Discrete Applied Mathematics 194 (2015) 23–36

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Bandwidth allocation in cellular networks with multiple interferences[☆]

Reuven Bar-Yehuda [a], Gleb Polevoy [b], Dror Rawitz [c,*]

[a] *Department of Computer Science, Technion, Haifa 32000, Israel*

[b] *Faculty of Engineering, Mathematics and Computer Science, TU Delft, P.O. Box 5031, 2600 GA Delft, The Netherlands*

[c] *Faculty of Engineering, Bar Ilan University, Ramat Gan 52900, Israel*

ARTICLE INFO

Article history:
Received 11 August 2013
Received in revised form 9 February 2015
Accepted 6 May 2015
Available online 29 May 2015

Keywords:
Approximation algorithms
Bandwidth allocation
Multiple interferences
Interval scheduling
Interval selection
Local ratio

ABSTRACT

We study the problem of bandwidth allocation with *multiple interferences*. In this problem the input consists of a set of users and a set of base stations. Each user has a list of requests, each consisting of a base station, a frequency demand, and a profit that may be gained by scheduling this request. The goal is to find a maximum profit set of user requests $\mathcal{S}$ that satisfies the following conditions: (i) $\mathcal{S}$ contains at most one request per user, (ii) the frequency sets allotted to requests in $\mathcal{S}$ that correspond to the same base station are pairwise non-intersecting, and (iii) the QoS received by any user at any frequency is reasonable according to an *interference model*. In this paper we consider two variants of bandwidth allocation with multiple interferences. In the first each request specifies a demand that can be satisfied by any subset of frequencies that is large enough. In the second each request specifies a specific frequency interval. Furthermore, we consider two interference models, *multiplicative* and *additive*. We show that these problems are extremely hard to approximate if the interferences depend on both the interfered and the interfering base stations. On the other hand, we provide constant factor approximation algorithms for both variants of bandwidth allocation with multiple interferences for the case where the interferences depend only on the interfering base stations. We also consider a restrictive special case that is closely related to the Knapsack problem. We show that this special case is NP-hard and that it admits an FPTAS.

1. Introduction

Allocating bandwidth while maintaining a reasonable level of quality of service (QoS) is an important part of managing a cellular network (see, e.g., [21,19]). In this paper we focus on the throughput maximization aspect of bandwidth allocation, namely we study the problem of assigning frequencies to base stations according to user requests so as to maximize system profit.

Several studies regarding cell planning of cellular networks have been performed. Roughly speaking, the cell planning problem is to find a minimum cost set of locations for base stations that will be able to satisfy all user demands. (See [2] for a survey on cell planning.) A client may be covered by several base stations, and in this case the transmission by a base station

☆ A preliminary version was presented at the 6th ACM SIGACT/SIGMOBILE International Workshop on Foundation of Mobile Computing (DIALM-POMC), 2010.

* Corresponding author.
E-mail addresses: reuven@cs.technion.ac.il (R. Bar-Yehuda), g.polevoy@tudelft.nl (G. Polevoy), dror.rawitz@biu.ac.il (D. Rawitz).

http://dx.doi.org/10.1016/j.dam.2015.05.013

may be interfered by the other base stations that transmit at the same frequency. Hence, one of the most interesting aspects of the cell planning problem is the *interference model*. In wideband systems all transmissions are spread over the entire bandwidth, namely the same spectrum is used in every cell, and therefore simultaneous transmissions may cause interferences. In the Orthogonal Frequency Division Multiplexing (OFDM) system transmissions within cells are kept orthogonal due to hopping in the time frequency grid, but inter-cell interferences are still possible, namely base station transmissions may interfere with other cells (see [24, Chapter 4]). Interference is typically modeled by an interference matrix that provides information about the impact of any base station on the transmission of any other base station in areas that are covered by both cells (see [11, Appendix 6B]).

Following Amzallag et al. [3,1] we consider two interference models for measuring quality of service. In the *additive interference model*, each pair of base stations i and k have an interference parameter $w(i, k)$ representing the interference to a user listening to i transmitting at the same frequency as k. We have a QoS upper bound on the interference, and the weights are normalized such that the upper bound is 1. A user is able to listen to base station i transmitting at frequency t only if the interference at t is at most 1, i.e., if $\sum_{k\in B(t)\setminus\{i\}} w(i, k) \leq 1$, where $B(t)$ denotes the set of base stations transmitting at t. In the *multiplicative interference model*, each pair of base stations i and k have an interference parameter $\sigma(i, k)$ representing the fraction of service that is lost by a user listening to i transmitting at the same frequency as k. If several base stations are transmitting at the same frequency t, then the remaining service received by a user listening to base station i at frequency t is $\prod_{k\in B(t)\setminus\{i\}} (1 - \sigma(i, k))$. A user is able to listen to base station i at frequency t only if the interference at t is at most δ, i.e., if $\prod_{k\in B(t)\setminus\{i\}} (1 - \sigma(i, k)) \geq 1 - \delta$.

We introduce the problem of bandwidth allocation with multiple interferences in which the goal is to maximize the throughput of an existing cellular network. In this problem the input consists of a set of m users, and a set of n base stations. Each user has a list of requests, each consisting of a base station, a frequency demand, and a profit that may be gained by scheduling this request. The goal is to find a maximum profit set of user requests that satisfies three types of constraints:

User Constraints: The solution contains at most one request per user.

Base Station Constraints: The frequency sets that are assigned to requests that correspond to the same base station are pairwise non-intersecting.

QoS Constraints: The quality of service received by any user at any frequency is reasonable according to some interference model.

We study two variants of bandwidth allocation with multiple interferences. The first is Frequency Allocation with Multiple Interferences (abbreviated, FAMI) where each request specifies a demand that can be satisfied by any subset of frequencies that is large enough. The second is the Interval Selection with Multiple Interferences problem (ISMI) where each request specifies a frequency interval (i.e., a consecutive set of frequencies). We consider both FAMI and ISMI with the two interferences models, the *multiplicative* and the *additive*. We note both problems can be defined such that user requests are for sets of time units or time intervals instead of sets of frequencies or frequency intervals. In this case the QoS constraints are per time unit.

ISMI is an extension of a single machine scheduling problem called Weighted Job Interval Selection Problem (WJISP). In WJISP the input consists of n weighted jobs, where a job is a set of (time) intervals on the real line. To schedule a job one of its intervals must be selected. A schedule of several jobs contains a set of non-intersecting job intervals. The goal is to find a maximum weight schedule. It is not hard to see that when there is only one base station (i.e., $n = 1$), the QoS constraints no longer matter, and therefore this special case of ISMI is equivalent to WJISP. Spieksma [23] proved that the unweighted version of WJISP (JISP) is APX-hard even for the case where there are at most two intervals per job (JISP2). Hence, ISMI is also APX-hard.

When interferences are uniform (i.e., all interferences are the same) the QoS constraints translate into a bound on the number of transmissions per frequency. Hence, ISMI with uniform interferences captures an extension of WJISP in which we are given a set of identical machines, and we have an upper bound on the number of machines that can work simultaneously. This can model an upper bound on the energy consumption or on the size of the crew operating the machines. We note that scheduling jobs on several machines can be modeled by WJISP by concatenating the schedules for the machines along the real line, and specifying the possible intervals for each job accordingly. However, in ISMI we cannot concatenate the schedules for the base stations due to the QoS constraints.

Our results. We show that both FAMI and ISMI are extremely hard to approximate if the interferences depend on both the interfered and the interfering base stations. More specifically, we show that these problems are not approximable within n^ε, for some $\varepsilon > 0$, unless P = NP, even for the very restrictive special case in which there is only one frequency, one request per user, one user request per base station and unit profits.

Given the hardness result we consider the case where the interferences depend only on the interfering base stations. First, we prove that FAMI is strongly NP-hard in both interference models. (Recall that ISMI is APX-hard even if $n = 1$ [23].) Assuming the additive interference model, we provide a 7-approximation algorithm for ISMI and a 12-approximation algorithm for FAMI. Our algorithms are based on the local ratio technique [9,6]. Our techniques can be used in the multiplicative interference model, but the approximation ratios increase by a factor of $(1 + \varepsilon)$. We note that our approximation ratios improve if all interferences are small. We also note that using the techniques of [6], our algorithm for ISMI can be modified to deal with frequency windows (see Section 4.3). The resulting approximation ratio is $7 + \varepsilon$, for every $\varepsilon > 0$.