

Scheduling Dynamic Parallel Workload of Mobile Devices with Access Guarantees

Antonio Fernández Anta, IMDEA Networks Institute
Dariusz R. Kowalski, University of Liverpool
Miguel A. Mosteiro, Pace University
Prudence W. H. Wong, University of Liverpool

We study a dynamic resource-allocation problem that arises in various parallel computing scenarios such as mobile cloud computing, cloud computing systems, Internet of Things systems and others. Generically, we model the architecture as client mobile devices and static base stations. Each client “arrives” to the system to upload data to base stations by radio transmissions, and then “leaves”. The problem, called *Station Assignment*, is to assign clients to stations so that every client uploads its data under some restrictions, including a target subset of stations, a maximum delay between transmissions, a volume of data to upload, and a maximum bandwidth for each station. We study the solvability of Station Assignment under an adversary that controls the arrival and departure of clients, limited to maximum rate and burstiness of such arrivals. We show upper and lower bounds on the rate and burstiness for various client arrival schedules and protocol classes. To the best of our knowledge, this is the first time that Station Assignment is studied under adversarial arrivals and departures.

CCS Concepts: • **Theory of computation** → **Parallel algorithms; Online algorithms; Approximation algorithms analysis;**

Additional Key Words and Phrases: Station Assignment, Mobile Cloud Computing, Radio Networks, Continuous Adversarial Dynamics, Internet of Things, Health Monitoring Systems

ACM Reference Format:

Antonio Fernández Anta, Dariusz R. Kowalski, Miguel A. Mosteiro, and Prudence W. H. Wong, 2018. Scheduling Dynamic Parallel Workload of Mobile Devices with Access Guarantees. *ACM Trans. Parallel Comput.*, , Article (March), 20 pages.
DOI: 0000001.0000001

1. INTRODUCTION

We study a dynamic allocation problem that arises in various parallel computing scenarios where minimum volume of data transferred with some regularity must be guaranteed. For instance, mobile access to cloud computing services through multiple devices (e.g. smart phones, tablets, etc.) that connect to the cloud through access points using radio communication. Also, cloud computing systems where virtual machines with various user-defined resource requirements must be allocated to real machines. Another example is Internet of Things subnetworks such as wearable health-monitoring systems where ambulatory patients carry physiological sensors and the data gathered must be periodically uploaded.

To comprise all the above applications, we consider a general model assuming a continuous arrival of client mobile devices who have to upload data to static base stations

Author’s addresses: Antonio Fernández Anta IMDEA Networks Institute, Madrid, Spain; Dariusz R. Kowalski and Prudence W. H. Wong University of Liverpool, Department of Computer Science, Liverpool, UK; Miguel A. Mosteiro Pace University, Seidenberg School of CSIS, New York, NY, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© ACM. 1539-9087/03-ART \$15.00

DOI: 0000001.0000001

via radio transmissions. Clients have a life interval and a fixed subset of stations to communicate. Mobility can then be modeled as multiple instances of the same client with subsequent life intervals and possibly different subsets of stations. Under this model, we study the *Station Assignment* problem of allocating clients to base stations so that every client transmits to some station in its subset, limited by maximum delays between transmissions and data bandwidths in the client and station side.

We show upper and lower bounds on the rate and burstiness of client arrivals for solvability of Station Assignment under various client arrival schedules and protocol classes. We assume the presence of an adversary that controls the arrival and departure of clients. The adversary is limited by two parameters that model the rate and burstiness of the arrival. We also study the connections of this problem with online load balancing and scheduling, usually studied using competitive analysis. To the best of our knowledge, this is the first study of Station Assignment under adversarial arrivals.

2. ADVERSARIAL MODEL AND PROBLEM DEFINITION

Model. We consider a Mobile Radio Network composed of a set S of base stations, or simply **stations** for short and a set C of **clients** that want to transmit packets to some station. Throughout we denote $n \triangleq |C|$ and $m \triangleq |S|$. The time is assumed to be slotted and the time domain is \mathbb{N} . Each time slot is long enough to transmit one packet.

Each client $c \in C$ has the following characterization.

- A **life interval**, which is the set $\tau_c = [a, b] \subseteq \mathbb{N}$ of consecutive slots in which c is active.
- A **group of stations**, which is the set $S_c \subseteq S$ of stations to which c may transmit packets.
- A **laxity** $w_c \in \mathbb{N}$, $0 < w_c \leq |\tau_c|$, such that $c \in C$ must transmit to some station in S_c at least once within every w_c consecutive time slots in τ_c .
- A **bandwidth** $b_c \in \mathbb{R}^+$ that models a resource requirement (such as frequency bandwidth).

On the other hand, each station $s \in S$ has the following characterization.

- A **bandwidth** $B_s \in \mathbb{R}^+$, which limits the sum of the bandwidth of the clients transmitting to s .

We refer to the set of stations (with their parameters) as the **system** and to the set of clients (with their parameters) as the **client arrival schedule**. The model described above is very rich and allows for many different combinations of system and client arrival schedule. In this work, we focus only on a subset of these combinations, which have been found to have enough variety of challenging cases to explore, leaving the rest for future work. Hence, we assume in the rest of the paper systems in which all stations have the same bandwidth B , and client arrival schedules in which all clients have the same laxity w .

To carry out a worst-case analysis, we consider adversarial client arrival schedules where the adversary is limited as follows. For any $C' \subseteq C$, let $S(C') = \bigcup_{c \in C'} S_c$. For a given pair of values $\rho > 0$ and $\beta \geq 0$ (that limit the rate and burstiness of the stations load, which in turn limits the arrival/departure of clients), we say that a client arrival schedule is (ρ, β) -**admissible** if the following conditions hold:

$$\forall C' \subseteq C : \forall T = [t, t'] \subseteq \mathbb{N} : \sum_{c \in C'} b_c \frac{|\tau_c \cap T|}{w} \leq |T| |S(C')| \rho B + \beta \quad (1)$$

$$\forall c \in C : b_c \leq B. \quad (2)$$

The first condition (1) restricts the load of the stations for any set of clients C' and any time interval T . In particular, given any C' and any T , the total bandwidth requested by the clients in C' (specifically, $\sum_{c \in C'} b_c |\tau_c \cap T|/w$) has to be no larger than a fraction ρ of the bandwidth that can be provided by the stations that can serve the clients in C' ($S(C')$) plus a constant term β (that allows for some burstiness). The second condition (2) imposes that the requested bandwidth b_c of each client must be no larger than the bandwidth B of each station. Naturally, if some client had a request of bandwidth larger than B it would be impossible to satisfy it. Adversarial methodology characterized as above is typically used for performing worst-case analysis of the considered problem [Borodin et al. 2001; Andrews et al. 2001; Blesa et al. 2009].

Problem. The **Station Assignment problem** is defined as follows. For a given system and admissible client arrival schedule, for each time slot $t \in \mathbb{N}$, schedule a set of clients to transmit to each station in t , so that

- (1) Each client $c \in C$ transmits to some station in S_c at least once within each w consecutive time slots in τ_c using a bandwidth b_c ;
- (2) For each station $s \in S$ the sum of the bandwidths of the clients transmitting to s in any time slot is at most B .

Protocols. We consider the following classes of protocols, commonly used in scheduling literature.

- A Station Assignment algorithm is called **irrevocable** if for each client c all the transmissions of c are to the same station s . We say that the algorithm **irrevocably assigns** the client c to station s .
- A Station Assignment algorithm is called **online** if the information about any client c is revealed to the algorithm only at the arrival time of c .
- A Station Assignment algorithm is called **improvident** if the algorithm does not know when a client will leave the system.

3. OUR RESULTS

The results presented in this work are summarized in Tables I and II. The tables are organized by the system characteristics (columns) and the rows are further subdivided by double lines into comparable settings for which upper and lower bounds are presented. The terms “distinct”, “identical”, and “any” in these tables refer to the values that the system characteristics can take. For instance, for b_c , “identical” means that all clients have the same bandwidth, “distinct” means that different clients have different bandwidth, and “any” means that it is not restricted. The bounds are on the adversary limitations ρ and β . Thus, lower bounds are for impossibility whereas upper bounds are for solvability.

We introduce the Station Assignment problem motivated by wireless networks. Our main contribution is a variety of separation results that expose the complexity of Station Assignment according to various model assumptions. Starting from an optimistic scenario where all clients have the same bandwidth and the same group of stations, we gradually remove assumptions making the model more pessimistic (hence, realistic), which gives insight on what the inherent algorithmic challenges of Station Assignment are. We note in Tables I and II that not all combinations of systems assumptions have been considered. That is, the question of whether other combinations could lead to more separations is open.

Specifically, we start considering adversarial client arrival schedules where all clients have the same group of stations and bandwidth. Then, Theorem 5.1 shows that for each $\beta > mwB((n/(mw))/\lceil n/(mw) \rceil - \rho)$, where $n \geq \lceil (mwB\rho + \beta)/B \rceil$, there exists a (ρ, β) -admissible client arrival schedule such that no Station Assignment algorithm

Table I. Summary of bounds on problem solvability for **offline** protocols.

b_c	S_c	arrival time	protocol class	β	ρ	Thm
identical	identical	identical	any	$> mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right)$ $n = \lceil \frac{mwB\rho + \beta}{B} \rceil$	$> \frac{n/(mw)}{\lceil n/(mw) \rceil}$	5.1
identical	identical	identical	even assignment	$\leq mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right)$	$\leq \frac{n/(mw)}{\lceil n/(mw) \rceil}$	5.3
distinct	identical	identical	any	—	$> 1/2$	5.4
distinct	identical	identical	any	$> mB(1/m + 1/2 - \rho)$	—	5.5
any	identical	identical	balance bandwidth	$< mwB(1/2 - \rho)$	$< 1/2$	5.6
distinct	distinct	identical	any	$> mwB(1/(mw) - \rho)$	$> 1/(mw)$	5.7
any	any	any	any	$\leq mwB(1/(mw) - \rho)$	$\leq 1/(mw)$	5.8

Table II. Summary of bounds on problem solvability for **online** protocols.

b_c	S_c	arrival time	τ_c	protocol class	β	ρ	Thm
1	distinct	distinct	open	irrevocable	≥ 0	$> \frac{1}{1+\ln m}$	6.1
1	distinct	distinct	open	irrevocable	$> mB \left(\frac{1}{\ln m} - \rho \right)$	$> \frac{1}{\ln m}$	6.3
1	distinct	distinct	distinct	irrevocable improvident randomized	$> mB \left(\frac{3}{\sqrt{2m}} - \rho \right)$	$> \frac{3}{\sqrt{2m}}$	6.4
1	distinct	distinct	distinct	irrevocable improvident deterministic	$> mB \left(\frac{1}{\sqrt{2m}} - \rho \right)$	$> \frac{1}{\sqrt{2m}}$	6.4
identical	any	any	open	irrevocable improvident	$\leq \rho B$	$< \frac{1}{2(m+1)}$	6.5
$\rho B \leq b < \frac{B}{m-1}$	any	any	open	irrevocable improvident	$\leq \rho B$	$\rho < \frac{1}{m-1}$	6.6

can solve the problem, even if all clients arrive simultaneously and have the same life interval. The intuition is the following. Leaving aside rounding (i.e., $\beta > mwB(1 - \rho)$), which is similar to ignoring the effect of one station poorly used, and noticing that mwB conveys a measure of the overall resources of which a ρ fraction may be “regularly” used, what Theorem 5.1 shows is that if the adversary is allowed to inject a burst larger than the resources available, the system is overloaded. Given that it must be $\beta \geq 0$, this lower bound for non-solvability implies also a lower bound of $\rho > (n/(mw))/\lceil n/(mw) \rceil$. Corollary 5.2 shows a stronger bound on β that holds for any positive ρ . Under the same conditions, Theorem 5.3 shows that the offline algorithm that distributes the clients evenly solves Station Assignment, for any (ρ, β) -admissible client arrival schedule that matches those bounds on β and ρ .

Then, we move to a class of client arrival schedules where clients may have different bandwidths, although the group of stations is still the same for all. In this scenario, Theorem 5.4 shows that, for each $\rho > 1/2$, there exists a (ρ, β) -admissible client arrival schedule such that no Station Assignment algorithm can solve the problem, even if all clients must arrive simultaneously. Changing the adversarial client arrival schedule slightly, Theorem 5.5 shows a bound of $\beta > mB(1/m + 1/2 - \rho)$ for the same conditions. This bound implies a bound on ρ as well, but it is subsumed by Theorem 5.4. Under the same conditions, Theorem 5.6 shows that an algorithm that (somehow) balances the station-bandwidth usage solves Station Assignment, for any (ρ, β) -admissible client arrival schedule such that $\beta < mwB(1/2 - \rho)$ and $\rho < 1/2$.

The last class of client arrival schedules we consider in our offline analysis does not restrict groups of stations or bandwidths. In Theorem 5.7, it is shown that, for each $\beta > mB(1/m - \rho)$, there exists a (ρ, β) -admissible client arrival schedule such that Station Assignment cannot be solved by any algorithm, even if all clients arrive at the same time. This result implies $\rho > 1/m$ because $\beta \geq 0$. Theorem 5.8 matches those bounds, showing that if $\rho \leq 1/(mw)$ and $\beta \leq mwB(1/(mw) - \rho)$ the Station Assignment problem is solvable.

Moving to online protocols (c.f., Table II), we prove in Theorem 6.3 that for any irrevocable algorithm, that is, algorithms where the station-client assignments are final, if $m > 1$, there is a client arrival schedule such that if $\beta > mB(1/\ln m - \rho)$ the Station Assignment problem is not solvable. The latter lower bound implies a lower bound of $\rho > 1/\ln m$ because $\beta \geq 0$. If the algorithm is additionally improvident, that is, the departure time of clients already in the system is not known in advance, by showing a reduction from Load Balancing [Azar et al. 1994], Theorem 6.4 shows lower bounds of $\beta > mB(3/\sqrt{2m} - \rho)$ and $\beta > mB(1/\sqrt{2m} - \rho)$ for randomized and deterministic algorithms respectively. Those bounds imply that if $\rho > 3/\sqrt{2m}$ and if $\rho > 1/\sqrt{2m}$ respectively the Station Assignment problem cannot be solved. Finally, Theorem 6.5 shows that, when all clients have the same bandwidth $b \leq B$ and do not depart, even if the groups of stations and arrival times are different, if $\rho < 1/(2(m+1))$ and $\beta \leq \rho B$ the algorithm that distributes clients evenly (restricted to groups of stations) solves Station Assignment. For a similar scenario but with restricted $\rho B \leq b < B/(m-1)$, a better bound of $\rho < 1/(m-1)$ is shown in Theorem 6.6.

We also show in Theorem 6.1 (not included in Table II) a lower bound on ρ for non-solvability with irrevocable algorithms that applies to systems with distinct station-bandwidths. Corollary 6.2 shows that instantiating Theorem 6.1 on a system where all stations have the same bandwidth B , the lower bound on ρ for non-solvability is $\rho > 1/(1 + \ln m)$.

4. RELATED WORK

Adversarial queuing was introduced in [Andrews et al. 2001; Borodin et al. 2001], applied to store-and-forward networks, to measure stability of buffers and packet latency of dynamically injected packets. Later, there were approaches to apply it in the context of wireless networks: modelled as time-varying channels [Andrews and Zhang 2005], radio channels with collisions [Chlebus et al. 2012; Anantharamu et al. 2009], or SINR networks [Kesselheim 2012]. A more detail competitive analysis of dynamic and stochastic traffic was performed in a single-hop radio channels with collisions [Binkowski et al. 2012]. The difference between this line of research and our work is that it considered simple packet forwarding requests without additional scheduling constraints.

In [Azar et al. 1994], Azar, Broder, and Karlin studied a load balancing problem where a set of tasks have to be assigned to a set of machines. In this work, tasks are

Table III. Competitive ratios of load balancing problem

	Unknown duration	Known duration	Permanent
Identical	$2 - o(1)$ [Graham 1966; Azar et al. 1994]	$2 - o(1)$ [Graham 1966; Azar et al. 1994]	$2 - \epsilon$ [Karger et al. 1996]
Related	$\Theta(1)$ [Azar et al. 1997]	$\Theta(1)$ [Azar et al. 1997]	$\Theta(1)$ [Azar et al. 1997]
Restricted	$O(\sqrt{m})$ [Azar et al. 1997] $\Omega(\sqrt{m})$ [Azar et al. 1994]	$O(\log mT)$ [Azar et al. 1997]	$\Theta(\log m)$ [Aspnes et al. 1993]

temporary (as opposed to permanent when tasks do not depart), and task arrivals and departures always occur in time. Each task has an associated weight that represents the load that processing of such a task adds to a machine. Additionally, each task has an associated subset of machines that may process the task (restricted assignment). Upon arrival, a task must be assigned to a machine immediately and cannot be transferred to another machine later. The machine starts processing the task immediately and continues until the task departs. An assignment algorithm selects a machine to assign each task upon arrival. In the online version the algorithm does not know future arrivals or departures, whereas an offline algorithm has complete knowledge. The cost of an assignment of a given input is the maximum load of the machines for such assignment.

The load balancing problem is related to our Station Assignment problem when the laxity is the same for all clients. One main difference is the objective function. In load balancing the load can be unbounded and the objective is to minimize the maximum load over the machine, while in the Station Assignment problem the maximum load of the server is bounded and the objective is to find criteria for a feasible schedule. As will shown later (Theorem 6.4), some special instance of the load balancing problem can be mapped to some instance of the Station Assignment problem. However, this mapping is not generic, and there is no evidence that an algorithm that works well for one objective would work for the other. Moreover, the fact that our adversary has two limitations also separates the problems. Indeed, allowing a larger burst (β) imposes a more stringent limitation on arrival rate (ρ) in the Station Assignment problem, while intuitively the restriction on the load in the load balancing problem corresponds only to one of these parameters, mainly the arrival rate.

For the load balancing problem, the authors in [Azar et al. 1994] study the competitive ratio of an online algorithm with respect to an offline one as the supremum over all inputs of the cost ratio. Specifically, for the greedy online algorithm that assigns each task to the least loaded machine, they show matching upper and lower bounds of $((3m)^{2/3}/2)(1 + o(1))$ on the competitive ratio, and a lower bound of $\Omega(\sqrt{m})$ for any deterministic or randomized algorithm. The lower bound is matched in [Azar et al. 1997]. Variants of the problem include relaxing the constraint such that the duration of a job is known on arrival (temporary) or the job never depart (permanent). Another direction of relaxation includes making all machines to be available for all jobs (identical or related). Table III gives a summary of the results.

In [Alon et al. 1997], Alon et al. studied a similar model for permanent tasks. They consider two cases: (i) the tasks have associated weights and can be assigned to any machine (unrestricted), (ii) the tasks have unit weights and can be assigned only to a subset of the machines (restricted). They provide an ϵ -approximation scheme for the L_p norm of the loads. Interestingly, for the restricted unit-weights model, they show that there exists an assignment that is optimal for all norms. For further references on dynamic online scheduling and load balancing, see the chapters [Pruhs et al. 2004; Azar 1996].

Other related problems include the so-called windows scheduling [Bar-Noy et al. 2007; Bar-Noy and Ladner 2003] and bin packing of unit-fraction items [Bar-Noy et al.

2007; Chan et al. 2008]. In the windows scheduling problem, items (cf. clients) with windows w_i have to be broadcasted on a channel periodically such that every two broadcasts of the same item cannot be more than w_i timeslots apart. In the context of our problem, this is analogous to the scenario where the bandwidth of a station and the bandwidth requirement of a client is the same. The objective has been however to minimize the number of channels (cf. stations). Windows scheduling has been studied as a restricted version of bin packing of unit-fraction items [Bar-Noy et al. 2007; Chan et al. 2008] where the size of the items is the reciprocal of the windows. The objective has been again to minimize the number of bins (cf. stations).

5. ANALYSIS OF OFFLINE PROTOCOLS

In this section, we study the impact of ρ and β on the offline solvability of Station Assignment.

5.1. Unique Group of Stations and Client Bandwidth

We start with a very optimistic scenario (for Station Assignment algorithms) where all clients have the same group of stations and the same bandwidth. We show a lower bound for non-solvability that holds even under those optimistic conditions. Given that $\beta \geq 0$ by definition, the bound obtained implies a lower bound on ρ .

THEOREM 5.1. *Given a system of m stations each with bandwidth B , even if all clients have the same group of stations and the same bandwidth, for any*

$$\beta > mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right),$$

where $n \geq \lceil (mwB\rho + \beta)/B \rceil$, $n \in \mathbb{Z}^+$, there exists a (ρ, β) -admissible client arrival schedule such that no algorithm can solve the Station Assignment problem, even if all clients have the same life interval.

PROOF. Consider a client arrival schedule of n clients, for any $n \geq \lceil (mwB\rho + \beta)/B \rceil$, $n \in \mathbb{Z}^+$, with the same bandwidth $b = (mwB\rho + \beta)/n$ and the same life interval of length w . Such schedule is (ρ, β) -admissible because, for any $n' \leq n$ and any subinterval T of the life interval of the clients (i.e., $|T| \leq w$), it holds $n'b \frac{|T|}{w} \leq nb \frac{|T|}{w} = (mwB\rho + \beta) \frac{|T|}{w} \leq m|T|B\rho + \beta$, and $b = \frac{mwB\rho + \beta}{n} \leq \frac{mwB\rho + \beta}{\lceil (mwB\rho + \beta)/B \rceil} \leq B$. However, by the pigeonhole principle, there is at least one station and one slot for which the sum of bandwidths of the clients assigned to the station in the slot is at least $\lceil n/(mw) \rceil b = \lceil n/(mw) \rceil (mwB\rho + \beta)/n$. Replacing $\beta > mwB((n/(mw))/\lceil n/(mw) \rceil - \rho)$, the latter is bigger than B . \square

Given that the client arrival schedule is adversarial, by choosing the station group to be a singleton in the above proof, that is $m = 1$, and the laxity $w = 1$, the lower bound obtained becomes $\beta > B(1 - \rho)$, which implies that if $\rho > 1$ the Station Assignment is not solvable. We assume that $\rho \leq 1$ throughout the rest of the paper. This result can also be used to show that, for some higher values of β , Station Assignment is not solvable for any $\rho > 0$.

COROLLARY 5.2. *Given a system of m stations each with bandwidth B , even if all clients have the same group of stations and the same bandwidth, if $\rho > 0$ and $\beta \geq nB/\lceil n/(mw) \rceil$, where $n = \lceil (mwB\rho + \beta)/B \rceil$, there exists a (ρ, β) -admissible client arrival schedule such that no algorithm can solve the Station Assignment problem, even if all clients have the same arrival time.*

PROOF. From Theorem 5.1, it is enough to prove that $\beta > mwB(n/(mw \lceil n/(mw) \rceil)) - \rho$, for $n = \lceil (mwB\rho + \beta)/B \rceil$ and $\rho > 0$. This holds if $\beta \geq mwB \frac{n/(mw)}{\lceil n/(mw) \rceil} = \frac{nB}{\lceil n/(mw) \rceil}$. \square

Now we show a matching upper bound for solvability in the same optimistic scenario. That is, all clients have the same group of stations and bandwidth.

THEOREM 5.3. *Given any (ρ, β) -admissible client arrival schedule of n clients, such that all clients have the same bandwidth, the same station group of size $m > 0$, and the same arrival time, if $\beta \leq mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right)$, the algorithm that assigns clients evenly among stations and intervals of w times slots solves the Station Assignment problem on any system of at least m stations each with bandwidth B .*

PROOF. Let b be the client bandwidth. In order to show the claim, it is enough to show it for the initial w time slots after the arrival of the clients, given that, if some client departs, the bandwidth usage of the assigned station is reduced. Note that the life interval of all clients is at least w , by the definition of laxity. Given that the assignment of clients is even, the station most used has at most $\lceil n/(mw) \rceil$ clients assigned per slot. Hence, in order to prove the claim, it is enough to prove $\lceil \frac{n}{mw} \rceil b \leq B$. Due to admissibility (Equation (1)) for w slots (i.e., $|T| = w$), we know that $nb \leq mwB\rho + \beta$. Replacing this bound on b , it is enough to show that $\lceil \frac{n}{mw} \rceil \frac{mwB\rho + \beta}{n} \leq B$. Replacing the bound on β , it can be seen that the inequality holds. \square

5.2. Unique Stations Group and Different Client Bandwidth

We now consider a less optimistic scenario where the client bandwidths may be different. Theorems 5.4 and 5.5 show lower bounds for non-solvability on ρ and β respectively.

THEOREM 5.4. *Given a system of m stations each with bandwidth B , even if all clients have the same station group, for any $\rho > 1/2$, there exists a (ρ, β) -admissible client arrival schedule such that no algorithm can solve the Station Assignment problem, even if all clients have the same life interval.*

PROOF. Consider a client arrival schedule of $mw + 1$ clients with the same station group S and the same life interval of length w . One of the clients, call it x , has bandwidth $b = (\rho - \delta)mwB$ for some value δ such that $1/2 < \delta < \rho$ and $\rho - 1/(mw) \leq \delta < (\rho mw - 1)/(mw - 1)$. Each of the remaining mw clients has bandwidth δB . Such schedule is (ρ, β) -admissible since, for any subset of $n \leq mw + 1$ clients that includes x , Equation (1) becomes $\forall T : |T| \leq w : ((n - 1)\delta B + (\rho - \delta)mwB) \frac{|T|}{w} \leq |T|m\rho B + \beta$, which is true because $n - 1 \leq mw$ and $\beta \geq 0$. On the other hand, if we consider the $n \leq mw$ clients that do not include x , Equation (1) becomes $\forall T : |T| \leq w : n\delta B \frac{|T|}{w} \leq |T|m\rho B + \beta$, which is true because $n \leq mw$, $\beta \geq 0$, and $\delta < \rho$. Finally, Equation (2) also holds because $\rho \leq 1$ and hence $\delta B < \rho B \leq B$, and $(\rho - \delta)mwB \leq B$ for $\delta \geq \rho - 1/(mw)$. However, given that there are $mw + 1$ clients, due to the pigeonhole principle two clients have to be assigned to the same slot of the same station. Then, there is a slot in some station such that the sum of the assigned clients is either $2\delta B > B$ or $\delta B + (\rho - \delta)mwB > B$ because $\delta < (\rho mw - 1)/(mw - 1)$. \square

The following theorem shows a lower bound on β for this scenario. The proof uses an adversarial client arrival schedule similar to the schedule used in the proof of Theorem 5.4.

THEOREM 5.5. *Given a system of m stations each with bandwidth B , even if all clients have the same station group, for any $\beta > mB(1/m + 1/2 - \rho)$, there exists a*

(ρ, β)-admissible client arrival schedule such that no algorithm can solve the Station Assignment problem, even if all clients have the same arrival time.

PROOF. Consider a client arrival schedule of $m + 1$ clients with the same station group S , laxity $w = 1$, and life interval $\tau_c \geq w$. One of the clients, call it x , has bandwidth B . Each of the remaining m clients has bandwidth δB , for some value $\delta > 1/2$. We show that for $\beta > mB(1/m + 1/2 - \rho)$ and $\rho > 0$ such schedule is (ρ, β) -admissible as follows. Given that all clients arrive at the same time and are active at least one slot, it is enough to consider the first slot upon arrival. This is true since, if some clients depart, the schedule is still admissible because the station group of all clients is S . Consider any subset of $n \leq m + 1$ clients that includes x . Then, Equation (1) becomes

$$(n - 1)\delta B + B \leq m\rho B + \beta$$

which is true because $n - 1 \leq m$ for $\beta = mB(1/m + \delta - \rho) > mB(1/m + 1/2 - \rho)$. On the other hand, if we consider $n \leq m$ clients that do not include x , Equation (1) becomes

$$n\delta B \leq m\rho B + \beta$$

which is true because $n \leq m$ and $\beta = mB(1/m + \delta - \rho) > mB(1/m + 1/2 - \rho)$. Finally, Equation (2) also holds because $\delta B < B$. However, given that there are $m + 1$ clients and $w = 1$, due to the pigeonhole principle two clients have to be assigned to the same slot of the same station. Then, there is a slot in some station such that the sum of the assigned clients is either $2\delta B > B$ or $\delta B + B > B$. Hence, the Station Assignment problem cannot be solved. \square

Now we show an upper bound for solvability for the same scenario. That is, the group of stations is unique among clients but the bandwidth may be different. This scenario is similar to packing items into mw bins and the major difference is that typically in bin packing the number of bins is not given but the objective to be minimized. First we present the algorithm that computes the assignment, and then we prove its properties.

ALGORITHM 1: Computation of Station Assignment for a set of clients that arrive simultaneously having all the same group of stations S . The transmissions schedule obtained corresponds to time slots in $[1, w]$, and after periodically with period w . Let \mathcal{A} be the output transmissions schedule implemented as a Boolean matrix of size $|C| \times |S| \times w$. Let $B_{s,t}(\mathcal{A})$ be the bandwidth usage on station s and time t in \mathcal{A} , and $B_{\max}(\mathcal{A})$ be the maximum bandwidth usage at any station and time slot in \mathcal{A} .

```

1 for each  $c \in C$  do
2   choose some station  $s \in S$  uniformly at random;
3   choose some time slot  $t \in [1, w]$  uniformly at random;
4    $\mathcal{A}_{c,s,t} \leftarrow true$ ;
5 while  $B_{\max}(\mathcal{A}) > B$  do
6   find station  $s$  and time slot  $t$  such that  $B_{s,t}(\mathcal{A}) = B_{\max}(\mathcal{A})$ ;
7   find stations  $s', s''$  and time slots  $t', t''$  such that  $B_{s',t'}(\mathcal{A}) + B_{s'',t''}(\mathcal{A}) \leq B$ ;
8   for each  $c \in C$  do
9     if  $\mathcal{A}_{c,s',t'} = true$  then
10       $\mathcal{A}_{c,s',t'} \leftarrow false$ ;
11       $\mathcal{A}_{c,s'',t''} \leftarrow true$ ;
12   find some client  $c$  such that  $\mathcal{A}_{c,s,t} = true$ ;
13    $\mathcal{A}_{c,s,t} = false$ ;
14    $\mathcal{A}_{c,s',t'} = true$ ;

```

THEOREM 5.6. *Algorithm 1 computes in polynomial time a Station Assignment for any (ρ, β) -admissible client arrival schedule such that all clients have the same arrival time, the same group of stations of size $m > 0$, each with bandwidth B , the same laxity w , and $\beta < mwB(1/2 - \rho)$. The transmission schedule of such assignment is periodic with period w .*

PROOF. Consider a (ρ, β) -admissible client arrival schedule where all clients have the same station group, arrive simultaneously, and all have laxity w . Let the time slot of arrival be labeled as 1. Algorithm 1 computes the assignment for the first interval of slots $[1, w]$. The assignment in all the subsequent intervals $[iw + 1, (i + 1)w]$, for each integer $i > 0$, is identical.

Algorithm 1 initially assigns each client at random to one of the m stations and one of the w slots (cf. Lines 1 to 4). Call such assignment \mathcal{A} . Then, as long as the maximum bandwidth used in \mathcal{A} in some slot t in some station s is above B , the algorithm finds two stations s', s'' and two time slots t', t'' such that the sum of bandwidth used in s', s'' in time slots t', t'' is at most B (cf. Line 7).

To see why s', s'' and t', t'' exist, assume they do not. Then, adding in pairs, the total bandwidth used throughout all stations and slots is at least $mwB/2$. But, according to Equation (1), the total bandwidth used must be at most $mw\rho B + \beta < mwB/2$. Which is a contradiction.

Then, in Lines 8 to 14, Algorithm 1 reassigns all clients from s', t' to s'', t'' (thus keeping the total bandwidth in s'' at time t'' at most B) and reassigns one client from s at time t to s' at time t' (which was left available). Such client exists since otherwise the client arrival schedule would violate Equation (2).

The polynomial running time follows by simple inspection of Algorithm 1. From Line 6 of the algorithm, when the algorithm stops it holds that $B_{\max} \leq B$. \square

A similar bound can be obtained if clients never depart, even if they arrive at different times.

5.3. Distinct Stations Group and Client Bandwidth

Now we consider the harshest scenario where clients may have different station groups and different bandwidths. Given that $\beta \geq 0$ by definition, the bound obtained implies that if $\rho > 1/m$ the problem is not solvable.

THEOREM 5.7. *Given a system of m stations each with bandwidth B , for each $\beta > mwB(1/(mw) - \rho)$, there exists a (ρ, β) -admissible client arrival schedule such that no algorithm can solve the Station Assignment problem, even if all clients have the same life interval.*

PROOF. Consider a client arrival schedule of $n + 1$ clients, where $n = amw$, for some integer $a \geq 1$, such that $n \geq (mwB\rho + \beta - B)/B$. The first n clients have a singleton station group so that, for each station $s_i, i = 1, 2, \dots, m$, the number of clients with station group $\{s_i\}$ is aw . The bandwidth of each of these n clients is $b = (mwB\rho + \beta - B)/n$. There is one additional client x with station group M and bandwidth B . All the $n + 1$ clients in the client arrival schedule have the same life interval of length $w \geq 1$. Such client arrival schedule is (ρ, β) -admissible because, for any subinterval T such that $|T| \leq w$, the total bandwidth of any subset of $n' \leq n + 1$ clients is, if x is included then $((n' - 1)b + B) \frac{|T|}{w} = ((n' - 1) \frac{mwB\rho + \beta - B}{n} + B) \frac{|T|}{w} \leq (mwB\rho + \beta) \frac{|T|}{w} \leq |T|mB\rho + \beta$. Otherwise, if x is not included, and hence $n' \leq n$, the total bandwidth is $n'b \frac{|T|}{w} = n' \frac{mwB\rho + \beta - B}{n} \frac{|T|}{w} = \frac{n'}{aw} B\rho |T| + n' \frac{\beta - B}{n} \frac{|T|}{w} \leq \left\lceil \frac{n'}{a} \right\rceil |T| B\rho + \beta$. Therefore, Equation (1) holds. Additionally, replacing the expression of n in b , it can be seen that $b \leq B$. Thus, Equation (2) holds for

all clients. However, for any assignment, there must be at least one slot of one station with bandwidth usage $B + ab = B + \frac{n}{mw}b = B + \frac{n}{mw} \frac{mwB\rho + \beta - B}{n} = B(1 + \rho) + \frac{\beta - B}{mw}$, which is bigger than B for $\beta > mwB(1/(mw) - \rho)$. \square

Now we show a matching upper bound for solvability for the same strict scenario. That is, both, the group of stations and bandwidth, may be different among clients. Let a *valid* assignment be one where each client c is scheduled to transmit to some station in S_c at least once in each sequence of w slots. The particular relation between β and ρ in this upper bound guarantees that, as long as the client arrival schedule is (ρ, β) -admissible, any valid Station Assignment solves the problem.

THEOREM 5.8. *For any system of $m > 1$ stations, each with the same bandwidth B , if the client arrival schedule is (ρ, β) -admissible with*

$$\beta \leq mwB \left(\frac{1}{mw} - \rho \right),$$

any valid assignment solves the Station Assignment problem.

PROOF. Consider an assignment of a given (ρ, β) -admissible client arrival schedule. Consider the set $C' \subseteq C$ of clients that are active at any given time step t in such assignment. Because the client arrival schedule is (ρ, β) -admissible, making $|T| = w$ in Equation (1) and using that $w \leq |\tau_c|$, it must be $\sum_{c \in C'} b_c \leq w|S(C')|\rho B + \beta \leq wm\rho B + \beta$. Replacing in the latter the upper bound on β , we have that $\sum_{c \in C'} b_c \leq B$. Thus, no station can have a bandwidth usage bigger than B . \square

6. ANALYSIS OF ONLINE PROTOCOLS

6.1. Lower Bounds for Non-Solvability

We show now conditions under which irrevocable algorithms do not solve the Station Assignment problem. Theorem 6.1 applies to a more general model where the station bandwidths may be different. The corollary that follows instantiates the result on a model where the station bandwidth is unique.

THEOREM 6.1. *For any system of m stations, where station s has bandwidth B_s , any $\beta \geq 0$, and for each irrevocable online algorithm \mathcal{A} , there is a station labeling $\{s_1, \dots, s_m\}$ and a (ρ, β) -client arrival schedule with $b_c = 1, \forall c \in C$, such that, if*

$$\rho > B_{s_m} / \left(B_{s_m} + \sum_{j=1}^{m-1} \left(\sum_{i=j}^m B_{s_i} - \max_{j \in [j, m]} B_{s_j} \right) \frac{1}{m-j+1} \prod_{k=2}^{m-j} \left(1 - \frac{1}{k} \right) \right),$$

\mathcal{A} cannot solve the Station Assignment problem.

PROOF. Consider the following (ρ, β) -client arrival schedule. $\forall c \in C : b_c = 1$. The life interval of all clients is open ended. That is, upon arrival, clients stay active forever. Clients arrive in batches. That is, groups of clients arrive simultaneously, rather than in separate time slots. A new batch of clients arrives after the previous batch has been irrevocably assigned by algorithm \mathcal{A} . Time is conceptually divided into **rounds**, which are enumerated sequentially as $1, 2, \dots, m$. A new round starts when a new batch of clients arrive. That is, clients will arrive in m groups, each group at a different time, and rounds are defined by those arrivals. The set of clients arriving at the beginning of round i is called C_i . All clients arriving in the same round have the same group of stations. Starting from the whole set of stations S in the first round, the group of stations for the following round (i.e., of the next batch of clients) has one station less. We say that such station was **removed**. (Observe that there are exactly m rounds

because after this number of rounds all stations have been removed.) For any station s , let $\gamma_s(i)$ be the number of clients assigned by \mathcal{A} to s up to the end of round i , and let the **available capacity** of station s at the end of round i be $w\rho B_s - \gamma_s(i)$. Then, for round $i + 1$ the station removed is some station s' with the largest available capacity at the end of round i , with ties broken arbitrarily. In other words, s' is not in the group of stations of clients arriving at the beginning of round $i + 1$. For the purpose of this analysis, we label the stations as follows. For any $0 < i < m$ the station removed in round i is labeled s_i , and the station left in round m is labeled s_m . Using this notation, for round i , the group of stations is $\forall c \in C_i : S_c = \{s_i, s_{i+1}, \dots, s_m\}$, and the number of clients injected in round i is

$$|C_i| = \begin{cases} w\rho \max_{j \in [1, m]} B_{s_j} & \text{if } i = 1 \\ w\rho (\max_{j \in [i, m]} B_{s_j} - \max_{j \in [i-1, m]} B_{s_j} + B_{s_{i-1}}) & \text{if } 2 \leq i \leq m \end{cases}$$

To show that this client arrival schedule is admissible, consider any set of clients $C' \subseteq C$. Let $B(C') = \sum_{s \in S(C')} B_s$. To show that Equation (1) holds, it is enough to show

$$\forall T = [t, t'] \subseteq \mathbb{N} : \sum_{c \in C'} \frac{|\tau_c \cap T|}{w} \leq |T|B(C')\rho + \beta$$

which is true if

$$|C'| \leq B(C')w\rho. \quad (3)$$

To show that the latter inequality holds, consider a partition $P = \{C'_1, C'_2, \dots, C'_m\}$ of C' such that, for each $i \in [1, m]$, it is $C'_i = C' \cap C_i$. That is, C'_i is the subset of clients in C' that arrived in round i . From the client arrival schedule, we know that $|C'_i| \leq |C_i|$. Additionally, for each round i such that $|C'_i| > 0$, we know that all stations s_j , $i \leq j \leq m$, are in the group of stations of clients in C'_i . Let $r = \min\{i : |C'_i| > 0\}$. Then, we know that $B(C') = \sum_{i=r}^m B_{s_i}$. After replacing in Equation (3), it is enough to show

$$\sum_{i: |C'_i| > 0} |C'_i| \leq w\rho \sum_{i=r}^m B_{s_i}$$

which is true if

$$\sum_{i=r}^m |C'_i| \leq w\rho \sum_{i=r}^m B_{s_i}.$$

If $r > 1$, it is enough to prove

$$\sum_{i=r}^m \left(\max_{j \in [i, m]} B_{s_j} - \max_{j \in [i-1, m]} B_{s_j} + B_{s_{i-1}} \right) \leq \sum_{i=r}^m B_{s_i}.$$

Expanding and cancelling, it is enough to prove

$$\left(- \max_{j \in [r-1, m]} B_{s_j} + B_{s_{r-1}} \right) + B_{s_r} + \dots + B_{s_m} \leq \sum_{i=r}^m B_{s_i}$$

which is true because $\max_{j \in [r-1, m]} B_{s_j} \geq B_{s_{r-1}}$. For $r = 1$, the derivation is similar but this term does not appear on the left-hand side.

We show now that \mathcal{A} does not solve the Station Assignment problem for this admissible client arrival schedule as follows. Let the total bandwidth of stations $\{s_i, \dots, s_m\}$ be $B_i = \sum_{j=i}^m B_{s_j}$. Let the total number of clients assigned to stations $\{s_{i+1}, \dots, s_m\}$ at the end of round i be $\Gamma_i = \sum_{j=i+1}^m \gamma_{s_j}(i)$. Recall that at the beginning of the first

round $|C_1|$ clients arrive, and that at the end of that round, after \mathcal{A} assigns these clients, the station s_1 is removed from the group of stations for the second round because it has the largest available capacity at the end of round 1. After the assignment of round 1 is done, the available capacity of s_1 is $w\rho B_{s_1} - \gamma_{s_1}(1)$, which, being the largest, is at least the average available capacity $(w\rho B_1 - |C_1|)/m$. Then, s_1 has $\gamma_{s_1}(1) \leq w\rho B_{s_1} - (w\rho B_1 - |C_1|)/m$ clients assigned. Therefore, for the second round, there are

$$\Gamma_1 = |C_1| - \gamma_{s_1}(1) \geq w\rho \left(\left(\max_{j \in [1, m]} B_{s_j} - B_{s_1} \right) \left(1 - \frac{1}{m} \right) + \frac{1}{m} \sum_{i=2}^m B_{s_i} \right)$$

clients that are already assigned to stations in $\{s_2, \dots, s_m\}$, and $|C_2|$ new clients arrive, yielding a total of $|C_2| + \Gamma_1$. After the assignment of round 2 is done, the available capacity of s_2 is the largest, and therefore it is at least the average available capacity. Thus, s_2 has

$$\gamma_{s_2}(1) + \gamma_{s_2}(2) \leq w\rho B_{s_2} - \frac{w\rho B_2 - \Gamma_1 - |C_2|}{m-1}$$

clients assigned. Therefore, for the third round, there are

$$\begin{aligned} \Gamma_2 &= |C_2| - \gamma_{s_2}(1) - \gamma_{s_2}(2) + \Gamma_1 \\ &\geq (|C_2| + \Gamma_1 - w\rho B_{s_2}) \left(1 - \frac{1}{m-1} \right) + \frac{w\rho \sum_{i=2}^m B_{s_i}}{m-1} \\ &\geq w\rho \left(\left(\max_{j \in [2, m]} B_{s_j} - B_{s_2} \right) \left(1 - \frac{1}{m-1} \right) \right. \\ &\quad \left. - \left(\max_{j \in [1, m]} B_{s_j} - B_{s_1} \right) \frac{1}{m} \left(1 - \frac{1}{m-1} \right) \right) \\ &\quad + \sum_{i=2}^m B_{s_i} \frac{1}{m} \left(1 - \frac{1}{m-1} \right) + \sum_{i=3}^m B_{s_i} \frac{1}{m-1} \end{aligned}$$

clients that are already assigned to stations in $\{s_3, \dots, s_m\}$, and $|C_3|$ new clients arrive, yielding a total of $|C_3| + \Gamma_2$. Proceeding similarly, for the m -th round, there are $\Gamma_{m-1} = |C_{m-1}| - \sum_{i=1}^{m-1} \gamma_{s_{m-1}}(i) + \Gamma_{m-2}$ clients that are already assigned to station s_m , and $|C_m|$ new clients arrive. Lower bounding Γ_{m-1} as before, it yields a total of

$$\begin{aligned} |C_m| + \Gamma_{m-1} &\geq \\ &w\rho \left(B_{s_m} + \sum_{j=1}^{m-1} \left(\sum_{i=j}^m B_{s_i} - \max_{j \in [j, m]} B_{s_j} \right) \frac{1}{m-j+1} \prod_{k=2}^{m-j} \left(1 - \frac{1}{k} \right) \right). \end{aligned}$$

However, s_m may receive at most wB_{s_m} transmissions within w steps. Thus, for any

$$\rho > \frac{B_{s_m}}{B_{s_m} + \sum_{j=1}^{m-1} \left(\sum_{i=j}^m B_{s_i} - \max_{j \in [j, m]} B_{s_j} \right) \frac{1}{m-j+1} \prod_{k=2}^{m-j} \left(1 - \frac{1}{k} \right)}$$

the Station Assignment problem cannot be solved. \square

COROLLARY 6.2. *For any system of m stations each with bandwidth B , and for each irrevocable algorithm \mathcal{A} , and for any $\rho > 1/(1 + \ln m)$ and $\beta \geq 0$, there is a (ρ, β) -client arrival schedule with $b_c = 1, \forall c \in C$, such that \mathcal{A} cannot solve the Station Assignment problem.*

PROOF. Replacing all bandwidths in the lower bound of ρ in Theorem 6.1 by B , we get $\rho > \left(\sum_{j=1}^m 1/j\right)^{-1} = H_m^{-1} > 1/(1 + \ln m)$. \square

Observe that for the above proof to work it is not needed that an irrevocable algorithm assigns a client to a station forever. It is enough that it assigns it for $m + w$ steps to reach the same result.

The following theorem for irrevocable algorithms relates β and ρ for the case where the bandwidth of all stations is the same and $m > 1$. Given that $\beta \geq 0$ by definition, the bound implies that if $\rho > 1/\ln m$, the Station Assignment problem is not solvable.

THEOREM 6.3. *For any system of $m > 1$ stations, such that all stations have the same bandwidth B , and for each irrevocable algorithm \mathcal{A} , there is a (ρ, β) -client arrival schedule with $b_c = 1, \forall c \in C$, and $w = 1$ such that, if $\beta > mB(1/\ln m - \rho)$, then \mathcal{A} cannot solve the Station Assignment problem.*

PROOF. Consider an adaptive adversary that decides which clients arrive according to the actions of \mathcal{A} . The adversarial client arrival schedule is the following. Let $w = 1$. For each client c , it is $b_c = 1$. The life interval of all clients is open ended. That is, upon arrival, clients stay active forever. Clients arrive in batches. A new batch of clients arrive after the previous batch has been irrevocably assigned by algorithm \mathcal{A} . Time is conceptually divided in m **rounds**, which are enumerated sequentially as $1, 2, \dots, m$. A new round starts when a new batch of clients arrive. The number of clients arriving in each round is $\rho B + \beta/m$. (We omit ceilings and floors throughout the proof for clarity.) All clients arriving in the same round i have the same group of stations S_i . Starting from the whole set of stations S in the first round, the group of stations for each new round is reduced by one station. We say that such station is **removed**. Thus, for round 1 the group of stations has size m , for round 2 the size is $m - 1$, and so on until round m when the group of stations has size 1. For any round $r > 1$, the station removed is the station with the smallest number of clients assigned.

First we notice that the client arrival schedule defined is (ρ, β) -admissible. For this purpose, it is enough to show that the property is preserved after each batch of arrivals. Consider any round $i = 1, \dots, m$. Let C_j be any subset of clients arriving in round $j = 1, \dots, i$ with group of stations S_j . Given that, by definition of the client arrival schedule, $\rho B + \beta/m$ clients arrive in each round, we know that $|C_j| \leq \rho B + \beta/m$. So, in Equation (1), the ρB term can be applied to the station removed in round j , and putting together all the β/m terms they add up to $i\beta/m \leq \beta$.

We show now that, with the above client arrival schedule, the sum of the bandwidths of the clients assigned to the station in S_m is more than B . Let the number of clients arriving in each round be called $X = \rho B + \beta/m$. In round 1 the overall number of clients is X . Given that the station removed is the one with the smallest number of clients, in round 2 the overall number of clients assigned to stations in S_2 is at least $X(1 - 1/m) + X$. Likewise, in round 3, the overall number of clients assigned to stations in S_3 is at least $((X(1 - 1/m) + X)(1 - 1/(m - 1)) + X$. Inductively, the number of clients assigned to the station in S_m is at least

$$\begin{aligned} & \left(\dots \left(\left(X \frac{m-1}{m} + X \right) \frac{m-2}{m-1} + X \right) \frac{m-3}{m-2} \dots \right) \frac{1}{2} + X = \\ & X \left(\frac{1}{m} + \frac{1}{m-1} + \dots + \frac{1}{2} + 1 \right) > X \ln m. \end{aligned}$$

That is, the total bandwidth of the clients assigned to the station in S_m is at least $\ln m(\rho B + \beta/m)$. Thus, if $\beta > mB(1/\ln m - \rho)$ the claim follows. \square

The following theorem shows that the restriction on ρ for solvability with irrevocable assignments is stronger for improvident algorithms. Theorem 6.4 shows that, for randomized online algorithms, if $\beta > mB(3/\sqrt{2m} - \rho)$ the Station Assignment problem is not solvable, and if $\beta > mB(1/\sqrt{2m} - \rho)$ the Station Assignment problem is not solvable online deterministically. Given that $\beta \geq 0$ by definition, the bound implies that if $\rho > 3/\sqrt{2m}$, or if $\rho > 1/\sqrt{2m}$ respectively, the problem is not solvable.

THEOREM 6.4. *For any set of m stations each with bandwidth B , the following holds, even if all clients have the same bandwidth:*

- (1) *For any $m \geq 5$ and $\beta > mB(3/\sqrt{2m} - \rho)$, there exists a (ρ, β) -admissible client arrival schedule such that no online irrevocable improvident randomized algorithm can solve Station Assignment.*
- (2) *For any $m \geq 1$ and $\beta > mB(1/\sqrt{2m} - \rho)$, there exists a (ρ, β) -admissible client arrival schedule with $w = 1$ such that no online irrevocable improvident deterministic algorithm can solve Station Assignment.*

PROOF. If $\beta > mB(1 - \rho)$, the claim follows from Theorem 5.1. So, for the rest of the proof we assume that $\beta \leq mB(1 - \rho)$.

For the Load Balancing problem, where computing tasks have to be assigned to servers, the proof of Theorem 3.3 in [Azar et al. 1994] shows a sequence of unit-weight tasks such that, the maximum (over the servers) off-line load at all times is 1, and the competitive ratio of any randomized irrevocable improvident algorithm is at least $(\sqrt{2m}/3)(1 + o(1))$. (The theorem is stated in asymptotic notation, but the bound obtained in the proof is the expression given here.) We reuse such adversary mapping tasks to clients, servers to stations and weights/loads to bandwidths. Let the bandwidth of such clients be instead $\rho B + \beta/m$ and the laxity $w = 1$. This client arrival schedule is (ρ, β) -admissible because (i) $\beta \leq mB(1 - \rho)$ and then $\rho B + \beta/m \leq B$, and (ii) the maximum off-line bandwidth at all times on each station is at most $\rho B + \beta/m$. However, the bandwidth used at some station is at least $(\sqrt{2m}/3)(\rho B + \beta/m)$, which is larger than B if $\beta > mB(3/\sqrt{2m} - \rho)$, which is feasible for $m > 9/2$. The same argument can be used for deterministic algorithms and competitive ratio of $\sqrt{2m}$. \square

6.2. Upper Bounds for Solvability

The following upper bound applies to a setting where the station bandwidth is unique and all laxities are $w = 1$, but the station group may be different for each client. First we present the algorithm that computes the assignment, and then we prove its properties.

ALGORITHM 2: Online computation of Station Assignment for a set of stations S with unique bandwidth B , and clients with unit laxity. C_s is the set of clients assigned to station s .

```

1 for each station  $s \in S$  do
2    $C_s \leftarrow \emptyset$ ;
3 for each arriving client  $c$  do
4   find a station  $s \in S_c$  such that  $\sum_{c' \in C_s} b_{c'} = \min_{s' \in S_c} \sum_{c' \in C_{s'}} b_{c'}$ ;
5    $C_s \leftarrow C_s \cup \{c\}$ ;

```

THEOREM 6.5. *For any system of m stations each with bandwidth B , if $\rho < 1/(2(m+1))$, $\beta \leq \rho B$, and all clients have the same bandwidth $b \leq B$, laxity $w = 1$, and never depart, Algorithm 2 solves the Station Assignment problem.*

PROOF. For the sake of contradiction, assume that some station gets overloaded. Consider the clients in the order they arrive and are assigned to stations. In particular, the first client \hat{c} whose assignment (at some time slot t) makes some station $s \in S$ to be overloaded. Let us assume, without loss of generality, that b divides exactly B ,¹ and denote $r \triangleq B/b$. Then, after the assignment of \hat{c} , we have that s has exactly $r + 1$ clients assigned, and the rest of stations no more than r clients assigned each. That is, the total number of clients assigned is at most $rm + 1$. We will show that in this scenario the actual number of clients has to be higher, reaching a contradiction.

Let \hat{C} be the set of $r + 1$ clients assigned to s up to \hat{c} . Recall that, for any $C' \subseteq \hat{C}$, it is $S(C') = \bigcup_{c \in C'} S_c$. Also, given that the bandwidth of all clients is the same value b , we have that $\sum_{c \in C'} b_c = |C'|b$. Replacing in Equation 1 and considering as time interval the time slot t , we have for any $C' \subseteq \hat{C}$,

$$\begin{aligned} |C'|b &\leq |S(C')|\rho B + \beta, \text{ given that } \beta \leq \rho B, \\ &\leq (|S(C')| + 1)\rho B. \end{aligned}$$

Hence, the following holds.

$$|S(C')| \geq \frac{|C'|}{\rho r} - 1, \text{ for any } C' \subseteq \hat{C}. \quad (4)$$

Consider a labeling of clients in $\hat{C} = \{c_0, c_1, c_2, \dots, c_r = \hat{c}\}$, where the subindex is the order in which the clients were assigned to s . Let us define subsets C_i containing the latest $r - i + 1$ clients assigned to s up to $c_r = \hat{c}$. That is,

$$C_i = \{c_i, \dots, c_r\}, \text{ for each } i = 0, \dots, r.$$

Then, from Equation 4 we have that

$$|S(C_i)| \geq \frac{r - i + 1}{\rho r} - 1, \text{ for each } i = 0, \dots, r. \quad (5)$$

The following observations will be used. Observe that $S(C_i) \subseteq S(C_{i-1})$, since $C_i \subset C_{i-1}$. The set of stations $S(C_i)$ contains by definition the stations to which at least one of the clients $c_j \in C_i$ could be assigned, where $j \geq i$. Also, notice that Algorithm 2 assigns c_j to the station in S_{c_j} with the largest available bandwidth, breaking ties arbitrarily. Then, when c_j (which is the $j + 1$ st client assigned to s) was processed in Line 4 of Algorithm 2, all stations in S_{c_j} had at least j clients already assigned. In order to lower bound the number of clients in all the stations after $\hat{c} = c_r$ is assigned to s , we will use the latter argument for each station in $S(\hat{C})$.

Let the number of clients in all stations after $\hat{c} = c_r$ is assigned to s be denoted as X . We associate to each station $s' \in S(\hat{C})$ the largest j , denoted $k(s')$, such that $s' \in S_{c_j}$. In other words, we associate to each station s' the maximum of the minimum number of clients that should have been assigned to s' , due to assignments to s . Then, we have the following lower bound.

$$X \geq 1 + \sum_{s' \in S(\hat{C})} k(s'). \quad (6)$$

¹Otherwise, the stations' bandwidth can be redefined to a value B' such that $B' = \lfloor B/r \rfloor b$.

Where the additive one corresponds to client \hat{c} . If we show that the latter sum is $\sum_{s' \in S(\hat{C})} k(s') > rm$ we reach a contradiction proving the theorem. So, in the rest of the proof we focus on proving the latter inequality.

First, we re-write the sum in Equation 8 grouping the stations that have the same $k(\cdot)$, as follows.

$$\begin{aligned}
X &\geq 1 + r|S(C_{\bar{r}})| + \sum_{i=1}^{r-1} i(|S(C_{\bar{i}})| - |S(C_{\bar{i+1}})|) \\
&= 1 + \sum_{i=1}^r |S(C_{\bar{i}})| \\
&\geq 1 + \sum_{i=1}^r \left(\frac{r-i+1}{\rho r} - 1 \right), \text{ from Equation 5,} \\
&= 1 + \frac{r(r+1)}{2\rho r} - r = 1 + \frac{r+1}{2\rho} - r.
\end{aligned}$$

Since $\rho < \frac{1}{2(m+1)} < \frac{r+1}{2(m+1)r}$, it is $X > rm + 1$, which is a contradiction. \square

THEOREM 6.6. *For any system of m stations each with bandwidth B , if $\beta \leq \rho B$, and all clients have the same bandwidth b such that $\rho B \leq b < B/(m-1)$, laxity $w = 1$, and never depart, Algorithm 2 solves the Station Assignment problem.*

PROOF. For the sake of contradiction, assume that some station gets overloaded. Consider the clients in the order they arrive and are assigned to stations. In particular, the first client \hat{c} whose assignment (at some time slot t) makes some station $s \in S$ to be overloaded. Let us assume, without loss of generality, that b divides exactly B ,² and denote $r \triangleq B/b$. Then, after the assignment of \hat{c} , we have that s has exactly $r + 1$ clients assigned, and the rest of stations no more than r clients assigned each. That is, the total number of clients assigned is at most $rm + 1$. We will show that in this scenario the actual number of clients has to be higher, reaching a contradiction.

Let \hat{C} be the set of $r + 1$ clients assigned to s up to \hat{c} . Recall that, for any $C' \subseteq \hat{C}$, it is $S(C') = \bigcup_{c \in C'} S_c$. The following property holds.

PROPERTY 6.7. $\forall C' \subseteq \hat{C} : |C'| \leq |S(C')|$.

That is, for each group of clients of cardinality x , there are at least x stations to which these clients can be assigned. To see this, consider a set $C' \subseteq \hat{C}$, the corresponding set $S(C')$, and considering as time interval the time slot t . Given that the bandwidth of all clients is the same, we have that $\sum_{c \in C'} b_c = |C'|b$. Replacing in Eq. 1 we have the following.

$$\begin{aligned}
|C'|b &\leq |S(C')|\rho B + \beta, \text{ given that } \beta < \rho B, \\
&< (|S(C')| + 1)\rho B, \text{ given that } b \geq \rho B, \\
&\leq (|S(C')| + 1)b,
\end{aligned}$$

which proves the property.

Consider a labeling of clients in $\hat{C} = \{c_0, c_1, c_2, \dots, c_r = \hat{c}\}$, where the subindex is the order in which the clients were assigned to s . Let us define subsets $C_{\bar{i}}$ containing the

²Otherwise, the stations' bandwidth can be redefined to a value B' such that $B' = \lfloor B/r \rfloor b$.

latest $r - i + 1$ clients assigned to s up to $c_r = \hat{c}$. That is,

$$C_{\bar{i}} = \{c_i, \dots, c_r\}, \text{ for each } i = 0, \dots, r.$$

Then, from Property 6.7 we have that

$$|S(C_{\bar{i}})| \geq r - i + 1, \text{ for each } i = 0, \dots, r. \quad (7)$$

The following observations will be used. Observe that $S(C_{\bar{i}}) \subseteq S(C_{\bar{i-1}})$, since $C_{\bar{i}} \subset C_{\bar{i-1}}$. The set of stations $S(C_{\bar{i}})$ contains by definition the stations to which at least one of the clients $c_j \in C_{\bar{i}}$ could be assigned, where $j \geq i$. Also, notice that Algorithm 2 assigns c_j to the station in S_{c_j} with the largest available bandwidth, breaking ties arbitrarily. Then, when c_j (which is the $j + 1$ st client assigned to s) was processed in Line 4 of Algorithm 2, all stations in S_{c_j} had at least j clients already assigned. In order to lower bound the number of clients in all the stations after $\hat{c} = c_r$ is assigned to s , we will use the latter argument for each station in $S(\hat{C})$.

Let the number of clients in all stations after $\hat{c} = c_r$ is assigned to s be denoted as X . We associate to each station $s' \in S(\hat{C})$ the largest j , denoted $k(s')$, such that $s' \in S_{c_j}$. In other words, we associate to each station s' the maximum of the minimum number of clients that should have been assigned to s' , due to assignments to s . Then, we have the following lower bound.

$$X \geq 1 + \sum_{s' \in S(\hat{C})} k(s'). \quad (8)$$

Where the additive one corresponds to client \hat{c} . If we show that the latter sum is $\sum_{s' \in S(\hat{C})} k(s') > rm$ we reach a contradiction proving the theorem. So, in the rest of the proof we focus on proving the latter inequality.

First, we re-write the sum in Equation 8 grouping the stations that have the same $k(\cdot)$, as follows.

$$\begin{aligned} X &\geq 1 + r|S(C_{\bar{r}})| + \sum_{i=1}^{r-1} i(|S(C_{\bar{i}})| - |S(C_{\bar{i+1}})|) \\ &= 1 + \sum_{i=1}^r |S(C_{\bar{i}})| \\ &\geq 1 + \sum_{i=1}^r (r - i + 1), \text{ from Equation 7,} \\ &= 1 + r(r + 1). \end{aligned}$$

Since $r + 1 = B/b + 1 > m$, it is $X > rm + 1$, which is a contradiction. \square

7. CONCLUSIONS

This paper presented worst-case (adversarial) analysis of scheduling periodic communication between base stations and mobile clients. We considered various classes of scheduling settings and protocols, and provided limitations on feasible mobility patterns given in the form of upper and lower bounds on client injection rates and burstiness.

The separation results obtained expose the dependency of the complexity of Station Assignment on model assumptions. The question of whether other combinations could lead to more separations is enticing. For instance, what are the bounds on the adversary limitations for offline protocols when the bandwidth is the same for all clients

but the group of stations is not. Or what is the impact in online protocols of having different client bandwidths if the group of stations is the same for all clients. Also, in our study, we assumed that all clients have the same laxity, but the additional combinatorial problem of having different laxities may yield more complexity separations. The answers to these questions are open.

In this paper we explored the solvability of Station Assignment, without evaluating the performance of the algorithms proposed. A future line of work could explore these algorithms (and other that may be proposed) in terms of some goodness parameter that would have to be defined. For instance, one could allow the algorithms to refuse service to some of the arriving clients, and measure their goodness by the volume of clients that are in fact granted service. The analyses may also explore the competitiveness of the algorithms, and their performance in simulated scenarios and real traces.

ACKNOWLEDGMENTS

This research was partially supported by the UK Royal Society International Exchanges 2017 Round 3 Grant #170293, Pace University SRC and Kenan Fund, the initiative Networks Sciences & Technologies (NeST) by School of EEE & CS, University of Liverpool, the Polish National Science Center (NCN) grant UMO-2017/25/B/ST6/02553, the Regional Government of Madrid (CM) grant Cloud4BigData (S2013/ICE-2894), cofunded by FSE & FEDER, the Spanish Ministry of Science, Innovation and Universities grant DiscoEdge (TIN2017-88749-R), the EU Commission H2020 project RECAP (RIA 732667), and the NSF of China grant 61520106005.

REFERENCES

- Noga Alon, Yossi Azar, Gerhard J. Woeginger, and Tal Yadid. 1997. Approximation Schemes for Scheduling. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 493–500.
- L. Anantharamu, Bogdan S. Chlebus, and Mariusz A. Rokicki. 2009. Adversarial multiple access channel with individual injection rates. In *Proceedings of the 13th International Conference on Principles of Distributed Systems (OPODIS) (LNCS 5923)*. Springer-Verlag, 174–188.
- Matthew Andrews, Baruch Awerbuch, Antonio Fernández, Frank Thomson Leighton, Zhiyong Liu, and Jon M. Kleinberg. 2001. Universal-stability results and performance bounds for greedy contention-resolution protocols. *J. ACM* 48, 1 (2001), 39–69.
- Matthew Andrews and Lisa Zhang. 2005. Scheduling over a time-varying user-dependent channel with applications to high-speed wireless data. *J. ACM* 52, 5 (2005), 809–834.
- James Aspnes, Yossi Azar, Amos Fiat, Serge A. Plotkin, and Orli Waarts. 1993. On-line load balancing with applications to machine scheduling and virtual circuit routing. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC)*. 623–631.
- Yossi Azar. 1996. On-line Load Balancing. In *Online Algorithms, The State of the Art (the book grow out of a Dagstuhl Seminar, June 1996) (Lecture Notes in Computer Science)*, Amos Fiat and Gerhard J. Woeginger (Eds.), Vol. 1442. Springer, 178–195. DOI : <http://dx.doi.org/10.1007/BFb0029569>
- Yossi Azar, Andrei Z. Broder, and Anna R. Karlin. 1994. On-Line Load Balancing. *Theor. Comput. Sci.* 130, 1 (1994), 73–84.
- Yossi Azar, Bala Kalyanasundaram, Serge A. Plotkin, Kirk Pruhs, and Orli Waarts. 1997. On-Line Load Balancing of Temporary Tasks. *Journal of Algorithms* 22, 1 (1997), 93–110.
- Amotz Bar-Noy and Richard E. Ladner. 2003. Windows Scheduling Problems for Broadcast Systems. *SIAM J. Comput.* 32, 4 (2003), 1091–1113. DOI : <http://dx.doi.org/10.1137/S009753970240447X>
- Amotz Bar-Noy, Richard E. Ladner, and Tami Tamir. 2007. Windows scheduling as a restricted version of bin packing. *ACM Trans. Algorithms* 3, 3 (2007), 28. DOI : <http://dx.doi.org/10.1145/1273340.1273344>
- Marcin Bienkowski, Tomasz Jurdzinski, Mirosław Korzeniowski, and Dariusz R. Kowalski. 2012. Distributed Online and Stochastic Queuing on a Multiple Access Channel. In *Proceedings of the 26th International Symposium on Distributed Computing (DISC)*. 121–135.
- Maria J. Blesa, Daniel Calzada, Antonio Fernández, Luis López, Andrés L. Martínez, Agustín Santos, Maria J. Serna, and Christopher Thraves. 2009. Adversarial Queueing Model for Continuous Network Dynamics. *Theory of Computing Systems* 44, 3 (2009), 304–331.
- Allan Borodin, Jon M. Kleinberg, Prabhakar Raghavan, Madhu Sudan, and David P. Williamson. 2001. Adversarial queueing theory. *J. ACM* 48, 1 (2001), 13–38.

- Joseph Wun-Tat Chan, Tak Wah Lam, and Prudence W. H. Wong. 2008. Dynamic Bin Packing of Unit Fractions Items. *Theoretical Computer Science* 409, 3 (2008), 172–206.
- Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. 2012. Adversarial Queuing on the Multiple Access Channel. *ACM Trans. Algorithms* 8, 1, Article 5 (Jan. 2012), 31 pages.
- R. L. Graham. 1966. Bounds on Multiprocessing Timing Anomalies. *The Bell System Technical Journal* 45 (1966), 1563–1581.
- David R. Karger, Steven J. Phillips, and Eric Torng. 1996. A Better Algorithm For an Ancient Scheduling Problem. *Journal of Algorithms* 20 (1996), 400–430.
- Thomas Kesselheim. 2012. Dynamic packet scheduling in wireless networks. In *Proceedings of the 31st ACM Symposium on Principles of Distributed Computing (PODC)*. 281–290.
- Kirk Pruhs, Jirí Sgall, and Eric Torng. 2004. Online Scheduling. In *Handbook of Scheduling - Algorithms, Models, and Performance Analysis.*, Joseph Y.-T. Leung (Ed.). Chapman and Hall/CRC, Chapter 15, pp. 15–1 to 15–41. <http://www.crcnetbase.com/isbn/978-1-58488-397-5>

Received ; revised ; accepted