

Transmit Beamforming for Physical-Layer Multicasting

Nicholas D. Sidiropoulos, *Senior Member, IEEE*, Timothy N. Davidson, *Member, IEEE*, and Zhi-Quan (Tom) Luo, *Senior Member, IEEE*

Abstract—This paper considers the problem of downlink transmit beamforming for wireless transmission and downstream precoding for digital subscriber wireline transmission, in the context of common information broadcasting or multicasting applications wherein channel state information (CSI) is available at the transmitter. Unlike the usual “blind” isotropic broadcasting scenario, the availability of CSI allows transmit optimization. A minimum transmission power criterion is adopted, subject to prescribed minimum received signal-to-noise ratios (SNRs) at each of the intended receivers. A related max–min SNR “fair” problem formulation is also considered subject to a transmitted power constraint. It is proven that both problems are NP-hard; however, suitable reformulation allows the successful application of semidefinite relaxation (SDR) techniques. SDR yields an approximate solution plus a bound on the optimum value of the associated cost/reward. SDR is motivated from a Lagrangian duality perspective, and its performance is assessed via pertinent simulations for the case of Rayleigh fading wireless channels. We find that SDR typically yields solutions that are within 3–4 dB of the optimum, which is often good enough in practice. In several scenarios, SDR generates exact solutions that meet the associated bound on the optimum value. This is illustrated using measured very-high-bit-rate Digital Subscriber line (VDSL) channel data, and far-field beamforming for a uniform linear transmit antenna array.

Index Terms—Broadcasting, convex optimization, downlink beamforming, minimization of total radiation power, multicasting, semidefinite programming, semidefinite relaxation (SDR), very-high-bit-rate Digital Subscriber line (VDSL) precoding.

I. INTRODUCTION

CONSIDER a transmitter that utilizes an antenna array to broadcast information to multiple radio receivers within a certain service area. The traditional approach to broadcasting is

Manuscript received November 1, 2004; revised May 11, 2005. The work of N. D. Sidiropoulos was supported in part by the U.S. ARO under ERO Contract N62558-03-C-0012, the E.U. under FP6 U-BROAD STREP # 506790, and the GSRT. The work of T. N. Davidson was supported in part by the Natural Sciences and Engineering Research Council of Canada and the Canada Research Chairs program. The work of Z.-Q. Luo was supported in part by the U.S. National Science Foundation, Grant no. DMS-0312416. An earlier version of part of this work appears in *Proceedings of the IEEE SAM 2004 workshop*, vol. 1, pp. 489–493, Sitges, Barcelona, Spain, July 18–21, 2004. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Daniel Fuhrman.

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

T. N. Davidson is with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, ON L8S 4K1, Canada (e-mail: davidson@mcmaster.ca).

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).

Digital Object Identifier 10.1109/TSP.2006.872578

to radiate transmission power isotropically, or with a fixed directional pattern. However, future digital video/audio/data broadcasting and multicasting applications are likely to be based on subscription to services; hence, it is plausible to assume that the transmitter can acquire channel state information (CSI) for all its intended receivers.

The goal of this paper is to develop efficient algorithms for the design of broadcasting schemes that exploit this channel information in order to provide better performance than the traditional approaches.

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

Our designs are initially developed for a wireless broadcast scenario in which each user employs a single receive antenna and the channel is modeled as being flat in frequency and quasi-static in time. However, the designs are also appropriate on a per-tone basis for orthogonal-frequency-division multiplexing (OFDM) and related multicarrier systems, and, as we will show, they can be generalized in a straightforward manner to single-carrier systems transmitting over frequency-selective channels. In addition to wireless systems, applications of the proposed methodology also appear in downstream multicast transmission for multicarrier and single-carrier digital subscriber line (DSL) systems. In this context, (linear) *precoding* of multiple DSL loops in the same binder that wish to subscribe to a common service (e.g., news feed, video-conference, or movie multicast) can be employed to improve QoS and/or reduce far-end crosstalk (FEXT) interference to other loops in the binder. In scenarios in which the customer-premise equipment (CPE) receivers are not physically co-located (as in residential service) or cannot be co-ordinated (legacy CPE), multiuser decoding of the downstream transmission is not feasible, while transmit precoding is viable. The most important difference between DSL and the wireless

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

It is interesting to note that, as of today, Internet multicasting (using IP's Multicast Backbone—MBone) is performed at the *network layer*, e.g., via packet-level flooding or spanning-tree access of the participant nodes and any intermediate nodes needed to access the participants. To complement that approach, what we advocate herein can be interpreted as judicious *physical layer multicasting*, that is, enabled by i) the availability of multiple transmitting elements; ii) exploiting opportunities for joint beamforming/precoding; and iii) the availability of CSI at the transmitting node or one of its proxies. This is a cross-layer optimization approach that exploits information available at the physical layer to reduce relay retransmissions at the network layer, thus providing congestion relief and QoS guarantees.

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

II. DATA MODEL AND PROBLEM STATEMENT

Consider a wireless scenario incorporating a single transmitter with N antenna elements and M receivers each with a single antenna. Let \mathbf{h}_i denote the $N \times 1$ complex vector that models the propagation loss and phase shift of the frequency-flat quasi-static channel from each transmit antenna to the receive antenna of user $i \in \{1, \dots, M\}$, and let \mathbf{w}^H denote the beamforming weight vector applied to the N transmitting antenna elements. If the signal to be transmitted is zero-mean and white with unit variance, and if the noise¹ at receiver i is zero-mean and white with variance σ_i^2 , then the receiver SNR for the i th user is $|\mathbf{w}^H \mathbf{h}_i|^2 / \sigma_i^2$. Let $\rho_{\min,i}$ be the prescribed minimum SNR for the i th user and define the normalized channel vectors $\tilde{\mathbf{h}}_i := \mathbf{h}_i / \sqrt{\rho_{\min,i} \sigma_i^2}$. Then $|\mathbf{w}^H \mathbf{h}_i|^2 / \sigma_i^2 \geq \rho_{\min,i} \Leftrightarrow |\mathbf{w}^H \tilde{\mathbf{h}}_i|^2 \geq 1$. Therefore, the design of the beamformer that minimizes the transmitted power, subject to (possibly different) constraints on the received SNR of each user, can be written as

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

We will denote an instance of problem \mathcal{Q} as $\mathcal{Q}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M)$, keeping in mind that $\tilde{\mathbf{h}}_i = \mathbf{h}_i / \sqrt{\rho_{\min,i} \sigma_i^2}$.

Remark 1: One could think of imposing the stricter constraints $\mathbf{w}^H \tilde{\mathbf{h}}_i = 1, \forall i$ in order to avoid the need for single-tap equalization at the receivers. However, we are interested in the practically important case of $M > N$, wherein the stricter constraints generically yield an overdetermined system of equations, and thus an infeasible problem. On the other hand, it is

easy to see that problem \mathcal{Q} is always feasible, provided of course that none of the channel vectors is identically zero.

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

Problem \mathcal{Q} is a quadratically constrained quadratic programming (QCQP) problem, but unfortunately the constraints are not convex.² Nonconvexity, per se, does not mean that the problem is difficult to solve; however, we have the following claim, whose proof can be found in Appendix I.

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

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

A. Review of Pertinent Prior Art

The above problem is reminiscent of some closely related problems. For $M = 1$, the optimum \mathbf{w} is a matched filter. When the scaled channel vectors $\tilde{\mathbf{h}}_i$ span a ball or ellipsoid about a “nominal” channel vector,³ the problem can be transformed *exactly* into a second-order cone program, and hence can be efficiently solved [13]. Unfortunately, this transformation cannot be employed in the case of finitely many channel vectors (intended receivers).

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

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

³This implies a continuum of intended receivers.

¹The noise may include unmodeled interference.

Transmit beamforming for the dissemination of common information to multiple users has been considered in the Ph.D. dissertation of Lopez [7, ch. 5]. Lopez proposed maximizing the sum of received SNRs, which is equivalent to maximizing the average SNR over all users. This formulation leads to a principal component computational problem for the optimum beamformer, which is relatively simple to solve. The drawback is that quality of service cannot be guaranteed to all users in this way. This is important, because the weakest user link determines the common information rate. Still, the work of Lopez is the closest in spirit to ours, and for this reason we will include the maximum average SNR approach in our performance evaluations in Section VIII (see Table V).

III. RELAXATION

Toward solving our problem, we first recast it as follows:

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

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

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

where now \mathbf{X} is an $N \times N$ complex matrix, and the inequality $\mathbf{X} \succeq \mathbf{0}$ means that the matrix \mathbf{X} is symmetric positive semidefinite. Note that, in the above *equivalent* formulation of our problem, the cost function is linear in \mathbf{X} ; the trace constraints are linear inequalities in \mathbf{X} , and the set of symmetric positive semidefinite matrices is convex; however, the rank constraint on \mathbf{X} is not convex.⁴ The important observation is that the above problem is in a form suitable for semidefinite relaxation (SDR) (see, e.g., [9] and references therein); that is, dropping the rank-one constraint, one obtains the relaxed problem

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

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

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

which is now expressed in a standard form used by SDP solvers, such as SeDuMi [11]. Here, \mathbf{I}_N is the identity matrix of size $N \times N$.

SDP problems can be efficiently solved using interior point methods, at a complexity cost that is at most $O((M + N^2)^{3.5})$ and is usually much less. SeDuMi [11] is a MATLAB implementation of modern interior point methods for SDP that is particularly efficient for up to moderate-sized problems, as is the case in our context. Typical run times for realistic choices of N and M are under 1/10 s, on a typical personal computer.

IV. ALGORITHM

Due to the relaxation, the matrix \mathbf{X}_{opt} obtained by solving the SDP in Problem \mathcal{Q}_r will not be rank one in general. If it is, then its principal component will be the optimal solution to the original problem. If not, then $\text{trace}(\mathbf{X}_{\text{opt}})$ is a lower bound on the power needed to satisfy the constraints. This comes from the fact that we have removed one of the original problem’s constraints. Researchers in optimization have recently developed ways of generating good solutions to the original problem, \mathcal{Q} , from \mathbf{X}_{opt} , [9], [12], [15]. This process is based on *randomization*: using \mathbf{X}_{opt} to generate a set of candidate weight vectors, $\{\mathbf{w}_\ell\}$, from which the “best” solution will be selected. We consider three methods for generating the \mathbf{w}_ℓ ’s, which have been designed so that their computational cost is negligible compared to that of computing \mathbf{X}_{opt} . (For consistency, the principal component is also included in the set of candidates.) In the first method (*randA*), we calculate the eigen-decomposition of $\mathbf{X}_{\text{opt}} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^H$ and choose \mathbf{w}_ℓ such that $\mathbf{w}_\ell = \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{e}_\ell$, where the elements of \mathbf{e}_ℓ are independent random variables, uniformly distributed on the unit circle in the complex plane; i.e., $[\mathbf{e}_\ell]_i = e^{j\theta_{\ell,i}}$, where the $\theta_{\ell,i}$ are independent and uniformly distributed on $[0, 2\pi)$. This ensures that $\mathbf{w}_\ell^H \mathbf{w}_\ell = \text{trace}(\mathbf{X}_{\text{opt}})$, irrespective of the particular realization of \mathbf{e}_ℓ . In the second method (*randB*), inspired by Tseng [12], we choose \mathbf{w}_ℓ such that $[\mathbf{w}_\ell]_i = \sqrt{[\mathbf{X}_{\text{opt}}]_{ii}} [\mathbf{e}_\ell]_i$, which ensures that $||[\mathbf{w}_\ell]_i||^2 = [\mathbf{X}_{\text{opt}}]_{ii}$. The third method (*randC*), motivated by successful applications in related QCQP problems [8], uses $\mathbf{w}_\ell = \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{v}_\ell$, where \mathbf{v}_ℓ is a vector of zero-mean, unit-variance complex circularly symmetric uncorrelated Gaussian random variables. This ensures that $E[\mathbf{w}_\ell \mathbf{w}_\ell^H] = \mathbf{X}_{\text{opt}}$ [8].

For both *randA* and *randB*, $||\mathbf{w}_\ell||_2^2 = \text{trace}(\mathbf{X}_{\text{opt}})$, and hence when $\text{rank}(\mathbf{X}_{\text{opt}}) > 1$, at least one of the constraints $|\mathbf{w}_\ell^H \mathbf{h}_i|^2 \geq c_i$ will be violated.⁵ However, a feasible weight vector can be found by simply scaling \mathbf{w}_ℓ so that all the constraints are satisfied. Under *randC*, $||\mathbf{w}_\ell||_2^2$ depends on the particular realization of \mathbf{v}_ℓ , but again the resulting \mathbf{w}_ℓ can be scaled to the minimum length necessary to satisfy the constraints. The “best” of these randomly generated weight vectors is the one that requires the smallest scaling. For convenience, we have summarized the algorithm in Table I, which includes a simple MATLAB interface to SeDuMi [11] for the solution of the semidefinite relaxation, \mathcal{Q}_r . We point out that we have not yet been able to obtain theoretical *a priori* bounds on the extent

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

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

TABLE I
BROADCAST QoS BEAMFORMING VIA SDR: ALGORITHM

- Solve the relaxed problem:

A simple MATLAB interface for SeDuMi is as follows:

```
% Input Data:
% H: N by M, columns are scaled channels  $\mathbf{h}_i/\sqrt{\rho_{\min,i}\sigma_i^2}$ 
% Output Data:
% Xopt: the solution to the SDR
vecQs = [];
for i=1:M,
    Qi = H(:,i)*H(:,i)';
    vecQs = [vecQs, vec(Qi.')];
end;
A = [-eye(M), vecQs.'];
b = ones(M,1);
c = [zeros(M,1); vec(eye(N))];
K.l=M; K.s=N; K.scomplex=1;
[x_opt, y_opt, info] = sedumi(A, b, c, K);
Xopt = mat(x_opt(M+1:end));
```

- Randomization:

Use randA, and/or randB, randC to generate the candidates, w_e11.

For each w_e11, find the most violated constraint.

Scale w_e11 so that that constraint is satisfied with equality.

Pick the w_e11 with the smallest norm.

of the suboptimality of solutions generated in this way, but our simulation results are quite encouraging.

V. MAX-MIN FAIR BEAMFORMING

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

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

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

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

It is clear that P is immaterial with respect to optimization; the solution scales up with \sqrt{P} , while the optimum value scales up with P . We will denote an instance of problem \mathcal{F} as $\mathcal{F}(\{\mathbf{h}_i/\sigma_i\}_{i=1}^M, P)$. Let \mathbf{w}_q be a solution to $\mathcal{Q}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M)$,

and P_q the associated minimum transmitted power. Consider $\mathcal{F}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M, P_q)$, that is

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

and let \mathbf{w}_f denote an optimal solution. Since \mathbf{w}_q already attains $|\mathbf{w}_q^H \tilde{\mathbf{h}}_i|^2 \geq 1, \forall i$, it follows that $|\mathbf{w}_f^H \tilde{\mathbf{h}}_i|^2 \geq 1, \forall i$. Hence, \mathbf{w}_f also satisfies the constraints of the QoS formulation, and at the same power as \mathbf{w}_q . It follows that \mathbf{w}_f is equivalent to \mathbf{w}_q . This shows Claim 2.

Claim 2: $\mathcal{F}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M, P)$ is equivalent to $\mathcal{Q}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M)$ up to scaling. In the special case that $\rho_{\min,i} = \rho_{\min}, \forall i$, we have that $\tilde{\mathbf{h}}_i = (\mathbf{h}_i/\sigma_i)/\sqrt{\rho_{\min}}, \forall i$, and hence $\mathcal{F}(\{\mathbf{h}_i/\sigma_i\}_{i=1}^M, P)$ is equivalent to $\mathcal{Q}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M)$ up to scaling.

Corollary 1: One way to solve the max-min fair problem $\mathcal{F}(\{\mathbf{h}_i/\sigma_i\}_{i=1}^M, P)$ is to solve the QoS problem $\mathcal{Q}(\{\mathbf{h}_i/\sigma_i\}_{i=1}^M)$, then scale the resulting solution to the desired power P . Conversely, scaling the solution of $\mathcal{F}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M, 1)$ yields a solution to $\mathcal{Q}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M)$, even in the case of unequal $\rho_{\min,i}$.

Remark 2: It is important not to lose sight of the fact that $\mathcal{F}(\{\mathbf{h}_i/\sigma_i\}_{i=1}^M, P)$ is not equivalent up to scaling to $\mathcal{Q}(\{\tilde{\mathbf{h}}_i\}_{i=1}^M)$ when the $\rho_{\min,i}$'s are unequal. This can be intuitively appreciated by noting that the max-min fair formulation aims to maximize the minimum received SNR, without regard to the individual SNR constraints. The QoS formulation, on the other hand, explicitly guarantees the prescribed minimum SNR level at each node.

From the above, and Claim 1, Claim 3 follows.

Claim 3: The max-min fair problem \mathcal{F} is NP-hard.

If the QoS problem could be solved *exactly*, there would have been no need for a separate algorithm for the max-min fair problem. However, we can only solve the QoS problem approximately (cf., randomization postprocessing of the generally higher rank solution). Due to this, it is of interest to develop a customized SDR algorithm directly for the max-min fair problem. Using the fact that $\mathbf{h}_i^H \mathbf{w} \mathbf{w}^H \mathbf{h}_i = \text{trace}(\mathbf{w} \mathbf{w}^H \mathbf{h}_i \mathbf{h}_i^H)$, and defining $\mathbf{Q}_i := \mathbf{h}_i \mathbf{h}_i^H / \sigma_i^2$, we recast the max-min fair problem as follows:

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

Dropping the rank constraint, we obtain the relaxation

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

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

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

TABLE II
BROADCAST MAX-MIN BEAMFORMING VIA SDR: ALGORITHM

```

• Solve the relaxed problem:
A suitable MATLAB interface for SeDuMi is as follows:
% Input Data:
% H: N by M, columns are scaled channels  $\mathbf{h}_i/\sigma_i$ 
% P: scalar, the total transmit power constraint
% Output Data:
% Xopt: the solution to the SDR
% t_opt: the minimum objective value of the SDR
vecQs = [];
for i=1:M,
    Qi = H(:,i)*H(:,i)';
    vecQs = [vecQs vec(Qi.')];
end;
A1=[-ones(M,1); 0];
A2=[-eye(M); zeros(1,M)];
A3=[vecQs.'; vec(eye(N)).'];
A=[A1 A2 A3];
b=[zeros(M,1); P];
c = [-1; zeros(M+N,1)];
K.l=M+1; K.s=N; K.scomplex=1;
[x_opt,y_opt,info]=sedumi(A,b,c,K);
Xopt=mat(x_opt(M+2:end));
t_opt = x_opt(1)

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

```

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

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

This problem is formatted for direct solution via SeDuMi [11]. Table II provides a suitable MATLAB interface for solving this relaxation. Postprocessing of the solution of the relaxed problem to approximate the solution of the original max-min-fair problem can be accomplished using randA, randB, and randC, but the selection criterion is different (see Table II).

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

$$C := \max_{\substack{\mathbf{X} \succeq \mathbf{0}, \\ \text{trace}(\mathbf{X}) \leq P}} \min_i \left\{ \log \left(1 + \frac{\mathbf{h}_i^H \mathbf{X} \mathbf{h}_i}{\sigma_i^2} \right) \right\}_{i=1}^M.$$

Alternatively, we can rewrite this max-min problem as

$$\begin{aligned} & \max t \\ & \text{subject to } \mathbf{X} \succeq \mathbf{0}, \text{trace}(\mathbf{X}) \leq P, \log \left(1 + \frac{\mathbf{h}_i^H \mathbf{X} \mathbf{h}_i}{\sigma_i^2} \right) \geq t, \quad \forall i. \end{aligned}$$

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

$$\begin{aligned} & \max \tilde{t} \\ & \text{subject to } \mathbf{X} \succeq \mathbf{0}, \text{trace}(\mathbf{X}) \leq P, \frac{\mathbf{h}_i^H \mathbf{X} \mathbf{h}_i}{\sigma_i^2} \geq \tilde{t}, \quad \forall i \end{aligned}$$

in the sense that they yield the same optimal transmit covariance matrix \mathbf{X} . The latter problem is identical to problem \mathcal{F}_r . In other words, the semidefinite relaxation of problem \mathcal{F} actually yields a transmit covariance matrix that achieves the maximum common information rate C . In a similar manner, we can argue that the rank-one transmit covariance matrix obtained from problem \mathcal{F} achieves the maximum common information rate under the restriction that beamforming is employed. However, the latter rate can be significantly lower than C for a large number of users [6]. Nonetheless, from a practical perspective, beamforming is attractive because it is simple to implement,⁶ requiring only a single standard additive white Gaussian noise (AWGN) channel encoder and decoder. In contrast, achieving the maximum common information rate C in general requires higher rank transmit covariance matrix \mathbf{X} . In that case, a weighted sum of multiple independent signals is transmitted from each antenna, with each independent signal requiring a separate AWGN channel encoder and decoder. Hence, the beamforming strategy considered in this paper trades off a potential reduction in the maximum common information rate for implementation simplicity.

VI. CASE OF FREQUENCY-SELECTIVE MULTIPATH

Although we have focused our attention so far on frequency-flat fading channels, the situation is quite similar in the case of spatial beamforming⁷ for common information transmission over frequency-selective (intersymbol interference) channels. Let $\mathbf{h}_i^{(\ell)}$ denote the ℓ th $N \times 1$ vector tap of the baseband-equivalent discrete-time impulse response of the multipath channel between the transmitter antenna array and

⁶A properly weighted common temporal signal is transmitted from each antenna.

⁷It is perhaps worth emphasizing that, while *space-time precoding* would generally be preferable from a performance point of view when the channels are time dispersive, we (continue to) consider *spatial beamforming* only in this section. This is motivated from a complexity point of view. Space-time multi-antenna precoding is an interesting topic for future research.

the (single) receive antenna of receiver i . Assume that delay spread is limited⁸ to L nonzero vector channel taps. Define the channel matrix for the i th receiver as

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

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

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

or, in vector form

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

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

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

where now $\bar{\mathbf{Q}}_i := \mathbf{H}_i \mathbf{H}_i^H / \sigma_i^2$ and is generally of higher rank than before, but otherwise things remain conceptually the same. In particular, the relaxations \mathcal{Q}_r and \mathcal{F}_r and the algorithms in Tables I and II can be employed as they were in the frequency-flat case—only the definition of the input matrices changes.

VII. INSIGHTS AFFORDED VIA DUALITY

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

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

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

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

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

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

and the dual problem is

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

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

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

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

which is a semidefinite program.

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

Theorem 1: [14, pp. 403–404] $\mathcal{D}(\mathcal{D}(\mathcal{P})) = \mathcal{R}(\mathcal{P})$ and $\beta(\mathcal{R}(\mathcal{P})) = \beta(\mathcal{D}(\mathcal{P}))$.

More specifically, Theorem 1 states that the dual of the dual of \mathcal{P} is the SDR of \mathcal{P} and that the optimal objective value of the SDR of \mathcal{P} is the same as the optimal objective value of the dual of \mathcal{P} . Hence, SDR yields the same lower bound on the optimal value of \mathcal{P} as that obtained from duality, and the associated gap between this bound and the optimal value is equal to the duality gap.

Theorem 1 along with Claim 2 directly yield the following corollary for the max–min–fair problem \mathcal{F} .

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

VIII. SIMULATION RESULTS

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

TABLE III

MC SIMULATION RESULTS FOR QoS BEAMFORMING: MEAN AND STANDARD DEVIATION OF UPPER BOUND ON POWER BOOST. EACH ELEMENT OF \mathbf{h}_i IS I.I.D. WITH A CIRCULARLY SYMMETRIC COMPLEX GAUSSIAN (RAYLEIGH) DISTRIBUTION OF VARIANCE 1. ALL THREE RANDOMIZATION TECHNIQUES (randA, randB, randC) ARE USED IN PARALLEL, FOR 1000 RANDOMIZATIONS EACH. $\rho_{\min,i}\sigma_i^2 = 1, \forall i$

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

TABLE IV

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

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

bound on the power boost required to satisfy the constraints. If our algorithm achieves a power boost of η , then the transmitted power is guaranteed to be within a factor η of that of the optimal solution $\beta(\mathcal{Q})$ and will often be closer.

A. Rayleigh Fading Wireless Channels

We consider the standard independent and identically distributed (i.i.d.) Rayleigh fading model described in the caption of Table III. That table summarizes the results obtained using the direct QoS relaxation algorithm in Table I ($\rho_{\min,i}\sigma_i^2 = 1, \forall i$) with all three randomization options (randA, randB, and randC) employed in parallel, for a fixed number of 1000 randomization samples each. Table IV summarizes results for the same scenario, except that $30 NM$ randomization samples are drawn for each randomization strategy—thus the number of randomizations grows linearly in the problem size. Note that, in many cases, our solutions are within 3–4 dB from the (generally optimistic) lower bound on transmitted power provided by SDR, and thus are guaranteed to be at most 3–4 dB away from optimal; this is often good enough from an engineering perspective. Comparing the corresponding entries in Tables III and IV, it is evident that switching from 1000 to $30 NM$ randomizations per channel realization only yields a minor performance improvement in the cases considered.

Table V summarizes our simulation results for max–min fair beamforming, using the direct algorithm in Table II ($\sigma_i^2 = 1, \forall i, P = 1$). Table V presents Monte Carlo averages for the upper bound on the minimum SNR (the optimum attained

TABLE V

MC SIMULATION RESULTS FOR MAX–MIN FAIR BEAMFORMING: AVERAGES FOR THE UPPER BOUND ON $\min_i \text{SNR}_i$, THE RELAXATION-ATTAINED $\min_i \text{SNR}_i$, THE $\min_i \text{SNR}_i$ ATTAINED BY MAXIMIZING AVERAGE SNR (ACROSS USERS), AND THE $\min_i \text{SNR}_i$ FOR THE CASE OF NO BEAMFORMING. THE RESULTS ARE AVERAGED OVER 1000 MONTE CARLO (MC) RUNS. FOR EACH MC RUN, THE ELEMENTS OF \mathbf{h}_i ARE INDEPENDENTLY REDRAWN FROM A CIRCULARLY SYMMETRIC COMPLEX GAUSSIAN DISTRIBUTION OF VARIANCE 1. $\sigma_i^2 = 1, \forall i, P = 1$. ALL THREE RANDOMIZATION TECHNIQUES (randA, randB, randC) ARE USED IN PARALLEL, FOR $30 NM$ RANDOMIZATIONS EACH

N/M	upper bound	SDR	Max Avg SNR	no BMF
4/8	1.05	0.94	0.25	0.12
4/16	0.73	0.51	0.11	0.06
8/16	1.43	0.86	0.14	0.06
8/32	1.07	0.45	0.06	0.03

in problem \mathcal{F}_r), the SDR-attained minimum SNR (after randomization), the minimum SNR attained by the maximum average SNR beamformer⁹ [7, ch. 5], and the minimum SNR for the case of no beamforming. For the latter, we have used $\mathbf{w} = (1/\sqrt{N})\mathbf{1}_{N \times 1}$, which fixes transmitted power to 1. Under the i.i.d. Rayleigh fading assumption, this is *equivalent* to selecting an arbitrary transmit antenna, allocating the entire power budget to it, and shutting off all others. To see this, note that the sum channel $(1/\sqrt{N})\mathbf{1}_{N \times 1}^T \mathbf{h}_i$ viewed by any particular receiver i will still be Rayleigh, of the same variance as the elements of \mathbf{h}_i . For this reason, we can view the beamforming vector $\mathbf{w} = (1/\sqrt{N})\mathbf{1}_{N \times 1}$ as corresponding to no beamforming at all. All three randomization options (randA, randB, and randC) were employed in parallel, for $30 NM$ samples each. It is satisfying to note that the SDR solution attains a significant fraction of the (possibly unattainable) upper bound. Furthermore, the SDR technique provides a substantial improvement in the average minimum SNR relative to no beamforming and to maximum average SNR beamforming [7, ch. 5]. Like SDR, maximum average SNR beamforming uses full CSI at the transmitter. However, it is generally not meaningful to compare designs produced under different objectives. Accordingly, the maximum average SNR beamforming results in Table V are only meant to convey an idea of how much QoS improvement SDR can provide over computationally simpler solutions that also exploit full CSI.

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

⁹This beamformer maximizes the average SNR for each channel matrix realization (Monte Carlo run), where the average is taken over the users. The resulting beamforming vector is also scaled to unit norm.




Fig. 1. Broadcast beamforming example using algorithm in Table I. Optimized beam pattern for $N = 8$ -element transmit ULA ($d/\lambda = 1/2$) and $M = 24$ downlink users, in six clusters of four users each. Clusters centered at $[-51, -31, -11, 11, 31, 51]^\circ$ with extent $\pm 2^\circ$. Channel vectors are Vandermonde, of element modulus 1. $\sigma_i^2 = \sigma^2, \forall i, \rho_{\min,i} = 1/\sigma^2, \forall i$ (here, σ_i^2 also models propagation loss, in addition to thermal noise). Symmetric lobes appear due to the inherent ULA ambiguity. randA, # post-SDR randomizations = 300. In this case, the solution is guaranteed to be within 0.1% of the optimum.

method per se may degrade as the problem size grows. In a related, but distinct, problem the quality of the SDR approximation degrades logarithmically in the problem size [10].

B. Far-Field Beamforming for a Uniform Linear Transmit Antenna Array

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

C. Measured VDSL Channels

In this section, we test the performance of our algorithms using measured VDSL channel data collected by France Telecom R&D as part of the EU-FP6 U-BROAD project # 506 790.

Gigabit VDSL technology for very short twisted copper loops (in the order of 100–500 m) is currently under development in the context of fiber to the basement (FTTB) or fiber to the curb/cabinet (FTTC) hybrid access solutions. Multiple-input multiple-output (MIMO) transmission modalities are an important component of gigabit VDSL. These so-called *vectoring* techniques rely on transmit precoding and/or multiuser detection to provide reliable communication at very high transmission rates [5]. Transmit precoding is particularly appealing when the targeted receivers are not physically co-located, or when legacy equipment is being used at the receive

site. In both cases, multiuser detection is not feasible. In this context, media streaming (e.g., news-feed, pay-per-view, or video-conferencing) may involve multiple recipients in the same binder.

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

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

We use IL and far-end FEXT measured data for S88 cable comprising 14 quads, i.e., 28 loops. The length of the cable is 300 m. For each channel, a log-frequency sweeping scheme was used to measure the I/Q components of the frequency response from 10 kHz to 30 MHz, yielding 801 complex samples per channel. Cubic spline complex interpolation was used to convert these samples to a linear frequency scale. We consider $17 N \times N$ channel matrices, with $N = 14$, in the frequency range 21.5 to 30 MHz. Insertion loss drops between -40 and -45 dB in this range of frequencies, while FEXT coupling is between -77 and -82 dB in the mean, with over 10-dB standard deviation and significant variation across frequency as well. For each channel matrix, we apply our max–min–fair beamforming algorithm with $\sigma_i^2 = \sigma^2, \forall i$, and $P/\sigma^2 = 1$. Fig. 2 shows the resulting plots of minimum received signal power, the associated relaxation/duality bound, and the minimum received signal power when no precoding is used. We observe that SDR can almost double the minimum received signal power relative to no precoding, and it often attains zero gap relative to the relaxation/duality bound. For shorter loops (e.g., 100 m), the situation is even more in favor of SDR, because then FEXT resembles near-end crosstalk (NEXT) and is relatively more pronounced.

D. Further Observations

1) *Comparison of the Two Relaxations:* We have shown theoretically that the two problem formulations (QoS, \mathcal{Q} , and max–min–fair, \mathcal{F}) are algorithmically equivalent, i.e., had we had an optimal algorithm that provides the exact solution to one, it could have also been used to obtain the exact solution to the other. What we have instead is two generally approximate algorithms, obtained by direct relaxation of the respective problems. The link between the two formulations can still be exploited. For example, we may obtain an approximate




Fig. 2. Transmit precoding for VDSL multicasting.

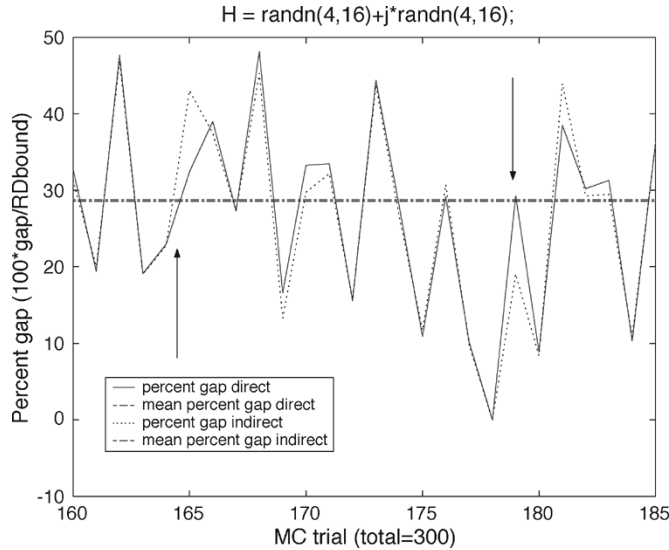


Fig. 3. Comparison of direct and indirect solutions to the max-min-fair problem.

solution to the max-min-fair problem by first running the QoS algorithm in Table I with all the $\rho_{\min,i} = 1$, then scaling the resulting solution to the desired power level P . Of course, we can also use the direct relaxation of the max-min-fair problem in Table II. Due to approximation, there is no *a priori* reason to expect that the two solutions will be identical, even in the mean.

In order to address this issue, we have compared the two strategies by means of Monte Carlo simulation. We chose $N = 4$, $M = 16$, $\sigma_i^2 = 1$, $\forall i$, and $P = 1$, and ran both algorithms for 300 i.i.d. Rayleigh fading channels. All three randomizations (randA, randB, randC) were employed in parallel, for 30 NM randomization samples each. For each channel, we recorded the percent gap (100 times the gap over the relaxation bound) of each algorithm. Fig. 3 shows a portion of the results, along with the mean percent gap attained by each algorithm (averaged over all 300 channels). By “direct” we refer to the algorithm in Table II, whereas by “indirect,” we refer to the algorithm in Table I with all $\rho_{\min,i} = 1$, followed by scaling.

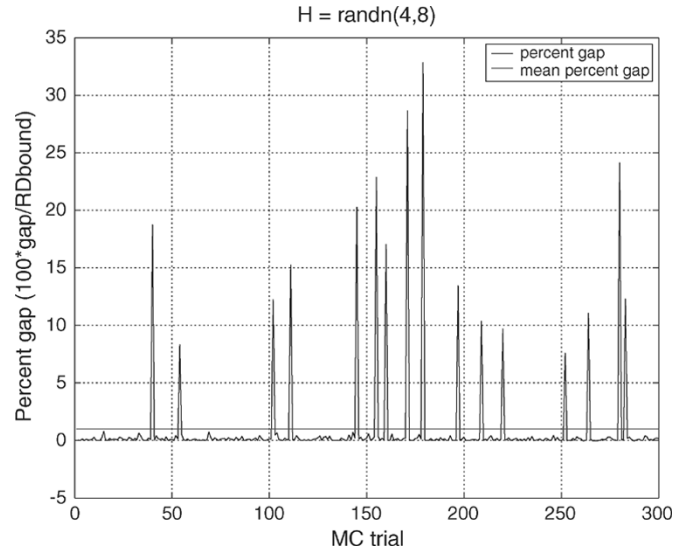


Fig. 4. Percent gap outcomes for 300 real Gaussian channel realizations.

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

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

It is evident that the gap statistics depend on the channel statistics. Interestingly, the gap statistics are far more favorable for *real* (as opposed to complex circular) i.i.d. Gaussian channels. This is illustrated in Fig. 4, using the QoS algorithm in Table I for $N = 4$, $M = 8$, $\rho_{\min,i}\sigma_i^2 = 1$, $\forall i$, and 300 real i.i.d. Gaussian channels. All three randomizations (randA, randB, randC) are employed in parallel, for 30 NM randomization samples each. For each channel, we recorded the percent gap (100 times the gap over the relaxation bound) of the algorithm in Table I. Observe that for about 95% of the channels the percent gap is down to numerical accuracy in this case. Contrast this situation with Fig. 5, which shows the respective results for complex circular Gaussian channel matrices—the difference is remarkable.

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




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




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

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

IX. CONCLUSION

We have taken a new look at the broadcasting/multicasting problem when channel state information is available at the transmitter. We have proposed two pertinent problem formulations: minimizing transmitted power under multiple minimum received power constraints, and maximizing the minimum received power subject to a bound on the transmitted power. We have shown that both formulations are NP-hard optimization problems; however, their solution can often be well approximated using semidefinite relaxation tools. We have explored the relationship between the two formulations and also insights

afforded by Lagrangian duality theory. In view of i) our extensive numerical experiments with simulated and measured data, verifying that semidefinite relaxation consistently yields good performance, ii) proof that the basic problem is NP-hard, and thus approximation is unavoidable, and iii) corroborating motivation provided by duality theory, we conclude that the approximate solutions provided herein offer useful designs across a broad range of applications.

It would be useful to analyze the duality gap for the problem at hand, for this would yield *a priori* bounds on the degree of suboptimality introduced by relaxation, as opposed to the *a posteriori* bound that we now have by virtue of Theorem 1. Our numerical results indicate that the degree of suboptimality is often acceptable in our intended applications. In an effort to understand the apparent success of the SDR approach (e.g., in the case where the channel vectors have nonnegative real and imaginary parts), one can consider the following simple linearly constrained convex quadratic program (QP) *restriction* of the QoS problem:

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

Notice that the feasible region of this problem is a subset of that of the original nonconvex (and NP-hard) QoS formulation \mathcal{Q} . Thus, $P^* \leq \bar{P}$, where P^* and \bar{P} denote the minimum beamforming power obtained from optimal solutions of \mathcal{Q} and \mathcal{Q}_s , respectively. We have recently shown [16] that the gap between P^* and \bar{P} is never more than $1/\cos^2(\alpha/2)$, where α is the maximum phase spread across the different users measured at each transmit antenna and is assumed to be less than π . Notice that the two cases where channel vectors i) are real and nonnegative or ii) have nonnegative real and imaginary parts correspond to $\alpha = 0$ and $\alpha \leq \pi/2$. Thus, \mathcal{Q}_s provides an exact solution in the first case and a factor of 2 approximation in the second case. These results indicate that problem \mathcal{Q} is well approximated by \mathcal{Q}_s if the phase spread α is small.

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

APPENDIX I
PROOF OF CLAIM 1

Before dealing with Claim 1 directly, we first consider the following restriction of the QoS problem \mathcal{Q} : the case when all \mathbf{h}_i are real, and optimization is over \mathbb{R}^N . We will show that¹⁰

$$\mathcal{S} : \quad \min_{\mathbf{x} \in \mathbb{R}^N} \mathbf{x}^T \mathbf{x}$$

$$\text{subject to: } |\mathbf{x}^T \mathbf{h}_m| \geq 1, \quad m \in \{1, \dots, M\}$$

contains

$$\mathcal{A} : \quad \min_{y_n \in \mathbb{R}} y_1^2 + \dots + y_N^2 + \left(\sum_{n=1}^N a_n y_n \right)^2$$

$$\text{subject to: } y_n^2 \geq 1, \quad n \in \{1, \dots, N\}$$

as a special case and that problem \mathcal{A} is at least as hard as the following problem:

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

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

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

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

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

where $\mathbf{a} := [a_1, \dots, a_N]^T$, and $\mathbf{Q} := \mathbf{I} + \mathbf{a}\mathbf{a}^T$. The matrix \mathbf{Q} is positive definite. Let $\mathbf{Q} = \mathbf{S}^T \mathbf{S}$, and $\mathbf{x} := \mathbf{S} \mathbf{y}$. Then $\mathbf{y}^T \mathbf{Q} \mathbf{y} = \mathbf{x}^T \mathbf{x}$, $\mathbf{y} = \mathbf{S}^{-1} \mathbf{x}$, and $|\mathbf{y}^T \mathbf{e}_n| \geq 1$ can be written as $|\mathbf{x}^T \mathbf{S}^{-T} \mathbf{e}_n| \geq 1$, or, with $\mathbf{h}_n := \mathbf{S}^{-T} \mathbf{e}_n$, as $|\mathbf{x}^T \mathbf{h}_n| \geq 1$. This shows that an arbitrary instance of problem \mathcal{A} can be transformed to a special instance of problem \mathcal{S} (with $M = N$). Thus, \mathcal{S} is at least as hard as \mathcal{A} , which is at least as hard as the partition problem. \square

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

$$\min_{\mathbf{w} \in \mathbb{C}^N} \quad \mathbf{w}^H \mathbf{w}$$

$$\text{subject to: } |\mathbf{w}^H \mathbf{h}_i| \geq 1, \quad i = 1, \dots, M. \quad (1)$$

Define the $N \times M$ matrix $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_M]$, and the $M \times 1$ vector \mathbf{z} , with $\mathbf{z}^H := \mathbf{w}^H \mathbf{H}$. Consider the case that $M \geq N$,

¹⁰We henceforth use \mathbf{h}_m to denote possibly scaled channel vectors, dropping the tilde for brevity.

and \mathbf{H} is full row-rank (N). Then $\mathbf{w}^H = \mathbf{z}^H \mathbf{H}^\dagger$, where $\mathbf{H}^\dagger = \mathbf{H}^H (\mathbf{H} \mathbf{H}^H)^{-1}$ denotes the right pseudoinverse of \mathbf{H} , and the problem in (1) is equivalent to

$$\min_{\mathbf{z} \in \mathbb{C}^M} \quad \mathbf{z}^H \mathbf{Q} \mathbf{z}$$

$$\text{subject to: } |z_k| \geq 1, \quad k = 1, \dots, M \quad (2)$$

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

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

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

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

Let us denote

$$\mathbf{a} := [a_1 \quad a_2 \quad \dots \quad a_P]^T$$

$$\mathbf{A} := \begin{bmatrix} -\mathbf{1}_P & \mathbf{I}_P & \mathbf{I}_P \\ -\frac{1}{2} \mathbf{1}_P^T \mathbf{a} & \mathbf{a}^T & \mathbf{0}_P^T \end{bmatrix}$$

$$\mathbf{Q} := \mathbf{A}^T \mathbf{A} + \mathbf{I}_M$$

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

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

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

it follows that

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

is equivalent to

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

The latter gives rise to a set of linear equations

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

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

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

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

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

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

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

Therefore, (5) is satisfied if and only if

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

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

ACKNOWLEDGMENT

The authors would like to thank Dr. M. Ouzzif and Dr. R. Tarafi of France Telecom R&D, who conducted the VDSL channel measurements under the auspices of the U-BROAD project.

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, Aug. 2001, ch. 18.
- [2] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [3] G. D. Forney, "Maximum likelihood sequence estimation of digital sequences in the presence of intersymbol interference," *IEEE Trans. Inf. Theory*, vol. 18, no. 3, pp. 363–378, May 1972.
- [4] M. R. Garey and D. S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-Completeness*. San Francisco, CA: Freeman, 1979.
- [5] G. Ginis and J. M. Cioffi, "Vectored transmission for digital subscriber line systems," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 5, pp. 1085–1104, Jun. 2002.
- [6] N. Jindal and A. Goldsmith, "Optimal power allocation for parallel gaussian broadcast channels with independent and common information," presented at the Int. Symp. Information Theory, Chicago, IL, Jun. 2004.

- [7] M. J. Lopez, "Multiplexing, scheduling, and multicasting strategies for antenna arrays in wireless networks," Ph.D. thesis, Dept. of Elect. Eng. and Comp. Sci., MIT, Cambridge, MA, 2002.
- [8] Z.-Q. Luo, "Lecture Notes for EE8950: Engineering Optimization," Univ. of Minnesota, Minneapolis, Spring 2004.
- [9] 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.
- [10] A. Nemirovski, C. Roos, and T. Terlaky, "On maximization of quadratic form over intersection of ellipsoids with common center," *Math. Program.*, ser. A, vol. 86, pp. 463–473, 1999.
- [11] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optim. Meth. Softw.*, vol. 11–12, pp. 625–653, 1999.
- [12] P. Tseng, "Further results on approximating nonconvex quadratic optimization by semidefinite programming relaxation," *SIAM J. Optim.*, vol. 14, no. 1, pp. 268–283, Jul. 2003.
- [13] S. A. Vorobyov, A. B. Gershman, and Z.-Q. Luo, "Robust adaptive beamforming using worst-case performance optimization: A solution to the signal mismatch problem," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 313–324, Feb. 2003.
- [14] 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.
- [15] S. Zhang, "Quadratic maximization and semidefinite relaxation," *Math. Program.*, ser. A, vol. 87, pp. 453–465, 2000.
- [16] Z.-Q. Luo, N. D. Sidiropoulos, P. Tseng, and S. Zhang, "Approximation bounds for quadratic optimization with homogeneous quadratic constraints," *SIAM J. Optim.*, Oct. 2005, submitted for publication.



Nicholas D. Sidiropoulos (M'92–SM'99) received the Diploma degree in electrical engineering from the Aristotelian University of Thessaloniki, Thessaloniki, Greece, in 1988 and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park (UMCP), in 1990 and 1992, respectively.

From 1988 to 1992, he was a Fulbright Fellow and a Research Assistant at the Institute for Systems Research (ISR), UMCP. 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. He was a Postdoctoral Fellow from 1994 to 1995 and Research Scientist from 1996 to 1997 at 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. He is currently a Professor with the Telecommunications Division of the Department of Electronic and Computer Engineering, 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 and multiway analysis. He is an active consultant for industry in the areas of frequency-hopping systems and signal processing for xDSL modems.

Dr. Sidiropoulos is a member of both the Signal Processing for Communications (SPCOM) and Sensor Array and Multichannel Signal Processing (SAM) Technical Committees of the IEEE Signal Processing Society and currently serves as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING. From 2000 to 2002, he also served as Associate Editor for the IEEE SIGNAL PROCESSING LETTERS. He received the NSF/CAREER award (Signal Processing Systems Program) in June 1998 and an IEEE Signal Processing Society Best Paper Award in 2001.



Timothy N. Davidson (M'96) received the B.Eng. (Hons. I) degree in electronic engineering from the University of Western Australia (UWA), Perth, in 1991 and the D.Phil. degree in Engineering Science from the University of Oxford, U.K., in 1995.

He is currently an Associate Professor in the Department of Electrical and Computer Engineering at McMaster University, Hamilton, ON, Canada, where he holds the (Tier II) Canada Research Chair in Communication Systems. His research interests are in signal processing, communications and control,

with current activity focused on signal processing for digital communication systems. He has held research positions at the Communications Research Laboratory at McMaster University, the Adaptive Signal Processing Laboratory at UWA, and the Australian Telecommunications Research Institute at Curtin University of Technology, Perth, Western Australia.

Dr. Davidson was awarded the 1991 J. A. Wood Memorial Prize [for "the most outstanding (UWA) graduand" in the pure and applied sciences] and the 1991 Rhodes Scholarship for Western Australia.



Zhi-Quan (Tom) Luo (SM'03) received the B.Sc. degree in mathematics from Peking University, China, in 1984 and the Ph.D. degree in operations research from the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, in 1989.

During the academic year of 1984 to 1985, he was with 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 has held the Canada Research Chair in Information Processing since 2001. Starting 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 at the 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 is a member of SIAM and 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. He is presently serving as an Associate Editor for several international journals, including the *SIAM Journal on Optimization* and *Mathematics of Operations Research*. He received an IEEE Signal Processing Society Best Paper Award in 2004.