

Max-stretch minimization on an edge-cloud platform

Anne Benoit, Redouane Elghazi, Yves Robert

▶ To cite this version:

Anne Benoit, Redouane Elghazi, Yves Robert. Max-stretch minimization on an edge-cloud platform. IPDPS 2021 - IEEE International Parallel and Distributed Processing Symposium, May 2021, Portland, Oregon, United States. pp.1-10, 10.1109/IPDPS49936.2021.00086 . hal-03509637

HAL Id: hal-03509637 https://inria.hal.science/hal-03509637

Submitted on 4 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Max-Stretch Minimization on an Edge-Cloud Platform

Anne Benoit LIP, ENS Lyon, France Anne.Benoit@ens-lyon.fr Redouane Elghazi LIP, ENS Lyon, France & FEMTO Univ. Franche Comté, Besançon, France Redouane.Elghazi@ens-lyon.fr Yves Robert LIP, ENS Lyon, France & UT Knoxville, TN, USA Yves.Robert@ens-lyon.fr

Abstract—We consider the problem of scheduling independent jobs that are generated by processing units at the edge of the network. These jobs can either be executed locally, or sent to a centralized cloud platform that can execute them at greater speed. Such edge-generated jobs may come from various applications, such as e-health, disaster recovery, autonomous vehicles or flying drones. The problem is to decide where and when to schedule each job, with the objective to minimize the maximum stretch incurred by any job. The stretch of a job is the ratio of the time spent by that job in the system, divided by the minimum time it could have taken if the job was alone in the system. We formalize the problem and explain the differences with other models that can be found in the literature. We prove that minimizing the max-stretch is NP-complete, even in the simpler instance with no release dates (all jobs are known in advance). This result comes from the proof that minimizing the max-stretch with homogeneous processors and without release dates is NPcomplete, a complexity problem that was left open before this work. We design several algorithms to propose efficient solutions to the general problem, and we conduct simulations based on real platform parameters to evaluate the performance of these algorithms.

I. INTRODUCTION

Edge-Cloud computing is a recent paradigm that is currently deployed by many vendors (see [20], [21], [22], [29] and many others). The idea is to execute some jobs in-situ, directly on the edge server where they originate from, thereby providing a flexible, cost-effective and decentralized solution that dramatically reduces the volume of data transfers. But some jobs may be too demanding in terms of computing power, or some edge servers may be overloaded because they launch too many jobs. This calls for coupling the edge server with a powerful cloud platform, where some jobs can be delegated whenever needed.

Deciding which jobs should be executed locally and which ones should be communicated (up and down) to the cloud platform is a challenging problem. The main focus of this paper is to lay the foundations for solving the instance of the problem dealing with response time as a single optimization criterion. Response time is a very important criterion, in particular for the users, and arguably the most important one in a real-time framework. But it is not the unique criterion to optimize on an edge-cloud platform: energy consumption and resource costs are important criteria too. However, we show that dealing with response-time is already an intricate problem, and we leave multi-objective optimization for further work.

The *response time* for a job [13], [30] (also called *flow time* in the scheduling literature [9]) is the time spent by

that job in the system, starting from its release date and up to final completion (including output transfers if any). The standard scheduling objective is to minimize the maximum response time over all jobs¹. However, maximum response time is not appropriate to ensure fairness, because giving the same response time for all jobs results in worst performance for short jobs, compared to the performance achieved for long ones. In order to provide a fair processing of jobs, job lengths should be taken into account. The *stretch* [12], [3] is the metric used to ensure fairness among jobs, and it is defined as the response time normalized by the job length. More precisely, the stretch is the response time divided by the best possible execution time for the job if it were alone on the platform. Therefore, a job stretch measures how the performance of a job is degraded compared to a system dedicated exclusively to this job. The *maximum stretch* is the maximum of the stretches of all jobs² and provides a measure of the responsiveness of the system: it quantifies the user expectation that the response time should be proportional to the load incurred by the execution of the job by the system. For an illustration, consider two jobs released at the same time, one lasting 1 hour and the other 10 hours. With a single processor, executing the long job first leads to a maximum stretch of 11, while executing the short job first leads to a maximum stretch of 1.1. Intuitively, the latter scheduling is more fair to users than the former. A word of caution: assume now that we have one edge processor and one cloud processor for the two jobs. For each of them, the time that it would spend in a dedicated system must be computed as the minimum of its execution times on the edge and on the cloud (including transfers), if it was the only job.

In this paper, we focus on the maximum stretch as the optimization metric, and we investigate how to minimize the maximum stretch of independent jobs executing on an edgecloud platform. This is an online scheduling problem, as the jobs have release dates that are unknown until the jobs are submitted. The first task is to design a realistic model for such a platform. In a nutshell, we propose a model with preemption and possible re-execution, but no migration: jobs are either executed on the edge processor where they originate from, or on a cloud processor, and can be dynamically re-assigned to

¹Minimizing the average response time of all jobs, or equivalently the sum of the response times of all jobs (also called *total flow time*), has been extensively studied too [9].

 $^{^{2}}$ Just as for response time, the average stretch has also been studied as a metric [5].

another resource. If this is case, the execution must then restart from scratch, thus the time spent executing the job until reassignment is lost. Note that migration would require some kind of checkpoint mechanism to be able to restart the job on another resource from its current state.

If a job is executed on the cloud, up and down transfers are accounted for. We allow for computation and communication overlap, and many communications can occur in parallel across edge and cloud processors. However, we enforce the one-port full-duplex communication model that states that any resource cannot be involved either in two sending operations, or in two receiving operations, at the same-time step, but sending and receiving simultaneously is allowed. Independent sender/receiver pairs can communicate in parallel. A given message may be preempted (and resumed later) if the sender or the receiver has a more urgent message to process (because of the release of a new job). This communication model [17], [19], [34] has been advocated as being much more accurate than the traditional macro-dataflow communication model of the scheduling literature, which allows a processor to, say, send 100 different messages in parallel without accounting for any bandwidth limitation.

After designing the model, we assess complexity results. The offline problem is shown to be NP-hard, even in the simpler instance with no release dates (all jobs are known in advance). This result comes from the proof that minimizing the max-stretch with homogeneous processors and without release dates is NP-complete, a problem whose complexity was left open before this work. Given these negative (but somewhat expected) results, we introduce several heuristics for the general online problem, and compare their performance through simulations. In summary, the main contributions of this paper are the following:

- A realistic model for scheduling jobs on edge-cloud platforms;
- Several complexity results for the offline problem, proving in particular the NP-hardness of an open problem;
- New heuristics and their experimental comparison for the online problem.

The rest of the paper is organized as follows. We first discuss related work in Section II. The detailed model of the edge-cloud platform is provided in Section III, along with the MINMAXSTRETCH-EDGECLOUD problem definition. We then prove the NP-completeness of MINMAXSTRETCH-EDGECLOUD in Section IV, and we design several heuristic algorithms to solve this problem in Section V. Finally, extensive simulations are presented in Section VI to evaluate and compare the algorithms. Conclusions and directions for future work are presented in Section VII.

II. RELATED WORK

Edge-cloud platforms: A range of applications of edge computing is given in [8], and these include e-health, disaster recovery, autonomous vehicles or flying drones. Although no specific model is given, a state of the art is established with a list of objectives and constraints that have been studied, such

as the delay, the bandwidth, the energy, QoS-assurance, etc. Some challenges of edge computing are also discussed in [23].

Some papers influenced our choices related to the model. For example, we chose a model that allows preemption but not migration. This was influenced by papers such as [1], where it is shown that allowing migration does not affect the efficiency of the algorithms that minimize the average flow time, whereas allowing preemption does. This is also proven for the minimization of the average stretch in [2].

We also reviewed papers specifically to design our edge/cloud model. For example, in [26], even though the concern is not exactly the same as ours, a model is proposed for mobile-edge computing systems. Our model is simpler than theirs because we do not consider the power consumption, whereas they do. Existing models include those from [32] and [31], where the focus is on the placement of data stream processing applications. However, this latter approach is application-specific, which our model aspires to be agnostic of the nature of the application.

Finally, we decided to consider heterogeneous edge processors but kept homogeneity inside the cloud platform. This decision is motivated by the fact that users can always request identical virtual machines on cloud platforms. On the theoretical side, dealing with heterogeneous processors might be an obstacle to finding competitive algorithms: in some cases [16], having heterogeneous processors invalidates the theoretical bounds of some usual scheduling algorithms. However, it is not difficult to extend our model with heterogeneous cloud processors, and all the algorithms can be modified in a straightforward manner to handle a fully heterogeneous edgecloud platform.

Communication model: Contention-aware task scheduling has been considered since many years in the scheduling literature [17]. In particular, [34] showed through simulations that taking contention into account is essential for the generation of accurate schedules. They investigate both end-point and network contention. Here end-point contention refers to the bounded multi-port model [18]: the volume of messages that can be sent/received by a given processor is bounded by the limited capacity of its network card. Network contention refers to the one-port model, which has been advocated by [7] because "current hardware and software do not easily enable multiple messages to be transmitted simultaneously." Even if non-blocking multi-threaded communication libraries allow for initiating multiple send and receive operations, all these operations "are eventually serialized by the single hardware port to the network." Experimental evidence of this fact has been reported by [33], which reports that asynchronous sends become serialized as soon as message sizes exceed a few megabytes. The one-port model fully accounts for end-point contention. Coupled with message preemption, it provides a very realistic model for communications.

Stretch: The stretch is a particular case of weighted response time [12] and has been studied extensively because of its fairness. Here is a quote from [5]: "Flow time measures the time that a job is in the system regardless of the service it

requests; the stretch measure relies on the intuition that a job that requires a long service time must be prepared to wait longer than jobs that require small service times." An interesting analogy is made in [14] between maximum stretch and distributive justice, defined as the perceived fairness in the way costs and rewards are shared within a group of individuals.

Stretch optimization was studied for independent jobs without preemption [4], bag-of-tasks applications [25], [11], multiple parallel task graphs [10], for sharing broadcast bandwidth between client requests [35], and also *k*-norm of stretch with one processor and preemption [27].

On the algorithmic side, [28] proposes an algorithm called Shortest Remaining Processing Time (SRPT), that is an approximation for the online case with the average stretch as an objective. It is a greedy approach. Then [3] proposes an algorithm for minimizing the maximum stretch in the online case of only one processor. This algorithm is optimal for the offline problem, and is Δ -competitive for the online problem, where Δ is the ratio between the longest and the shortest job. It means that the result of this algorithm is at most Δ times the optimal result. Another version of this algorithm in given in [4], with a better time complexity but similar bounds on the resulting stretch.

III. MODEL

We introduce the framework in Section III-A, we survey schedule constraints in Section III-B, and we work out an example in Section III-C.

A. Framework

We consider a two-level platform, with P^c homogeneous processors in a cloud, and P^e edge computing units. The cloud processors have a speed normalized to 1, while edge computing units run at a slower pace. Hence, the *j*-th edge computing unit, with $1 \leq j \leq P^e$, is operating at a speed $s_j \leq 1$.

The application consists in n independent jobs J_1, \ldots, J_n , where each job is issued from an edge computing unit. Job J_i is generated by the edge computing unit o_i (with $1 \le o_i \le P^e$), and this edge computing unit must obtain the result of job J_i , either by executing J_i locally, or by delegating the job to the cloud, and getting the result back. Parameters for job J_i are as follows:

- o_i is the origin processor on the edge ;
- w_i is the amount of work required to complete the job;
- r_i is the release date;
- up_i and dn_i are the communication times required to send the job to the cloud and get the result back (up-link/downlink communications).

The execution time of job J_i hence depends whether it is executed directly at the edge processing unit where the job originates from, or whether it is sent for an execution on the cloud. If the job is executed on the edge, it takes a time $t_i^e = \frac{w_i}{s_{o_i}}$, since the edge processing unit o_i operates at speed s_{o_i} . Computation may be faster on the cloud, but we must then account for communication time, hence an execution on the cloud takes a time $t_i^c = up_i + w_i + dn_i$ (send the job up to the cloud, compute at speed 1, and get the result back down). In both cases, the execution of the job can be preempted, which means that we can interrupt its execution (to schedule another job), and resume it at a later time. Once started on a given resource, the execution cannot migrate to another resource, but re-execution from scratch is possible (and the time spent up to re-execution is lost).

Let C_i denote the time at which the execution of J_i is completed. Then, the *stretch* S_i of job J_i is defined by:

$$S_i = \frac{C_i - r_i}{\min(t_i^e, t_i^c)}.$$

Indeed, job J_i is released at time r_i and spends a time $C_i - r_i$ in the system, while it could have taken a time $\min(t_i^e, t_i^e)$ in the best case of a dedicated system. Hence, the stretch of each job is equal to one if the job is executed with the minimum possible time, and we want the maximum stretch (over all jobs) to be as close to one as possible. Overall, the objective function is to minimize the maximum stretch, which is defined as $\max_{1 \le i \le n} S_i$.

While several jobs may be competing for cloud resources, several constraints must be enforced within a schedule of jobs. We overlap computations and communications, and we consider communication channels to be full-duplex between an edge processor and a cloud processor, hence it is possible to perform in parallel a computation, an uplink communication, and a downlink communication. However, communications involving a common processor must be sequentialized: if two jobs originating from the same edge processing unit are delegated to the cloud, we cannot perform the two uplink (and then downlink) communications in parallel. Similarly, if two jobs (that may come from different edge processing units) are sent to the same cloud processor, their communications (uplink and then downlink) must be sequentialized.

Furthermore, communications, just as computations, are preemptive, which means that we can interrupt a communication (for instance, to schedule a communication for a smaller job), and resume the interrupted communication at a later time.

Optimization problem: The goal is to find a schedule that respects all constraints, with the aim of minimizing the maximum stretch. This problem is denoted as MINMAXSTRETCH-EDGECLOUD.

B. Schedules

For each job J_i , we first need to decide where it is executed: alloc(i) = 0 if the job is executed on its local edge processor o_i , otherwise we set alloc(i) = k, where k is the cloud processor on which J_i is executed $(1 \le k \le P^c)$. Formally, a schedule consists in sets of disjoint execution intervals E_i for each job (recall that we allow preemption), and disjoint uplink/downlink communication intervals for jobs executed on the cloud $(alloc(i) \ne 0), U_i(o_i, alloc(i))$ and $D_i(alloc(i), o_i)$.

Several constraints must be ensured for the schedule to be valid. In particular, for any two jobs $J_i, J_{i'}$ (with

 $1 \le i, i' \le n$), executed on the same processor, i.e., alloc(i) = $alloc(i') \neq 0$ (cloud), or alloc(i) = alloc(i') = 0 and $o_i = o_{i'}$ (edge), all their execution intervals must be disjoint:

$$\forall I \in E_i, \forall I' \in E_{i'}, \ I \cap I' = \emptyset.$$

We must also ensure that communications are serialized, i.e., for any two sets $U_i(j,k)$ and $U_{i'}(j',k')$ (resp. $D_i(k,j)$ and $D_{i'}(k',j')$), if j = j' or k = k', then the communication originates from or targets the same processor, and hence all intervals must be pairwise disjoint. Finally, we must ensure that all computations and communications are done, in the correct order. Given a set of intervals E, we denote by $\min(E)$ (resp. max(E)) the smallest (resp. largest) extremity of all intervals in E. Hence, for job J_i $(1 \le i \le n)$:

- If alloc(i) = 0, then $\sum_{I \in E_i} |I| \ge \frac{w_i}{s_{o_i}}$;
- Otherwise,

 - $\sum_{I \in E_i} |I| \ge w_i;$ $\sum_{I \in U_i(o_i, alloc(i))} |I| \ge up_i;$ $\sum_{I \in D_i(alloc(i), o_i)} |I| \ge dn_i;$
 - we must complete uplink communications before starting computation: $\max(U_i(o_i, alloc(i))) \leq \min(E_i);$
 - computation we must complete – and before starting downlink communications: $max(E_i)$ \leq $\min(D_i(alloc(i), o_i)).$

C. Example

An example of a valid execution with one edge processor and one cloud processor is depicted in Figure 1, where J_1, J_4 and J_6 are executed on the edge, while J_2, J_3 and J_5 are sent to the cloud. Job parameters are as follows, where the speed of the edge processor is $\frac{1}{3}$:

- J_1 : $r_1 = 0$, $w_1 = 1$, $up_1 = dn_1 = 5$;
- J_2 : $r_2 = 0$, $w_2 = 4$, $up_2 = dn_2 = 2$;
- J_3 : $r_3 = 3$, $w_3 = 2$, $up_3 = 2$, $dn_3 = 1$;
- J_4 : $r_4 = 5$, $w_4 = 4/3$, $up_4 = dn_4 = 5$;
- J_5 : $r_5 = 5$, $w_5 = 2$, $up_5 = 2$, $dn_5 = 1$;
- $J_6: r_6 = 6, w_6 = 1/3, up_6 = dn_6 = 5.$

The execution intervals, as well as uplink/downlink communication intervals, are depicted. At time-step 6, we compute at the same time on the cloud, on the edge, and we perform an uplink communication and a downlink communication. Also, this is the time when J_6 preempts job J_4 , since it can be executed locally within one time unit, while J_4 is a longer job and will be delayed only by one time unit.

One can check, by an exhaustive search at all possible schedules, that this one is optimal. Indeed, jobs J_1 and J_6 run at their minimum possible time $w_i \times 3$ on the edge, while



executing them on the cloud would have cost at least 10 units of communication time, hence they have a stretch of 1. On the other hand, J_2 would take a time 12 on the edge, and it is sent to the cloud where the total time is $up_2 + w_2 + dn_2 = 8$, hence also a stretch of 1. Next, consider J_3 and J_5 , which have the same characteristics: they can be executed in 6 units of time on the edge, but 5 on the cloud. Here, we execute them on the cloud, where they both experience one time-step of delay (waiting for the processor to be available), hence they have a stretch of $\frac{6}{5}$, and they could not have a better stretch on the edge. Finally, job J_4 takes 5 units of time on the edge, while its minimum time is 4, because it is preempted at time-step 6 to execute the shorter job J_6 that would greatly impact the stretch if it was delayed. On the cloud, it would have taken a time at least $10 + \frac{4}{3}$, hence much greater. Its stretch is $\frac{5}{4}$, which determines the maximum stretch for this example.

Throughout this example, we see that decisions are more difficult to take when there is no knowledge about jobs that will be released in the future. For instance, one could schedule job J_3 either on the edge or on the cloud, since it would complete in both cases within time 9, but depending on the jobs that come next (computation-intensive vs communicationintensive), one decision would be better than the other. The online case is the problem where jobs are not known in advance, while in the offline case, all job parameters are known in advance. In the following, we prove that the problem is difficult even in the offline case, and then we derive heuristics to address the general online problem in Section V.

IV. PROBLEM COMPLEXITY

This section is devoted to assess the study of the complexity of MINMAXSTRETCH-EDGECLOUD in the offline case. To shorten notations, we write MMSECO instead of MINMAXSTRETCH-EDGECLOUD-OFFLINE throughout this section, and MMSECO-DEC is the corresponding decision problem (is it possible to achieve a target maximum stretch?). First, we prove that MMSECO-DEC is in NP in Section IV-A. Then, we derive two main complexity results, that remain true even without release dates (consider that all jobs are released at time 0). The results are the following:

- For a fixed number of processors, MMSECO-DEC is NP-complete in the weak sense (Section IV-B).
- For a variable number of processors, MMSECO-DEC is NP-complete in the strong sense (Section IV-C).

A. NP membership

Lemma 1. MMSECO-DEC is in NP.

Proof. A solution to MMSECO-DEC is a sequence of events, where an event is defined with the following list:

- a job J_i being released;
- a job J_i finishing its uplink communication;
- a job J_i finishing its execution;
- a job J_i finishing its downlink communication.

In between two events, there is no reason to reconsider any decision taken, nor to preempt any computation or computation. There are only two events for the jobs executed on edge processors, and four for those executed on a cloud processor, so there are at most 4n events. Therefore, a schedule can be represented in polynomial space: at each of these 4n timesteps, there might be at most one event per processor. The validity of the solution can then be verified in polynomial time by checking each constraint for a valid schedule, as detailed in Section III-B. Therefore, MMSECO-DEC is in NP.

B. Weak NP-completeness with two processors

We are now ready to prove the NP-completeness of MMSECO-DEC. We first focus on minimizing the *maximum* stretch for a fixed number of homogeneous processors, and without release dates (hence forgetting about the cloud-edge system, called MMSH problem), and the associated decision problem MMSH-DEC. The execution of job J_i hence takes w_i time units, and there are no communications. To the best of our knowledge, the complexity of MMSH has not been established yet, and we prove its NP-completeness. In a second step, we show that it is easy to create an instance of MMSECO-DEC from a general instance of MMSH-DEC, hence proving the NP-completeness of MMSECO-DEC.

As a preliminary, we derive an interesting result about the optimal order of jobs in a schedule for MMSH: with a single processor, jobs should be ordered from the shortest to the longest (i.e., by non-decreasing w_i 's).

Lemma 2. For MMSH with a single processor, there always exists a schedule that minimizes the maximum stretch by ordering the jobs from the shortest to the longest, and without preemption.

Proof. First, we observe that preemption is not useful when all jobs have the same release date (which is 0 here). To see this, assume we execute a sequence with preemption $(J_i^{(1)}, J_j^{(1)}, J_i^{(2)})$ of three job chunks, where J_i and J_j are two different jobs. Here, $J_i^{(1)}$ and $J_i^{(2)}$ are two chunks of job J_i , and $J_j^{(1)}$ is one chunk of job J_j . Swapping the first two chunks leads to the sequence $(J_j^{(1)}, J_i^{(1)}, J_i^{(2)})$. Because J_i and J_j have the same release date, this swap is valid: it does not change the stretch of J_i and can only decrease or leave unchanged the stretch of J_j . After a finite number of such swaps, we have a non-preemptive schedule whose stretch does not exceed that of the original schedule with preemption.

Then, to prove that the jobs can be ordered as stated, we consider a schedule that does not order the jobs from the shortest to the longest, and we construct another schedule with fewer mis-orderings, and which has a maximum stretch that is smaller than or equal to the original maximum stretch. Let (k_1, \ldots, k_n) be (the indices of) the ordering of the jobs in the initial schedule, meaning that we execute J_{k_1} first and J_{k_n} last. By assumption, there exists *i* such that $w_{k_i} > w_{k_{i+1}}$ (i.e., job J_{k_i} is longer than job $J_{k_{i+1}}$). We swap these two jobs, which is possible because all jobs have same release dates, and which does not change the stretch of the other jobs.

Since we consider the maximum stretch of the schedule, let us compare the maximum between the stretches of the two jobs before and after the swap. If X is the starting time of the earlier job, then the values that we consider are:

$$S_b = \max\left(\frac{X + w_{k_i}}{w_{k_i}}, \frac{X + w_{k_i} + w_{k_{i+1}}}{w_{k_{i+1}}}\right)$$
 before the swap;

$$S_a = \max\left(\frac{X + w_{k_{i+1}}}{w_{k_{i+1}}}, \frac{X + w_{k_{i+1}} + w_{k_i}}{w_{k_i}}\right) \text{ after the swap.}$$

As we have $\frac{X+w_{k_{i+1}}}{w_{k_{i+1}}} < S_b$ and $\frac{X+w_{k_{i+1}}+w_{k_i}}{w_{k_i}} < S_b$, we indeed get $S_a < S_b$, so the stretch after the swap is at most the stretch before the swap. We can repeat the operation until there are no mis-orderings anymore, and we get a schedule that orders the jobs from the shortest to the longest and has a maximum stretch less than or equal to that of the initial schedule.

Theorem 1. MMSH-DEC with two homogeneous processors is NP-complete in the weak sense.

Proof. It is easy to see that MMSH-DEC is in NP, since it is a simpler case of MMSECO-DEC (see Section IV-A). We now prove that it is NP-complete.

Let I_1 be an instance of 2-PARTITION-EQ [15], which is a special case of 2-PARTITION where the partitions must have equal cardinality: Given 2n integers $\{a_1, \ldots, a_{2n}\}$ with $2S = \sum_{i=1}^{2n} a_i$, does there exist a subset $I \in \{1, \ldots, 2n\}$ with |I| = n and $\sum_{i \in I} a_i = S$?

We create an instance I_2 of MMSH-DEC with two processors and 2n + 2 jobs, such that:

- for $1 \leq i \leq 2n$, the execution time of job J_i is $w_i = nS + a_i$;
- the execution time of the two additional jobs J_{2n+1} and J_{2n+2} is $w_{2n+1} = w_{2n+2} = (n+1)S$.

We ask whether it is possible to achieve a stretch of $\frac{n^2+n+2}{n+1}$.

Note that, building upon Lemma 2, we know that each processor will sort its jobs by non-decreasing execution time in an optimal solution, so a schedule is characterized only by a partition P_1 , P_2 of the jobs. We now show that I_1 has a solution if and only if I_2 has a solution.

Consider first that I₁ has a solution, I. We consider the schedule with P₁ containing jobs Ji with i ∈ I and job J_{2n+1}, while P₂ contains all other jobs (Ji, i ∉ I and J_{2n+2}). We prove that this schedule has a max-stretch equal to n+1/n+1.

First, note that the sum of job weights in each set is:

$$\sum_{i \in I} (nS + a_i) + (n+1)S = n^2 S + S + (n+1)S = (n^2 + n + 2)S$$

since there are exactly n jobs in set |I| (and the same number of jobs for P_2).

By Lemma 2, job J_{2n+1} (resp. J_{2n+2}) is executed last in P_1 (resp. P_2), since $w_{2n+1} = w_{2n+2} > w_i$ for all $1 \le i \le 2n$. Therefore, these additional jobs both complete at time $(n^2 + n + 2)S$, and they have a stretch $\frac{n^2 + n + 2}{n+1}$.

For any other job J_i , the stretch is $S_i = \frac{knS+x_i}{nS+a_i}$, where there are $k \leq n$ jobs before J_i (including J_i) with a

total weight knS + X, where $X = \sum_{j \in Pred(i)} a_j$. Here, Pred(i) denotes the indices of the k jobs that precede J_i , and we have $X \leq S$. We have the following inequalities: (i) $\frac{knS+x}{nS+a_i} \leq \frac{(n^2+1)S}{nS+a_i} \leq \frac{n^2+1}{n}$ (constraints on k and X), and (ii) $\frac{n^2+1}{n} \leq \frac{n^2+n+2}{n+1}$ (for any $n \geq 1$).

Therefore, job J_i has a stretch $S_i \leq \frac{n^2+n+2}{n+1}$, meaning that the max-stretch of the schedule is $\frac{n^2+n+2}{n+1}$, and hence I_2 has a solution.

• Consider now that I_2 has a solution: Let P_1 , P_2 be a schedule with max-stretch $\frac{n^2+n+2}{n+1}$ or less. We prove that P_1 , P_2 is a partition of the jobs into two sets of equal sum and equal size, with $J_{2n+1} \in P_1$ and $J_{2n+2} \in P_2$. We first prove that each processor finishes its whole execution at a time t_{finish} such that $t_{finish} \leq (n^2+n+2)S$. Let J_i be the last job to finish, its stretch is $S_i = \frac{t_{finish}}{w_i}$. Since $w_i \leq (n+1)S$ for all jobs, $S_i \geq \frac{t_{finish}}{(n+1)S}$. Hence, since P_1, P_2 is a solution to I_2 , it means that $S_i \leq \frac{n^2+n+2}{n+1}$ and therefore $t_{finish} \leq (n^2+n+2)S$. Furthermore, since the sum of all execution times is $2(n^2+n+2)S$, this means that both processors finish at exactly time $t_{finish} = (n^2+n+2)S$.

Now, we specifically look at the two jobs that finish their execution at time $(n^2 + n + 2)S$, one on each processor. We prove that these two jobs are J_{2n+1} and J_{2n+2} . For their stretch to be at most $\frac{n^2+n+2}{n+1}$, they need to have an execution time $w_i \ge (n+1)S$. The only jobs for which this constraint holds are J_{2n+1} and J_{2n+2} . So, we indeed have J_{2n+1} and J_{2n+2} that are scheduled on different processors, say J_{2n+1} in P_1 and J_{2n+2} in P_2 (the two jobs are identical).

Let us now consider the jobs in P_1 , excluding the large job J_{2n+1} . Their total weight is hence $t_{finish} - (n+1)S = n \times nS + S$, which means that there are exactly n jobs (the a_i 's are small in comparison with nS). Since their sum is $n^2S + S$, we obtain that $\sum_{i \in P_1 \setminus \{2n+1\}} a_i = S$, and the set of jobs from P_1 is a solution to 2-PARTITION-EQ, and I_1 has a solution.

Overall, I_1 has a solution if and only if I_2 has a solution, and the construction of I_2 from I_1 can obviously be done in polynomial time, hence MMSH-DEC is NP-complete. This NP-completeness result is established in the weak sense since 2-PARTITION-EQ can be solved by a pseudo-polynomial algorithm.

We can now derive the result for the original cloud-edge problem, by performing a simple reduction from MMSH-DEC.

Corollary 1. MMSECO-DEC is NP-complete in the weak sense, even with one edge processor, one cloud processor, and no release dates.

Proof. By Lemma 1, MMSECO-DEC is in NP. We perform a reduction from MMSH-DEC with two processors, where $P^e = P^c = 1$, the speed of the edge processor is set to 1, and all communication costs of jobs are set to $up_i = dn_i = 0$.

The problem is then exactly equivalent to the original instance, which concludes the proof. \Box

C. Strong NP-completeness for a variable number of processors

Theorem 2. MMSH-DEC with a variable number of homogeneous processors is NP-complete in the strong sense.

Proof. We already showed that MMSH-DEC is in NP when proving Theorem 1. We now prove that it is strongly NP-hard, i.e., it remains NP-hard even if the input parameters are bounded above by a polynomial in p, where p is the number of processors.

Let I_1 be an instance of 3-PARTITION [15]: Given 3nintegers $\{a_1, \ldots, a_{3n}\}$ whose sum is $nB = \sum_{i=1}^{3n} a_i$ and with $\frac{B}{4} < a_i < \frac{B}{2}$ for $1 \le i \le 3n$, does there exist a partition of $\{1, \ldots, 3n\}$ into n subsets S_1, \ldots, S_n such that $\sum_{i \in S_i} a_i = B$ for $1 \le i \le n$? We create an instance I_2 of MMSH-DEC with n processors and 4n jobs as follows:

- For $1 \le i \le 3n$, the execution time of job J_i is $w_i = a_i$;
- For 3n + 1 ≤ i ≤ 4n, the execution time of job J_i is w_i = ^B/₂; these n additional jobs are hence larger than any other jobs;
- The bound on the max-stretch is set to 3.

We now prove that I_1 has a solution if and only if I_2 has a solution, and the solution of I_2 always maps exactly one of the additional large jobs on each processor.

- Consider first that I_1 has a solution, $\{S_1, \ldots, S_n\}$, with $\sum_{i \in S_i} a_i = B$. Given the constraints on the a_i 's, there are exactly three elements per set S_i . We consider the schedule with jobs from S_i and one additional job on processor i, hence four jobs per processor. As already stated, jobs are scheduled from the shortest to the longest, and therefore, the k-th job always has a stretch that is at most k. Hence, the first three jobs have a stretch at most 3 (equal to three if the three jobs have a weight $\frac{B}{3}$, and then the third job has a stretch $\frac{B}{B/3} = 3$). The last job finishes at time 3B/2 and it has a length B/2, hence it has a stretch of 3. Finally, each processor has a stretch of 3, hence we have constructed a solution to I_2 .
- Consider now that I₂ has a solution, that is a partition of jobs onto processors, {P₁,..., P_n}, such that the maxstretch is at most 3. Since the max-stretch is 3 and all jobs are of length at most ^B/₂, no job can finish after ^{3B}/₂ and only a job of length B/2 can finish at ^{3B}/₂ (a smaller job would lead to a larger stretch). Furthermore, the total work is ^{3nB}/₂ and there are n processors, so each processor finishes exactly at time ^{3B}/₂, with one of the n large jobs. Consider S_i = P_i \ {k | k > 3n}: S_i consists in three jobs of sum B (since all job weights are between B/4 and B/2), hence the S_i's are a solution to I₁.

To conclude, I_1 has a solution if and only if I_2 has a solution, and we have performed in polynomial time the reduction from 3-PARTITION that is NP-hard in the strong sense, therefore MMSH-DEC is NP-complete is the strong

sense when considering a variable number of homogeneous processors. $\hfill \Box$

Corollary 2. MMSECO-DEC with an arbitrary number of processors is strongly NP-complete.

Proof. By Lemma 1, MMSECO-DEC is in NP. We perform a reduction from MMSH-DEC with p processors, where $P^e = 1$, $P^c = p - 1$, the speed of the edge processor is set to 1, and all communication costs of jobs are set to $up_i = dn_i = 0$. The problem is then exactly equivalent to the original instance, since all jobs originate from the only edge processor, and they can be scheduled either on the edge processors (cloud and edge) behave similarly; the platform is hence fully homogeneous.

V. Algorithms

Since MINMAXSTRETCH-EDGECLOUD is NP-complete even in the offline case, we design general heuristic algorithms to address the problem in the most general online setting, with P^e different edge computing units generating jobs, that can be processed either locally, or on one of the P^c cloud processors.

We first propose a strategy that does not exploit the cloud platform, but executes all jobs locally following a competitive strategy (Section V-A). We then design a simple greedy algorithm that greedily decides whether to execute a job locally or on the cloud (Section V-B). Finally, we revisit efficient algorithms from the literature to adapt them to the edge-cloud setting, building upon the shortest remaining processing time of a job in Section V-C, and on an earliest deadline first strategy in Section V-D. Please note that running examples for the algorithms are described in the companion research report [6], but omitted here due to lack of space.

The algorithms are event-based, i.e., we reconsider decisions only when an event occurs, so that they all remain of polynomial cost. There are at most 4n events, corresponding to the time-steps at which job J_i , for $1 \le i \le n$,

- 1) is released at edge processor o_i (event $\mathscr{E}_r(i)$);
- 2) completes its execution on processor alloc(i) (event $\mathscr{E}_f(i)$);
- 3) completes an uplink communication (event $\mathscr{E}_{up}(i)$);
- 4) completes a downlink communication (event $\mathscr{E}_{dn}(i)$).

A. Baseline: Edge-Only strategy

We first consider a simple heuristic, *Edge-Only*, that does not use the cloud platform, but where all jobs are executed locally on the edge. This might actually be a good strategy when edge processing units have a good processing speed and/or when communications are costly.

Since all edge processors are independent with this strategy, the problem consists in minimizing the max-stretch on one processor, and this is done independently at each edge processing unit. We use the Stretch-So-Far Earliest-Deadline-First algorithm, designed by Michael Bender et al. [3]. Indeed, this algorithm is Δ -competitive with a single processor, where Δ is the ratio between the longest and the shortest job. Note that we have to modify the algorithm of [3] to account for the edge-cloud framework: when computing the stretch of a job, we consider its potential execution time on the cloud. We will compare heuristics exploiting the cloud with this edgeonly strategy in Section VI.

B. Greedy

We design a greedy heuristic, called *Greedy*, that greedily schedules first the job that would currently achieve the highest stretch, as it is the job that might impact most the maximum stretch. Decisions are taken only when an event occurs (new job released, job completion, completion of an uplink or downlink communication).

The algorithm works as follows: at each event, as long as there are available resources, we compute for each job the minimum stretch that can be achieved using an available resource immediately. We select the job that maximizes this value, and execute the job on the resource (edge or cloud) on which it achieves the minimum stretch.

The complexity of the routine that runs at each event is $O(n(1 + P^c))$: compute the stretch for each job, for either a local execution, or the execution on one of the P^c cloud computing units. Since there are at most 4n events, the worst case complexity is $O(n^2P^c)$. Of course, in an online setting, we expect the number of jobs that are simultaneously processed in the system to be far less than n. Hence, the average complexity is expected to be much lower.

C. Shortest Remaining Processing Time

The *SRPT* heuristic builds on the classical Shortest Remaining Processing Time strategy, which assigns to a processing unit the job that it can finish the earliest, so that shortest jobs will be executed with little delay. Again, decisions are taken at each event.

The incentive for the use of this algorithm is that it is a O(1)-competitive algorithm for the average stretch problem, as shown in [28]. However, this is not the case for the maximum stretch: a long job could be blocked by short jobs for an arbitrary long duration. Note that migration is not allowed, but full re-execution is possible, thus a job that has been preempted by another job might start again (from scratch) on another processor, if it becomes the job that will finish first on that other processor.

The algorithm takes new decisions at each event. As long as we still have a processor on which we can execute a job, we choose the job that can be completed the earliest and execute it on the processor that can complete it the earliest, and we remove this job and processor from the list of available jobs/processors. Due to lack of space, the algorithm is detailed in the companion research report [6].

The complexity of the routine that runs at each event is in $O(n(P^e + P^c))$, and since there are at most 4n events, the worst case complexity is $O(n^2(P^e + P^c))$.

D. Stretch-so-far EDF

For one processor in the offline case, the algorithm in [4] finds the optimal maximum stretch in polynomial time. This

algorithm was extended to the online case, and obtained an algorithm that is Δ -competitive, where Δ is the ratio between the longest and the shortest job. The core idea of this algorithm consists in giving deadlines to jobs, based on an estimate of the maximum stretch. This estimate can change over the course of the algorithm and is computed by finding the lowest stretch we can get, through a binary search on the stretch. More precisely, these deadlines are computed as $d_i = r_i + S_e \times t_i$, where S_e is the stretch estimate. This estimate is computed as S_e = $\alpha \times S^*_{\leq k},$ where $S^*_{\leq k}$ is the optimal stretch of the jobs that have already been released (computed through binary search), and α is a parameter of the algorithm. Then, the algorithm schedules the jobs by taking the earliest of these deadlines first. With a well chosen α , this algorithm has a competitive ratio that is a function of $\Delta = \frac{w_{max}}{w_{min}}$. More specifically, for $\alpha = 1$, it is Δ -competitive, which is optimal up to a constant factor. However, the result can be better if Δ is known in advance. This algorithm is called Stretch-so-far Earliest-Deadline-First, since it is the algorithm Earliest Deadline First, with, as input, deadlines derived from the optimal stretch so far.

In the original problem, this algorithm relies on the fact that *Earliest Deadline First (EDF)* is optimal with a single processor. However, this is not the case anymore when we consider the MINMAXSTRETCH-EDGECLOUD problem, since we need to account for uplink and downlink communications for each job. We provide an example of the non-optimality of *EDF* with two jobs and one cloud processor:

- J_1 with $up_1 = 3$, $w_1 = 1$, $dn_1 = 0$, and deadline $d_1 = 5$;
- J_2 with $up_1 = 1$, $w_2 = 3$, $dn_2 = 0$, and deadline $d_2 = 6$.

EDF establishes that the job with the highest priority is J_1 , since $d_1 < d_2$. Therefore, it first sends J_1 to the cloud, and hence J_2 starts its uplink communication at time 3 and completes at time $3 + 1 + 3 = 7 > d_2$. However, by executing J_2 first, J_1 is able to also fit its deadline (uplink communication starting at time 1 while J_2 is executed on the cloud, and execution of J_1 at time 4).

Still, we propose an adaptation of this algorithm for MINMAXSTRETCH-EDGECLOUD. However, while *EDF* tells us which job to execute, we now also have to decide on which processor we execute the job with the highest priority. We decide to execute the job on the processor that minimizes the stretch of this job (or, equivalently, its finishing time). This algorithm is called *SSF-EDF*.

At each event corresponding to the release of a job (event $\mathscr{E}_r(i)$), we compute deadlines for each of the jobs currently on the platform, accounting for the remaining working time for the job if part of it has already been executed. These deadlines are computed using a target stretch and $\alpha = 1$, and a binary search is done to try different stretches. Indeed, the goal is to find the best possible achievable stretch at the current time, but as already pointed out, the *EDF* strategy is no longer optimal and we may not get the optimal stretch. In some cases, a stretch might not be achievable, while a smaller stretch would have been possible with another scheduling strategy.

Given a set of deadlines (and hence a target stretch), we assign the job with the smallest deadline on the processor

where it completes the earliest (either the origin edge processor of the job, or one of the P^c cloud processors), and then iterate on other jobs sorted by non-decreasing deadlines. If all deadlines can be respected, a smaller stretch might be considered, and when the binary search is done, we update remaining processing and communication times, similarly to *SRPT*, with the assignment returned by *SSF-EDF* until the next event.

The complexity of *SSF-EDF* algorithm is $O(n(1 + P^c))$ at each event, as for the greedy (compute for each job its shortest execution time locally or on one of the cloud processors), but we should re-execute the routine as many times as needed by the binary search. So the actual worst case complexity that we get is $O(n^2 P^c \log(\frac{1}{\epsilon}))$, where ϵ is the relative precision we want to keep on the estimate of the stretch (the binary search takes a time $\log(\frac{1}{\epsilon})$).

VI. SIMULATIONS

In order to test the efficiency of the proposed algorithms and to compare them with the *Edge-Only* strategy, we implemented and simulated them in different scenarios, using parameters from real edge-cloud platforms. All heuristics are implemented in C++ and the code downloaded github.com/Redouane-Elghazi/ be at can Max-Stretch-Minimization-on-an-Edge-Cloud-Platform.git for reproducibility purpose.

A. Simulation setup

We consider two main problem instances: instances with random scenarios to experiment with the various parameters, and instances with parameters coming from [24], inspired by a real-life situation.

Random instances: We first consider a platform with 20 cloud processors, 10 slow edge processors with speed s = 0.1, and 10 fast edge processors with speed s = 0.5. The jobs are generated using a uniform distribution for the execution and communication times, as well as the release date and the origin processor. Both execution and communication times follow the same distribution. The parameters of the distribution for execution, through the notion of Communication/Computation-Ratio (CCR). More precisely, both distributions are chosen so that the ratio between their expected values is equal to some value determined in advance. We consider CCRs ranging from 0.1 (compute-intensive scenario) to 10 (communication-intensive scenario).

Kang instances: These instances are meant to be closer to a real life situation, inspired from [24]. They contain different types of edge processors depending on whether their computational unit is a GPU or a CPU, and their communication channel is 3G, LTE, or Wi-Fi. Then, the jobs are created with an execution time and a communication time that is directly inferred from the type of computational unit and the type of communication channel. More precisely, the values are chosen as follows (expected values and speeds according to [24]):



Fig. 2: Maximum stretch, from left to right, (a) varying the CCR on random instances; (b) varying the load on random instances; (c) varying the number of jobs n on the Kang instance with 20 edge processors and 10 cloud processors; (d) varying n on the Kang instance with 100 edge processors and 10 cloud processors.

- the execution time follows a normal distribution with mean 6 and relative standard deviation ¹/₃;
- the uplink communication time follows a normal distribution with mean t and relative standard deviation $\frac{1}{3}$, where t = 95 with Wi-Fi, t = 180 with LTE, and t = 870 with 3G;
- the downlink communication time is 0 for all jobs, as the place of delivery is not relevant in this specific case.

Also, the speed of an edge processor is $\frac{6}{81}$ if the processor computes on a GPU, and $\frac{6}{382}$ for CPUs.

Release dates and load: With both instance types, the distribution of the release dates is chosen to control the load on edge processors, i.e., the average number of jobs originating from the edge processor that are simultaneously in the system. Hence, for a load ℓ , the maximum release date is set to $\sum_{\substack{i=1 \\ l \\ j=1}}^{n} w_i$: the sum of the work over the aggregated speed is the average execution time using all processors; dividing this ratio by, say, $\ell = 0.1$, augments release times by a factor ten, thereby decreasing the load accordingly. By default, we consider a load of 5% ($\ell = 0.05$), hence the processors are able to cope with the jobs before new jobs arrive in the system, but we still have some interesting collisions on both the edge and the cloud.

B. Simulation results

For each experiment, we plot the maximum stretch, with and without the *Edge-Only* heuristic, since it often leads to much higher and out-of-range stretch values. We further report the execution time of each algorithm (time to compute the schedule). All plots are available in the companion research report [6]. We present only a subset of them and give general trends in this paper. Each point corresponds to the average of 1000 instances created with the same parameters.

Random instances: We first run the heuristics on random instances (see Figure 2(a)), where the number of jobs is set to n = 4000, and the CCR varies from 0.1 to 10. This illustrates the fact that the new heuristics obtain much better stretches compared to *Edge-Only* for small values of CCR, because it is then very useful to send jobs to the cloud. *SSF-EDF* is the best in all scenarios, with *SRPT* being very close. Their stretch exceeds two only for the largest values of CCR. *Greedy* is slightly behind, with stretches going up to 3.5.

Also, we investigate scenarios with a higher load, with a load ℓ up to 2 (recall that the default load was $\ell = 0.05$ in Figure 2(a)). However, we do not report *Edge-Only*, which becomes too costly since all jobs compete on the edge. We focus on the instance with a CCR of 1 and n = 4000 jobs, see Figure 2(b). Again, *SSF-EDF* is clearly the best, and this is even more striking when the load increases, where the stretch of *SRPT* and *Greedy* drastically increases while it stays under three for *SSF-EDF*. It is also interesting to note that *Greedy* sometimes becomes better than *SRPT* with an increased load.

Kang instances: The CCR is dictated by the platform parameters, hence we vary the number of jobs, in two different scenarios with 20 or 100 edge processors, while keeping ten cloud processors (see Figures 2(c) and (d)). Again, *SSF-EDF*, closely followed by *SRPT*, is clearly the best heuristic, and *Edge-Only* cannot keep up when the number of jobs increases (see [6]). Note that with more edge processors, and hence more competition for cloud resources, *Greedy* gets a stretch that is close to the ones achieved by *SRPT* and *SSF-EDF*.

Execution times: Detailed execution times are available in [6]; overall, heuristics do not exceed a few seconds, with *SRPT* being much faster to execute than *SSF-EDF* and *Edge-Only. Greedy* is similar to *SRPT* with low loads, but its execution time drastically increases with the load. In general, execution times increase with n or with the load, but remain constant with the CCR. For the Kang instances, execution times are much higher in the scenario with 100 edge processors (up to 16 seconds for *SSF-EDF*).

Summary: Overall, we conclude that it is always a better option to choose SRPT over Greedy for lightly loaded systems, since it always leads to smaller stretches, and has approximately the same time complexity. Greedy however may outperform SRPT when the system is heavily loaded, but then also becomes more costly. Then, choosing between SRPT and SSF-EDF is a matter of trade-off between efficiency and execution time, since SSF-EDF is more costly than SRPT, but also more efficient. Edge-Only is a costly solution that does not exploit the cloud, and a comparison to this strategy highlights the importance of using cloud resources when available, in particular when communication costs are not too important.

VII. CONCLUSION

We have tackled the problem of scheduling independent jobs on an edge-cloud platform, with the goal of minimizing the maximum stretch of jobs. Jobs are produced by edge computing units, and they can be processed either locally (at a slow speed), or delegated to a cloud (at the price of communications to pay). We have designed a general model, accounting for a realistic communication model, and formalizing the scheduling constraints in this particular setting. While establishing problem complexity, we have encountered the problem of minimizing the max-stretch without release dates on a homogeneous platform, whose complexity was left open in the literature; we have proven that it is NP-hard, filling a gap in fundamental scheduling complexity results. The NPhardness of the problem in the edge-cloud setting is then a corollary of this result.

We designed heuristic algorithms to address the problem in the general online setting, and we assessed the performance of these algorithms through simulations, using parameters coming from real-world edge-cloud platforms. The algorithms delegating jobs to the cloud achieve much better stretches than a competitor *Edge-Only* approach, in particular when communication costs are not too high. The proposed *SSF-EDF* algorithm always achieves very good stretches. *SRPT* is also is very efficient, and computes its schedules more rapidly, hence it might be an interesting alternative to *SSF-EDF*.

As future work, it would be very interesting to derive theoretical bounds for the proposed online algorithms (competitive results), for instance for some specific job distributions. Furthermore, we would like to investigate a more complicated framework where cloud processors are not available full-time: in practice, a realistic but intricate framework is to consider that cloud processors may be dynamically requested by other applications at certain time intervals.

Acknowledgments: We thank Denis Trystram for several fruitful discussions on edge computing.

REFERENCES

- B. Awerbuch, Y. Azar, S. Leonardi, and O. Regev. Minimizing the flow time without migration. SIAM J. on Computing, 31(5):1370–1382, 2002.
- [2] L. Becchetti, S. Leonardi, and S. Muthukrishnan. Average stretch without migration. J. Comput. Syst. Sci., 68:80–95, 02 2004.
- [3] M. A. Bender, S. Chakrabarti, and S. Muthukrishnan. Flow and stretch metrics for scheduling continuous job streams. In *Proc. of the 9th ACM-SIAM Symp. on Discrete Algorithms*, SODA'98, page 270–279, 1998.
- [4] M. A. Bender, S. Muthukrishnan, and R. Rajaraman. Improved algorithms for stretch scheduling. In *Proc. of the 13th Annual ACM-SIAM Symp. on Discrete Algorithms*, SODA'02, page 762–771, USA, 2002.
- [5] M. A. Bender, S. Muthukrishnan, and R. Rajaraman. Approximation algorithms for average stretch scheduling. *Journal of Scheduling*, 7:195– 222, 2004.
- [6] A. Benoit, R. Elghazi, and Y. Robert. Max-Stretch Minimization on an Edge-Cloud Platform. Research report RR-9369, Inria, 2020. Available at hal.inria.fr/hal-02972296v1.
- [7] P. Bhat, C. Raghavendra, and V. Prasanna. Efficient collective communication in distributed heterogeneous systems. *Journal of Parallel and Distributed Computing*, 63:251–263, 2003.
- [8] A. Brogi, S. Forti, C. Guerrero, and I. Lera. How to place your apps in the fog: State of the art and open challenges. *CoRR*, 1901.05717, 2019.
- [9] P. Brucker. Scheduling Algorithms. Springer-Verlag, 2004.

- [10] H. Casanova, F. Desprez, and F. Suter. Minimizing stretch and makespan of multiple parallel task graphs via malleable allocations. In 39th International Conference on Parallel Processing, pages 71–80, 2010.
- [11] J. Celaya and L. Marchal. A Fair Decentralized Scheduler for Bag-of-Tasks Applications on Desktop Grids. In 10th IEEE/ACM Int. Conf. on Cluster, Cloud and Grid Computing, pages 538–541, 2010.
- [12] C. Chekuri and S. Khanna. Approximation schemes for preemptive weighted flow time. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02, page 297–305, 2002.
- [13] P. Chrétienne, E. G. Coffman Jr., J. K. Lenstra, and Z. Liu, editors. Scheduling Theory and its Applications. John Wiley and Sons, 1995.
- [14] V. Gama Pinheiro. The management of multiple submissions in parallel systems : the fair scheduling approach. Theses, Université de Grenoble & Universidade de São Paulo, 2014. Available at https: //tel.archives-ouvertes.fr/tel-01677743.
- [15] M. R. Garey and D. S. Johnson. Computers and Intractability, a Guide to the Theory of NP-Completeness. W.H. Freeman and Company, 1979.
- [16] A. Gupta, S. Im, R. Krishnaswamy, B. Moseley, and K. Pruhs. Scheduling heterogeneous processors isn't as easy as you think. In *Proc. of the* 23rd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'12, page 1242–1253, USA, 2012.
- [17] L. Hollermann, T. S. Hsu, D. R. Lopez, and K. Vertanen. Scheduling problems in a practical allocation model. *J. Combinatorial Optimization*, 1(2):129–149, 1997.
- [18] B. Hong and V. Prasanna. Distributed adaptive task allocation in heterogeneous computing environments to maximize throughput. In *Int. Parallel and Distributed Processing Symposium IPDPS*'2004, 2004.
- [19] T. S. Hsu, J. C. Lee, D. R. Lopez, and W. A. Royce. Task allocation on a network of processors. *IEEE Trans. Computers*, 49(12):1339–1353, 2000.
- [20] IBM. Cloud at the Edge, 2019. https://www.ibm.com/cloud/blog/ cloud-at-the-edge.
- [21] Infradata. What is Edge-Cloud Computing, 2020. https://www.infradata. com/resources/what-is-edge-cloud/.
- [22] Intel. Edge-Cloud and Edge Servers, 2020. https://www.intel.com/ content/www/us/en/edge-computing/edge-cloud.html.
- [23] C. Jiang, T. Fan, H. Gao, W. Shi, L. Liu, C. Cérin, and J. Wan. Energy aware edge computing: A survey. *Computer Communications*, 151, 02 2020.
- [24] Y. Kang, J. Hauswald, C. Gao, A. Rovinski, T. Mudge, J. Mars, and L. Tang. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. In 22nd Int. conf. on Architectural Support for Prog. Languages and Operating Systems (ASPLOS), pages 615–629, 2017.
- [25] A. Legrand, A. Su, and F. Vivien. Minimizing the stretch when scheduling flows of divisible requests. *Journal of Scheduling*, 11:381– 404, 2008.
- [26] J. Liu, Y. Mao, J. Zhang, and K. B. Letaief. Delay-optimal computation task scheduling for mobile-edge computing systems. *CoRR*, abs/1604.07525, 2016.
- [27] B. Moseley, K. Pruhs, and C. Stein. The complexity of scheduling for p-norms of flow and stretch. In *Proc. IPCO'13*, 2013.
- [28] S. Muthukrishnan, R. Rajaraman, A. Shaheen, and J. E. Gehrke. Online scheduling to minimize average stretch. In 40th Annual Symposium on Foundations of Computer Science, pages 433–443, 1999.
- [29] Nokia. Edge-Cloud, 2020. https://www.nokia.com/networks/solutions/ edge-cloud/.
- [30] Y. Robert and F. Vivien, editors. *Introduction to Scheduling*. Chapman and Hall/CRC Press, 2009.
- [31] F. Rodrigo De Souza, A. Da Silva Veith, M. Dias de Assuncao, and E. Caron. Scalable Joint Optimization of Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Infrastructure. In 18th Int. Conf. on Service Oriented Computing (ICSOC), Dubai, 2020.
- [32] F. Rodrigo De Souza, M. Dias de Assuncao, E. Caron, and A. da Silva Veith. An Optimal Model for Optimizing the Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Computing. In SBAC-PAD'20, Porto, Portugal, Sept. 2020.
- [33] T. Saif and M. Parashar. Understanding the behavior and performance of non-blocking communications in MPI. In *Proceedings of Euro-Par* 2004: Parallel Processing, LNCS 3149, pages 173–182, 2004.
- [34] O. Sinnen and L. Sousa. Communication contention in task scheduling. IEEE Trans. on Parallel and Distributed Systems, 16(6):503–515, 2005.
- [35] Y. Wu and G. Cao. Stretch-optimal scheduling for on-demand data broadcasts. In Proc. Tenth Int. Conf. on Computer Communications and Networks, pages 500–504, 2001.