxation of the QCQP problem in the Lagrangian class. Dropping the rank-one constraints yields a semidefinite programming (SDP) problem, which can be efficiently solved using modern interior point methods [2]:

$$\min_{\{\mathbf{W}_m \in \mathbb{C}^{N \times N}, \{\mathbf{S}_i \in \mathbb{R}^{2 \times 2}\}_{i \in \mathcal{G}_m}\}_{m=1}^G} \phi \left( \{\mathbf{W}_m, \{\mathbf{S}_i\}_{i \in \mathcal{G}_m}\}_{m=1}^G \right)$$

$$:= \epsilon \sum_{m=1}^G \text{Tr}(\mathbf{W}_m) + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} \text{Tr}(\mathbf{1}_{2 \times 2} \mathbf{S}_i) \quad (17)$$

$$\text{subject to: } \sum_{m=1}^G \text{Tr}(\mathbf{W}_m) \leq P, \quad (18)$$

$$\frac{\text{Tr}(\mathbf{H}_i \mathbf{W}_m) + \delta^{-1} \text{Tr}(\mathbf{1}_{2 \times 2} \mathbf{S}_i)}{\sum_{\ell \neq m} \text{Tr}(\mathbf{H}_i \mathbf{W}_\ell) + \sigma_i^2} \geq c_i, \quad \forall i \in \mathcal{G}_m, \forall m, \quad (19)$$

$$\mathbf{W}_m \geq 0, \quad \forall m, \quad (20)$$

$$\mathbf{S}_i \geq 0, \mathbf{S}_i(1,1) = \mathbf{S}_i(2,2) = 1, \quad \forall i \in \mathcal{G}_m, \forall m. \quad (21)$$

*Remark 1:* Being a relaxation of (12)–(16), the problem in (17)–(21) is also feasible when  $\epsilon$  and  $\delta$  are chosen as per Claim 2.

Solution of the relaxed problem in (17)–(21) in general only provides a lower bound on the cost of an optimal solution to the original problem in (12)–(16). If a particular solution of (17)–(21) consists of rank-one  $\mathbf{W}_m$ ,  $\forall m$  and rank-one  $\mathbf{S}_i$ ,  $\forall i$ , then it is also a solution (12)–(16); although this situation does happen in practice, it does not always happen. What is needed is a way to “convert” an optimal solution of (17)–(21) into a good feasible solution of (12)–(16). This step is presently an art guided by partial results. A commonly used technique is *Gaussian randomization*, which can generate provably good approximate solutions in certain cases (e.g., see [10] which is applicable in the special case of an isolated multicast scenario, albeit not accounting for admission control). In other cases reasonable heuristics are often used, and the overall solution is benchmarked against the relaxation lower bound<sup>5</sup> and/or exhaustive search when applicable—in particular for small problem instances.

In our present context, the following algorithm seems to work best in practice, among numerous options we have tried:

---

**Algorithm 1: Multicast Membership Deflation by Relaxation (MDR)**

---

- 1)  $\mathcal{U} \leftarrow \{1, \dots, K\}$ .
- 2)  $\mathcal{G}_m \leftarrow \mathcal{G}_m \cap \mathcal{U}, \forall m$ .
- 3) Solve (41)–(45), and let  $\check{\mathbf{W}}_m$  denote the resulting transmit covariance matrices.
- 4) For each  $m$  such that  $\mathcal{G}_m \neq \emptyset$ , extract the principal component of  $\check{\mathbf{W}}_m$ , and scale it to power  $\text{Tr}(\check{\mathbf{W}}_m)$ ; i.e., set  $\check{\mathbf{w}}_m := \sqrt{\text{Tr}(\check{\mathbf{W}}_m)} \check{\mathbf{u}}_m$ , where  $\check{\mathbf{u}}_m$  is the unit-norm principal component of  $\check{\mathbf{W}}_m$ .
- 5) For all  $m$  such that  $\mathcal{G}_m \neq \emptyset$ , and each  $i \in \mathcal{G}_m$ , check whether

$$\frac{|\check{\mathbf{w}}_m^H \mathbf{h}_i|^2}{\sum_{\ell \in \mathcal{G}_m, \ell \neq m} |\check{\mathbf{w}}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i$$

holds. If this is true for all  $i \in \mathcal{G}_m$ , and all  $m$  such that  $\mathcal{G}_m \neq \emptyset$ , stop (a feasible solution has been found); else pick the user with largest gap to its target SINR (smallest attained SINR if all the SINR targets are equal), remove from  $\mathcal{U}$  and go to Step 2).

---

<sup>5</sup>We know that an optimal solution of the original NP-hard problem can never have lower cost than that attained by the possibly higher-rank solution of the corresponding rank relaxation.

### A. Implementation Complexity

The worst-case complexity of MDR is  $O(K^{4.5} \log(1/\epsilon))$ , where  $\epsilon$  is the required relative accuracy of the duality gap at termination [13], [18].

Note that it is possible to further reformulate (17)–(21) to reduce complexity by a factor of two, as shown in the Appendix.

## IV. SPECIAL CASE: SINGLE MULTICAST

The special case of a single (isolated) multicast group ( $G = 1$ ) is of particular interest, due to on-going standardization activity in the context of UMTS-LTE. This has been considered in [17], but without regard to infeasibility/admission control issues. In the case of a single multicast, assuming that  $\mathbf{h}_i \neq \mathbf{0}, \forall i$ , infeasibility arises only due to the transmit power constraint—there is no interference from other groups. An alternative scenario where the same problem arises is when the beamformers of the different groups are sequentially optimized, and interference is clumped together with the additive noise terms. Letting  $\mathcal{U}$  denote the set of potential subscribers, the two-stage formulation in (4)–(8) reduces to the following problem:

$$S_o = \text{argmax}_{S \subseteq \mathcal{U}, \mathbf{w} \in \mathbb{C}^N} |S| \quad (22)$$

$$\text{subject to: } \|\mathbf{w}\|_2^2 \leq P \quad (23)$$

$$\frac{|\mathbf{w}^H \mathbf{h}_i|^2}{\sigma_i^2} \geq c_i, \quad \forall i \in S \quad (24)$$

and given  $S_o$

$$\min_{\mathbf{w} \in \mathbb{C}^N} \|\mathbf{w}\|_2^2 \quad (25)$$

$$\text{subject to: } \frac{|\mathbf{w}^H \mathbf{h}_i|^2}{\sigma_i^2} \geq c_i, \quad \forall i \in S_o. \quad (26)$$

The single-stage reformulation in (9)–(11) reduces to

$$\min_{\mathbf{w} \in \mathbb{C}^N, \{s_i \in \{-1, +1\}\}_{i \in \mathcal{U}}} \epsilon \|\mathbf{w}\|_2^2 + (1 - \epsilon) \sum_{i \in \mathcal{U}} (s_i + 1)^2 \quad (27)$$

$$\text{subject to: } \|\mathbf{w}\|_2^2 \leq P \quad (28)$$

$$\frac{|\mathbf{w}^H \mathbf{h}_i|^2 + \delta^{-1} (s_i + 1)^2}{\sigma_i^2} \geq c_i, \quad \forall i \in \mathcal{U} \quad (29)$$

where  $\delta \leq \min_{i \in \mathcal{U}} 4c_i^{-1}/\sigma_i^2$ ,  $\epsilon < 1/(P/4 + 1)$ , as per Claim 2. The final optimization problem after “squaring” of variables and rank relaxation is the following SDP:

$$\min_{\mathbf{W} \in \mathbb{C}^{N \times N}, \{\mathbf{S}_i \in \mathbb{R}^{2 \times 2}\}_{i \in \mathcal{U}}} \epsilon \text{Tr}(\mathbf{W}) + (1 - \epsilon) \sum_{i \in \mathcal{U}} \text{Tr}(\mathbf{1}_{2 \times 2} \mathbf{S}_i) \quad (30)$$

$$\text{subject to: } \text{Tr}(\mathbf{W}) \leq P, \quad (31)$$

$$\frac{\text{Tr}(\mathbf{H}_i \mathbf{W}) + \delta^{-1} \text{Tr}(\mathbf{1}_{2 \times 2} \mathbf{S}_i)}{\sigma_i^2} \geq c_i, \quad \forall i \in \mathcal{U} \quad (32)$$

$$\mathbf{W} \geq 0 \quad (33)$$

$$\mathbf{S}_i \geq 0, \mathbf{S}_i(1,1) = \mathbf{S}_i(2,2) = 1, \quad \forall i \in \mathcal{U}. \quad (34)$$

The overall convex approximation approach for  $G = 1$  entails a trimmed-down version of MDR, listed below for clarity.

---

**Algorithm 2: Single Group MDR**


---

- 1)  $\mathcal{U} \leftarrow \{1, \dots, K\}$ .
  - 2) Solve the relaxed problem (30)–(34), and let  $\tilde{\mathbf{W}}$  denote the resulting transmit covariance matrix.
  - 3)  $\tilde{\mathbf{w}}$  = principal component of  $\tilde{\mathbf{W}}$ , scaled to power  $Tr(\tilde{\mathbf{W}})$ .
  - 4) For each  $i \in \mathcal{U}$ , check whether  $|\tilde{\mathbf{w}}^H \mathbf{h}_i|^2 / \sigma_i^2 \geq c_i$ . If true  $\forall i \in \mathcal{U}$ , stop (feasible solution has been found); else pick user with largest gap to its target SNR, remove from  $\mathcal{U}$  and go to Step 2).
- 

Extensive experiments with measured channels (cf. Section VII) indicate that MDR provides very satisfactory performance (about one user less than optimum, on average, in our experiments) at a moderate average and worst-case complexity that grow gracefully with the problem size. Still, MDR takes order of second to execute on a typical PC, and, perhaps more important, it is a batch algorithm that solves the problem from scratch every time. Even though warm-start options can be readily envisioned and could be used to track small variations inexpensively, a practitioner would naturally prefer a solution that works well at very low complexity—an adaptive filtering-type algorithm, preferably.

Is it possible to have good performance at really low complexity for an NP-hard problem? This seems *a priori* highly unlikely. Yet Lozano recently proposed an intriguing adaptive filtering algorithm that apparently works well in limited but realistic experiments in the context of UMTS-LTE [9]. In the following section, we take a closer look at Lozano’s algorithm.

#### V. LOZANO’S ALGORITHM

In the sequel, let  $\mathbf{H}_i$  denote either  $\mathbf{h}_i \mathbf{h}_i^H$  or its expectation, depending on whether instantaneous CSI-T or long-term CSI-T is assumed. When instantaneous CSI-T is available,  $\mathbf{w}^H \mathbf{H}_i \mathbf{w} / \sigma_i^2$  is the instantaneous SNR at receiver  $i$  for weight vector  $\mathbf{w}$ ; with long-term CSI-T,  $\mathbf{w}^H \mathbf{H}_i \mathbf{w} / \sigma_i^2$  is the expected SNR at receiver  $i$  for weight vector  $\mathbf{w}$ . Lozano’s algorithm is a very simple alternating gradient iteration: at each step it sorts users according to presently attained SNRs, discards the users with the poorest SNRs, and makes a gradient step in the direction of the *weakest retained* user. The choice of users to drop is based on either a fixed SNR threshold, or (better) a fixed number/percentage of users to keep at each iteration. Dropped users participate as candidates in the next iteration. Lozano’s algorithm can be summarized as follows.

---

**Algorithm 3: Lozano’s Iteration for Single-Group Multicast [9]**


---

- 1) Initialize:  $\mathbf{w} = [1 \ 0 \ \dots \ 0]^T$ .
  - 2) Compute  $\text{SNR}_i(t-1) = (P/\sigma_i^2) \mathbf{w}_{t-1}^H \mathbf{H}_i \mathbf{w}_{t-1}$ ,  $\forall i \in \mathcal{U}$ .
  - 3) Sort  $\text{SNR}_i(t-1)$ ,  $\forall i \in \mathcal{U}$ .
  - 4) Drop a fixed proportion of users with lowest attained SNRs.
  - 5) Find weakest link among remaining ones ( $\rightarrow k$ ).
  - 6) Take a step in its direction:  $\mathbf{w}_t = \mathbf{w}_{t-1} + \mu \mathbf{H}_k \mathbf{w}_{t-1}$ ; then  $\mathbf{w}_t = \mathbf{w}_t / \|\mathbf{w}_t\|_2$ .
  - 7) Repeat until no significant change in minimum SNR.
- 

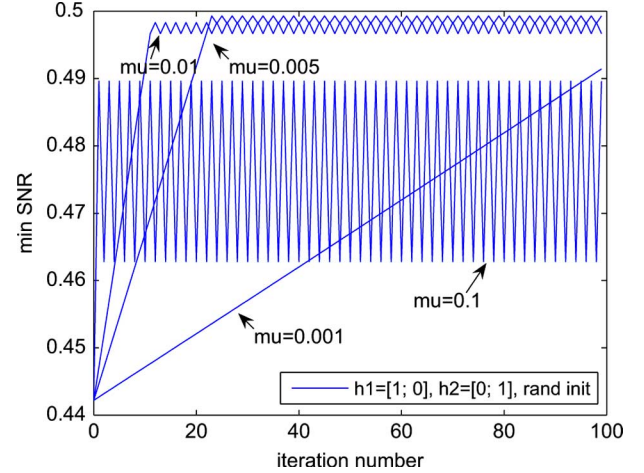


Fig. 1. Illustration of the effect of  $\mu$  on Lozano’s algorithm for a contrived but instructive two-user scenario.

The ingenuity of this algorithm lies precisely on giving up on the weakest links—to focus on the remaining, more promising ones. Performance is far worse when rejection is not employed. The choice of number/percentage of users to keep at each iteration has a significant impact on performance. The choice of step-size parameter  $\mu$  is also important. Unfortunately, theoretical analysis of Lozano’s algorithm appears difficult, precisely due to the rejection step; but even empirical choice of parameters is difficult, as illustrated next.

#### A. A Closer Look

Despite its conceptual simplicity, Lozano’s algorithm exhibits intricate convergence behavior. Consider the following contrived but instructive scenario: there are  $N = 2$  transmit antennas and  $K = 2$  users, with channels  $\mathbf{h}_1 = [1 \ 0]^T$  and  $\mathbf{h}_2 = [0 \ 1]^T$  (each user only listens to a single transmit antenna). Let  $\sigma_1 = \sigma_2 = P = 1$ . If both users should be served, the optimal solution is  $\mathbf{w} = (1/\sqrt{2})[1 \ 1]^T$ , attaining an SNR of  $1/2$  for each user. Lozano’s algorithm initialized with  $\mathbf{w}_0 = \mathbf{h}_1$  (say, because it was previously serving only user 1, and now user 2 comes into the system) has a fixed point at  $\mathbf{h}_1$ , which is in the null space of  $\mathbf{H}_2$ —thus  $\mathbf{w}_t = \mathbf{w}_0$  for all  $t \geq 0$  and all  $\mu$ , implying that user 2 is simply shut off from the system. This shows that the algorithm can converge to a suboptimal and unfair solution.

A small perturbation of either  $\mathbf{h}_2$  or  $\mathbf{w}_0$  takes the algorithm away from this undesirable fixed point; for small enough  $\mu$ , the iterates typically approach the optimum solution, albeit slowly. Beyond the usual speed—misadjustment trade-off, however, in this simple scenario Lozano’s algorithm typically exhibits limit cycle behavior when randomly initialized. Choosing a smaller  $\mu$  helps reduce the magnitude of the oscillation, but naturally reduces the speed of adaptation, as illustrated in Fig. 1. Limit cycles are particularly annoying because they make it hard to select an appropriate tolerance threshold.

Evidently, Lozano’s algorithm may fail to converge, or converge to a suboptimal or unfair solution, and is sensitive with respect to initialization, problem instantiation, and the choice

of parameters. These issues do arise in realistic scenarios, however the algorithm performs considerably better, on average, than what the above example may suggest, and its simplicity is certainly appealing. It therefore makes sense to consider modifying it to mitigate its two major shortcomings (sensitivity to initialization, and the potential for limit cycle behavior) while maintaining its simplicity.

## VI. IMPROVING LOZANO'S ALGORITHM

We propose the following two simple but worthwhile improvements to Lozano's algorithm.

- Lopez [8] considered multicast beamforming from the viewpoint of maximizing *average* SNR. While *minimum* SNR is what determines the common multicast rate [17], the average SNR solution can serve as a reasonable starting point for further improvement via adaptive algorithms. Lopez has shown [8] that maximizing the average SNR in a multicast context reduces to determining the principal eigenvector of the (normalized) channel correlation matrix. In the case of instantaneous CSI-T, this is defined as  $\sum_{k=1}^K \mathbf{h}_k \mathbf{h}_k^H / \sigma_k^2 = \mathbf{H} \mathbf{H}^H$ , where  $\mathbf{H} := [\mathbf{h}_1 / \sigma_1, \dots, \mathbf{h}_K / \sigma_K]$ ; whereas for long-term CSI-T it is  $\sum_{k=1}^K E[\mathbf{h}_k \mathbf{h}_k^H / \sigma_k^2]$ . Finding the principal eigenvector can be accomplished via adaptive algorithms (e.g., based on the power method), which can also be employed in tracking mode. As a result, the overall solution remains simple and adaptive in nature. We call this algorithm LLI, for *Lozano with Lopez Initialization*. The only difference between LLI and Lozano's original algorithm is in the initialization—step 1), where the average SNR beamformer of Lopez is used in LLI.

- A simple and effective way to suppress limit cycle behavior is to damp  $\mu$  according to a predefined back-off schedule. This should be balanced against our primary objective, which is to find a good solution. Aggressively damping  $\mu$  limits how much of the search space we can explore.

The weight update of Lozano's algorithm can be interpreted as taking a step in the direction of the *local* subgradient of  $\min_{i \in \mathcal{A}_t} \text{SNR}_i(t-1)$ , where the minimum is taken over the currently active user set  $\mathcal{A}_t$ . There are two difficulties here: this is not a subgradient in the usual sense, because it is only a local, not a global under-estimator of  $\min_{i \in \mathcal{A}_t} \text{SNR}_i(t-1)$ ; and  $\mathcal{A}_t$  can change as iterations progress. These difficulties arise because we are dealing with a non-convex and NP-hard problem. For convex problems it is known that subgradient optimization using a step-size sequence  $\mu_t$  such that  $\sum_{t=1}^{\infty} \mu_t = \infty$  but  $\sum_{t=1}^{\infty} \mu_t^2 < \infty$  (e.g.,  $\mu_t = 1/t, t > 0$ ) yields an algorithm that converges to the optimum. In our context, the choice of back-off schedule is not obvious. We have tried the following options ( $\mu_0 = 1$  in all cases):

- 1)  $\mu_t = 1/t$ ;
- 2)  $\mu_t = 1/\text{floor}(t/10)$  (thus  $\mu_t$  is reduced every 10 iterations);
- 3)  $\mu_t = \mu_{t-1}/2$  (exponential back-off);
- 4)  $\mu_t = \mu_{t-1}/2$  if  $t \bmod 10 = 0$ , else  $\mu_t = \mu_{t-1}$  (exponential back-off every 10 iterations);

- 5)  $\mu_t = \mu_{t-1}/t$  (even more aggressive);
- 6)  $\mu_t = \mu_{t-1}/(t/10)$  if  $t \bmod 10 = 0$ , else  $\mu_t = \mu_{t-1}$  (same as the previous one but back-off every 10 iterations).

We have tested these options (with Lopez initialization) in *extensive* experiments with simulated and measured channel data (cf. simulations section). Options 1, 2, and 6 performed equally well, whereas 3, 4, and 5 were worse. Between 1, 2, and 6, option 6 was two orders of magnitude faster than the other two. We therefore settled on option 6, in which  $\mu$  is aggressively damped every ten iterations. This is of course *ad-hoc*, but so is the overall algorithm—and it is hard to argue with something simple that works very well in practice, as we will show in our experiments. We call this variant dLLI (for *damped* LLI). dLLI differs from Lozano's original algorithm in that it uses Lopez initialization in Step 1), and step-size back-off option 6 for the weight update in Step 6).

*Remark 2:* Lozano's algorithm and LLI use a fixed step-size, thus having the potential to track changes in the operational environment—e.g., due to users coming in and out of the system and/or user mobility. The use of a vanishing step-size, on the other hand, implies that dLLI *per se* is not capable of tracking. In our particular context, however, channel vector updates will generally be infrequent (relative to the down-link signaling rate), and users will drop or attempt to join at an even slower time scale. In between such updates, we have to solve a static problem. Each static problem can be solved with dLLI, using either the average SNR beamformer (Lopez) or the beamvector computed for the previous “slot” (problem instance) as initialization. Which initialization is best will depend on the type of update (e.g., new user or updated channel vector for existing user, and user mobility). Since dLLI is a cheap algorithm, the pragmatic approach would be to run two parallel iterations with both initializations and choose the best in the end.

## VII. EXPERIMENTS

The problems we are aiming to solve are important in practice, albeit NP-hard. MDR is only a well-motivated approximation, and the adaptive algorithms (Lozano's, LLI, dLLI) are merely common-sense engineering. We have conducted extensive and carefully designed experiments to assess the performance of all algorithms using measured channel data. We further tested all algorithms in simulations with i.i.d. Rayleigh channels. We do not report the i.i.d. Rayleigh results for brevity, but note that these were consistent with those obtained using measured channels.

Approximate solutions should ideally be compared to exact (optimal) ones to assess the quality of approximation. Unfortunately, this is not possible in our context, because even the single-group version of the problem without admission control is NP-hard. Still, for fixed SINR targets, we can enumerate over all possible subsets of users using the potentially higher-rank SDP relaxation in [6] to test each subset. This will be referred as ENUM in the sequel, and it yields an upper bound on the number of users that can be served (and a lower bound on the power required to serve them) under the given SINR and power constraints. This bound is the tightest that can be obtained via

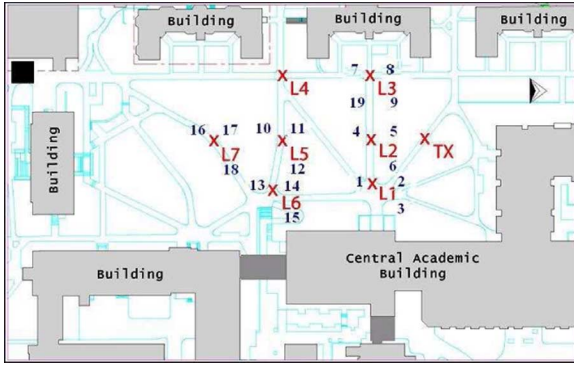


Fig. 2. Quad: Outdoor measurement scenario from <http://www.ece.ualberta.ca/~mimo/>.



Fig. 3. 2nd Floor of ECERF: Indoor office measurement scenario from <http://www.ece.ualberta.ca/~mimo/>.

duality theory, but remains optimistic in general, because it allows for higher-rank transmit covariances (beamforming corresponds to rank-one transmit covariance).

If after testing all subsets ENUM returns a set of transmit covariances which are all rank-one, then this set is an exact (optimal) solution of the original NP-hard problem in (4)–(6), and thus ENUM yields the ultimate benchmark. This is because it is not possible to serve any more users in this case, even using higher-rank transmit covariances—this possibility has already been tested during ENUM. This is very important, because it happens in the vast majority of cases considered in our experiments. Only in rare cases does ENUM return higher-rank solutions, as we will see in the sequel. The drawback of ENUM is of course its exponential complexity in the number of users,  $K$ , which makes it prohibitive for  $K$  over 10–12 on a current PC.

Two different kinds of wireless scenarios are considered for both single multicast and multiple multicast groups: measured outdoor channel data, and measured indoor channel data. Measured channel data were downloaded from the iCORE HCDC Laboratory, University of Alberta, at <http://www.ece.ualberta.ca/~mimo/> (see also [5]). The outdoor scenario (“Quad”) is illustrated in Fig. 2 and described in [12] and [13]. The indoor scenario (“2nd Floor ECERF”) is illustrated in Fig. 3, and is briefly described next. In both Figures, Tx denotes the (four-element) transmit antenna array location, whereas the numbers denote the positions of each user’s single receive antenna. Data selection and preprocessing follows [12], [13].

“2nd Floor ECERF” is a typical office environment. The floor includes many small offices, divided by thin wooden plates with embedded windows. There are many small corridors as well. The whole room is mainly used by staff and professors of the University of Alberta. The transmitter is placed at the reception area, where many people walked into during measurements. Both the transmitter and the receivers are fixed; positions where measurements have been taken are marked in the floor plan. The distance of the transmitter to a concrete wall, which is covered by wooden plates, is less than 0.5 m. The main corridor, in which locations 1, 2, and 6 are marked, has a width of 2.5 m and a height of 4 m to a concrete ceiling. Locations 1 and 2 are about 18 and 34 m, respectively, from the transmitter. Location 6 is halfway between the transmitter and location 2. Locations 3, 4, and 5 have a distance of 23, 19, and 10 m, respectively, to the main corridor. No measurements are available for location 7. Both the transmitter and the receivers are equipped with antenna arrays, each comprising four vertically polarized dipoles, spaced  $\lambda/2$  ( $\approx 16$  cm) apart. As described in [13], at each Rx Location, nine different measurements were taken by shifting the Rx antenna array on a  $3 \times 3$  square grid with  $\lambda/4$  spacing. Each measurement contains about 100  $4 \times 4$  channel snapshots, recorded 3 per second. We used measurements corresponding to the locations that are marked in Fig. 3.

We next present experiments for the case of multiple co-channel multicast groups, in which the simpler adaptive algorithms are not applicable; the case of an isolated multicast follows.

#### A. Multiple Co-Channel Multicast Groups

In experiments 1 and 2, we consider the case of  $G = 3$  co-channel multicast groups. The number of transmit antennas is  $N = 4$ , while the number of single-antenna receivers is  $K = 12$ , in all cases. We use instantaneous channel vectors (rank-one channel covariance matrices). The reported results are averages over 30 temporal channel snapshots, spanning 30 s.

Experiment 1 concerns the “Quad” measured outdoor scenario. Users are split in the following three groups of four users each:  $\mathcal{G}_1 = \{1, 3, 13, 15\}$ ,  $\mathcal{G}_2 = \{4, 6, 10, 16\}$ ,  $\mathcal{G}_3 = \{7, 8, 9, 19\}$ , see Fig. 2. The remaining parameters are as follows:  $P = 10^3$ ;  $\sigma_i^2 = \sigma^2 = 1$ ,  $c_i = c$ ,  $\forall i$ ; for MDR,  $\epsilon = 10^{-9} < 1/(P/4+1)$ , and  $\delta = (2c^{-1})/(P \max_m \|\mathbf{h}_m\|_2^2 + \sigma^2)$ . Performance of MDR is compared to that of ENUM. The detailed results are reported in Table I. For ease of visualization, Fig. 4 is a plot of the average number of users served, as a function of target SINR in dB, for both ENUM and MDR. It is important to note here that ENUM returned only rank-one transmit covariance matrices in 95% of the cases considered. Only in the rest 5% of the cases were higher-rank covariance matrices returned by ENUM, and in all these cases it was possible to serve only one additional user using these higher-rank covariances, and this required significant excess power. In this experiment, MDR serves on average about half a user less than ENUM when the target SINR is high, and about one and a half user less when the target SINR is low.

Experiment 2 concerns the “2nd Floor ECERF” measured indoor scenario. Users are split in the three groups

TABLE I

EXPERIMENT 1: STATIONARY OUTDOOR, THREE MULTICAST GROUPS; MONTE-CARLO RESULTS (30 MEASURED CHANNEL SNAPSHOTS):  $N = 4$  TX ANT.,  $K = 12$  USERS (ALL DEPICTED IN FIG. 2),  $P = 1000$ ;  $\sigma_k^2 = \sigma^2 = 1$ ,  $c_k = c$ ,  $\lambda_k = 1, \forall k$ ;  $\epsilon = 10^{-9} < 1/(P/4 + 1)$ ,  $\delta = (2c^{-1})/(P \max_m \|\mathbf{h}_m\|_2^2 + \sigma^2)$ . ENTRIES FORMATTED AS  $n @ p/30$  MEAN THAT  $n$  USERS (OR GROUPS) ARE SERVED IN  $p$  OUT OF 30 CASES. ALL POWERS REPORTED IN LINEAR SCALE.

QoS target	Alg	# users served	# Groups	Avg Min Tx Power	Max Min Tx.Power	Avg Time
3	ENUM	12 @ 10/30 11 @ 19/30 10 @ 1/30	3 @ 30/30	70	224	2960 s
3	D-SDR	12 @ 2/30 11 @ 6/30 10 @ 13/30 9 @ 4/30 8 @ 5/30	3 @ 29/30 2 @ 1/30	26	88	0.95 s
5	ENUM	11 @ 24/30 10 @ 6/30	3 @ 30/30	185	749	1064 s
5	D-SDR	11 @ 6/30 10 @ 22/30 9 @ 2/30	3 @ 30/30	122	846	0.88 s
10	ENUM	10 @ 6/30 9 @ 24/30	3 @ 30/30	146	774	535 s
10	D-SDR	10 @ 3/30 9 @ 14/30 8 @ 13/30	3 @ 24/30 2 @ 6/30	123	803	0.99 s
15	ENUM	9 @ 16/30 8 @ 14/30	3 @ 16/30 2 @ 14/30	414	950	352 s
15	D-SDR	9 @ 1/30 8 @ 26/30 7 @ 3/30	3 @ 3/30 2 @ 27/30	153	447	0.79 s

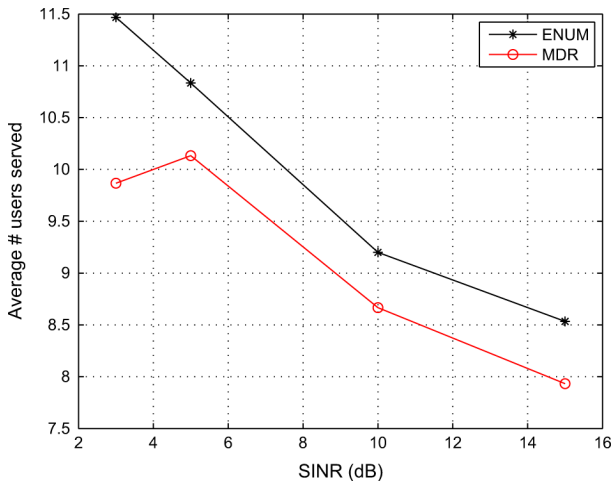


Fig. 4. Experiment 1 (Outdoor): Average number of users served versus target SINR for 30 measured channel snapshots.




Fig. 5. Experiment 2 (Indoor): Average number of users served versus target SINR for 30 measured channel snapshots.

$\mathcal{G}_1 = \{1, 2, 3, 13\}$ ,  $\mathcal{G}_2 = \{5, 14, 6, 16\}$ ,  $\mathcal{G}_3 = \{4, 15, 9, 8\}$ , see Fig. 3. Detailed results are reported in Table II, and summary plots in Fig. 5. The parameters are the same as in Experiment 1. In this indoor scenario, MDR serves on average about one user less than ENUM, which returns higher-rank covariance matrices in 9% of cases.

Summarizing the multigroup multicast experiments, MDR appears to work well in the cases considered, albeit the gap to ENUM is not as small as in the multiuser SDMA downlink case considered in [13]. This is natural, because multicasting is a much harder problem—even its plain-vanilla version is NP-hard. The gap in MDR performance relative to ENUM should be considered in light of the associated complexities: MDR terminates in under 1 second in all cases considered, whereas ENUM takes 5–50 minutes for  $K = 12$ , on a current

PC (see Tables I and II, for more detailed execution time results).

It is also worthwhile to note that ENUM indeed yields the exact solution of the original NP-hard problem in the vast majority of cases considered—which was rather unexpected. It appears that *subset selection diversity* plays a role here: among the many possibilities of choosing a subset of users of fixed cardinality, there is typically one for which the optimal beamforming problem is “easy”—rank relaxation is not a relaxation after all; see also [3], where again a suboptimal solution operates close to the optimal one. In the few cases that ENUM returned a higher-rank solution, this served exactly one additional user, and required significant excess power to do so. The conclusion is that ENUM indeed yields a tight upper bound on the achievable performance.



TABLE II

EXPERIMENT 2: STATIONARY INDOOR, THREE MULTICAST GROUPS; MONTE-CARLO RESULTS (30 MEASURED CHANNEL SNAPSHOTS):  $N = 4$  TX ANT.,  $K = 12$  USERS (ALL DEPICTED IN FIG. 3),  $P = 1000$ ;  $\sigma_k^2 = \sigma^2 = 1$ ,  $c_k = c$ ,  $\lambda_k = 1$ ,  $\forall k$ ;  $\epsilon = 10^{-9} < 1/(P/4 + 1)$ ,  $\delta = (2c^{-1})/(P \max_m \|\mathbf{h}_m\|_2^2 + \sigma^2)$ . ENTRIES FORMATTED AS  $n$  @  $p/30$  MEAN THAT  $n$  USERS (OR GROUPS) ARE SERVED IN  $p$  OUT OF 30 CASES. ALL POWERS REPORTED IN LINEAR SCALE

QoS target	Alg	# users served	# Groups	Avg Min Tx Power	Max Min Tx.Power	Avg Time
3	ENUM	12 @ 12/30 11 @ 18/30	3 @ 30/30	76	432	479 s
3	MDR	12 @ 4/30 11 @ 7/30 10 @ 11/30 9 @ 7/30 8 @ 1/30	3 @ 30/30	51	422	0.40 s
5	ENUM	11 @ 14/30 10 @ 16/30	3 @ 30/30	101	388	308 s
5	MDR	11 @ 1/30 10 @ 15/30 9 @ 11/30 8 @ 3/30	3 @ 30/30	145	997	0.43 s
10	ENUM	10 @ 1/30 9 @ 29/30	3 @ 30/30	194	983	383 s
10	MDR	10 @ 15/30 8 @ 14/30 6 @ 1/30	3 @ 21/30 2 @ 9/30	338	981	0.45 s
15	ENUM	8 @ 30/30	3 @ 29/30 2 @ 1/30	296	692	330 s
1 15	MDR	8 @ 7/30 7 @ 17/30 6 @ 6/30	3 @ 2/30 2 @ 28/30	226	752	0.53 s

### B. Single Multicast

We now turn to single group multicasting ( $G = 1$ ), and compare ENUM, MDR, Lozano's algorithm, LLI, and dLLI.

The three adaptive algorithms (Lozano's, LLI, dLLI) fix coverage (number of users served) and attempt to maximize the minimum SNR among those served under the transmit power constraint. MDR, on the other hand, attempts to maximize coverage subject to received SNR and transmit power constraints, while ENUM yields a (usually tight) upper bound on the number of users that can be served under the same constraints. A meaningful way to compare all algorithms is via the respective minimum SNR—coverage curves, parameterized by transmit power  $P$ . This is analogous to the use of the Receiver Operating Characteristic (ROC) to compare different detectors.

We again consider two different wireless scenarios: measured outdoor, and measured indoor (i.i.d. Rayleigh simulations were also conducted, yielding consistent results, but these are omitted for brevity). In addition to instantaneous CSI-T, we also considered long-term CSI-T. For the latter, we estimated channel correlation matrices by averaging over 30 temporal snapshots; i.e., with  $\mathbf{h}_{i,n}$  denoting the channel from the transmit antenna array to receiver  $i$  at time  $n \in \{1, 2, \dots, 30\}$ , we used  $\hat{\mathbf{H}}_i := (1/30) \sum_{n=1}^{30} \mathbf{h}_{i,n} \mathbf{h}_{i,n}^H$  in place of  $\mathbf{H}_i$  for all algorithms.

In all experiments with a single multicast group, ENUM yielded the optimum rank-one solution of the original NP-hard problem in all cases except for those that correspond to full coverage (i.e., the received signal power requirements are low enough to ensure that everyone can be served—there is no need for admission control). This is interesting in itself, and it also suggests that ENUM is a tight upper bound in all cases where admission control is active.

The parameters used in the experiments were as follows:  $N = 4$ ;  $K = 10$  in experiments 3, 4, and  $K = 12$  for experiments

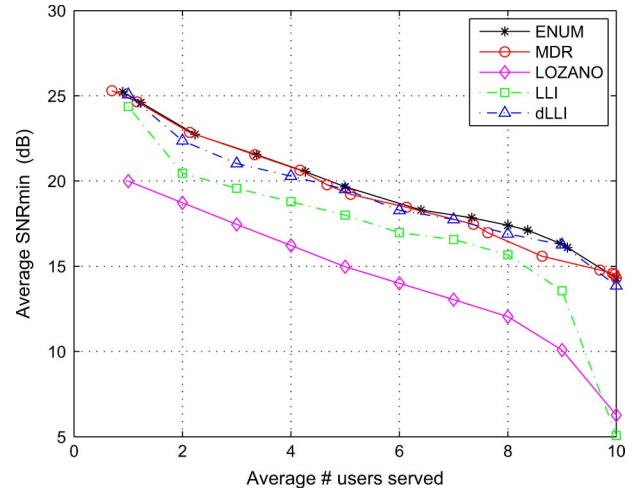


Fig. 6. Experiment 3 (Outdoor, single multicast, instantaneous CSI-T): Average minimum SINR versus average number of users served over 30 measured channel snapshots.

5, 6;  $P = 30$ ;  $c_i = c$  and  $\sigma_i^2 = \sigma^2 = 1$ ,  $\forall i$ ; for MDR,  $\epsilon = 10^{-10}$ ,  $\delta < \min_{i \in \mathcal{U}} 4c^{-1}/\sigma_i^2$ . For Lozano's and LLI algorithm,  $\mu = 10^{-3}$ . For dLLI,  $\mu_0 = 1$  and back-off schedule number 6 is used. For all three adaptive algorithms, convergence is declared when the change in minimum SNR drops below  $10^{-3}$ . These parameters were empirically tuned to optimize the performance of each algorithm.

Experiments 3 and 4 concern the 'Quad' measured outdoor scenario. Measurements corresponding to positions 1, 3, 4, 6, 7, 9, 12, 13, 15, and 17 in Fig. 2 were selected for the respective ten users. Results for instantaneous CSI-T (experiment 3) are summarized in Fig. 6, while those for long-term CSI-T (experiment 4) in Fig. 7.

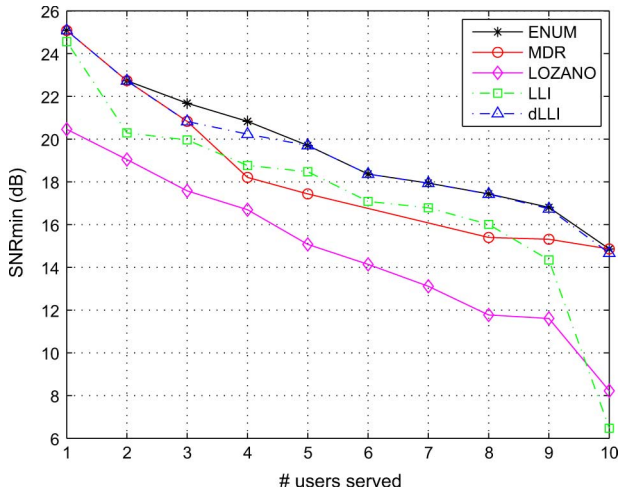


Fig. 7. Experiment 4 (Outdoor, single multicast, long-term CSI-T): Minimum SINR versus number of users served.

In Fig. 6 (instantaneous CSI-T), MDR and dLLI perform very close to the optimum, while Lozano’s algorithm is far behind. Specifically, the average coverage gap of Lozano’s compared to MDR and dLLI is up to five users (50%) for a given average minimum SNR, while the average minimum SNR gap is up to 5 dB for a given average coverage. LLI’s curve falls between Lozano’s and MDR; dLLI significantly improves the performance of LLI. Unlike MDR and dLLI, the other two adaptive algorithms (Lozano’s and LLI) do a poor job for full coverage. MDR and dLLI are on par performance-wise (MDR is somewhat better at lower coverage, dLLI at higher coverage); but dLLI is faster.

The situation is different for long-term CSI-T, as illustrated in Fig. 7. Here dLLI is close to optimal throughout the coverage range *and* clearly outperforms the rest, including MDR.

Indoor measurements (“2nd Floor ECERF”) were used for the next two experiments. Measurements corresponding to positions 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, and 12 in Fig. 3 were selected for the respective 12 users. Instantaneous CSI-T is used in experiment 5, whereas long-term CSI-T is used in experiment 6. As can be seen in Fig. 8, MDR almost coincides performance-wise with ENUM, dLLI performs very close to MDR, while LLI keeps trailing both MDR and dLLI, by a significant margin. This picture changes (again) when long-term CSI-T is considered, in Fig. 9: dLLI clearly outperforms all other algorithms, including MDR; but this time, dLLI is not as close to optimal as in the outdoors case.

Summarizing the insights obtained from single-group experiments, MDR and dLLI emerge as the clear winners (in light of the fact that ENUM is prohibitively complex for realistic values of  $K$ ). Performance-wise, MDR is somewhat better for instantaneous CSI-T (rank-one channel correlations), especially in the higher SNR/lower coverage regime, whereas dLLI is clearly preferable for long-term CSI-T (higher rank channel correlations), and is generally close to MDR even for instantaneous CSI-T. The proposed modifications of Lozano’s algorithm (LLI, dLLI) are simple, yet significantly boost performance. Complexity-wise, Lozano’s algorithm together with




Fig. 8. Experiment 5 (Indoor, single multicast, instantaneous CSI-T): Average minimum SINR versus average number of users served over 30 measured channel snapshots.




Fig. 9. Experiment 6 (Indoor, single multicast, long-term CSI-T): Minimum SINR versus number of users served.

LLI are the fastest ones, requiring typically  $10^{-3}$  to  $10^{-2}$  seconds to terminate, per problem instance. LLI is slightly faster than Lozano due to its better initialization, but their run-times remain in the same order of magnitude. dLLI is somewhat slower, requiring on average  $10^{-2}$  seconds. MDR follows next, requiring from  $10^{-1}$  up to 1.5 s to run. dLLI is more than an order of magnitude faster than MDR in all cases considered. ENUM takes from 2 to 7 min per problem instance.

VIII. CONCLUSION

We have considered the problem of joint multicast beamforming and admission control for cellular and indoor/outdoor wireless networks. The objective is to serve as many potential subscribers as possible at or above their prescribed SINR, and minimize the power required to serve them. The problem is unfortunately NP-hard, but we have shown that it is possible to design approximate solutions of acceptable complexity.

Two distinct approaches have been developed: one extending our earlier work for the multiuser SDMA downlink; the other building upon Lozano’s alternating gradient iteration. The

former (MDR) is a batch algorithm based on convex approximation, and handles multiple co-channel multicast groups. The latter (dLLI) is an adaptive filtering-type algorithm that is restricted to a single multicast group. These algorithms were thoroughly tested in experiments using measured indoor and outdoor wireless channel data. These experiments indicate that MDR and dLLI are generally good and sometimes remarkably good low-complexity approximations for the problem at hand.

For a single multicast group, dLLI has better or comparable performance relative to MDR, in all cases considered. For long-term CSI-T, dLLI is the clear winner. Long-term CSI-T is more realistic in a cellular context, due to mobility and the desire to limit signaling overhead. The simplicity of dLLI is also very appealing. Taken together, these factors swing the verdict in favor of dLLI, at least for cellular applications and a single multicast group (multiple multicasts can be served via frequency- or time-division multiplexing, but this is generally not spectrally efficient). When multiple co-channel multicasts are considered, and/or in fixed wireless applications where instantaneous CSI-T is available, MDR is the method of choice.

#### APPENDIX

##### *Proof of Claim 2:*

*Feasibility:* Using the Cauchy–Schwartz inequality, it is easy to show that  $\mathbf{w}_m = \mathbf{0}$ ,  $\forall m$ ,  $s_i = 1$ ,  $\forall i$  is always feasible for (9)–(11), provided

$$\delta \leq \min_{m \in \{1, \dots, G\}} \min_{i \in \mathcal{G}_m} \frac{4c_i^{-1}}{P \max_{n \in \{1, \dots, K\}} \|\mathbf{h}_n\|_2^2 + \sigma_i^2}. \quad (35)$$

*Optimality:* We next show that under the additional condition

$$\epsilon < \frac{1}{\frac{P}{4} + 1} \quad (36)$$

the single-stage reformulation in (9)–(11) is equivalent to the two-stage problem in (4)–(8). The proof is by contradiction. Let  $\{\check{\mathbf{w}}_m, \{\check{s}_i \in \{-1, +1\}\}_{i \in \mathcal{G}_m}\}_{m=1}^G$  be a solution of (9)–(11), and let  $\{\tilde{\mathbf{w}}, \{\tilde{s}_i \in \{-1, +1\}\}_{i \in \mathcal{G}_m}\}_{m=1}^G$  denote a feasible alternative [that satisfies (10)–(11)] for which

$$\sum_{m=1}^G \sum_{i \in \mathcal{G}_m} 1(\check{s}_i = -1) > \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} 1(\tilde{s}_i = -1)$$

where  $1(\cdot)$  stands for the indicator function. It follows that

$$\sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\check{s}_i + 1)^2 \leq \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\tilde{s}_i + 1)^2 - 4$$

so

$$\begin{aligned} \epsilon \sum_{m=1}^G \|\check{\mathbf{w}}_m\|_2^2 + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\check{s}_i + 1)^2 \\ \leq \epsilon P + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\tilde{s}_i + 1)^2 - (1 - \epsilon)4. \end{aligned}$$

Now,  $\epsilon < 1/(P/4 + 1) \Leftrightarrow \epsilon P - (1 - \epsilon)4 < 0$ , therefore

$$\begin{aligned} \epsilon \sum_{m=1}^G \|\check{\mathbf{w}}_m\|_2^2 + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\check{s}_i + 1)^2 \\ < (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\tilde{s}_i + 1)^2 \\ \leq \epsilon \sum_{m=1}^G \|\check{\mathbf{w}}_m\|_2^2 + (1 - \epsilon) \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (\tilde{s}_i + 1)^2 \end{aligned}$$

which is a contradiction.

This shows that serving more than  $\sum_{m=1}^G \sum_{i \in \mathcal{G}_m} 1(\check{s}_i = -1)$  users is impossible. It remains to show that users whose  $\check{s}_i = -1$  are served by  $\{\check{\mathbf{w}}_m \in \mathbb{C}^N\}_{m=1}^G$  at minimum power.

For this, notice that joint optimality of  $\{\check{\mathbf{w}}_m, \{\check{s}_i \in \{-1, +1\}\}_{i \in \mathcal{G}_m}\}_{m=1}^G$  implies conditional optimality of  $\{\check{\mathbf{w}}_m\}_{m=1}^G$  given  $\{\check{s}_i\}$ —for otherwise we would again have a contradiction. For given  $\{\check{s}_i\}$ , and denoting  $S_o := \cup_i |_{\check{s}_i = -1}$ , note that  $\mathcal{G}_m \cap S_o = \emptyset \Rightarrow \check{\mathbf{w}}_m = \mathbf{0}_{N \times 1}$ . This comes from the cost function and constraints in (9)–(11): if the constraints for a certain subset of users are satisfied via the slack variables, allocating power to these users is wasteful (increases the cost function) and simply adds interference to other users. It follows that, conditioned on  $\{\check{s}_i\}$ , the  $\{\check{\mathbf{w}}_m\}_{m=1}^G$  should minimize  $\epsilon \sum_{m | \mathcal{G}_m \cap S_o \neq \emptyset} \|\check{\mathbf{w}}_m\|_2^2 + \text{constant}$  under  $\sum_{m | \mathcal{G}_m \cap S_o \neq \emptyset} \|\check{\mathbf{w}}_m\|_2^2 \leq P$ , and

$$\frac{|\check{\mathbf{w}}_m^H \mathbf{h}_i|^2}{\sum_{\ell | \mathcal{G}_\ell \cap S_o \neq \emptyset, \ell \neq m} |\check{\mathbf{w}}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i, \quad \forall i \in \mathcal{G}_m \cap S_o, \quad \forall m. \quad (37)$$

Finally, note that if there are multiple solutions of (4)–(6), i.e., if the maximal subset of users that can be served is not unique, then (9)–(11) will automatically pick a maximal subset requiring minimal total power—for otherwise a contradiction would emerge: cf. (9), and note that the second term is only a function of the number of users served. This completes the proof. ■

*Further Simplifications:* The reformulated problem in (9)–(11) can be further simplified as follows:

$$\begin{aligned} \min_{\{\mathbf{w}_m \in \mathbb{C}^N, \{s_i \in \{-1, +1\}\}_{i \in \mathcal{G}_m}\}_{m=1}^G} \mathcal{J}(\{\mathbf{w}_m, \{s_i\}_{i \in \mathcal{G}_m}\}_{m=1}^G) \\ := \epsilon \sum_{m=1}^G \|\mathbf{w}_m\|_2^2 + (1 - \epsilon)2 \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (s_i + 1) \end{aligned} \quad (38)$$

$$\text{subject to: } \sum_{m=1}^G \|\mathbf{w}_m\|_2^2 \leq P, \quad (39)$$

$$\frac{|\mathbf{w}_m^H \mathbf{h}_i|^2 + \delta^{-1}2(s_i + 1)}{\sum_{\ell \neq m} |\mathbf{w}_\ell^H \mathbf{h}_i|^2 + \sigma_i^2} \geq c_i, \quad \forall i \in \mathcal{G}_m, \quad \forall m, \ell \in \{1, \dots, G\} \quad (40)$$

since for  $s_i \in \{-1, +1\}$  it holds that  $(s_i + 1)^2 = s_i^2 + 2s_i + 1 = 2(s_i + 1)$ . This way, the quadratic penalty term in the objective and the constraints can be equivalently replaced by the linear

term  $2(s_i + 1)$ . In a similar vein, the semidefinite relaxation in (17)–(21) can be equivalently rewritten as

$$\begin{aligned} & \min_{\{\mathbf{W}_m \in \mathbb{C}^{N \times N}, \{s_i \in \mathbb{R}\}_{i \in \mathcal{G}_m}\}_{m=1}^G} \phi \left( \{\mathbf{W}_m, \{s_i\}_{i \in \mathcal{G}_m}\}_{m=1}^G \right) \\ & = \epsilon \sum_{m=1}^G \text{Tr}(\mathbf{W}_m) + (1 - \epsilon) 2 \sum_{m=1}^G \sum_{i \in \mathcal{G}_m} (s_i + 1) \end{aligned} \quad (41)$$

$$\text{subject to: } \sum_{m=1}^G \text{Tr}(\mathbf{W}_m) \leq P, \quad (42)$$

$$\frac{\text{Tr}(\mathbf{H}_i \mathbf{W}_m) + \delta^{-1} 2(s_i + 1)}{\sum_{\ell \neq m} \text{Tr}(\mathbf{H}_i \mathbf{W}_\ell) + \sigma_i^2} \geq c_i, \quad \forall i \in \mathcal{G}_m, \forall m, \ell \in \{1, \dots, G\} \quad (43)$$

$$\mathbf{W}_m \geq 0, \quad \forall m, \quad (44)$$

$$-1 \leq s_i \leq 1, \quad \forall i \in \mathcal{G}_m, \forall m. \quad (45)$$

which avoids the introduction of  $2 \times 2$  positive semidefinite matrix variables  $\mathbf{S}_i$ , using instead scalar variables. This helps speed up computations, by roughly a factor of two.

REFERENCES

[1] 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 Press, 2001, ch. 18.

[2] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004 [Online]. Available: <http://www.stanford.edu/~boyd/cvxbook.html>

[3] G. Dimic and N. D. Sidiropoulos, "On downlink beamforming with greedy user selection: Performance analysis and a simple new algorithm," *IEEE Trans. Signal Process.*, vol. 53, no. 10, pp. 3857–3868, Oct. 2005.

[4] F.-R. Farrokhi, K. R. Liu, and L. Tassiulas, "Transmit beamforming and power control for cellular wireless systems," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1437–1450, Oct. 1998.

[5] 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, p. 10, 2006.

[6] E. Karipidis, N. D. Sidiropoulos, and Z.-Q. Luo, "Quality of service and max-min-fair transmit beamforming to multiple co-channel multicast groups," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1268–1279, Mar. 2008.

[7] 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.

[8] M. J. Lopez, "Multiplexing, scheduling, and multicasting strategies for antenna arrays in wireless networks," Ph.D. dissertation, Elect. Eng. and Comp. Sci. Dept., MIT, Cambridge, MA, 2002.

[9] A. Lozano, "Long-term transmit beamforming for wireless multicasting," presented at the Int. Conf. Acoustics, Speech, Signal Processing (ICASSP), Honolulu, HI, Apr. 15–20, 2007.

[10] Z.-Q. Luo, N. D. Sidiropoulos, P. Tseng, and S. Zhang, "Approximation bounds for quadratic optimization with homogeneous quadratic constraints," *SIAM J. Optimiz.*, vol. 18, no. 1, pp. 1–28, 2007.

[11] W.-K. Ma, T. N. Davidson, K. M. Wong, Z.-Q. Luo, and P.-C. Ching, "Quasi-ML multiuser detection using semi-definite relaxation with application to synchronous cdma," *IEEE Trans. Signal Process.*, vol. 50, no. 4, pp. 912–922, Apr. 2002.

[12] E. Matskani, N. D. Sidiropoulos, Z.-Q. Luo, and L. Tassiulas, "A second-order cone deflation approach to joint multiuser downlink beamforming and admission control," presented at the Proc. SPAWC, Helsinki, Finland, Jun. 17–20, 2007.

[13] E. Matskani, N. D. Sidiropoulos, Z.-Q. Luo, and L. Tassiulas, "Convex approximation techniques for joint multiuser downlink beamforming and admission control," *IEEE Trans. Wireless Commun.*, vol. 7, no. 7, pp. 2682–2693, Jul. 2008.

[14] E. Matskani, N. D. Sidiropoulos, Z.-Q. Luo, and L. Tassiulas, "Joint multicast beamforming and admission control," presented at the IEEE CAMSAP, St. Thomas, U.S. Virgin Islands, Dec. 12–14, 2007.

[15] E. Matskani, N. D. Sidiropoulos, and L. Tassiulas, "A comparison of multicast beamforming algorithms for UMTS-LTE," presented at the IEEE Int. Conf. Acoustics, Speech, Signal Processing (ICASSP), Las Vegas, NV, Apr. 12–14, 2008.

[16] "Long term evolution (LTE): A technical overview," Technical White Paper, Motorola Inc., Schaumburg, IL [Online]. Available: <http://business.motorola.com/experiencelte/pdf/LTE%20Technical%20Overview.pdf>

[17] N. D. Sidiropoulos, T. N. Davidson, and Z.-Q. Luo, "Transmit beamforming for physical layer multicasting," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pt. 1, pp. 2239–2251, Jun. 2006.

[18] Y. Ye, *Interior Point Algorithms: Theory and Analysis*. New York: Wiley, 1997.

[19] H. Wolkowicz, "Relaxations of Q2P," in *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*, H. Wolkowicz, R. Saigal, and L. Venberghe, Eds. Boston, MA: Kluwer, 2000, ch. 13.4.



**Evaggelia Matskani** received the Diploma degree in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2005, and the M.S. degree in electronic and computer engineering from the Technical University of Crete, Greece, in 2007, where she is currently working towards the Ph.D. degree.


Her research interests are in signal processing for communications, with emphasis on convex optimization and cross-layer design of wireless networks.



**Nicholas D. Sidiropoulos** (F'09) received the Diploma degree from the Aristotelian University of Thessaloniki, Greece, in 1988, and the M.Sc. and Ph.D. degrees from the University of Maryland at College Park (UMCP) in 1990 and 1992, respectively, all in electrical engineering.

He has been a Postdoctoral Fellow (1994–1995) and Research Scientist (1996–1997) at the Institute for Systems Research, UMCP, and has held positions as Assistant Professor, Department of Electrical Engineering, University of Virginia-Charlottesville (1997–1999), and Associate Professor, Department of Electrical and Computer Engineering, University of Minnesota—Minneapolis (2000–2002). Since 2002, he has been a Professor in the Department of Electronic and Computer Engineering at the Technical University of Crete, Chania-Crete, Greece, and Adjunct Professor at the University of Minnesota. His current research interests are primarily in signal processing for communications, convex approximation of NP-hard problems, cross-layer resource allocation for wireless networks, and multiway analysis.

Prof. Sidiropoulos received the U.S. NSF CAREER Award in 1998, the IEEE Signal Processing Society Best Paper Award in 2001 and in 2007, and is currently a Distinguished Lecturer of the IEEE SP Society (2008–2009). He has served as Chair (2007–2008) of the Signal Processing for Communications and Networking Technical Committee (SPCOM-TC) of the IEEE SP Society (2004–2009), Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING (2000–2006) and the IEEE SIGNAL PROCESSING LETTERS (2000–2002), and currently serves on the editorial board of *IEEE Signal Processing Magazine*.



**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 academic year 1984–1985, 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 the *SIAM Journal on Optimization* and *Mathematics of Operations Research*.



**Leandros Tassiulas** (F'07) was born in Katerini, Greece, in 1965. He received the Diploma degree in electrical engineering from the Aristotelian University of Thessaloniki, Thessaloniki, Greece, in 1987, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1989 and 1991, respectively.

He is a Professor in the Department of Computer and Telecommunications Engineering, University of Thessaly, Greece, and a Research Professor in the Department of Electrical and Computer Engineering

and the Institute for Systems Research, University of Maryland, College Park, since 2001. He has held positions as Assistant Professor at Polytechnic University New York (1991–1995), Assistant and Associate Professor at University of Maryland, College Park (1995–2001), and Professor at University of Ioannina, Greece (1999–2001). His research interests are in the field of computer and communication networks with emphasis on fundamental mathematical models, architectures and protocols of wireless systems, sensor networks, high-speed internet and satellite communications.

Dr. Tassiulas received a National Science Foundation (NSF) Research Initiation Award in 1992, an NSF CAREER Award in 1995, an Office of Naval Research (ONR) Young Investigator Award in 1997, a Bodosaki Foundation Award in 1999, the IEEE INFOCOM '94 Best Paper Award, and the inaugural INFOCOM Achievement award (2007) for “fundamental contributions to resource allocation in communication networks.”