Dynamic Power Allocation Under Arbitrary Varying Channels – The Multi-User Case

Niv Buchbinder
Microsoft Research, New England
Cambridge, MA 02142
Email: nivbuchb@microsoft.com

Liane Lewin-Eytan
IBM Haifa Research Laboratory
Haifa 31905, Israel
Email: lianel@il.ibm.com

Ishai Menache
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Email: ishai@mit.edu

Joseph (Seffi) Naor
Computer Science Department
Technion, Israel Institute of Technology
Email: naor@cs.technion.ac.il

Ariel Orda
Department of Electrical Engineering
Technion, Israel Institute of Technology
Email: ariel@ee.technion.ac.il

***Abstract*—We consider the power control problem in a time-slotted wireless channel, shared by a finite number of mobiles that transmit to a common base station. The channel between each mobile and the base station is time varying, and the system objective is to maximize the overall data throughput. It is assumed that each transmitter has a limited power budget, to be sequentially divided during the lifetime of the battery. We deviate from the classic work in this area, by considering a realistic scenario where the channel quality of each mobile changes *arbitrarily* from one transmission to the other.**

Assuming first that each mobile is aware of the channel quality of all other mobiles, we propose an *online* power-allocation algorithm, and prove its optimality under mild assumptions. We then indicate how to implement the algorithm when only local state information is available, requiring minimal communication overhead. Notably, the competitive ratio of our algorithm (nearly) matches the bound obtained for the (much simpler) single-transmitter case [2], albeit requiring significantly different algorithmic solutions.

I. Introduction

A. Background and Motivation

The multiple access nature of wireless networks introduces a fundamentally different resource allocation problem as compared to wired networks, which often provide a dedicated channel for each user. Indeed, the shared nature of the wireless domain implies that the rate obtained by a user depends not only on its own transmission parameters, but also on the transmission characteristics of other users. Specifically, when considering an uplink scenario, the rate that a user can sustain depends on the transmission power of *all* users, and at the same time on their respective channel qualities (or gains). Naturally, the channel gains are time-varying (an effect known as *channel fading*), making the rate allocation task (as determined by the power assignment) a *dynamic control* problem. Adequate algorithmic solutions for power allocation are crucial, especially for limited-battery devices, for which inefficient use of power might be devastating.

Much research has been devoted within the information theory community to the study of the optimal power allocation problem in the face of varying channel conditions. The most commonly studied objective is maximizing the *sum-rate* of the system, under an average power constraint per user (see [6] for a detailed survey). An underlying assumption in this line of research is that the channel gain variation between mobile users and a base station satisfies a probabilistic rule, known to all parties. In practice, however, such information on the channel-gain distribution may not be available, hence calling for adaptive schemes for estimating it. Even worse, the probability rule governing the channel state processes might change over time, due to non-stationary network elements affecting transmission quality (e.g., mobility, line of sight, etc.). Accordingly, our goal is to investigate how well can power be allocated in an uplink scenario under *arbitrarily varying* channel conditions for each and every mobile.

Our model consists of multiple mobiles transmitting to a single base station over a time-slotted wireless channel. The channel between each mobile and the base station is arbitrarily time varying, and the system objective is to maximize the overall data throughput. A distinctive property of our model is that transmitters have a limited battery that can be recharged only occasionally. Hence, instead of satisfying a long-term power average constraint (as traditionally considered in related research), each transmitter has to be aware of its actual remaining energy. Due to the arbitrarily changing channel conditions, we study the problem within the framework of *online computation* [3], with the objective of devising online power-allocation algorithms with proven bounds on the competitive ratio. A second objective is to establish lower bounds on the worst-case performance of *any* online algorithm operating under arbitrarily varying channel conditions, hence providing a benchmark for the quality of our proposed solutions.

The technological relevance of our work lies, for example, in ad-hoc and sensor networks, where the battery of each mobile is limited and can be charged only occasionally (e.g., by solar energy). Sensors that are required to send informative data, may do so in a relatively slow pace, with the objective of maximizing their overall throughput. Due to the low rate of transmission, the assumption of arbitrary channel conditions is commensurate with the unknown changes (e.g., environmental) that take place between subsequent transmissions.

B. Related Work

The information theory community has considered the case where the transmitter and the receiver operate with incomplete knowledge of the probability law governing the channel over which a transmission takes place. The focus of the related research is in designing encoders and decoders which achieve reliable communication and in analyzing the capacity over such channels (see [4] for a survey). The problem that we consider here is fundamentally different, as we concentrate on the *sequential* power-allocation problem, constrained by limited power budgets.

Recently, there has been growing interest in *jamming games* (e.g., [1]), where a malicious adversary, equipped with its own power budget, aims at deteriorating system performance by allocating its own power (affecting the throughput of other users) in a harmful way. Our study differs from the jamming game framework in several respects. Our focus is on worst-case analysis, rather than on the notion of equilibrium between "equal" players. For such settings, in which decisions are made on an arbitrary input pattern revealed piece-by-piece, the methodology of *competitive analysis* [3] provides a framework for the systematic design of algorithmic solutions, as well as for establishing worst-case performance bounds. Online methods have gained prominence in solving algorithmic problems in a variety of networking domains, e.g., network switches and buffers, call admission control, and scheduling.

Due to the finite time-horizon assumption, the mathematical formulation of our *offline* power control problem shows much similarity with the well studied problem of assigning transmission powers to orthogonal frequency bands in a general-topology wireless network. This problem is studied as a centralized optimization problem as well as in a non-cooperative game theoretic setting (see, e.g., [7], [8], [11], [10], [9] and the references therein). The centralized optimization problem is known to be NP-hard [8], further highlighting the algorithmic challenge in our problem, since the online decision requirement imposes additional difficulty.

In recent work, the online power allocation problem for the *single* transmitter case [2] was considered. The multiuser case studied here uses some basic observations from [2]; yet, due to its higher complexity, it requires new formulations and algorithmic solutions.

C. Contribution and Paper Organization

To the best of our knowledge, this is the first paper studying the problem of power allocation in a multi-transmitter environment, under dynamically varying channel quality, through the methodology of online (competitive) analysis. We address two scenarios which correspond to different capabilities of the transmitter in terms of the feedback it obtains from the channel. The first one, the *fixed channel gain* scenario, corresponds to binary feedback, i.e., the channel quality information provided to a transmitter in each time slot is either "reception" or "no reception". We start by characterizing some fundamental properties of an optimal (*offline*) solution. Then, we provide an *online* algorithm for which we prove, under mild assumptions, that its worst-case performance (i.e., competitive ratio) is at most a *constant factor* away from an (offline) optimum. We then turn to the general scenario, where the information received from the channel is the gain value, rather than just a binary flag. For this case, we provide an online algorithm whose competitive ratio is on the same order as previously obtained for the single transmitter case [2].

The above results are obtained under the full information assumption, i.e., each transmitter is aware of the gain of all other transmitters in each time slot. We get around this assumption and provide a distributed scheme allowing the employment of the above algorithms in the practical case of "local" information, where each transmitter is aware of only its own gain value. We complement our work with a simulation study, where we validate our online algorithm for the fixed channel gain scenario, and examine the effect of certain parameters on its performance. In certain cases, we observe that the performance of our online algorithm is much better than the worst-case bounds.

The paper is organized as follows. The channel and transmitter are modeled in Section II, followed by a formulation of the problem. Section III addresses the fixed channel gain ("binary feedback") case, whereas the general case and the local information perspective are treated in Sections IV and V. Section VI presents the simulation study and discusses its results. Finally, conclusions appear in Section VII.

II. The Model

A. The Channel Model

We consider a (single cell) CDMA-like system with a finite set of transmitters $\{1, \ldots, n\}$, transmitting to a single base station (receiver) over a common bandwidth of W hertz. The channel between the transmitters and the receiver is modeled as a frequently-flat fading channel with additive white Gaussian noise. Specifically, at each time t, the received signal $y(t)$ is given by

$$y(t) = \sum_j \sqrt{\tilde{h}_j(t)} x_j(t) + z(t), \tag{1}$$

where $x_j(t)$ and $\tilde{h}_j(t) \geq 0$ are the transmitted signal and channel gain (state) for the jth user, and $z(t)$ is an additive white Gaussian noise with power spectral density $N_0/2$. The sequence of channel gains $\tilde{\mathbf{h}}(t) = \big(\tilde{h}_1(t), \ldots, \tilde{h}_n(t)\big)$ is modeled as a block-fading process [6], so that for $i = 1, 2, \ldots$,

$$\tilde{\mathbf{h}}(t) = \tilde{\mathbf{h}}^i = \big(\tilde{h}_1^i, \ldots, \tilde{h}_n^i\big), \quad \text{for all } t \in [iL, (i+1)L),$$

where $\tilde{h}_j^i$ is the gain of user j at time slot i, and L is the length of each time slot.

A distinctive feature of our model is that the process $\{\tilde{\mathbf{h}}^i\}$ evolves *arbitrarily*, i.e., without an underlying probability rule. At the beginning of every time slot i, each transmitter j receives information h_j^i regarding its current channel gain $\tilde{h}_j^i$.

This information is passed through a finite lossless feedback link with capacity of C bits per second[1]. The information $h_j^i \in \{q_0 = 0, q_1, \ldots, q_M\} \subset \mathcal{R}_+$ is a quantized version of the actual gain $\tilde{h}_j^i$, so that if $\tilde{h}_j^i \in [q_m, q_{m+1})$ (where $q_{M+1} \equiv \infty$), then $h_j^i = q_m$. Throughout the paper, we use the notation $h_{\min} = q_1$ for the smallest (nonzero) quantized gain, and $h_{\max} = q_M$ for the maximum one. To simplify the exposition, we shall henceforth refer to h_j^i as the channel gain of transmitter j at time-slot i.

The transmitter observes h_j^i and can accordingly decide whether to transmit or not, and also adapt its transmission power. We assume that there is no retransmission mechanism, so that each transmission arrives at the base station with very high probability. The basic measure that determines the instantaneous throughput is the received signal to noise ratio (SINR), given by

$$\text{SINR}_j^i = \frac{h_j^i p_j^i}{N_0 + \sum_{m \neq j} h_m^i p_m^i}, \tag{2}$$

where p_j^i is the transmission power of transmitter j at time i, and W is normalized to one. Note that in CDMA systems, the signals of other users are often treated as interfering noise signals (i.e., there is no interference cancellation at the receiver), hence the direct dependence of throughput on SINR, as we elaborate below.

Let $U(\text{SINR}_j^i)$ be the instantaneous rate which transmitter j can reliably transmit at time slot i. In this paper, we shall consider the function $U(\text{SINR}_j^i) = \log(1+\text{SINR}_j^i)$, which can be interpreted as being proportional to the Shannon capacity of user j, if we make the simplifying assumption that the noise plus the interference of all other users constitute an independent Gaussian noise.

B. The Optimization Problem

We assume that each transmitter j has an initial power budget of P_j that can be divided between different time slots. Further, transmitters can recharge their battery to their initial power; however, due to practical limitations, a period of T time slots elapses between consecutive battery charges.

The objective of the transmitters is to cooperatively maximize the total throughput of the system, subject to the constraints described above. We often refer to the total network throughput achieved as *profit*. The general optimization problem is thus the following:

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{T} \sum_{j=1}^{n} U(\text{SINR}_j^i) \qquad (3) \\ \text{subject to} \quad & \sum_{i=1}^{T} p_j^i \leq P_j, \quad \forall j \in \{1, \ldots, n\} \end{aligned}$$

where SINR_i^j is given in (2).

We emphasize that since the channel gain sequence $\mathbf{h}^i = \left(h_1^i, \ldots, h_n^i\right)$ is not known a-priori, we pose (3) as an *online optimization problem*, where in each time slot i a new channel gain sequence $\mathbf{h}^i$ is revealed. We consider two different scenarios, corresponding to the information available to the transmitters when making their power-allocation decision.

1) *Global information:* At every time slot i, each transmitter is aware of the gain of all transmitters (e.g., the gains are broadcasted to all transmitters). That is, the vector $\mathbf{h}^i = \left(h_1^i, \ldots, h_n^i\right)$ is known to all transmitters.
2) *Local information:* At every time slot i, each transmitter j is aware only of its own gain h_j^i.

In the bulk of the paper, we focus on the global information case. We describe a distributed and practical way to deal with the local information case in Section V.

[1]E.g., each feedback link may employ a different frequency band in a Frequency Division Duplex (FDD) system, which does not suffer from fading.

III. The Fixed Channel Gain Problem

We study here a version of our power allocation problem in which the channel gain value given to a user (in each time-slot) can be either 0 or a fixed value $\bar{h}_0$. This scenario corresponds to binary feedback, i.e., the channel quality information provided to a transmitter in each time slot is either "reception" or "no reception". Our main result is an online algorithm with constant competitive factor for this case. We later show that this online algorithm can be used in a black box fashion to design an online algorithm with competitive ratio $O\left(\log \frac{h_{\max}}{h_{\min}}\right)$ for the general online multi-user power allocation problem.

Let $F^i \subseteq \{1, \ldots, n\}$ denote the set of users that obtain a gain of $\bar{h}_0$ at time slot i. We refer to F^i as the set of users that can transmit (or "allowed" to transmit) at time slot i (since other users will obtain a zero throughput if they decide to transmit). The fixed channel gain optimization problem is thus:

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{T} \sum_{j \in F^i} \log\left(1 + \frac{p_j^i \cdot h_0}{1 + \sum_{m \neq j} p_m^i h_0}\right) \\ \text{subject to} \quad & \sum_{i=1}^{T} p_j^i \leq P_j, \quad \forall j \in \{1, \ldots, n\}, \end{aligned}$$

where $h_0 = \bar{h}_0 / N_0$.

We now point out the relation between our (offline) optimization problem and the well-studied power allocation problem in multiuser frequency selective environments with no fading (see, e.g., [8]). In our setting, users share a *single* band, whereas the channel quality of each one changes over time. By mapping time-slots to bands it is straightforward to reduce our (offline) optimization problem, stated above, to the problem investigated in [8]. This reduction implies that under mild assumptions, the optimal solution of our power allocation problem in the fixed channel gain scenario is a TDMA solution, where in each slot only a single user is allowed to transmit. Specifically, the assumption in our model (as follows from [8]) is that each user transmits during at least two time slots. Since this is a very reasonable assumption, we thus compare our online algorithm for the fixed channel case to an optimal TDMA solution.

A. Characterizing an Optimal TDMA Solution

In this section we characterize an optimal TDMA solution in the fixed gain case. In a TDMA solution only one user is allowed to transmit at each time slot. Thus, for each slot i some user from the set F^i is chosen to transmit. Let T_j be the number of slots during which user j transmits. The user may split its budget P_j between these slots in order to maximize its total throughput. Since the instantaneous throughput is concave in the transmission power, it follows by Jansen's inequality that the total throughput is maximized when the budget is equally divided between the T_j slots. The total profit is therefore

$$T_j \log\left(1 + \frac{P_j \cdot h_0}{T_j}\right). \tag{4}$$

The question is thus how to divide the time slots between the users. Consider an allocation of slots to users. Let $T_j^r = T_j / P_j$ be the ratio between the number of slots given to user j and its budget. Arrange the users $\{1, \ldots, n\}$ in a non-increasing order with respect to T_j^r. Let $T_{(j)}^r$ be the jth highest ratio in the allocation. Thus, we get an allocation vector $\left(T_{(n)}^r, T_{(n-1)}^r, \ldots, T_{(1)}^r\right)$ such that $T_{(n)}^r \leq T_{(n-1)}^r \leq \ldots \leq T_{(1)}^r$. It is then possible to compare any two allocations by comparing them coordinate after coordinate, starting from $T_{(n)}^r$ and stoping whenever a coordinate in one vector is strictly larger than a coordinate in the other vector. The allocation that is the best among all allocations with respect to this comparison procedure is referred to as lexicographically maximal. It is also known as the *max-min fair allocation*. The next theorem characterizes the optimal allocation as lexicographically maximal among all allocation vectors.

Theorem 1: Consider the fixed channel gain problem. The optimal allocation vector $\mathbf{T}^* = (T_{(n)}^r, T_{(n-1)}^r, \ldots, T_{(1)}^r)$ for that problem is a max-min fair allocation vector; that is, $\mathbf{T}^*$ is a lexicographically maximal allocation.

Proof: Assume by contradiction that $\mathbf{T}^*$ is not a max-min fair allocation. Then, there are two users j and k such that $T_j^r + \frac{1}{P_j} \leq T_k^r - \frac{1}{P_k}$, and there exists a slot which is allocated to k, yet it and can be allocated to j. That is, if we take a slot from k and give it to j, the ratio of k is still larger than the ratio of j. Suppose we take a slot from k and give it to j. Then, using Equality (4) we get that the total gain after this change increases, that is,

$$\begin{aligned}(T_j+1)\log\left(1+\frac{P_j \cdot h_0}{T_j+1}\right) + (T_k-1)\log\left(1+\frac{P_k \cdot h_0}{T_k-1}\right) \\ > T_j \log\left(1+\frac{P_j \cdot h_0}{T_j}\right) + T_k \log\left(1+\frac{P_k \cdot h_0}{T_k}\right).\end{aligned}$$

To see this, consider the function $f(x) = x \cdot \log(1 + \frac{h_0}{x})$. For this function, $f''(x) = \frac{h_0}{x+h_0}(\frac{1}{x+h_0} - \frac{1}{x}) \leq 0$, for any $x > 0$. Therefore, this is a concave function for any $x > 0$, and so for $x < y$, $f'(x) > f'(y)$. Next, using this fact we get that,

$$\begin{aligned}&(T_j+1)\log\left(1+\frac{P_j \cdot h_0}{T_j+1}\right) - T_j \log\left(1+\frac{P_j \cdot h_0}{T_j}\right)\\ =\ & P_j\left(\frac{(T_j+1)}{P_j}\log\left(1+\frac{P_j \cdot h_0}{T_j+1}\right) - \frac{T_j}{P_j}\log\left(1+\frac{P_j \cdot h_0}{T_j}\right)\right)\\ =\ & P_j\left(f\left(\frac{T_j}{P_j}+\frac{1}{P_j}\right) - f\left(\frac{T_j}{P_j}\right)\right)\\ \geq\ & P_j \cdot \frac{1}{P_j} \cdot f'\left(\frac{T_j}{P_j}+\frac{1}{P_j}\right) = f'\left(\frac{T_j}{P_j}+\frac{1}{P_j}\right).\end{aligned}$$

On the other hand,

$$\begin{aligned}&T_k \log\left(1+\frac{P_k \cdot h_0}{T_k}\right) - (T_k-1)\log\left(1+\frac{P_k \cdot h_0}{T_k-1}\right)\\ =\ & P_k\left(\frac{T_k}{P_k}\log\left(1+\frac{P_k \cdot h_0}{T_k}\right) - \frac{T_k-1}{P_k}\log\left(1+\frac{P_k \cdot h_0}{T_k-1}\right)\right)\\ =\ & P_k\left(f\left(\frac{T_k}{P_k}\right) - f\left(\frac{T_k}{P_k}-\frac{1}{P_k}\right)\right)\\ \leq\ & P_k \frac{1}{P_k} f'\left(\frac{T_k}{P_k}-\frac{1}{P_k}\right) = f'\left(\frac{T_k}{P_k}-\frac{1}{P_k}\right).\end{aligned}$$

Since $T_j^r + \frac{1}{P_j} \leq T_k^r - \frac{1}{P_k}$ and using concavity again we get the desired result. Thus, unless $\mathbf{T}^*$ is a max-min fair allocation, the total profit can be improved, proving the theorem. ∎

For the sake of analysis let us consider a *fractional* max-min fair allocation. In a fractional allocation, a slot can be fractionally divided between several users, where the sum of the allocated fractions of each slot adds up to 1. The total gain of user j is still considered (for analysis purpose) as $T_j \log\left(1 + \frac{P \cdot h_0}{T_j}\right)$, although now T_j can be non-integral. We denote by $T_{(j)}^*$ the allocation of the user with the jth highest ratio $T_{(j)}^r$ according to a max-min fair fractional allocation. It is easy to verify that the proof of Theorem 1 still goes through when considering a fractional max-min fair allocation, and thus the best TDMA solution satisfies the following inequality with respect to a fractional optimal solution.

$$OPT \leq \sum_{j=1}^{n} T_{(j)}^* \log\left(1 + \frac{P_j \cdot h_0}{T_{(j)}^*}\right).$$

From this point on we assume, for simplicity, that all budgets are equal to some value P. In this case we are only concerned about the number of slots assigned to each user, rather than the ratio of number of slots to the budget. This simplifies a bit our notation and presentation. We later remark on how to extend all of our results to the case of arbitrary budgets.

When considering a fractional max-min fair solution, we have the following nice property.

Fact 1: In every fractional max-min fair solution, each user j is allocated the same total number of slots T_j^*.

One way of proving this fact is that a max-min allocation is in particular a market equilibrium allocation in which equilibrium utilities are unique. See, e.g. [5, Theorem 5.1]. This motivates the following definition. For any j let us define $S_j^* = \left\{k \mid T_k^* \geq T_{(j)}^*\right\}$, where $T_{(j)}^*$ is the number of slots the user with the jth highest allocation received in a max-min fair fractional allocation. The set S_j^* consists of all the users that get at least $T^*_{(j)}$ slots in the fractional max-min fair allocation. Note that the size of S^*_j is at least j since there are at least j users that got at least $T^*_{(j)}$ in the max-min fair solution, but there can be more than j users in S^*_j in the case where more than one user got exactly $T^*_{(j)}$ slots. Let W^*_j be the set of slots that are fractionally allocated to users in S^*_j. The following simple claim states that the slots in W^*_j cannot be allocated to users outside the set S^*_j. That is, each user outside the set S^*_j cannot transmit over the slots in W^*_j.

Lemma 1: Any slot in W^*_j cannot be allocated to users outside S^*_j. Therefore, each slot in W^*_j is fully allocated to only users in S^*_j.

Proof: Assume to the contrary that there exists a slot $w \in W^*_j$ that can be allocated to some user that is not in S^*_j. This user is allocated strictly less than $T^*_{(j)}$. By the definition of W^*_j there exists a user in S^*_j who is allocated a positive fraction w. Therefore, it is possible to decrease w by some $\epsilon > 0$ and allocate ϵ to a user outside S^*_j. This contradicts the fact that the allocation is a max-min fair allocation. ■

We prove another useful lemma connecting a fractional max-min fair allocation that allocates the slots fractionally to users and an integral allocation that can only allocate slots in an indivisible way.

Lemma 2: Let $\mathbf{T}^*$ be a fractional max-min allocation. Then there exists an integral allocation of the slots such that each user j gets at least $\lfloor T^*_j \rfloor$ slots. Furthermore, for any j there exists an integral allocation of the slots in W^*_j to the users in S^*_j that allocates to each user in S^*_j exactly $\lfloor T^*_{(j)} \rfloor$ slots.

Proof: We prove the first part of the claim. Given the values T^*_j, we construct a flow network between the users and the slots by adding a super-source s connected to all the users with infinite capacity, and a super-sink t with directed edges from all slots having unit capacity. We add an edge between each user j and slot w if the user can transmit in this slot. The capacity of these edges is also infinity. It can easily be verified that any flow in this graph immediately translates to a feasible (fractional) allocation of slots to the users and vice versa. Next, we put a lower bound capacity of $\lfloor T^*_j \rfloor$ on the edge to user j which bounds from below the amount of flow to user j. Since the max-min fair allocation satisfies these lower bounds we know that there exists a feasible flow that satisfies all upper and lower bounds on the capacities. Next, we compute a maximum flow in this graph. It is well known that for any flow problem, if there is a feasible fractional solution, then there is also a feasible integral solution, assuming capacities (lower and upper bounds) are integral. Thus, we get an integral feasible flow that satisfies the constraints.

To get the the second claim, we build a subgraph containing only the users in S^*_j and the slots in W^*_j. We then remove integral flow from users that got more than $\lfloor T^*_{(j)} \rfloor$ units of flow until each user gets exactly that amount. ■

B. The Online Algorithm

We now present our online algorithm for allocating slots to users and analyze its performance.

Algorithm Balance: *When a new time slot is available, allocate it to a user j that can transmit in this slot and has gotten the least number of slots so far. Break ties arbitrarily.*

The power allocation of user j is done as follows. Assume slot i is allocated to user j. At this point, j must determine the amount of power to be invested in slot i. We proceed as described in [2]. The algorithm guesses the length t_j of the sequence, starting from 1, and doubling it each time the current length of the online sequence turns out to be longer than the guess. In case the channel gain is a fixed value h_0, the power allocation is determined as follows. For a sequence length t, the algorithm invests in each time slot a power equal to $(\frac{P \cdot h_0}{h_0 \cdot t \cdot c})^{1/2}$, where $c = \frac{2}{(\sqrt{2}-1)^2}$. That is, when at time t the algorithm realizes that its current guess is wrong, it updates its guess to be $2t$, and works with this value until time $2t$. At time $2t$ the algorithm will again update its guess to $4t$, etc. The algorithm continues to invest some power until the sequence length becomes longer than P. After that point, the algorithm no longer invests any power, and thus does not make any additional profit.

Analysis. For each user j, let T^B_j be the number of slots allocated to the user in the online assignment. As in the allocation of the optimal solution, we sort the users by the number of slots allocated to them. Let $T^B_{(j)}$ be the number of slots the jth highest user received in the allocation of Algorithm Balance. Below is a technical lemma that proves a lower bound on the number of slots the online assignment allocates inside any subset S^*_j.

Lemma 3: For every j:

$$\sum_{k \in S^*_j} \min\{T^B_k, \lfloor T^*_{(j)} \rfloor\} \geq \frac{1}{2} |S^*_j| \lfloor T^*_{(j)} \rfloor.$$

Proof: By Lemma 2 there exists an integral allocation of a subset of the slots in W^*_j that allocates to each user exactly $\lfloor T^*_{(j)} \rfloor$ slots. For each user $k \in S^*_j$ let $W^*_{j,k}$ be the slots that were allocated to this user in this allocation (so $|W^*_{j,k}| = \lfloor T^*_{(j)} \rfloor$). For each user $k \in S^*_j$ let T^B_k be the number of slots allocated to user k by the online algorithm.

Next, we define a subset $W'^*_{j,k} \subseteq W^*_{j,k}$. A slot $w \in W^*_{j,k}$ belongs to $W'^*_{j,k}$ if it was allocated by the online assignment to a user which had at least $\lfloor T^*_{(j)} \rfloor$ at the time of the allocation. Our claim is that $\min\{T^B_k, \lfloor T^*_{(j)} \rfloor\} \geq |W'^*_{j,k}|$.

If $|W'^*_{j,k}| = 0$, the claim holds trivially. Otherwise, there exists a slot in $W^*_{j,k}$ that was allocated by the online algorithm to some user that had at least $\lfloor T^*_{(j)} \rfloor$ slots (at the time of allocation). In this case the slot could be allocated to user k and so it must be that user k also already received at least $\lfloor T^*_{(j)} \rfloor$ slots. Therefore, $T^B_k \geq \lfloor T^*_{(j)} \rfloor$. Since $|W'^*_{j,k}| \leq |W^*_{j,k}| = \lfloor T^*_{(j)} \rfloor$, the claim holds.

This means that:

$$\sum_{k \in S^*_j} \min\{T^B_k, \lfloor T^*_{(j)} \rfloor\} \geq \sum_{k \in S^*_j} |W'^*_{j,k}|. \tag{5}$$

We also claim that,

$$\sum_{k\in S_j^*}\min\{T_k^B,\lfloor T_{(j)}^*\rfloor\}\geq\sum_{k\in S_j^*}\left(|W_{j,k}^*|-|W'^*_{j,k}|\right). \quad (6)$$

This follows since each slot in $W_{j,k}^*$ that is not in $W'^*_{j,k}$ was allocated by the online assignment to some user that had strictly less than $\lfloor T_{(j)}^*\rfloor$ slots at the time of the assignment. Since by Lemma 1 the slots in W^* can only be allocated to users in S_j^* then this slot could only be allocated to users inside S_j^*. The LHS actually counts all slots that was allocated to users in S_j^* when the users had less than $\lfloor T_{(j)}^*\rfloor$ therefore such a slot is also counted in the LHS and the inequality holds. Summing up Inequality (5) and Inequality (6) we get that:

$$2\cdot\sum_{k\in S_j^*}\min\{T_k^B,\lfloor T_{(j)}^*\rfloor\}\geq\sum_{k\in S_j^*}|W_{j,k}^*|=|S_j^*|\cdot\lfloor T_{(j)}^*\rfloor$$

which proves the claim. ■

Next, we prove that the number of slots the jth highest user received in the online allocation is not too small. Specifically, we make a connection between the number of slots the jth highest user receives in our allocation and the number of slots the $4j$th highest user got in an optimal max-min fair allocation.

Lemma 4: For any $1\leq j\leq n$: $T_{(j)}^B\geq\frac{\lfloor T_{(4j)}^*\rfloor}{3}$, where for each $j>n$, $T_{(j)}^*=0$.

Proof: Consider some value j. Consider the set of users in S_{4j}^*. By Lemma 3 we know that:

$$\sum_{k\in S_{4j}^*}\min\{T_k^B,\lfloor T_{(4j)}^*\rfloor\}\geq\frac{1}{2}|S_{4j}^*|\lfloor T_{(4j)}^*\rfloor.$$

Consider some $1\leq r\leq|S_{4j}^*|$ and let $T_{(r)}^B$ be the user that received the rth highest number of slots in the online allocation out of the set of users in S_{4j}^*. We get that:

$$\begin{aligned}&(r-1)\lfloor T_{(4j)}^*\rfloor+\left(|S_{4j}^*|-r+1\right)T_{(r)}^B\\&\geq\sum_{k\in S_{4j}^*}\min\{T_k^B,\lfloor T_{(4j)}^*\rfloor\}\geq\frac{1}{2}|S_{4j}^*|\lfloor T_{(4j)}^*\rfloor.\end{aligned}$$

Setting $r=|S_{4j}^*|/4+1\geq j$ we get that:

$$\frac{|S_{4j}^*|}{4}\lfloor T_{(4j)}^*\rfloor+\frac{3|S_{4j}^*|}{4}T_{(|S_{4j}^*|/4)}^B\geq\frac{1}{2}|S_{4j}^*|\lfloor T_{(4j)}^*\rfloor.$$

Simplifying, we get: $T_{(|S_{4j}^*|/4)}^B\geq\frac{1}{3}\lfloor T_{(4j)}^*\rfloor$. Since $|S_{4j}^*|/4+1\geq j$ we get that $T_{(j)}^B$ is at least that amount. Note that we consider here the user that received the jth highest number of slots in the online allocation out of the set of users in S_{4j}^*. Thus, the user that received the jth highest number of slots in the online allocation out of all users can only be better. ■

Finally, the next lemma bounds from below the optimal solution that can be obtained if the users split their budget equally using the allocation of the online algorithm.

Lemma 5:

$$\sum_{j=1}^{n}T_{(j)}^B\cdot\log\left(1+\frac{P\cdot h_0}{T_{(j)}^B}\right)\geq\frac{1}{48}\cdot\sum_{j=4}^{n}T_{(j)}^*\cdot\log\left(1+\frac{P\cdot h_0}{T_{(j)}^*}\right),$$

where T^* is the integral optimal max-min fair allocation.

Proof: The proof follows from the following:

$$\begin{aligned}&\sum_{j=1}^{n}T_{(j)}^B\cdot\log(1+\frac{P\cdot h_0}{T_{(j)}^B}) && (7)\\
\geq\;& T_{(1)}^B\cdot\log(1+\frac{P\cdot h_0}{T_{(1)}^B})+\sum_{j=2}^{n}\frac{\lfloor T_{(4j)}^*\rfloor}{3}\cdot\log(1+\frac{3P\cdot h_0}{\lfloor T_{(4j)}^*\rfloor}) && (8)\\
\geq\;& T_{(1)}^B\cdot\log(1+\frac{P\cdot h_0}{T_{(1)}^B})+\frac{1}{3}\sum_{j=2}^{n}\lfloor T_{(4j)}^*\rfloor\cdot\log(1+\frac{P\cdot h_0}{\lfloor T_{(4j)}^*\rfloor})\\
\geq\;& T_{(1)}^*\cdot\log(1+\frac{P\cdot h_0}{T_{(1)}^*})+\frac{1}{3}\sum_{j=2}^{n}\lfloor T_{(4j)}^*\rfloor\cdot\log(1+\frac{P\cdot h_0}{\lfloor T_{(4j)}^*\rfloor}) && (9)\\
\geq\;& T_{(1)}^*\cdot\log(1+\frac{P\cdot h_0}{T_{(1)}^*})+\frac{1}{12}\sum_{j=8}^{n}\lfloor T_{(j)}^*\rfloor\cdot\log(1+\frac{P\cdot h_0}{\lfloor T_{(j)}^*\rfloor}) && (10)\\
\geq\;& \frac{1}{12}\sum_{j=1}^{n}\lfloor T_{(j)}^*\rfloor\cdot\log(1+\frac{P\cdot h_0}{\lfloor T_{(j)}^*\rfloor}). && (11)\end{aligned}$$

Inequality (8) follows by applying Lemma 4. Inequality (9) follows since the max-min fair allocation minimizes the largest coordinate in the allocation vector. Inequalities (10) and (11) follow by the fact that $T_{(j+1)}^*\leq T_{(j)}^*$, and thus $(T_{(j+1)}^*+T_{(j+2)}^*+T_{(j+3)}^*+T_{(j+4)}^*)/4\leq T_{(j)}^*$.

The claim almost proves the lemma. The only problem is that we take the floor of the values $T_{(j)}^*$. This is not a problem if $T_{(j)}^*\geq 1$, but it is problematic in case $T_{(j)}^*<1$. For the case where $T_{(j)}^*\geq 1$, the following equation holds:

$$\begin{aligned}&\sum_{j=1}^{n}T_{(j)}^B\cdot\log(1+\frac{P\cdot h_0}{T_{(j)}^B})\\
&\geq\frac{1}{12}\sum_{j=1}^{n}\lfloor T_{(j)}^*\rfloor\cdot\log(1+\frac{P\cdot h_0}{\lfloor T_{(j)}^*\rfloor})\\
&\geq\frac{1}{24}\sum_{j|T_{(j)}^*\geq 1}T_{(j)}^*\cdot\log(1+\frac{P\cdot h_0}{T_{(j)}^*}). \quad (12)\end{aligned}$$

Next, consider the group of users that got strictly less than one slot in a fractional max-min fair allocation. We assume that each user got strictly more than 0, since otherwise it cannot transmit at all and we can remove it from the solution. We remark that a scenario in which the gain of such users dominates the gain of the optimum is not very interesting, since it means in general that the number of slots is much smaller than the number of users. Nevertheless, we analyze this case for completeness.

Suppose there are k users that were allocated less than one slot by the fractional max-min allocation. Consider the set of slots that were assigned to them. The number of such slots is $j\leq k$, and none of the other slots can go to these users. Therefore, the total profit gained from these users in an integral allocation is at most $j\log(1+P\cdot h_0)$, and there is an integral allocation that gives 1 slot to j out of the k users. Next, applying the proof of Lemma 3 on these slots, we get that the number of users that got at least 1 slot in the online allocation is at least $j/2$. Therefore, the online algorithm gains at least $\frac{j}{2}\log(1+P\cdot h_0)$. If the gain of the optimal solution from users with $T^*_{(j)}\geq 1$ is more than $1/2$ of the total gain then the total gain of the online solution is at least $1/48$ of the optimum (by Inequality (12)). Otherwise, we are at least $1/4$ competitive, and we are done. ■

We are now ready to prove our main theorem. Before doing so we state a previous result from [2]. In [2], the online power allocation problem of each individual user was analyzed, and it was proved that each user is able to obtain at least a constant fraction of its optimal value. We use this fact along with the previous lemma to establish the following theorem.

Theorem 2: The competitive ratio of Algorithm Balance is $O(1)$.

Proof: Given the slots allocation $\mathbf{T}^B$ of Algorithm Balance, we denote by $\Pi^B_j(x)$ the profit achieved by user j in Algorithm Balance, and by $\Pi^*_j(x)$ the profit achieved by the optimal power allocation of user j, given an allocation of x slots to user j. Recall that T^B_j is the number of slots that Algorithm Balance allocated to user j, and T^*_j is the number of slots allocated to j in the optimal solution. We denote by Π^B the total profit of Algorithm Balance, and by Π^* the total profit of the optimal solution. Then, it holds that

$$\Pi^B = \sum_{j=1}^{n}\Pi^B_j(T^B_j)\geq\sum_{j=1}^{n}\frac{\Pi^*_j(T^B_j)}{c} \qquad (13)$$

$$\geq \frac{1}{48c}\sum_{j=1}^{n}\Pi^*_j(T^*_j)=O(1)\cdot\Pi^*. \qquad (14)$$

The value c is the competitive ratio proven for each user in [2]. Inequality (13) follows from [2], and inequality (14) follows from Lemma 5. ■

C. Arbitrary Power Budgets

Our online algorithm for the fixed channel gain problem, where the power budgets of the users can be arbitrary, is a weighted version of Algorithm Balance, proceeding as follows. Denote by $T^r_j(i-)$ the ratio of the number of slots allocated to user j up to (not including) slot i and its budget.

Algorithm Weighted Balance: When a new time slot is available, allocate the slot to a user j that can transmit in this slot and minimizes $T^r_j(i-)+(1/P_j)$. Break ties arbitrarily. Each user then allocates its power over the slots it is allocated to in a similar way as described in Section III-B.

It is easily seen that the lemmas in Sections III-A and III-B hold for the case of arbitrary budgets, using a simple reduction that splits each user to several "users" with equal budgets. Note that we assume here that all budgets are products of the same factor. We thus get that the competitive ratio of algorithm Weighted Balance is $O(1)$.

IV. The General Case

In this section we design a randomized algorithm for the general case where in each time slot i, each user j is given a channel gain h^i_j within the range $[h_{\min},h_{\max}]$. For ease of notation, we denote $h=h_{\min}$ and $H=h_{\max}$. We partition the range of gain values $[h,H]$ into $\log\frac{H}{h}$ levels, where the ℓth level contains gain values in the range $[2^{\ell-1}h,2^{\ell}h)$, and $\ell\in\{1,2,3,\ldots,\lceil\log\frac{H}{h}\rceil\}$. For each gain sequence $\mathbf{h}^i$ given at time slot i, we refer to users for which the gain value h^i_j is in the range $[2^{\ell-1}h,2^{\ell}h)$ as belonging to the same level.

The general idea of our algorithm is as follows. We choose uniformly in random a gain level δ from the range $\{1,2,3,\ldots,\lceil\log\frac{H}{h}\rceil\}$. For each slot i, we consider only the users belonging to level δ, and ignore all other users in the sequence $\mathbf{h}^i$. The sets of users considered for each slot define an instance of the power allocation problem where the ratio between the maximum and minimum possible gain values is at most 2, that is, the channel gain value is nearly fixed.

Our randomized algorithm thus works as follows. After choosing at random a gain level δ, we reduce our problem to the fixed channel gain case by ascribing all values in the range of level δ to the lower gain value $2^{\delta-1}h$. This means that the competitive ratio of the fixed channel gain drops by at most a factor of 2. Then, we apply Algorithm Balance to the set of users belonging to level δ in each slot.

Theorem 3: Given a γ-competitive online algorithm to the power allocation problem with fixed channel gain value, the expected competitive ratio the randomized algorithm is $O(\gamma\log\frac{H}{h})$.

Proof: We denote by $\mathcal{OPT}$ the optimal off-line algorithm for the general multi-user power allocation problem, and by OPT its profit. For each level ℓ, we denote by OPT_ℓ the profit obtained by $\mathcal{OPT}$ from users transmitting over a channel with gain values in the range $[2^{\ell-1}h,2^{\ell}h)$. Let $\Pi(\mathcal{OPT}_\ell)$ be the profit of the optimal off-line algorithm when applied to the ℓth level instance only (ignoring in each slot users belonging to other levels).

We denote $L=\log\frac{H}{h}$. As the level δ is chosen uniformly in random, the expected profit Π of our algorithm is bounded as follows:

$$\Pi\geq\sum_{\ell=1}^{\lceil\log\frac{H}{h}\rceil}\frac{1}{L}\cdot\frac{\Pi(\mathcal{OPT}_\ell)}{\gamma}\geq\sum_{\ell=1}^{\lceil\log\frac{H}{h}\rceil}\frac{1}{L}\cdot\frac{OPT_\ell}{\gamma}=\frac{OPT}{\gamma L}.$$

■

Since Algorithm Balance is a constant-competitive online algorithm for the fixed channel gain, we get a randomized algorithm with competitive ratio $O(\log\frac{H}{h})$ for the general online multi-user power allocation problem.

Theorem 4: Our online algorithm is $O(\log\frac{H}{h})$-competitive for the general multi-user power allocation problem.

V. The Local Information Case

We turn to examine the local information case, where at each time slot i, each transmitter j is aware only of its own gain value h^i_j. We present a distributed framework, where each user can learn the state of the system, and save it as local information. Each user takes an individual transmission decision, based on its observation of the state of the system and on the full knowledge of the algorithm.

Initially, all users agree on a level δ from the range $\{1,2,3,\ldots,\lceil\log\frac{H}{h}\rceil\}$. For each slot i, only users with channel gain values belonging to level δ (that is, gain values in the range $[2^{\delta-1}h, 2^{\delta}h)$) are candidates for transmission. We denote this set of users by S_i. In each slot, only a single user, out of this set, performs a transmission. We show how the identity of the transmitting user is set in each slot, in a fully distributed manner.

Each user is assigned with a different ID. During the initial state of the system, each user transmits its ID along with its initial power. We assume these transmissions are performed sequentially, so that all users get and save the data. At the beginning of each slot i, we allocate a small time interval σ^i. Using time division multiplexing, σ^i is divided into n time fractions, where time fraction σ^i_j is assigned to user j. Now, we add the following rule to our online algorithm. At the beginning of each slot i, each user j transmits during time fraction σ^i_j its ID along with a binary value TRUE/FALSE indicating whether its gain value h^i_j belongs to level δ (that is, whether $j \in F^i$). As it is the only user transmitting at this time, all other users receive its transmission and save the data.

Following this rule, each user can keep track of the number of slots allocated to the other users. That is, before slot i, each user j knows the total number of slots $T_m(i-)$ that were allocated to each user $m \neq j$ up to slot i. As after σ^i, each user knows which are the users in set S_i, then each user also knows which is the user $j \in F^i$ that minimizes $T^r_j(i-)+1/P_j$. Only this user will perform the transmission. Ties can be broken by choosing the user with minimal ID.

The appeal of the above scheme is in the overall message complexity, where each user is required to send only a *single* binary signal in each time slot.

VI. Simulations Study

The objective here is to validate our suggested online algorithm, and to examine the effect of certain parameters on its performance. We first describe some heuristic enhancements added to the online algorithm, and explicitly specify the algorithm used in the performed experiments. Then, we turn to describe the experiments and their respective results.

A. Heuristics Enhancements

We add several enhancements to the online algorithms described in Sections III-B and III-C. In each slot i, we denote by $F^i(\ell)$ the set of users with channel gain value belonging to level ℓ. As mentioned, our algorithm chooses at random a level δ, and allows transmissions only to users with gain value belonging to level δ. The main natural improvement is the following. Instead of considering in each slot i the set $F^i(\delta)$, we consider the set $F^i(\geq \delta)$ that contains the users that got in slot i a channel gain value belonging to one of the levels $\{\delta, \delta+1, \ldots, \lceil \log \frac{H}{h} \rceil\}$. That is, $F^i(\geq \delta)$ consists of users j such that $h^i_j \in [2^{\delta-1}h, H]$. Among these users, we allocate slot i to user $j \in F^i(\geq \delta)$ that minimizes $T^r_j(i-)+1/P_j$. In addition, in case there is a slot i where $F^i(\geq \delta) = \emptyset$, then a single user j with highest gain value $h^i_j \in [h, 2^{\delta-1}h)$ is allowed to transmit. We note that following these improvements, our assumption according to which each user is allocated at least two time slots in the optimal solution (see Section III), becomes even more evident.

We also apply the enhancements specified in [2] for the power allocation performed by each user over its allocated slots. Moreover, in practice, different gain values in level δ are not considered as equal, and thus the ratio between different gain values in this level is at most 2. The constant c is thus set to $c = \frac{2\lambda}{(\sqrt{2}-1)^2} = \frac{4}{(\sqrt{2}-1)^2}$. In addition, for a sequence length T_j, user j invests in each time slot a power equal to $(\frac{P_j}{h(\delta)\cdot T_j \cdot c})^{1/2}$, where $h(\delta) = 2^{\delta-1}h$. The algorithm used for our numerical experiments is specified below (Algorithm 1).

Algorithm 1 Online Algorithm for the Multi-User Power Allocation Problem in the Fixed Channel Gain Scenario

Initialization: choose uniformly in random a value δ from $\{1, \ldots, \lceil \log \frac{H}{h} \rceil\}$.
set $\forall j$, $(P_j)' = P_j$, current length guess $T_j = 1$ and current length $\tau_j = 0$.
set $c = \frac{4}{(\sqrt{2}-1)^2}$.

Given a new channel gain sequence $\mathbf{h}^i = (h^i_1, \ldots, h^i_n)$:

$F^i(\geq \delta)$ is the set of users with gain values belonging to levels $\geq \delta$ in slot i;
if $F^i(\geq \delta) \neq \emptyset$ **then**
  choose user $j \in F^i(\geq \delta)$ with minimum ratio $(\tau_j+1)/P_j$, such that $(P_j)' > 0$;
  $k = j$;
else
  if (($F^i(\geq \delta) = \emptyset$) OR (there is no user $j \in F^i(\geq \delta)$ with $(P_j)' > 0$)) **then**
    choose user m with highest gain value h^i_m, such that $(P_m)' > 0$;
    $k = m$;
  end if
end if
power to be invested in current slot is $p^i_k = \left(\frac{P_k}{h(\delta)\cdot T_k \cdot c}\right)^{1/2}$;
invest at current slot i a power of $\min\{(P_k)', p^i_k\}$;
remaining power of user k is $(P_k)' = \max\{0, (P_k)' - p^i_k\}$;
current sequence length of user k is $\tau_k = \tau_k + 1$;
if (length τ_k is equal to length guess T_k) **then**
  double the length guess $T_k = 2T_k$.
end if

B. Experimental Results

We perform simulations comparing an optimal TDMA solution to our online balance algorithm in case of a fixed channel gain. Comparing our algorithm to the optimal solution in the general case would be possible once the latter is fully characterized, which at this stage is an open problem.

We study the effect of several parameters on performance. First, we examine the influence of the power budgets given initially to the users on the throughput ratio between the optimal (off-line) and online algorithms. Accordingly, we fix the number of users (100) and length of time frame (10000). It is seen in Figure 1 that the performance bound stabilizes around 2.2, which is much better than the constant ratio guaranteed by our theoretical analysis (Equation (14)).

Next, we examine the influence of the number of users n on the throughput ratio between the optimal (off-line) and online algorithms. Accordingly, we fix the initial budgets (1000) and length of time frame (10000). Observe that the performance ratio improves with the number of users (Figure 2). The intuitive explanation for this phenomenon is that most of the time slots can be used for useful transmissions, as in the optimal solution.

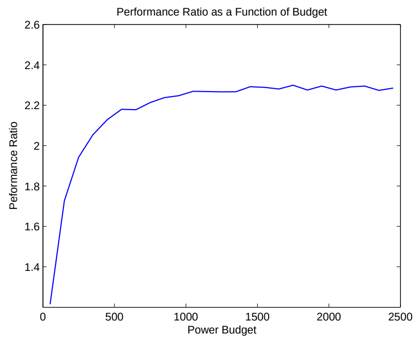

Fig. 1. The ratio between the optimal (offline) power allocation and the online allocation, as a function of the power budget. Results are averaged over 10 runs.

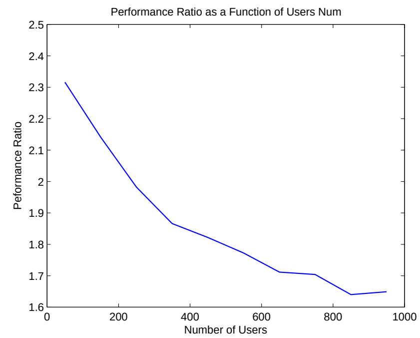

Fig. 2. The ratio between the optimal (offline) power allocation and the online allocation, as a function of the number of users. Results are averaged over 10 runs.

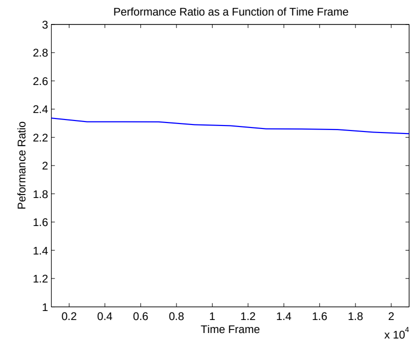

Fig. 3. The ratio between the optimal (offline) power allocation and the online allocation, as a function of the time frame length. Results are averaged over 10 runs.

Finally, we examine the influence of the time frame length T on the throughput ratio between the optimal (off-line) and online algorithms. Accordingly, we fix initial budgets (1000) and number of users (100). Observe that the performance ratio decreases slightly between 2.4 and 2.2, i.e., is almost fixed (Figure 3). This may follow from the concavity of the logarithmic objective function, reaching saturation from relatively low values of T (given that the total powers are held fixed).

Summarizing our results, we have observed that the online algorithm performs even better than the theoretical guarantees. In addition, it seems that performance is mostly affected by the number of users, but stays almost fixed as a function of the other system parameters.

VII. Conclusion

We considered the problem of power allocation under dynamic channel quality, in a multi-transmitter environment, within the framework of online computation. We addressed both a “fixed gain” case, where transmitters are provided with binary information on the channel quality, and the general case, where the precise gain values are provided. For both cases, we established online power-allocation algorithms with proven worst-case performance bounds. We then designed a distributed scheme that allows to implement these algorithms in a practical setting of local information. We complemented our work with a simulation study, where we validated our suggested online algorithm for the fixed channel gain case, and observed that our online algorithm performs even better than the theoretical bound.

Our framework can be extended in several different aspects. An immediate research direction would be to consider general network topologies (for example, multiple users transmitting to multiple base-stations); instead of having per-user gains as in our case, the gains would in general correspond to the instantaneous pairwise interference between any two users (hence involving a larger parameter space). Another interesting topic is to go beyond the CDMA-like cell studied here, and consider Information-Theoretic bounds and the corresponding capacity region under arbitrary fading channels.

At a higher level, it is of great interest to consider our model under a *noncooperative* framework, where mobiles adjust their transmission power in order to optimize their individual rate. From a game-theoretic perspective, the combination of multiple time periods, arbitrary channel conditions, and possibly incomplete information (e.g., regarding the gains of other users or their budget), creates a very complicated and challenging domain, whose analysis may require novel tools and solution concepts.

References

[1] E. Altman, K. Avrachenkov and A. Garnaev, “A jamming game in wireless networks with transmission cost”, *Network Control and Optimization, 1st EuroFGI International Conf., NET-COOP*, pp. 1–12, 2007.

[2] N. Buchbinder, L. Lewin-Eytan, I. Menache, S. Naor and A. Orda, “Dynamic power allocation under arbitrary varying channels – an online approach”, *Proc. of INFOCOM*, 2009.

[3] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis*. Cambridge University Press, 1998.

[4] A. Lapidoth and P. Narayan, “Reliable communication under channel uncertainty,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2148–2177, 1998.

[5] N. Nisan, T. Roughgarden, E. Tardos, and Vijay V. Vazirani, Algorithmic Game Theory, Cambridge University Press, 2007.

[6] E. Biglieri, J. Proakis, and S. Shamai, “Fading channels: information-theoretic and communications aspects,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2619–2692, 1998.

[7] R. Etkin, A. Parekh, and D. Tse. Spectrum sharing for unlicensed bands. *IEEE J. on Selected Areas in Communications*, 25(3):517–528, 2007.

[8] S. Hayashi and Z. Q. Luo. Spectrum management for interference-limited multiuser communication systems. *IEEE Transactions on Information Theory*, 55(3):1153–1175, 2009.

[9] J. Huang, R.A. Berry, and M.L. Honig. Spectrum sharing with distributed interference compensation. In *New Frontiers in Dynamic Spectrum Access Networks, DySPAN*, pages 88–93, 2005.

[10] Z. Q. Luo and J. S. Pang. Analysis of iterative waterfilling algorithm for multiuser power control in digital subscriber lines. *EURASIP Journal on Applied Signal Processing*, pages 1–10, 2009.

[11] G. Scutari, D. P. Palomar, and S. Barbarossa. Asynchronous iterative water-filling for gaussian frequency-selective interference channels. *IEEE Transactions on Information Theory*, 54(7):2868–2878, 2008.

Acknowledgement

The research of Seffi Naor is supported by ISF grant 1366/07. The research of Ishai Menache was supported by a Marie Curie International Fellowship within the 7th European Community Framework Programme.