

Station Assignment with Applications to Sensing [★]

Antonio Fernández Anta¹, Dariusz R. Kowalski²,
Miguel A. Mosteiro³, and Prudence W. H. Wong²

¹ Institute IMDEA Networks, Madrid, Spain,
antonio.fernandez@imdea.org

² University of Liverpool, Department of Computer Science, Liverpool, UK,
{D.Kowalski,pwong}@liverpool.ac.uk

³ Kean University, Department of Computer Science, Union, NJ, USA,
& Univ. Rey Juan Carlos, GSyC, Madrid, Spain,
mmosteir@kean.edu

Abstract. We study an allocation problem that arises in various scenarios. For instance, a health monitoring system where ambulatory patients carry sensors that must periodically upload physiological data. Another example is participatory sensing, where communities of mobile device users upload periodically information about their environment. We assume that devices or sensors (generically called *clients*) join and leave the system continuously, and they must upload/download data to static devices (or *base stations*), via radio transmissions. The mobility of clients, the limited range of transmission, and the possibly ephemeral nature of the clients are modeled by characterizing each client with a *life interval* and a *stations group*, so that different clients may or may not coincide in time and/or stations to connect. The intrinsically shared nature of the access to base stations is modeled by introducing a maximum *station bandwidth* that is shared among its connected clients, a client *laxity*, which bounds the maximum time that an active client is not transmitting to some base station, and a *client bandwidth*, which bounds the minimum bandwidth that a client requires in each transmission. Under the model described, we study the problem of assigning clients to base stations so that every client transmits to some station in its group, limited by laxities and bandwidths. We call this problem the *Station Assignment* problem. We study the impact of the rate and burstiness of the arrival of clients on the solvability of Station Assignment. To carry out a worst-case analysis we use a typical adversarial methodology: 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 the burstiness of the stations load (hence, limiting the rate and burstiness of the client arrivals). Specifically, we show upper and lower bounds on the rate and burstiness of the arrival for various client arrival schedules and protocol classes. The problem has connections with Load Balancing and Scheduling, usually studied using competitive analysis. To the best of our knowledge, this is the first time that the Station Assignment problem is studied under adversarial arrivals.

Keywords: Station Assignment, Periodic Sensing, Health Monitoring Systems, Participatory Sensing, Continuous Adversarial Dynamics.

1 Introduction

We study a dynamic allocation problem that arises in various scenarios where data sensed using mobile devices has to be gathered using one of many static access points available. Examples include *wearable health-monitoring systems*, where ambulatory patients carry physiological sensors, and the data gathered must be periodically uploaded, and *participatory sensing*, where communities of mobile device users upload periodically information about their environment. We call this problem *Station Assignment*.

[★] This work was supported in part by the National Science Foundation (CCF-0937829, CCF-1114930), the Comunidad de Madrid (S2009TIC-1692), the Spanish MICINN/MINECO (TEC2011-29688-C02-01), and the National Natural Science Foundation of China (61020106002).

We study the Station Assignment problem assuming a continuous arrival of mobile devices, called *clients*, who have to upload data to static devices, called *base stations*, via radio transmissions. The mobility of clients, the limited range of transmission, and the possibly limited time that they will be uploading to a given station, are modeled by assigning to clients a *life interval* and a *stations group*. That is, as clients move, the set of stations that they could possibly upload data may change. We model this event as a client that departs and a new client that arrives. The intrinsically shared nature of the base stations is modeled by introducing a maximum *station bandwidth* that a station can share among its connected clients, a client *laxity*, which bounds the maximum time that an active client is not transmitting to some base station, and a *client bandwidth*, which bounds the minimum bandwidth that a client requires in each transmission.

Under the described model, we study the problem of assigning clients to base stations so that every client transmits to some station in its group, limited by laxities and bandwidths. We study the impact of the frequency of client arrivals on the solvability of Station Assignment. Specifically, we show upper and lower bounds on the rate and burstiness of the client arrivals for solvability of Station Assignment under various client arrival schedules and protocol classes. To carry out a worst-case analysis, we use a typical adversarial methodology: 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 time when the Station Assignment problem is studied 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 *stations group*, 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 . In this work we assume the laxity to be some value w , which is the same for all clients.
- 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 . In this work we assume the station bandwidth to be some value B , which is the same for all stations.

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*.

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 [11,3,10].

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 1 and 2. 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. Lower bounds are for impossibility whereas upper bounds are for solvability.

We study offline Station Assignment under various model assumptions, starting from a more optimistic one where all clients have the same bandwidth and the same stations group, and removing gradually assumptions making the model more pessimistic and, hence, realistic. Studying different models gives insight on what the inherent challenges of Station Assignment are, c.f., Table 1.

We start considering adversarial client arrival schedules where all clients have the same stations group and bandwidth. Then, Theorem 1 shows that for each $\beta > mwB ((n/(mw)) / \lceil n/(mw) \rceil - \rho)$, where $n = \lceil (mwB\rho + \beta) / B \rceil$, there exists a (ρ, β) -admissible client arrival schedule such that no Station Assignment algorithm can solve the problem, even if all clients arrive simultaneously and have the same life interval. Given that it must be $\beta \geq 0$, this lower bound for non-solvability implies also a lower bound of $\rho >$

b_c	S_c	arrival time	offline protocol class	β	ρ	Theorem
identical	identical	identical	any	$> mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right)$ $n = \lceil \frac{mwB\rho + \beta}{B} \rceil$	$> \frac{n/(mw)}{\lceil n/(mw) \rceil}$	1
identical	identical	identical	even assignment	$\leq mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right)$	$\leq \frac{n/(mw)}{\lceil n/(mw) \rceil}$	2
distinct	identical	identical	any	—	$> 1/2$	3
distinct	identical	identical	any	$> mB(1/m + 1/2 - \rho)$	—	4
any	identical	identical	balance station-bandwidth usage	$< mwB(1/2 - \rho)$	$< 1/2$	5
distinct	distinct	identical	any	$> mwB(1/(mw) - \rho)$	$> 1/(mw)$	6
any	any	any	any	$\leq mwB(1/(mw) - \rho)$	$\leq 1/(mw)$	7

Table 1. Summary of bounds on problem solvability for offline protocols.

b_c	S_c	arrival time	τ_c	online protocol class	β	ρ	Theorem
1	distinct	distinct	open	irrevocable	$> mB \left(\frac{1}{\ln m} - \rho \right)$	$> \frac{1}{\ln m}$	9
1	distinct	distinct	distinct	irrevocable improvident randomized	$> mB \left(\frac{3}{\sqrt{2m}} - \rho \right)$	$> \frac{3}{\sqrt{2m}}$	10
1	distinct	distinct	distinct	irrevocable improvident deterministic	$> mB \left(\frac{1}{\sqrt{2m}} - \rho \right)$	$> \frac{1}{\sqrt{2m}}$	10
$b \geq \rho B$	any	any	open	irrevocable improvident	$< \rho B$	$\leq \frac{1}{1 + \sqrt{2m}}$	11

Table 2. Summary of bounds on problem solvability for online protocols.

$(n/(mw))/\lceil n/(mw) \rceil$. Corollary 1 shows a stronger bound on β that holds for any positive ρ . Under the same conditions, Theorem 2 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 stations group is still the same for all. In this scenario, Theorem 3 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 4 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 3. Under the same conditions, Theorem 5 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 station groups or bandwidths. Theorem 6 shows that, for each $\beta > mB(1/m - \rho)$ (which implies $\rho > 1/m$ because $\beta \geq 0$) 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. Theorem 7 matches those bounds, showing that if $\rho \leq 1/(mw)$ and $\beta \leq mwB(1/(mw) - \rho)$ the Station Assignment problem is solvable using any algorithm for *any* client arrival schedule.

Moving to online protocols, c.f., Table 2, by showing a reduction from Load Balancing [7], we prove in Theorem 9 that for any irrevocable algorithm, that is, algorithms where the station-client assignments are final, there is a client arrival schedule such that if $\beta > mB(1/\ln m - \rho)$ the Station Assignment problem is not solvable. Again, the 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, then Theorem 10 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 11 shows that, when all clients have the same bandwidth $b \geq \rho B$ and do not depart, even if the station groups and arrival times are different, if $\rho \leq 1/(1 + \sqrt{2m})$ and $\beta < \rho B$ the algorithm that distributes clients evenly (restricted to station group) solves Station Assignment.

We also show in Theorem 8 (not included in Table 2) a lower bound on ρ for non-solvability with irrevocable algorithms that applies to systems with distinct station-bandwidths. Corollary 2 shows that instantiating Theorem 8 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 [3,11], 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 [4], radio channels with collisions [12,2], or SINR networks [15]. In a single-hop radio channels with collisions, more detail competitive analysis of dynamic and stochastic traffic was performed [9]. The difference between this line of research and our work is that it considered simple packet forwarding requests without additional scheduling constraints.

In [7], Azar, Broder, and Karlin studied a load balancing problem where a set of tasks that arrive and depart in time (temporary, as opposed to permanent when tasks do not depart) have to be assigned to a set of machines. Each task has an associated weight that represents the load that the processing of such 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 authors 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 [8]. 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 3 gives a summary of the results.

In [1], 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

	Unknown duration	Known duration	Permanent
Identical	$2 - o(1)$ [13,7]	$2 - o(1)$ [13,7]	$2 - \epsilon$ [14]
Related	$\Theta(1)$ [8]	$\Theta(1)$ [8]	$\Theta(1)$ [8]
Restricted	$O(\sqrt{m})$ [8] $\Omega(\sqrt{m})$ [7]	$O(\log mT)$ [8]	$\Theta(\log m)$ [5]

Table 3. Competitive ratios of load balancing problem

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 [16,6].

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 Identical Stations Group and Client Bandwidth

We start with a very optimistic scenario (for Station Assignment algorithms) where all clients have the same stations group 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 1. *Given a system of m stations each with bandwidth B , even if all clients must have the same stations group 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 must 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 $\beta > mwB \left((n/(mw))/\lceil n/(mw) \rceil - \rho \right)$, 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 1. *Given a system of m stations each with bandwidth B , even if all clients must have the same stations group 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 must have the same arrival time.*

Proof. Let $\rho > 0$, from Theorem 1 it is enough to prove the claim that $\beta > mwB \left(\frac{n/(mw)}{\lceil n/(mw) \rceil} - \rho \right)$, where $n = \lceil (mwB\rho + \beta)/B \rceil$. 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 stations group and bandwidth.

Theorem 2. *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 Identical Stations Group and Distinct Client Bandwidth

We now consider a less optimistic scenario where the client bandwidths may be different. Theorems 3 and 4 show lower bounds for non-solvability on ρ and β respectively.

Theorem 3. *Given a system of m stations each with bandwidth B , even if all clients must 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 must 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 because, 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 $\delta B < \rho B \leq B$ because $\rho \leq 1$, 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 3. The details are left to the full version of this paper for brevity.

Theorem 4. *Given a system of m stations each with bandwidth B , even if all clients must 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 must have the same arrival time.

Now we show an upper bound for solvability for the same scenario. That is, the stations group is identical among clients but the bandwidth may be different.

Theorem 5. *Given any (ρ, β) -admissible client arrival schedule, such that all clients have the same station group of size $m > 0$ and the same arrival time, if $\beta < mwB(1/2 - \rho)$, there exists a polynomial time algorithm that computes an assignment of clients to stations that solves the Station Assignment problem on any system of at least m stations each with bandwidth B . 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 clients arrival be labeled as 1. We will focus on the first interval of slots $[1, w]$. Notice that all clients that arrive at time 1 stay active during such interval, given that by definition $\forall c \in C : w_c \leq |\tau_c|$. We will show how to assign each client to one station and one slot within this window, so that no station is overloaded in any slot. The assignment in all the subsequent intervals $[iw + 1, (i + 1)w]$, for each integer $i > 0$, is identical. Let \mathcal{A} be any initial assignment of each client to one of the m stations and one of the w slots. Let B_{\max} be the maximum bandwidth used in \mathcal{A} in any slot, and let (s, i) be some station-slot pair with such bandwidth usage in the assignment \mathcal{A} . If $B_{\max} \leq B$, we are done. Otherwise, given that $B_{\max} > B$, station s has more than one client assigned in slot i , since otherwise the client arrival schedule would violate Equation (2). If the sum of the bandwidth used on some pairs (s', j) and (s'', k) is at most B , consider another assignment \mathcal{A}' where the clients assigned to (s', j) and (s'', k) in \mathcal{A} are now all assigned to (s', j) , and the clients assigned to (s, i) in \mathcal{A} are now split between (s, i) and (s'', k) . Repeat the procedure above until the sum of bandwidth usage in each two station-slot pairs is at least B , or $B_{\max} \leq B$. In the latter case we are done. Otherwise, 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. \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 6. *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 must 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 $n'b\frac{|T|}{w} = n'\frac{mwB\rho + \beta - B}{n}\frac{|T|}{w} = \frac{n'}{aw}B\rho|T| + n'\frac{\beta - B}{n}\frac{|T|}{w} \leq \left\lceil \frac{n'}{aw} \right\rceil |T|B\rho + \beta$. 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 $B + ab = B + \frac{n}{m}b = B + \frac{n}{mw}\frac{mwB\rho + \beta - B}{n} = B(1 + \rho) + \frac{\beta - B}{mw}$, 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 stations group and bandwidth, may be different among clients.

Theorem 7. *Given any (ρ, β) -admissible client arrival schedule, if $\beta \leq mwB(1/(mw) - \rho)$, the Station Assignment problem can be solved on any system of m stations each with bandwidth B .*

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 $\sum_{c \in C'} b_c \leq w|S(C')|\rho B + \beta \leq w m \rho B + \beta$. 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

In this section, we present bounds for irrevocable improvident online protocols.

6.1 Lower Bounds for Non-Solvability

We show now that irrevocable algorithms do not always solve the problem. Theorem 8 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 identical. The proof is left to the full version of this paper for brevity.

Theorem 8. *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 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.*

Corollary 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 such that \mathcal{A} cannot solve the Station Assignment problem.*

Proof. Replacing all bandwidths in the lower bound of ρ in Theorem 8 by B , we get $\rho > \left(\sum_{j=1}^m \frac{1}{j} \right)^{-1} = H_m^{-1} > \frac{1}{1 + \ln m}$. \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. Given that $\beta \geq 0$ by definition, the bound implies that if $\rho > 1/\ln m$, the Station Assignment problem is not solvable.

Theorem 9. *For any system of m 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 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 the clients that 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 stations group S_i . Starting from the whole set of stations S in the first round, the stations group for each new round is reduced by one station. We say that such station is **removed**. Thus, for round 1 the stations group has size m , for round 2 the size is $m - 1$, and so on until round m when the stations group 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 with stations group S_j , for $j = 1, \dots, i$. 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 $\left(\dots \left((X \frac{m-1}{m} + X) \frac{m-2}{m-1} + X \right) \frac{m-3}{m-2} \dots \right) \frac{1}{2} + X = X \left(\frac{1}{m} + \frac{1}{m-1} + \dots + \frac{1}{2} + 1 \right) > X \ln m$. 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 10 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 10. *For any set of m stations each with bandwidth B , the following holds, even if all clients must 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 such that no online irrevocable improvident deterministic algorithm can solve Station Assignment.*

Proof. If $\beta > mB(1 - \rho)$, the claim follows from Theorem 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 [7] 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 theorem applies to a setting where the station bandwidth is identical, but the station group may be different for each client.

Theorem 11. *For any system of m stations each with bandwidth B , there exists an online algorithm, such that if $\rho \leq 1/(1 + \sqrt{2m})$, $\beta < \rho B$, all clients have the same bandwidth $b \geq \rho B$ and laxity $w = 1$, and never depart, the Station Assignment problem is solved.*

Proof. Let S be the set of stations in the system and, for any subset of stations $S' \subseteq S$, let C' be the set of clients $C' = \{c | S_c = S'\}$. Using that $b \geq \rho B$ and $\beta < \rho B$, the following properties arise from admissibility.

Property 1. $\forall S' \subseteq S : |C'| \leq |S'|$.

That is, for each station group of x stations, there are at most x clients with that station group.

Property 2. $\forall S'' = \{S' | S' \subseteq S\} : |\cup_{S' \in S''} C'| \leq |\cup_{S' \in S''} S'|$.

That is, for any set of station groups, the maximum number of clients with those station groups is at most the size of the union of those groups.

Consider the online algorithm that, for each client c , assigns c to the station $s \in S_c$ with the largest available bandwidth, breaking ties arbitrarily. We show that, under the assumptions of the theorem, this algorithm solves the problem. For the sake of contradiction, assume that some station s_i is overloaded. That is, s_i has some integer number k of clients assigned such that $k > 1/\rho$. We show that then the number of clients in the system must be more than m , which is not possible according to property 1.

We compute a lower bound on the number of clients that should be in the system in order to have more than $1/\rho$ clients in s_i . For clarity, we label the clients assigned to s_i in the order in which they were assigned. Client 1 is the first one and, hence, does not require any other clients to be in the system before. For each client $c = 2, 3, \dots, k$, we identify clients that must have been assigned before c to some station. We *allocate* some of those clients to each c . In order to avoid over-counting, sometimes we may *reallocate* some clients, so that each client in the system is allocated to at most one client.

Assume that, for each client $c \in [2, k]$, we can allocate $c - 1$ “new” clients. Then, overall, we will have $1 + 2 + \dots + k - 1$ allocated clients which yields a lower bound of $k(k - 1)/2 > m$ clients in the system proving the claim. The details of the allocation procedure follow.

For each client $c = k, k - 1, \dots, 2$ in s_i , we know that there must be at least $c - 1$ clients in each station in S_c , because the algorithm distributes clients evenly in S_c . If the clients in one or more of the stations in S_c have not been allocated yet, we choose one of those stations arbitrarily and allocate the $c - 1$ clients assigned

to that station to c . If the clients in all stations in S_c have been already allocated to some client, assume that there is at least one client $c' \in [c + 1, k]$ such that the clients assigned to some station $s_j \in S_{c'}$ have not been allocated. Then, we reallocate $c - 1$ clients from c' to c , and we allocate the $c' - 1$ clients in s_j to c' .

We show now that if the latter assumption is false, property 2 has been violated. For the sake of contradiction, assume that, at the point of allocating clients for some client c , the clients in all stations in $\cup_{c'=c+1}^k S_{c'}$ have been already allocated to some client in $[c + 1, k]$. This implies that $k - c = |\cup_{c'=c+1}^k S_{c'}|$. Thus, if $S_c \subseteq \cup_{c'=c+1}^k S_{c'}$, property 2 is violated.

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 obtained variety of results is a promising starting point for further study of more complex scheduling settings in the proposed mobility model, including the settings motivated by sensor and local wireless network applications.

References

1. Noga Alon, Yossi Azar, Gerhard J. Woeginger, and Tal Yadid. Approximation schemes for scheduling. In *SODA*, pages 493–500, 1997.
2. L. Anantharamu, Bogdan S. Chlebus, and Mariusz A. Rokicki. Adversarial multiple access channel with individual injection rates. In *Proceedings of the 13th International Conference on Principles of Distributed Systems (OPODIS)*, LNCS 5923, pages 174–188. Springer-Verlag, 2009.
3. Matthew Andrews, Baruch Awerbuch, Antonio Fernández, Frank Thomson Leighton, Zhiyong Liu, and Jon M. Kleinberg. Universal-stability results and performance bounds for greedy contention-resolution protocols. *Journal of the ACM*, 48(1):39–69, 2001.
4. Matthew Andrews and Lisa Zhang. Scheduling over a time-varying user-dependent channel with applications to high-speed wireless data. *Journal of the ACM*, 52(5):809–834, 2005.
5. James Aspnes, Yossi Azar, Amos Fiat, Serge A. Plotkin, and Orli Waarts. On-line load balancing with applications to machine scheduling and virtual circuit routing. In *STOC*, pages 623–631, 1993.
6. Yossi Azar. On-line load balancing. In *Online Algorithms*, pages 178–195, 1996.
7. Yossi Azar, Andrei Z. Broder, and Anna R. Karlin. On-line load balancing. *Theor. Comput. Sci.*, 130(1):73–84, 1994.
8. Yossi Azar, Bala Kalyanasundaram, Serge A. Plotkin, Kirk Pruhs, and Orli Waarts. On-line load balancing of temporary tasks. *J. Algorithms*, 22(1):93–110, 1997.
9. Marcin Bienkowski, Tomasz Jurdzinski, Mirosław Korzeniowski, and Dariusz R. Kowalski. Distributed online and stochastic queuing on a multiple access channel. In *DISC*, pages 121–135, 2012.
10. Maria J. Blesa, Daniel Calzada, Antonio Fernández, Luis López, Andrés L. Martínez, Agustín Santos, Maria J. Serna, and Christopher Thraves. Adversarial queueing model for continuous network dynamics. *Theory Comput. Syst.*, 44(3):304–331, 2009.
11. Allan Borodin, Jon M. Kleinberg, Prabhakar Raghavan, Madhu Sudan, and David P. Williamson. Adversarial queueing theory. *Journal of the ACM*, 48(1):13–38, 2001.
12. Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Adversarial queueing on the multiple-access channel. In *Proceedings of the 25th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 92–101, 2006.
13. R. L. Graham. Bounds on multiprocessing timing anomalies. *The Bell System Technical Journal*, 45:1563–1581, 1966.
14. David R. Karger, Steven J. Phillips, and Eric Torng. A better algorithm for an ancient scheduling problem. *Journal of Algorithms*, 20:400–430, 1996.
15. Thomas Kesselheim. Dynamic packet scheduling in wireless networks. In *PODC*, pages 281–290, 2012.
16. K. Pruhs, J. Sgall, and E. Torng. Online scheduling. In J. Leung, editor, *Handbook of Scheduling: Algorithms, Models and Performance Analysis (Chapter 15)*, pages 15–1–15–41. CRC Press, 2004.