On Tractability Aspects of Optimal Resource Allocation in OFDMA Systems

Di Yuan, Jingon Joung, Chin Keong Ho and Sumei Sun

Linköping University Post Print

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

On Tractability Aspects of Optimal Resource Allocation in OFDMA Systems

Di Yuan, *Member, IEEE,* Jingon Joung, *Member, IEEE,* Chin Keong Ho, *Member, IEEE,* and
Sumei Sun, *Senior Member, IEEE*

Abstract

Joint channel and rate allocation with power minimization in orthogonal frequency-division multiple access (OFDMA) has attracted extensive attention. Most of the research has dealt with the development of sub-optimal but low-complexity algorithms. In this paper, the contributions comprise new insights from revisiting tractability aspects of computing the optimum solution. Previous complexity analyses have been limited by assumptions of fixed power on each subcarrier, or power-rate functions that locally grow arbitrarily fast. The analysis under the former assumption does not generalize to problem tractability with variable power, whereas the latter assumption prohibits the result from being applicable to well-behaved power-rate functions. As the first contribution, we overcome the previous limitations by rigorously proving the problem's NP-hardness for the representative logarithmic rate function. Next, we extend the proof to reach a much stronger result, namely that the problem remains NP-hard, even if the channels allocated to each user is restricted to a consecutive block with given size. We also prove that, under these restrictions, there is a special case with polynomial-time tractability. Then, we treat the problem class where the channels can be partitioned into an arbitrarily large but constant number of groups, each having uniform gain for every individual user. For this problem class, we present a polynomial-time algorithm and provide its optimality guarantee. In addition, we prove that the recognition of this class is polynomial-time solvable.

Index Terms

orthogonal frequency-division multiple access, resource allocation, tractability.

I. Introduction

In orthogonal frequency division multiple access (OFDMA) systems, resource allocation amounts to finding the optimal assignment of subcarriers to users, and, for each user, the allocation of rate or power over the assigned subset of subcarriers. In this paper, we focus on the problem of minimizing the total transmit power, subject to delivering specified data rates, by power assignment and subcarrier allocation. The popularity of OFDMA for wireless communications has led to an intense research effort in this area. Most previous work has focused on heuristic and thus sub-optimal solutions, e.g., [1]–[8], and the references therein. For example, in [1], a two-stage algorithm is proposed. First, the bandwidth-assignment-based-on-SNR (BABS) algorithm decides the number of subcarriers that each user gets. Second, the rate-craving-greedy (RCG) algorithm or the amplitude-craving-greedy (ACG) algorithm is used to assign specific subcarriers and power levels. The RCG and ACG algorithms are greedy approaches that assign subcarriers based on the maximum rate and the maximum effective channel, respectively. In [2], the successive user integration (SUSI) algorithm is proposed which proceeds in two stages. After selecting an initial candidate, the SUSI algorithm performs a local search algorithm, i.e., it searches around neighboring solution candidates, by continuously exploring if there is a net decrease in the transmission power by re-assigning one subcarrier from user to another user. The algorithm stops when a local optimum solution is reached. The complexity of the SUSI algorithm depends on the number of explorations and thus cannot be easily pre-determined nor bounded. For global optimality, a branch-and-bound algorithm is developed in [9].

Di Yuan is with Department of Science and Technology, Linköping University, Sweden (e-mail: diyua@itn.liu.se).

Jingon Joung, Chin Keong Ho, and Sumei Sun are with Institute for Infocomm Research (I^2R), A*STAR, Singapore (e-mails: jgjoung; hock; sunsm}@i2r.a-star.edu.sg).

We address a complementary but fundamental aspect of the resource allocation problem: To what extent is it tractable? In contrast to the significant amount of research on algorithms, investigations along the line of tractability analysis have been few [10], [11]. The edge of fundamental understanding of problem tractability is formed in respect of the following limitations. In [10], the power on each subcarrier is assumed to be given, i.e., it can be arbitrarily set to any fixed value for the purpose of proving complexity. Thus the result does not apply to the problem where powers are optimization variables. In [11], the resource allocation problem is shown to be NP-hard for one particular type of rate function $r_{mn}(P) \in \mathcal{C}_{\text{inc}}$, where $r_{mn}(P)$ is the rate as a function of power P for the mth user and nth subcarrier, and $\mathcal{C}_{\text{inc}}$ is the set of all increasing functions such that $r_{mn}(0) = 0$. By this result, there exists some (but possibly ill-behaved[1]) function in $\mathcal{C}_{\text{inc}}$ for which the problem is NP-hard. However, the result does not carry over to well-behaved subclasses of functions in $\mathcal{C}_{\text{inc}}$. In fact, for any $r_{mn}(P) \in \mathcal{C}_{\text{linear}}$, where $\mathcal{C}_{\text{linear}}$ is the set of linear functions, the problem is solvable in polynomial time (see Section IV). Thus, whereas dealing with $\mathcal{C}_{\text{inc}}$ is in general intractable for the problem in question, there exists some subclass in $\mathcal{C}_{\text{inc}}$ that admits global optimality at low complexity.

For the (wide) class of functions in $\mathcal{C}_{\text{inc}}$ but not in $\mathcal{C}_{\text{linear}}$, the tractability of the resource allocation problem remains unknown and thus calls for investigation. The aspect is of most relevance for the class of increasing and concave functions $\mathcal{C}_{\text{concave}}$. Indeed, it is very commonly assumed that the rate function is given by $r_{mn}(P) = \log(1 + g_{mn}P)$, with $g_{mn} > 0$. For the hierarchy $\mathcal{C}_{\text{linear}} \subset \mathcal{C}_{\text{concave}} \subset \mathcal{C}_{\text{inc}}$, the result in [11] applies to $\mathcal{C}_{\text{inc}} \setminus \mathcal{C}_{\text{concave}}$. We investigate the complexity for the representative logarithmic function case of $\mathcal{C}_{\text{concave}}$, and, if the problem is NP-hard for the function, examine to what extent restrictions on input structure (e.g., assuming the number of subcarrier for each user is given in the input) will admit better tractability. The specific contributions are as follows.

- We rigorously prove that for the representative logarithmic rate function $r_{mn}(P) = \log(1 + g_{mn}P)$ with $g_{mn} > 0$, the resource allocation problem is NP-hard. It follows that in the hierarchy $\mathcal{C}_{\text{linear}} \subset \mathcal{C}_{\text{concave}} \subset \mathcal{C}_{\text{inc}}$, the hardness result holds not only for some function in $\mathcal{C}_{\text{inc}} \setminus \mathcal{C}_{\text{concave}}$, but also for at least one representative function in $\mathcal{C}_{\text{concave}}$. Moreover, we prove that the problem is tractable for all functions in $\mathcal{C}_{\text{linear}}$. The findings lead to a refinement of the result of tractability.
- We extend the NP-hardness analysis to arrive at a much stronger result. Namely, the problem remains NP-hard, even if the following two restrictions are *jointly* imposed: 1) the subcarriers allocated to each user form a consecutive block, and 2) the block size is given as part of the problem's input. We also prove that, with these two restrictions, there is a special case admitting polynomial-time tractability.
- We identify a tractable problem subclass, with the structure that the channels can be partitioned into an *arbitrarily many but constant* number of groups, possibly with varying group size, and the channels within each group have uniform gain for each user (but may differ by users). The original problem is, in fact, equivalent to having the number of groups equal to the number of subcarriers. The tractable subclass goes beyond the logarithmic rate function – the result holds as long as the single-user resource allocation is tractable. Moreover, we prove that recognizing this problem class is tractable as well.

The remainder of the paper is organized as follows. The system model is given in Section II. After briefly discussing the significance of this work in Section III, we provide and prove the base result of NP-hardness in Section IV. Section V is devoted to the problem's tractability with restrictions on channel allocation. In Section VI, we consider the problem class with structured channel gain and prove its tractability. In addition, we prove that recognizing this problem class is computable in polynomial time. Conclusions are given in Section VII.

II. SYSTEM MODEL

Consider an OFDMA system with M users and N subcarriers. In this paper, the terms subcarrier and channel are used interchangeably. For convenience, we define sets $\mathcal{M} = \{1, \ldots, M\}$ and $\mathcal{N} = \{1, \ldots, N\}$. Notation $r_{mn}(P_{mn})$ is reserved for the rate as a function of the transmission power $P_{mn} \geq 0$ for the mth user and nth subcarrier. The inverse function r_{mn}^{-1}, returning the power for supporting a given rate, is denoted by f_{mn}. The required rate of user m is denoted by R_m.

---

[1]Indeed, the proof in [11] relies on a power function (i.e., the inverse of the rate function) growing arbitrarily fast for arbitrarily small rate increase, meaning that the function is not locally Lipschitz continuous.

The forthcoming analysis focuses on the representative increasing, concave rate function $r_{mn}(P_{mn}) = \log_2(1 + g_{mn}P_{mn})$, where $g_{mn} > 0$ represents the channel gain (normalized such that the noise variance is one). Hence, the corresponding signal-to-noise ratio (SNR) is $g_{mn}P_{mn}$. Some of the tractability results generalize to $\mathcal{C}_{inc}$; we will explicitly mention the generalization when it applies.

The optimization problem is to minimize the total power by joint channel and rate allocation, such that the users' required rate targets are met. The problem is formulated formally below. Throughout the rest of the paper, we refer to the problem as minimum-power channel allocation (MPCA).

Input: User set $\mathcal{M} = \{1, \ldots, M\}$ and channel set $\mathcal{N} = \{1, \ldots, N\}$ with $M \leq N$, positive channel gain $g_{mn}, m \in \mathcal{M}, n \in \mathcal{N}$, and positive rate targets $R_m, m \in \mathcal{M}$.

Output: A channel partitioning represented by $(\mathcal{N}_1, \ldots, \mathcal{N}_M)$, where $\mathcal{N}_m \subset \mathcal{N}, m \in \mathcal{M}, \mathcal{N}_{m_1} \cap \mathcal{N}_{m_2} = \emptyset$, for all $m_1, m_2 \in \mathcal{M}, m_1 \neq m_2$, and non-negative power $P_{mn}, m \in \mathcal{M}, n \in \mathcal{N}_m$, such that $\sum_{n \in \mathcal{N}_m} \log_2(1 + P_{mn}g_{mn}) \geq R_m, m \in \mathcal{M}$ and the total power $\sum_{m \in \mathcal{M}} \sum_{n \in \mathcal{N}_m} P_{mn}$ is minimized.

By the rate function, there is a unique mapping between rate allocation and power expenditure. In addition, at optimum, the rate may be zero on some of the allocated channels of a user. Given the channel allocation, power optimization is determined by solving the single-user rate allocation problem by water-filling [12] in linear time (e.g., [2]).

The combinatorial nature of MPCA stems from the fact that the users can not share any subcarrier, and the core of problem-solving is channel partitioning. For other types of wireless access, this characterization may not be present. Consider for example the MAC channel model, where transmissions take place concurrently on a common channel. The transmissions are coupled via signal-to-interference-and-noise (SINR) ratio. In this case, there is no channel assignment, and a data rate corresponds to an SINR threshold. Characterizing the achievable SINR performance region from an information theoretical perspective, as well as determining transmission power levels for given SINR thresholds, fall into a problem domain that is commonly known as optimal power control. This topic has been extensively studied. Here, we refer to the seminal results in [13]–[15]. In particular, as power control is defined by a system of linear equations of the power variables, the optimization problem of minimizing the total power subject to SINR threshold is effectively solved by linear programming (LP) with polynomial-time complexity. For OFDMA, in contrast, channel assignment is inherently a discrete-choice problem, for which the tractability aspect warrants a thorough investigation.

MPCA can be modeled by means of the following integer programming formulation. This formulation contains two sets of variables to represent power and channel assignment, respectively. The former is continuous. The latter is binary as channel assignment consists of discrete-choice decisions.

$$
\begin{aligned}
P_{mn} &= \text{Transmission power allocated to user } m \text{ on} \\
&\quad \text{channel } n. \\
x_{mn} &= \begin{cases} 1, & \text{if channel } n \text{ is assigned to user } m, \\ 0, & \text{otherwise.} \end{cases}
\end{aligned}
$$

$$\min \ \sum_{m=1}^{M} \sum_{n=1}^{N} P_{mn} \tag{1a}$$

$$\text{s. t.} \ \sum_{m=1}^{M} x_{mn} \leq 1, \ \forall n \in \mathcal{N}, \tag{1b}$$

$$\sum_{n=1}^{N} \log_2(1 + P_{mn}g_{mn}) \geq R_m, \ \forall m \in \mathcal{M}, \tag{1c}$$

$$P_{mn} \leq P_{mn}^{U} x_{mn}, \ \forall m \in \mathcal{M}, \forall n \in \mathcal{N}, \tag{1d}$$

$$P_{mn} \geq 0, \ \forall m \in \mathcal{M}, \forall n \in \mathcal{N}, \tag{1e}$$

$$x_{mn} \in \{0, 1\}, \ \forall m \in \mathcal{M}, \forall n \in \mathcal{N}. \tag{1f}$$

In the formulation, (1a) represents the objective of minimizing the total power. Constraints (1b) ensure that a

channel is assigned to at most one user. Constraints (1c) model the requirement of satisfying the rate target of the users. Next, inequalities (1d) connect the two sets of variables. In the right-hand side of (1d), P_{mn}^U represents an upper bound on the power of user m on channel n. It suffices to set $P_{mn} = \left(2^{R_m} - 1\right) g_{mn}^{-1}$, i.e., the power required to meet the rate target if channel n is the only channel assigned to user m. In (1d), x_{mn} acts as a switch. If $x_{mn} = 1$, P_{mn} is allowed to be a value being less than or equal to the upper bound, otherwise P_{mn} is forced to be zero.

III. SIGNIFICANCE OF PROBLEM TRACTABILITY

Our work focuses on the computational tractability/complexity of finding the global optimum of MPCA. In this context, we do not restrict ourselves to an exhaustive search with exponential complexity. Tractability means whether or not one shall expect *any* polynomial-time algorithm guaranteeing global optimality. It is known that some combinatorial optimization problems (e.g., bipartite matching) can be solved to global optimality with low complexity (i.e., in polynomial time), without resorting to an exhaustive search. Hence, it is fundamentally important to show if a given problem can be solved in polynomial time without restriction to any specific solution algorithm (even if it appears exceedingly difficult at the first sight). Our key contributions is to prove if a low complexity algorithm can be expected for a given problem scenario. If so, we give such a method along with proving its polynomial-time complexity; if not, we provide a NP-hardness proof, implying that no low-complexity algorithm can be expected regardless of how clever algorithms can be designed, unless P=NP.

To further motivate the investigation of tractability in view of the current literature, two remarks are noteworthy. First, an NP-hard problem may become tractable by imposing restrictions to the structure of its input parameters (e.g., the shape of the objective function). An example is the traveling salesman problem (TSP) having a cost function with the so called Klyaus-matrix structure (meaning that inversed triangular distance inequality holds). In this case, TSP is solvable in polynomial time [16]. As the second remark, an NP-hard problem may become tractable by removing some of its constraints (and hence enlarging the solution space). For example, minimum spanning tree (MST) with constrained node degree is NP-hard, but becomes easily-solved if this constraint is removed. Thus, a problem that is in general hard may become tractable, if a restriction or assumption on input is added (the TSP example), or removed (the MST example). This insight demonstrates the importance of examining to what extent a tractability result is applicable.

MPCA appears intuitively difficult, as it contains binary decision variables for channel allocation. Having binary variables, however, does not provide any evidence of problem hardness, but rather implies that MPCA falls into the domain of combinatorial optimization. Within this domain, some problems do admit polynomial-time algorithms guaranteeing global optimality. Classical examples of well-solved combinatorial optimization problems include shortest path, MST, maximum matching, minimum-cost assignment, and uncapacitated lot-sizing [18], [19]. Moreover, that MPCA can be represented by an integer formulation (Section II) does not imply its hardness either, despite the general NP-hardness of integer programming. Indeed, the aforementioned four problems can all be represented by integer programming formulations, yet they are tractable in polynomial time. Thus, without a thorough investigation, no conclusion can be made on the tractability of MPCA.

The above discussion highlights the significance of studying the tractability of MPCA. Furthermore, for an NP-hard problem, tractability analysis includes identifying if the complexity reduces for input that is more structured than the general case (cf. the tractable case of TSP). For MPCA, for example, we prove that it is in fact polynomially solvable by a reduction to the well-solved minimum-weight bipartite matching problem, if the power-rate function would be linear. Details of our analysis for this case are provided in the next section (see Theorem 3).

To the best of our knowledge, no formal analysis other than the results in [10], [11] is available. The study in [10] formalizes the NP-hardness result with the assumption that the powers on all channels are given. This is equivalent to introducing constraints fixing the power values. By the second remark above, the result does not answer the tractability if these constraints are removed (that is, the original problem with variable power). Indeed, for any linear rate function as well as the problem class in Section VI, the NP-hardness proof in [10] remains valid for fixed power, but these problem classes with variable power are solvable in polynomial time. For the analysis in [11], the proof requires unbounded power growth for arbitrarily small rate increase. Specifically, for the power function f, $f\left(n + \frac{1}{k}\right)$ is k times higher than $f(n)$, for arbitrary positive integers n and k. This

assumption of ill-behaved power-rate function excludes not only $f(x) = 2^x - 1$ (the inverse of the logarithmic rate function), but also all locally Lipschitz continuous functions. Recall that a (not necessarily continuous) function f is locally Lipschitz continuous, if for any x, there exists a small real number ϵ and an arbitrarily large but constant real number D, such that $|f(x + \epsilon) - f(x)| \leq D\epsilon$, that is, the growth of the function is bounded when the change in the input diminishes. Clearly, f is not locally Lipschitz continuous, if $f\left(n + \frac{1}{k}\right)$ increases by a factor k over $f(n)$ for any n and arbitrarily large k. Hence, by the previous remark of the impact of cost function on tractability, the tractability under more well-behaved functions calls for investigation.

In discussion of complexity results of combinatorial optimization problems, the following convention of interpretation is often used. NP-hardness of a combinatorial problem for a function class means that the problem is NP-hard for *at least one* though not necessarily all functions of the class. On the other hand, if a problem is polynomial-time tractable for a function class, then the tractability result must hold for *all* functions of the class. Thus, the result in [11] implies MPCA's hardness for $\mathcal{C}_{\text{inc}}$ (but not $\mathcal{C}_{\text{concave}}$ or $\mathcal{C}_{\text{linear}}$).

In view of the above, the significance of proving the hardness of the representative logarithmic function can be observed from the aforementioned relation $\mathcal{C}_{\text{linear}} \subset \mathcal{C}_{\text{concave}} \subset \mathcal{C}_{\text{inc}}$. By theory of computational complexity, an NP-hardness result carries over from the left to the right in the given sequence of function types, but not in the opposite direction. For example, if we can prove NP-hardness for one selected rate function $r \in \mathcal{C}_{\text{concave}}$, then the problem is also NP-hard for $\mathcal{C}_{\text{inc}}$, because the latter is more general. However, the reverse is not true – the validity of NP-hardness does not carry over from the right to the left. In other words, an NP-hardness result derived for a function $r \in \mathcal{C}_{\text{inc}}$ but not in $\mathcal{C}_{\text{concave}}$, i.e., $r \in \mathcal{C}_{\text{inc}} \setminus \mathcal{C}_{\text{concave}}$, does not prove NP-hardness for functions in $\mathcal{C}_{\text{concave}}$. Thus without further analysis, the tractability for $\mathcal{C}_{\text{concave}}$ remains open. Similarly, a hardness proof for $r \in \mathcal{C}_{\text{concave}} \setminus \mathcal{C}_{\text{linear}}$ (next Section) does not tell the complexity if the function type is restricted to $\mathcal{C}_{\text{linear}}$. As was mentioned earlier, we will rigorously prove that MPCA is polynomial-time solvable for $\mathcal{C}_{\text{linear}}$.

As is detailed in the above discussion of [11], the hardness result in the reference is derived for a function shape that is not locally Lipschitz continuous, that is, a function in $\mathcal{C}_{\text{inc}} \setminus \mathcal{C}_{\text{concave}}$. Therefore the result does not apply to more structured/restricted function types ($\mathcal{C}_{\text{concave}}$ or $\mathcal{C}_{\text{linear}}$). This observation shows the significance of our analysis in contributing to sharpening the tractability boundary for MPCA.

IV. TRACTABILITY OF MPCA: BASE RESULTS

We provide the tractability results that overcome the limitations of the currently available analysis in two aspects. First, we present a rigorous proof of the problem's NP-hardness with the representative logarithmic rate function. Second, in the next section, we extend the proof to reach a much stronger result, stating that the problem remains NP-hard even with two strong restrictions on channel allocation.

Theorem 1. *MPCA, as defined in Section II, is NP-hard.*

Proof: As the proof is mathematically intricate and its length requires the use of several lemmas, we outline the basic idea here and defer the details to Appendix A. The proof uses a reduction from the well-known 3-satisfiability (3-SAT) problem. Two groups of users are defined. At optimum, each user in the first group either uses one channel of superior channel gain, or splits the rate on three inferior channels, but not both. This corresponds to the true/false value assignment in 3-SAT. We rigorously prove that the optimal power for this group of users is a constant, while the optimal power for the second user group gives the correct answer to 3-SAT. ∎

For MPCA, the solution output consists of binary variables (for channel allocation) as well as continuous variables (for power assignment). In complexity theory of combinatorial optimization [17], finding the exact optimal variable values is referred to as the problem's *optimization version*. To conclude the NP-hardness of a problem, the standard method is to consider its *decision version* instead of the optimization version, and to prove that the decision version is NP-complete. By complexity theory [17], [21], an optimization problem is NP-hard, if its decision version is NP-complete. Therefore Appendix A addresses the decision version of MPCA. For decision problems, the solution output is a yes/no answer. For MPCA, the decision version amounts to determining the existence of a channel and power allocation, such that the total power does not exceed a given value. Note that the decision version is equivalent to the optimization version, as far as polynomial-time tractability is concerned – If the decision version can be solved in polynomial time, then the optimization version can be solved in
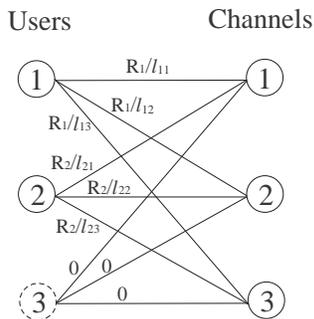
Fig. 1: An illustration of reduction of MPCA with linear rate function to minimum-weight bipartite matching for $M = 2$ (with user three as a dummy user) and $N = 3$.

polynomial time as well via bi-section search. For further details, we refer to two classical books in the area [17], [21].

Corollary 2. *MPCA remains NP-hard, even if the rate requirements of the users are uniform.*

Proof: Follows immediately from the equal-rate values used in the proof of Theorem 1. ∎

Earlier in this section, it was claimed that MPCA with any linear rate function is tractable. Even if this case is not much of practical interest, it is instructive in showing the importance of input assumption on problem tractability.

Theorem 3. *MPCA with linear rate function $r_{mn}(P_{mn}) = l_{mn}P_{mn}$, where $l_{mn} \geq 0$, $m \in \mathcal{M}, n \in \mathcal{N}$, is solvable in polynomial time.*

Proof: Since the rate (and hence power) function is linear, it follows that for any user m, it is optimal to allocate the entire rate R_m to a single channel. Specifically, denoting by $\mathcal{N}_m$ the channel set allocated to m, the minimum power required for rate R_m is given by $\min_{n \in \mathcal{N}_m} R_m/l_{mn}$. Because of this structure at optimum, MPCA reduces to pairing the M users with M out of the N channels. Hence the problem is equivalent to a minimum-weight bipartite matching problem (also known as minimum-cost assignment [18]) in a graph with node sets $\mathcal{M} \cup \{M+1, \dots, N\}$ and $\mathcal{N}$; the former represents the augmentation of $\mathcal{M}$ by $M - N$ artificial users. For edge (m, n), with $m \in \mathcal{M}$ and $n \in \mathcal{N}$, the cost is R_m/l_{mn}. All edges adjacent to artificial users have zero cost. An illustration of the transformation for $M = 2$ and $N = 3$ is provided in Fig. 1. By construction, the size of the graph is of $O(N^2)$. Because minimum-weight bipartite matching is polynomial-time solvable, the theorem follows. ∎

Remark: In light of Theorem 1, Corollary 2, and Theorem 3, we remark on the significance of assumption of input on tractability, by revisiting the result provided in [10]. In this reference, the result is proven for fixed power, that is, the power of each channel is part of problem input. Under this condition, [10] provides an elegant hardness proof of a reduction from the number partitioning problem, by setting specific power values on the channels. As long as power is fixed, a line-by-line copy of the proof in [10] remains valid even if the underlying MPCA rate function is linear. A similar observation applies to the tractable problem class that will be detailed in Section VI. In conclusion, for the original MPCA problem having power allocation as part of the output (for which the key assumption of [10] does not apply), our analysis provides new insights in the tractability rather than contradicting the previous results. □

V. TRACTABILITY WITH RESTRICTIONS ON CHANNEL ALLOCATION

Consider imposing jointly two restrictions to channel allocation. First, the number of channels to be allocated to each user is given. Tractability under this restriction is of significance to the two-phase OFDMA resource allocation (see [1]) that determines the number of channels per user in phase one, followed by channel allocation in phase two. The second restriction is the use of consecutive channels, that is, the channels of every user must be consecutive in the sequence $1, \dots, N$; this channel-adjacency constraint has been considered in, for example,

[20]. We prove that MPCA remains NP-hard even with these two seemingly strong restrictions, although there is a special case admitting polynomial-time tractability.

Theorem 4. *MPCA remains NP-hard, even if the number of channels allocated to each user, i.e., the cardinality of $\mathcal{N}_m, m = 1, \ldots, M$, is given in the input, and $\mathcal{N}_m, m = 1, \ldots, M$ must contain consecutive elements in the channel sequence $1, \ldots, N$.*

Proof: As we prove the result via an augmentation of reduction proof of in Appendix A, the details are deferred to Appendix B. In brief, the proof is built upon introducing a set of additional channels in the MPCA instance in Appendix A, in a way such that, at optimum, the channels allocated to each user are consecutive and the corresponding cardinality is known. ∎

Consider further narrowing down the problem to the case where $\mathcal{N}$ consists of M uniform-sized subsets of consecutive channels, and N is a multiple of M. For this case, the problem is tractable, as proven below.

Theorem 5. *If N is divisible by M, and it is restricted to allocate exactly $\frac{N}{M}$ consecutive channels to each user, then MPCA is tractable with time complexity $O\left(\max\left\{M^3, MN\right\}\right)$.*

Proof: We prove the result by a polynomial-time transformation to the well-solved matching problem in a bipartite graph. For the problem setting in question, channel partition is unique – the M subsets created by the partition are $\{1, \ldots, N/M\}$, $\{N/M + 1, \ldots, 2N/M\}$, $\ldots$, $\{N - N/M + 1, \ldots, N\}$, each containing N/M channels. The bipartite graph has $2M$ nodes representing the users and the subsets of channels. For each pair of the two node groups, say user m and channel subset S, there is an edge of which the cost equals the power of meeting the user rate using the channels in S. This cost is computed in $O\left(\frac{N}{M}\right)$ time (by single-user rate allocation). The complexity of computing all the edge costs is hence $O(MN)$. Maximum-weighted matching in a bipartite graph of V nodes and E edges is solved in $O\left(V^2 \log V + VE\right)$ time [18]. In our case, the second term is dominating, giving a time complexity of $O\left(M^3\right)$, which completes the proof. ∎

Theorem 5 provides a generalization of the trivial case of $M = N$. For $M = N$, solving MPCA amounts to finding an optimal matching with complexity $O\left(M^3\right)$. Moreover, from the proof, it follows that the analysis generalizes to any rate function in $\mathcal{C}_{\text{inc}}$, except that the complexity of single-use rate allocation has to be accounted for accordingly. The observation yields the following corollary.

Corollary 6. *The time required for computing the optimum to the MPCA problem class in Theorem 5 is of $O\left(\max\left\{M^3, M^2 T\left(\frac{N}{M}\right)\right\}\right)$, where $T\left(\frac{N}{M}\right)$ denotes the time complexity of optimal rate allocation of a single user on $\frac{N}{M}$ channels.*

VI. A TRACTABLE PROBLEM CLASS

Denote by K a (possibly large) fixed positive integer independent of M or N. Let $\mathcal{K} = \{1, 2, \cdots, K\}$. Consider MPCA in which the channels can be partitioned into (at most) K groups, such that within each channel group, the channel gains are uniform for every user, but remain non-uniform over the users. That is, from a user's perspective, the problem is characterized by K groups of channels rather than N distinct channels. Thus, we can write $\mathcal{N} = \bigcup_{k \in \mathcal{K}} \mathcal{N}_k$, such that for any channel $n \in \mathcal{N}_k$, the channel gain depends only on the user index m, i.e., $g_{mn} = g_m$. Equivalently, the rate functions belong to the subclass that satisfy $r_{mn} = r_m, n \in \mathcal{N}_k, k \in \mathcal{K}, m \in \mathcal{M}$. Note that the channel groups may vary in size, and the channel gain still differs by users within each group. In the sequel, we refer to the problem class as K-MPCA. The problem class is justified by scenarios with K distinct bands and channel difference is overwhelmingly contributed by the separation of the bands in the spectrum, whereas the subcarriers with each band are considered invariant for each user.

Consider 1-MPCA. The problem structure is significantly simpler than the general case. Namely, the optimization decision is no longer which, but how many channels each user should use. It follows that, for any subset of users $\mathcal{M}' \subset \mathcal{M}$, the optimum allocation of h channels (with $h \geq |\mathcal{M}'|$) among the users in $\mathcal{M}'$ is independent of the channel allocation of the remaining users. Thus 1-MPCA exhibits an optimal substructure, i.e., a part of optimal solution is also optimal for that part of the problem. The observation leads to a dynamic programming line of argument for problem-solving. As proven below, the solution strategy guarantees optimality in polynomial time.

Theorem 7. *Global optimum of* 1-*MPCA can be computed by dynamic programming in* $O\left(MN^2\right)$ *time.*

Proof: Consider the partial problem of optimally allocating h channels to the users in $\{1, \ldots, m\}$ with $h \geq m$, and $c_m(h)$ the corresponding optimum power. Clearly, at the optimum of this subproblem, the number of channels of user m is an integer in the set $\{1, 2, \ldots, h - m + 1\}$. (The upper bound $h - m + 1$ corresponds to having $m - 1$ channels left for the other $m - 1$ users.) Allocating $k \in \{1, 2, \ldots, h - m + 1\}$ channels to user m, the power equals $p_m^k + c_{m-1}(h - k)$, where p_m^k denotes the power of user m with k channels. This gives the following recursive formula for computing the optimal number of channels for user m.

$$c_m(h) = \min_{k=1,\ldots,h-m+1} \left\{ p_m^k + c_{m-1}(h - k) \right\} \tag{2}$$

We arrange the values $c_m(h)$ for $m = 1, \ldots, M$ and $h = 1, \ldots, N$ in an $M \times N$ matrix. Entries corresponding to infeasible solutions are called invalid, and their values are denoted by ∞. In the matrix, $c_m(h) = \infty$ for all entries where $h < m$, or $h > N - M + m$. We compute the valid entries as follows. For the first row, computing the entries $c_1(1), \ldots, c_1(N - M + 1)$ in the given order are straightforward, and each entry requires $O(1)$ computing time. Next, entries $c_m(m)$, i.e., one channel per user for the first m users, are calculated in $O(M)$ time for $m = 1, \ldots, M$. The bulk of the computation calculates the remaining entries row by row, starting from row two. For row m, the computations follow the order $c_m(m + 1), \ldots, c_m(N - M + m)$. Each of these entries is calculated using formula (2). For the valid entries of a row, the total number of comparisons that are used for computing the next row is $1 + \cdots + N - M + 1$. Hence the complexity for computing row m for any $m = 2, \ldots, M$, is of $O\left(N^2 + M^2\right)$. As there are m rows, the time complexity is of $O\left(MN^2 + M^3\right)$. However, because $M \leq N$, the second term can be discarded, leading to $O\left(MN^2\right)$ as the overall complexity result. The last entry computed, $c_M(N)$, gives the optimal allocation of the N channels to the M users and hence solves 1-MPCA. In parallel, the solution is stored in a second matrix of same size. Solution recording clearly has lower complexity than $O\left(MN^2\right)$, and the theorem follows. ∎

In the following theorem, we generalize the dynamic programming concept to any positive integer K. The generalized algorithms are able to solve K-MPCA to global optimality.

Theorem 8. *Global optimum of* K-*MPCA can be computed by dynamic programming in* $O\left(MN^{2K}\right)$ *time.*

Proof: Let $N_j = |\mathcal{N}_j|, j = 1, \ldots, K$. For the first m users, denote by $c_m(h_1, \ldots, h_K)$ the optimum power of allocating h_j channels of group, where $h_j \leq N_j$, $j = 1, \ldots, K$. Denote by $p_m^{(k_1, \ldots, k_K)}$ the power for user m, if it is allocated k_j channels of channel group j, $j = 1, \ldots, K$. We introduce the convention that $p_m^{(0,\ldots,0)} = \infty$ for convenience. By enumerating user m's allocation of channels of the K groups, we obtain the following recursion formula for $c_m(h_1, \ldots, h_K)$.

$$c_m(h_1, \ldots, h_K) = \min_{k_j \in \{0,\ldots,\min\{h_j, N-M+1\}\}, j=1,\ldots,K} \left\{ p_m^{(k_1,\ldots,k_K)} + c_{m-1}(h_1 - k_1, \ldots, h_K - k_K) \right\} \tag{3}$$

Extending the algorithm in the proof of Theorem 7, the corresponding matrix for K-MPCA has dimension $M \prod_{j=1,\ldots,K} N^j$, which does not exceed $O\left(MN^K\right)$. To compute an entry, there are no more than $O\left(N^K\right)$ calculations (including addition and comparison) using (3). For each calculation, the time required to compute $p_m^{(k_1,\ldots,k_K)}$ is linear[2] in K. Since K is a constant, this computation does not add to the complexity. These observations lead to the overall complexity of $O\left(MN^{2K}\right)$. Finally, entry $c_M\left(N^1, \ldots, N^K\right)$ is clearly the optimum to K-MPCA. The proof is complete by observing that, similar to 1-MPCA, recording the channel allocation solution does not form the computational bottleneck. ∎

Algorithm 1 provides a description of the dynamic programming algorithm for 1-MPCA. The adaptation to K-MPCA is straightforward. The input consists of M, N, the target rate vector $\mathbf{R} = (R_1, \ldots, R_M)$, the channel gain vector $\mathbf{g} = (g_1, \ldots, g_M)$, and power matrix $\mathbf{Q}$, where $Q(m, k) = \left(2^{\frac{R_m}{k}} - 1\right) g_m^{-1}$, $m = 1, \ldots, M, k = 1, \ldots, N$. The "$\infty$"-symbol denotes a large power value that is guaranteed to be non-optimal for any user; we can use

---

[2]This result follows by directly applying the single-user rate assignment to K channel groups, where the channels in each group have uniform gain.

Algorithm 1 Dynamic programming algorithm for 1-MPCA.

Input: M, N, $\mathbf{R}$, $\mathbf{g}$, $\mathbf{Q}$
Ouput: $\mathbf{C}$ and $\mathbf{S}$ (two matrices of size $M \times N$)

1: **for** $m = 1 : M$ **do**
2:    **for** $k = 1 : N$ **do**
3:       $C(m, k) \leftarrow \infty$
4:    **end for**
5: **end for**
6: **for** $m = 1 : M$ **do**
7:    $C(m, m) \leftarrow Q(m, 1)$
8: **end for**
9: **for** $k = 1 : N - M + 1$ **do**
10:    $C(1, k) \leftarrow Q(m, k)$
11: **end for**
12: **for** $m = 2 : M$ **do**
13:    **for** $k = m + 1 : N - M + m$ **do**
14:       $C(m, k) \leftarrow \min_{h=1,\ldots,k-m+1}(Q(m, h)$
15:              $+C(m - 1, h - k))$
16:       $S(m, k) \leftarrow \arg\min_{h=1,\ldots,k-m+1}(Q(m, h)$
17:              $+C(m - 1, k - h))$
18:    **end for**
19: **end for**
20: **return** $\mathbf{C}$, $\mathbf{S}$

---

for example $1 + \left(2^{R_m} - 1\right) g_m^{-1}$ in algorithm implementation. In the output, $C(M, N)$ is the optimal objective function value. The $\mathbf{S}$ matrix contains the solution, namely, $S(M, N)$ is the optimal number of channels for user M, $S(M - 1, N - S(M, N))$ is the corresponding number for user $M - 1$, and backtracking in this manner in descending order of users leads to the complete optimum solution. Note that the complexity of $O\left(MN^2\right)$ originates from Lines 12-19 in the algorithm description.

Remark: The polynomial-time tractability of K-MPCA holds only if K is not dependent on M or N. In fact, the general setting of MPCA is equivalent to N-MPCA, i.e., N channel groups with single channel each. The dynamic programming algorithm remains applicable for $K = N$. From Theorem 3, however, the algorithm corresponds to enumerating the solution space, and the running time is exponential. $\square$

Having concluded the tractability of K-MPCA, a natural question to ask next is whether or not the identification of the problem class is tractable as well. Theorem 9 states that this is indeed the case.

Theorem 9. *Recognizing K-MPCA can be performed in $O(MN)$ time for $K = 1$, and in $O\left(MN^2\right)$ time for any $K \geq 2$.*

Proof: For $K = 1$, identifying the problem class simply amounts to verifying, for each user, whether or not all the N channels are of the same gain; this immediately leads to the $O(MN)$ time complexity result. For $K \geq 2$, we construct a graph $\mathcal{G}$. The graph has N nodes, each representing a channel in $\mathcal{N}$. Consider two arbitrary channels $n_1, n_2 \in \mathcal{N}, n_1 \neq n_2$. If the two channels have the same gain for each of the users, i.e., $g_{mn_1} = g_{mn_2}, \forall m \in \mathcal{M}$, we denote it by $n_1 \cong n_2$. Checking whether or not this is the case runs obviously in $O(M)$ time. If $n_1 \cong n_2$, we add edge (n_1, n_2) to $\mathcal{G}$. Doing so for all unordered channel pairs has time complexity $O\left(MN^2\right)$. Next, note that the equivalence relation of channels is transitive, i.e., if $n_1 \cong n_2$ and $n_2 \cong n_3$, then $n_1 \cong n_3$. Hence channels that are equivalent for all users form a clique in $\mathcal{G}$, whereas channels that differ in gain for at least one user are not connected in $\mathcal{G}$. Consequently the number of strongly connected components in $\mathcal{G}$ equals the number of channel groups, each of which contains channels being equivalent for

Algorithm 2 Algorithm for identifying how many channel groups are created by partitioning $\mathcal{N}$, such that within each channel group the gain is uniform for every individual user.

Input: M, N, $\mathbf{g}$
Ouput: $K \in [1, N]$

1: $\mathcal{V} \leftarrow \mathcal{N}$
2: $\mathcal{E} \leftarrow \emptyset$
3: **for** $n_1 = 1 : N$ **do**
4:    **for** $n_2 = 1 : N$ **do**
5:       $s \leftarrow true$
6:       $m \leftarrow 1$
7:       **repeat**
8:          **if** $g(m, n_1) \neq g(m, n_2)$ **then**
9:             $s \leftarrow false$
10:         **end if**
11:         $m \leftarrow m + 1$
12:       **until** $s = false$ or $m = M$
13:       **if** $s = true$ **then**
14:          $\mathcal{E} = \mathcal{E} \cup (n_1, n_2)$
15:       **end if**
16:    **end for**
17: **end for**
18: $\mathcal{G} \leftarrow (\mathcal{V}, \mathcal{E})$
19: $K \leftarrow \text{Tarjan}(\mathcal{G})$
20: **return** K

TABLE I: Optimality gap (%) of the ACG heuristic [1] for K-MPCA with K=2 and 3.

	2-MPCA	3-MPCA
$(M, N) = (3, 8)$	9.29	10.15
$(M, N) = (20, 50)$	33.73	36.52

any user. Identifying the number of strongly connected components requires no more than $O\left(N^2\right)$ time for $\mathcal{G}$. Therefore the bottleneck lies in the $O\left(MN^2\right)$ complexity of obtaining the graph, and the theorem follows. ∎

In Algorithm 2, the procedure of identifying the value of K in K-MPCA is presented in algorithmic form. After constructing graph $\mathcal{G}$, the well-known Tarjan's algorithm of graph search [23] is called (at Line 19) to return the number of strongly connected components of $\mathcal{G}$. This gives the correct number of channel groups. The $O\left(MN^2\right)$ complexity of the algorithm originates from the for-loops starting at Lines 3 and 4, and the repeat-construction starting at Line 7.

Remark: The tractability results of this section are not restricted to the specific rate/power function defined in the section of system model. The problem class remains tractable (although the overall complexity may grow) for any function in $\mathcal{C}_{\text{inc}}$, as long as the function admits polynomial-time rate allocation of single user. □

We end the section by providing results of evaluating a previously proposed heuristic algorithm ACG [1], using instances of K-MPCA with K=2 and 3. The polynomial-time computational complexity of our dynamic programming algorithm for reaching global optimality is provided in Theorem 8. We consider a small-scale system with $(M, N) = (3, 8)$ and a large-scale system with $(M, N) = (20, 50)$. Uniform target rates are used, with $R_m = 1, \forall m = 1, \ldots, M$. The results are summarized in Table I. Each of the values in the table represents the average performance over 100 instances.

From the results in Table I, one observes that the heuristic leads to an optimality gap ranging approximately from 6% to 38%, for K-MPCA that is solved to global optimum in polynomial time. The gap exists for both the

small and large systems. Moreover, the gap has a small but noticeable increase when K goes from 2 to 3, and becomes significantly higher when the problem size grows. Thus general-purpose heuristic may not be able to utilize the structure of the polynomial-time tractable problem class. These observations clarify the significance of our theoretical analysis of identifying tractable problem classes of MPCA.

VII. Conclusions

We have considered the OFDMA resource allocation problem of minimizing the total power of channel allocation, so as to satisfy some rate constraints. Although it has been known that assuming the most general (and ill-behaved) increasing rate functions leads to NP-hard problems, we have shown that the same conclusion holds even if we restrict the class to increasing and concave rate functions. Interestingly, the problem admits a polynomial-time solution if the rate function is an increasing linear function. Hence, progress in the fundamental understanding on the tractability of the problem is made in the following sense: we have sharpened the boundary of tractability to between increasing concave and increasing linear rate functions. Finally, we have also identified specific cases when the problem remains NP-hard, or admits polynomial-time solutions, under various restrictions. Our contributions provide justification to the heuristic approaches that were undertaken for MPCA, and, at the same time, lead to new insights in problem structure that heuristics do not utilize, namely scenarios that do admit global optimality with low computational complexity.

APPENDIX A
PROOF OF THEOREM 1

There is no doubt that MPCA is in NP. The NP-hardness proof uses a polynomial-time reduction from the 3-satisfiability (3-SAT) problem that is NP-complete [21]. A 3-SAT instance consists of a number of boolean variables, and a set of clauses each consisting of a disjunction of exactly three literals. A literal is either a variable or its negation. The output is a yes/no answer to whether or not there is an assignment of boolean values to the variables, such that all the clauses become true. Denote by v and w the numbers of variables and clauses, respectively. For any binary variable z, its negation is denoted by $\hat{z}$. For the proof, we consider 3-SAT where each variable and its negation together appear at most 4 times in the clauses. Note that 3-SAT remains NP-complete with this restriction [22]. Without loss of generality, we assume that each variable z appears in at least one clause, and the same holds for its negation $\hat{z}$, because otherwise the optimal value of the variable becomes known, and the variable can be discarded. Hence the total number of occurrences of each literal in the clauses is between one and three.

We construct an MPCA instance with $M = 2v + w$ and $N = 7v + w$. We categorize the users and channels into groups, and, for convenience, name the groups based on their roles in the proof. The users consist in $2v$ literal users and w clause users. The channels are composed by three groups: v super-channels, $6v$ literal channels, and w auxiliary channels. The rate target $R_m = 1.0, \forall m \in \mathcal{M}$.

Problem reduction is illustrated in Fig. 2. For each binary variable z, three identical literal channels, denoted by z, z' and z'', are defined. A similar construction is done for $\hat{z}$. For this group of six literal channels, one super-channel is defined. We introduce two literal users for the seven channels. One user has channel gain g_l on the three literal channels z, z', and z'', and the other, complementary literal user has channel gain g_l on the remaining three literal channels. Both users have channel gain g_s on the super-channel. See Fig. 2a. Next, recall that each literal appears at most three times in the clauses in the 3-SAT instance. In the proof, for any binary variable z appearing t times in total in the clauses, with $1 \leq t \leq 3$, the occurrences are represented by any t elements in $\{z, z', z''\}$ in any order. A similar representation is performed for the negation $\hat{z}$. For each clause, we introduce one clause user with gain g_c on the channels corresponding to the original literals in the clause. In addition, one auxiliary channel is defined per clause user with channel gain g_a. See Fig. 2b. We set $g_s = g_c = 1$, $g_a = \frac{1}{0.9w + 0.1}, g_l = \frac{g_a}{26} = \frac{1}{26 \cdot (0.9w + 0.1)}$. For the user-channel combinations other than those specified, the channel gain is $g_\epsilon = \frac{1}{53w}$. For each user, we refer to the four channels with gain higher than g_ϵ as valid channels, and the other $7v + w - 4$ channels with gain g_ϵ as invalid channels. From the construction, clearly the reduction is polynomial.

We provide several lemmas characterizing the optimum to the MPCA instance. The first three lemmas use the following optimality conditions of single-user rate allocation (e.g., [2]). First, for any user, the derivatives

(a) Users and channels for a variable.

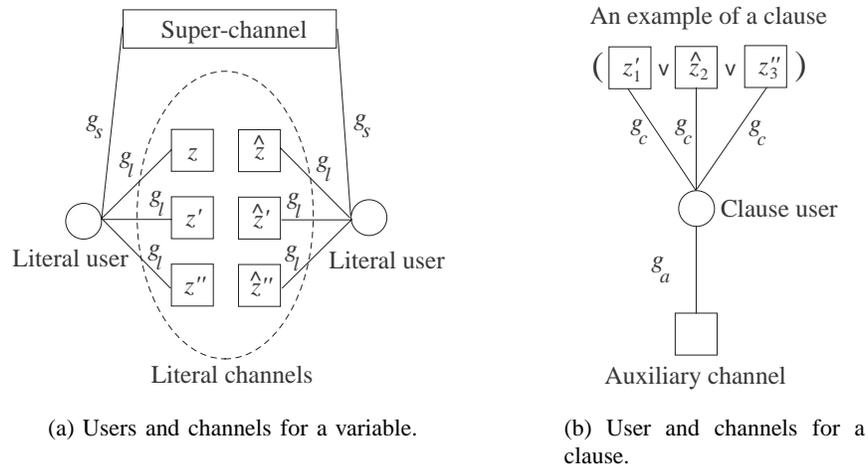(b) User and channels for a clause.

Fig. 2: An illustration of problem reduction.

of the power function, evaluated at the allocated rates, are equal on all channels with positive rates. Second, for channels not used, the function derivatives at zero rate are strictly higher than those of the used channels. For $f(x) = \frac{2^x - 1}{g}$, where x is the rate allocated and g is the channel gain, the derivative $f'(x) = \ln(2)\frac{2^x}{g}$.

Lemma 10. *There is an optimum allocation in which no user is allocated any invalid channel.*

Proof: Suppose that at optimum a clause user m_1 is allocated at least one invalid channel. Assume m_1 is also allocated any valid channel, then the invalid channels carry zero rate, because putting the entire rate of 1.0 on the auxiliary channel, the function derivative is at most $\ln(2) \cdot \frac{2}{g_a} = 2\ln(2)(0.9w + 0.1)$, whereas the function derivative for any invalid channel, at zero rate, is $\ln(2) \cdot 53w > 2\ln(2)(0.9w + 0.1)$. Thus the invalid channels can be eliminated from the allocation of m_1. Assume now the auxiliary channel is allocated to another user, say m_2. By construction, the auxiliary channel of m_1 is invalid for m_2. Consider re-allocating any invalid channel of m_1 to m_2, and allocating the auxiliary channel to m_1. Clearly, the total power will not increase. At this stage, the remaining invalid channels allocated to m_1 carry zero flow. Repeating the argument, we obtain an optimal allocation in which no clause user is allocated any invalid channel.

Let m_1 be any literal user and suppose it is allocated one or more invalid channels at optimum. If m_1 is allocated any of its four valid channels, the function derivative at rate 1.0 is at most $\ln(2) \cdot 52(0.9w + 0.1) < \ln(2) \cdot 53w$, and therefore the invalid channels carry zero rate and can be removed from m_1's allocation. Assume therefore all four valid channels are allocated to other users. Consider any literal channel of m_1, and suppose it is allocated to user m_2. For m_2, this literal channel is an invalid one. Swapping the allocation of the literal channel and any invalid channel currently allocated to m_1, the total power will not grow, and the remaining invalid channels allocated to m_1 can be released. The lemma follows from applying the procedure repeatedly. ∎

Lemma 11. *If a literal user is allocated its super-channel in the optimal solution, then none of the three literal channels is allocated to the same user.*

Proof: Putting the entire rate of 1.0 on the super-channel, the derivative value is $2\ln(2)$. For any literal channel, the derivative at zero rate is $\frac{\ln(2)}{g_l}$. That $g_l = \frac{1}{26 \cdot (0.9w + 0.1)}$ and $w \geq 1$ lead to $\frac{\ln(2)}{g_l} > 2\ln(2)$, and the result follows. ∎

Lemma 12. *If a literal user is not allocated its super-channel in the optimal solution, then the user is allocated all the three literal channels.*

Proof: By the assumption of the lemma and Lemma 10, the literal user in question is allocated one, two, or all three of its literal channels, with total power $f_1 = \frac{1}{g_l}$, $f_2 = 2 \cdot \frac{2^{1/2} - 1}{g_l}$, and $f_3 = 3 \cdot \frac{2^{1/3} - 1}{g_l}$, respectively. Note that, in the latter two cases, it is optimal to split the rate evenly because of the identical gain values. Clearly, $f_3 < f_2 < f_1$.

We prove that cases one and two are not optimal. Suppose that, at optimum, the literal user is allocated two of the literal channels. Then the remaining literal channel is used to carry a strictly positive amount of flow of a clause user. Consider modifying the solution by allocating all the three literal channels to the literal user, and letting the clause user use the auxiliary channel only. The power saving for the literal user is exactly $f_2 - f_3 > \frac{0.04}{g_l} > \frac{1}{g_a}$, whereas the power increase for the clause user is less than $\frac{1}{g_a}$. This contradicts the optimality assumption. Hence case two is not optimal. Since $f_1 - f_2 > f_2 - f_3$, a similar argument applies to case one, and the result follows. ∎

By Lemmas 11–12, at optimum, a literal user will use either the super-channel only, or all the three literal channels. Hence, for any literal in the 3-SAT instance, either none or all of the corresponding three literal channels become blocked for the clause users. Consequently, there is a unique mapping between a true/false variable assignment in the 3-SAT instance and the availability of literal channels to the clause users in the MPCA instance. The total power consumption of all the literal users equals exactly $v + 78v \cdot (2^{1/3} - 1)(0.9w + 0.1)$ at optimum. In the remainder of the proof, we concentrate on the power consumption of the clause users.

Lemma 13. *If every clause user is allocated at least one of the three literal channels corresponding to the literals in the clause in the 3-SAT instance, then the total power for all clause users is at most w.*

Proof: Allocating the entire rate of 1.0 to one literal channel gives a power consumption of $\frac{2^1 - 1}{g_c} = 1$ for the clause user. As there are w clause users, the lemma follows. ∎

Lemma 14. *If at least one clause user is not allocated any of its three literal channels, the total power for all clause users is strictly higher than w.*

Proof: By the assumption, at least one clause user is allocated the auxiliary channel only, with power $\frac{1}{g_a}$. Each of the other $w - 1$ clause users is allocated at most four channels. Assuming the availability of all four channels and setting $g_a = g_c$ leads to an under-estimation of the power consumption. The under-estimation has a total power of $4(w-1)\frac{2^{1/4} - 1}{g_c} + \frac{1}{g_a} = 4(w-1)(2^{1/4} - 1) + (0.9w + 0.1) > 0.4w - 0.4 + 0.9w + 0.1 = 1.3w - 0.3 \geq w$. ∎

By Lemmas 13-14, the optimum power for the clause users is at most w if and only if the answer is yes to the 3-SAT instance. Thus the recognition version of MPCA is NP-complete, and its optimization version is NP-hard.

APPENDIX B
PROOF OF THEOREM 4

In the proof of Theorem 1, a 3-SAT instance of v variables and w clauses is reduced to an MPCA instance with $2v + w$ users and $7v + w$ channels, such that no user will be allocated more than three channels at optimum. We make an augmentation by adding $2v$ channels, which we refer to as dummy channels. The channel gain of the dummy channels equals g_ϵ (defined in the proof of Theorem 1) for all users. After the augmentation, there is a total of $9v + w$ channels, organized in three blocks. The sequence of channels is as follows. The first block has $3v$ channels, including the v super-channels and the $2v$ dummy channels, in a sequence of v chunks of 3 channels each. Each chunk is composed by one super-channel and two dummy channels. The next block has the $6v$ literal channels, with v chunks having 6 channels each. Every chunk corresponds to a binary variable z in the 3-SAT instance, and the six literal channels appear in the order $z, z', z'', \hat{z}, \hat{z}'$ and $\hat{z}''$. The third block contains the w auxiliary channels. Consider the resulting MPCA instance with the restriction that, for the literal and clause users, respectively, the numbers of channels allocated per user are three and one. In addition, the channel allocation must be consecutive in the given sequence.

To prove the hardness result, consider first a relaxation of the problem, in which the two restrictions of channel allocation are ignored for the literal users. For this relaxation, it is clear that Lemmas 10-12 remain valid. By Lemma 10 and the signal-channel restriction of the clause users, each of these users will be allocated one of the four valid channels. Obviously, the optimum is to allocate one literal channel, or the auxiliary channel if all the three literal channels are allocated to literal users. Thus, as long as at least one literal channel is available to every clause user, the result of Lemma 13 holds. In addition, the validity of Lemma 14 obviously remains. Therefore the optimum to the problem relaxation provides the correct answer to the 3-SAT instance.

By Lemma 12, in the optimum of the relaxed problem, each literal user is either allocated its three consecutive literal channels, or the super-channel. In the latter case, we modify the solution by allocating the two dummy channels accompanying the super-channel, without changing the total power. After the modification, the allocation satisfies the cardinality requirement and restriction of using consecutive channels for all literal users, and the theorem follows.

ACKNOWLEDGMENTS

REFERENCES

[1] D. Kivanc, G. Li, and H. Liu, "Computationally efficient bandwidth allocation and power control for OFDMA," *IEEE Trans. Wireless Commun.*, vol. 2, no. 6, pp. 1150–1158, Nov. 2003.

[2] A. Feiten, R. Mathar, and M. Reyer. "Rate and power allocation for multiuser OFDM: an effective heuristic verified by branch-and-bound," *IEEE Trans. Wireless Commun.*, vol. 7, no.1, pp. 60–64, Jan. 2008.

[3] C. Y. Wong, R. S. Cheng, K. B. Letaief, and R. D. Murch, "Multiuser OFDM with adaptive subcarrier, bit, and power allocation," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 10, pp. 1747–1758, Oct. 1999.

[4] W. Rhee and J. M. Cioffi, "Increase in capacity of multiuser OFDM system using dynamic subchannel allocation," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, Tokyo, Japan, May 2000, pp. 1085–1089.

[5] M. Bohge, J. Gross, A. Wolisz, and M. Meyer, "Dynamic resource allocation in OFDM systems: An overview of cross-layer optimization principles and techniques," *IEEE Network Mag.*, vol. 21, no 1, pp. 53–59, Jan. 2007.

[6] S. Sadr, A. Anpalagan, and K. Raahemifar, "Radio resource allocation algorithms for the downlink of multiuser OFDM communication systems," *IEEE Commun. Surveys & Tutorials*, vol. 11, no. 3, pp. 92–106, Third Quarter 2009.

[7] K. Seong, M. Mohseni, and J. M. Cioffi, "Optimal resource allocation for OFDMA downlink systems," in *Proc. IEEE International Symposium on Infor. Theory (ISIT)*, Seattle, USA, Sep. 2006. pp. 1394–1398.

[8] T. Liu, C. Yang, and L.-L. Yang, "A low-complexity subcarrier-power allocation scheme for frequency-division multiple-access systems," *IEEE Trans. Wireless Commun.*, vol. 5. no. 5, pp. 1571–1576, May 2010.

[9] Y. B. Lin, T. H. Chiu, and Y. T. Su, "Optimal and near-optimal resource allocation algorithms for OFDMA networks," *IEEE Trans. Wireless Commun.*, vol. 8, no. 8, pp. 4066–4077, Aug. 2009.

[10] J. Gross and M. Bohge, "Dynamic mechanisms in OFDM wireless systems: a survey on mathematical and system engineering contributions," *Technical Report*, TKN-06-001, Technische University of Berlin, 2006.

[11] P.-H. Huang, Y. Gai, B. Krishnamachari, and A. Sridharan, "Subcarrier allocation in multiuser OFDM systems: complexity and approximability," in *Proc. IEEE Wireless Commun. Network Conf. (WCNC)*, Sydney, Australia, Apr. 2010.

[12] J. A. C. Bingham, "Multicarrier modulation for data transmission: an idea whose time has come," *IEEE Commun. Mag.*, vol. 28, no. 5, pp. 5–14, May 1990.

[13] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Trans. Veh. Technol.*, vol. 42, no. 4, pp. 641–646, 1993.

[14] R. Yates, "A framework for uplink power control in cellular radio systems," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 7, 1995.

[15] J. Zander, "Performance of optimum transmitter power control in cellular radio systems,", *IEEE Trans. Veh. Technol.*, vol. 41, no. 1, pp. 57–62, 1992.

[16] R. E. Burkard, V. G. Deineko, R. van Dal, J. A. A. van der Veen, and G. J. Woeginger. "Well-solvable special cases of the traveling salesman problem: a survey," *SIAM Review*, vol. 40, no. 3, pp. 496–546, Sep. 1998.

[17] C. Papadimitriou. *Computational Complexity*, Addison Wesley, 1994.

[18] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice-Hall, 1993.

[19] L. A. Wolsey, *Integer Programming*, John Wiley & Sons, 1998.

[20] I. C.Wong, O. Oteri, and W. McCoy, "Optimal resource allocation in uplink SC-FDMA systems," *IEEE Trans. Wireless. Commun.*, vol. 8, pp. 2161–2165, May 2009.

[21] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.

[22] C. A. Tovey, "A simplified NP-complete satisfiability problem," *Discrete Applied Mathematics*, vol. 8, no.1 pp. 85–89, Apr. 1984.

[23] R. E. Tarjan, "Depth-first search and linear graph algorithms", *SIAM Journal on Computing*, vol. 1, pp. 146–160, 1972.

Di Yuan received his MSc degree in computer science and engineering, and PhD degree in operations research at Linköping Institute of Technology in 1996 and 2001, respectively. At present he is full professor in telecommunications at the Department of Science and Technology, Linköping University, and head of a research group in mobile telecommunications. He has been guest professor at the Technical University of Milan (Politecnico di Milano), Italy, and visiting researcher at Ericsson Research, Sweden. He has been in the management committee of four European Cooperation in field of Scientific and Technical Research (COST) actions, and invited lecturer of European Network of Excellence EuroNF. He received a Best Paper Award at IEEE ICC 2012. He is an area editor of the Computer Networks journal. His research interests span design, analysis, and resource optimization of telecommunication systems, with over 90 publications. His current research mainly addresses network optimization of 4G systems.

Jingon Joung (S'04–M'07) received B.S. degree in electrical engineering from Yonsei University, Seoul, Korea, in 2001, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Korea Advance Science and Technology (KAIST), Daejeon, Korea, in 2003 and 2007, respectively.

From March 2007 to August 2008, he was a Postdoctoral Research Scientist with the Department of Electrical Engineering, KAIST. From April 2007 to August 2008, he worked as a Commissioned Researcher at Lumicomm, Inc., Daejeon, Korea. From October 2008 to September 2009, he was a Postdoctoral Fellow with the Department of Electrical Engineering at the University of California, Los Angeles (UCLA). Since October 2009, he has been working as a Scientist at Institute for Infocomm Research (I^2R), Singapore. His current research has focused on the study of energy efficient wireless communications.

Dr. Joung was the recipient of a Gold Prize at the Intel-ITRC Student Paper Contest in 2006.
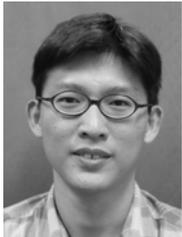
Chin Keong Ho received the B.Eng. (First-Class Hons.) and M. Eng degrees from the Department of Electrical Engineering, National University of Singapore in 1999 and 2001, respectively. Since August 2000, he has been with Institute for Infocomm Research (I2R), A*STAR, Singapore. He took leave in 2004-2007 to obtain his Ph.D. degree at the Eindhoven University of Technology, Eindhoven, The Netherlands. Concurrently during his leave, he conducted research work in Philips Research. Prior to his PhD leave, Chin Keong has worked on signal processing aspects of multi-carrier and multi-antenna communications. His current research interest lies in cooperative and adaptive wireless communications, energy efficient communications, and communications under energy harvesting constraints.

Sumei Sun (SM'12) obtained the B.Sc.(Honours) Degree from Peking University, China, the M.Eng Degree from Nanyang Technological University, and Ph.D Degree from National University of Singapore. She has been with Institute for Infocomm Research (formerly Centre for Wireless Communications) since 1995 and she is currently Head of Modulation & Coding Dept, developing physical layer-related solutions for next-generation communication systems. Her recent research interests are in energy efficient multiuser cooperative MIMO systems, joint source-channel processing for wireless multimedia communications, and wireless transceiver design.

Dr Sun has served as the TPC Chair of 12th IEEE International Conference on Communications in 2010 (ICCS 2010), General Co-Chair of 7th and 8th IEEE Vehicular Technology Society Asia Pacific Wireless Communications Symposium (APWCS), and Track Co-Chair of Transmission Technologies, IEEE VTC 2012 Spring. She has also an Associate Editor of IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, and an Editor of IEEE WIRELESS COMMUNICATION LETTERS. She is co-recipient of IEEE PIMRC'2005 Best Paper Award.