

# Scheduling over Scenarios on Two Machines\*

Esteban Feuerstein<sup>1</sup>, Alberto Marchetti-Spaccamela<sup>2</sup>, Frans Schalekamp<sup>3</sup>,  
René Sitters<sup>4,5</sup>, Suzanne van der Ster<sup>4</sup>, Leen Stougie<sup>4,5</sup>, and Anke van Zuylen<sup>3</sup>

<sup>1</sup> Departamento de Computación, FCEyN, UBA, Buenos Aires, Argentina  
efeuerst@dc.uba.ar

<sup>2</sup> Sapienza University of Rome, Italy  
alberto@dis.uniroma1.it

<sup>3</sup> College of William and Mary, Department of Mathematics,  
Williamsburg, VA 23185, USA  
{frans,anke}@wm.edu

<sup>4</sup> Vrije Universiteit Amsterdam, The Netherlands  
{r.a.sitters,suzanne.vander.ster,l.stougie}@vu.nl

<sup>5</sup> CWI Amsterdam, The Netherlands  
{r.a.sitters,stougie}@cwi.nl

**Abstract.** We consider scheduling problems over scenarios where the goal is to find a single assignment of the jobs to the machines which performs well over all possible scenarios. Each scenario is a subset of jobs that must be executed in that scenario and all scenarios are given explicitly. The two objectives that we consider are minimizing the maximum makespan over all scenarios and minimizing the sum of the makespans of all scenarios. For both versions, we give several approximation algorithms and lower bounds on their approximability. With this research into optimization problems over scenarios, we have opened a new and rich field of interesting problems.

**Keywords:** job scheduling, makespan minimization, scenarios, approximation.

## 1 Introduction

We consider optimization problems *over scenarios* where the goal is to find a single solution that performs well for each scenario in a given set of scenarios. In particular, we consider the scheduling problem where the objective function is the makespan: we are given a set  $J$  of jobs, each with a processing time, and a set of scenarios; each scenario is specified by a subset of jobs in  $J$  that must be executed in that scenario. Our goal is to find an assignment of jobs to machines that *is the same for all scenarios* and optimizes a function of the *makespan*, i.e., the completion time of the last completed job, *over all scenarios*.

---

\* This work was partially supported by EU-IRSES grant EUSACOU and Tinbergen Institute. EF was partially supported by projects PRH PICT 2009-119 and UBACYT 20020090100149. AvZ was partially supported by Suzann Wilson Matthews summer research award.

The two objectives that we consider are minimizing the *maximum* makespan over all scenarios and minimizing the *sum* of the makespan of all scenarios. We note that when the input contains only a single scenario, both versions of the problem reduce to the usual makespan minimization problem.

As an example, suppose that  $J$  contains three jobs, numbered 1, 2, and 3, that must be executed on two machines; the processing time of job 1 is 2 while the processing time of jobs 2 and 3 is 1. There are three scenarios  $S_1 = \{1, 2, 3\}$  and  $S_2 = S_3 = \{2, 3\}$ . Assigning job 1 to the first machine and jobs 2 and 3 to the second machine minimizes the maximum makespan over all scenarios, while assigning jobs 1 and 2 to the first machine and job 3 to the second one minimizes the sum of the makespans of all scenarios.

The more egalitarian objective function of minimizing the maximum makespan over all scenarios fits in the framework of *robust* optimization, where usually not so much a finite set of scenarios is explicitly given, as in our problem, but ranges for values of input parameters (see [3]). We will refer to this objective as the Min-Max objective. The more *utilitarian* objective function of minimizing the sum of the makespans of all scenarios fits in the framework of *a priori* optimization, though a priori optimization has so far only been introduced as a problem where the scenarios are random objects and the objective is to minimize the expected objective value. In that sense, minimizing the sum of makespans could be seen as the a priori problem with a uniform discrete distribution over a finite set of fully specified scenarios. In general, the deterministic problem of optimizing over a finite set of scenarios can be seen as an alternative to the stochastic a priori setting [13], in case a limited number of likely scenarios exists. We refer to this objective as the MinSum objective.

In an indirect way, combinatorial optimization problems over scenarios with the MinSum objective have appeared as the first-stage problem in a boosted sampling approach to two-stage stochastic optimization problems [11]. In [11], scenarios are defined within a so-called black box, meaning that they can only be learnt by sampling. From the black box, a finite set of scenarios is sampled, giving rise to a deterministic optimization problem over the drawn set of scenarios, in which a single solution needs to be found, that minimizes the sum of the objective values for the drawn scenarios. In this sense some results on combinatorial optimization problems over scenarios have appeared, like VERTEX COVER, STEINER TREE and UNCAPACITATED FACILITY LOCATION [11].

Modeling optimization problems over a finite set of given scenarios yields a rich source of interesting new combinatorial optimization problems, which are in general harder than their single-scenario versions. Specifically, almost any single-scenario scheduling problem has an interesting multi-scenario variation. As mentioned before, in this paper we focus, as a first example, on minimizing the maximum makespan over all scenarios and minimizing the sum (or, equivalently, the average) of the makespan of all scenarios.

The specific setting of the scheduling problem over scenarios appears in situations where jobs have to be performed by skilled machines (workers), and some investment is required to attain the skill for a particular job. In such situations,

one should decide on an assignment of all possible jobs to the workers, such that the workers can train for the jobs assigned to them ahead of time. The problem then is to assign jobs (specializations) to machines (workers), so that the workload of a machine for any scenario of jobs, from a set of scenarios likely to occur, is minimized. Examples of such a setting are assignment of clients to lawyers, households to power sources, compile-time assignment of computational tasks to processors. In most of such situations, the robust version of the problem with the MinMax objective is rather plausible, especially in situations where a set of likely scenarios to hedge against can be specified upfront.

Another motivation, though a bit indirect, comes from distributed information retrieval: in a term-partitioned index, it is good to allocate to the same processor terms appearing frequently together in queries, so as to minimize the communication cost (to solve an intersection query between two terms that reside in different processors, one of the posting lists must be sent to the processor holding the other). But this goal must be complemented with that of balancing the load, as it is not viable to put all the terms in the same processor. Therefore, it is necessary to divide “clusters” of commonly co-occurring terms among the processors, trying to balance the load. Naturally, queries appear sequentially over time and are not known a priori. One could, as an approximation, optimize considering as input the more likely scenarios. The partition must indeed be done a priori, because lists must be assigned to processors a priori.

To the best of our knowledge, this problem has not been considered in the literature. An a priori version of scheduling with stochastic scenarios has been studied in [4,5], albeit not from an approximation theory point of view, but merely presenting experimental results, and with the scheduling objective of minimizing the sum of completion times of all the jobs per scenario.

We now give a formal definition of the two problems we consider. We restrict ourselves to the case of two machines. We are given a set of jobs  $J$  with for each job  $j \in J$  a processing times  $p_j$ , and a set of  $k$  scenarios  $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$ , where each scenario  $S_i \in \mathcal{S}$  is a subset of  $J$ . In each scenario, we are interested in minimizing the makespan, but we are restricted to finding a solution, i.e., an assignment of the jobs to the machines, that applies to every one of the scenarios. Clearly, a solution that is good for one scenario may be bad for another. This gives rise to specifying objectives that reflect the trade-off between the various scenarios. In this paper we define the following two versions of the problem.

- **MM2** Assign the jobs in  $J$  to two machines in such a way that the maximum makespan over the given scenarios is minimized. In other words, if we denote the makespan of a subset  $S \subseteq J$  of jobs by  $p(S) = \sum_{j \in S} p_j$ , we are looking for a partition  $A, \bar{A}$  of  $J$ , that minimizes  $\max_{i=1,2,\dots,k} \max\{p(A \cap S_i), p(\bar{A} \cap S_i)\}$ .
- **SM2** Assign the jobs to the machines such that the sum of the makespans of the given scenarios is minimized. Using the notation just introduced, we are seeking a partition  $A, \bar{A}$  of  $J$ , that minimizes  $\sum_{i=1}^k \max\{p(A \cap S_i), p(\bar{A} \cap S_i)\}$ .

For both objective functions, the problems are NP-hard, since the single-scenario version is NP-hard. However, the single-scenario version is only weakly NP-hard

for 2 machines and an FPTAS exists [1], whereas the problems defined here are strongly NP-hard. We will give various approximability and inapproximability results for several different versions of the problem depending on restrictions of the input. In particular, the special cases that we consider are the following:

1.  $p_j = 1 \forall j \in J$ , that is, the case where all processing times are unitary;
2.  $|S_i| \leq r \forall S_i \in \mathcal{S}$ , that is, the case where the number of jobs in each scenario is bounded by a constant;
3.  $k = |\mathcal{S}|$  is constant, that is, the case that the number of scenarios is a constant.

In Section 2, we study the problem MM2; we show that the problem cannot be approximated to within a ratio of  $2 - \varepsilon$  already in the case where  $p_j = 1$  and a ratio of  $3/2$  if  $|S_i| \leq 3$  and  $p_j = 1$ . On the positive side, we give a polynomial-time algorithm for the version in which every scenario contains 2 jobs. If  $k$ , the number of scenarios, is constant then there exists a PTAS; for an arbitrary number of scenarios, a  $O(\log^2 k)$  approximation ratio exists. The latter two results are a consequence of an observed direct relation to the so-called VECTOR SCHEDULING problem (see Section 2 for its definition) and results of [6].

In Section 3, we study problem SM2. We prove inapproximability within 1.0196 assuming  $P \neq NP$ , and within 1.0404 under the *Unique Games Conjecture* [15]. On the positive side, we present a  $3/2$ -approximate randomized algorithm. For instances with scenarios of size at most 3, we use a reduction to MAX CUT to obtain a 1.12144-approximation algorithm. For scenarios of size at most  $r$ , we present a reduction to WEIGHTED MAX NOT-ALL-EQUAL  $r$ -SAT and use this to obtain better-than- $3/2$ -approximations for problems where the scenario sizes are not larger than 4.

Some thoughts about related problems, and ideas for future research are contained in a concluding section.

## 2 Minimizing Maximum Makespan

We obtain inapproximability of MM2 using a recent result [2] on the hardness of HYPERGRAPH BALANCING: given a hypergraph find a 2-coloring of the vertices such as to minimize over all hyperedges the discrepancy between the number of vertices of the two colors.

**Theorem 1.** *It is NP-hard to approximate MM2 with unitary jobs within ratio  $2 - \epsilon$ .*

This is remarkable since, trivially, any solution, for any job sizes, is 2-approximate (since we consider the problem for two machines only). In the full version of this paper [8] we prove a hardness bound of  $3/2$  when  $|S_i| \leq 3$  and  $p_j = 1$ . This result completes the hardness characterization.

We now show that, if the number of jobs per scenario is 2, then the problem is solvable in polynomial time.

**Theorem 2.** *MM2 with  $|S_i| = 2$  for all  $S_i \in \mathcal{S}$  can be solved in time  $O(|\mathcal{S}| \log |\mathcal{S}|)$ .*

*Proof.* We create a graph with a vertex for each job and connect by an edge the jobs that appear together in a scenario. We define the weight of edge  $(j, k)$  to be  $p_j + p_k$ , i.e., the sum of the processing times of the jobs associated to the incident vertices. Note that a solution for the a priori scheduling problem is a partitioning of the job set, and can be associated with a coloring of the vertices in this graph problem with two color classes. The objective value is then equal to the maximum of the highest weight of any monochromatic edge and the largest processing time of any job.

In other words, we should find a 2-coloring of the vertices of this graph, such that the maximum weight of a monochromatic edge is minimized. A lowest weight edge in any odd cycle gives a lower bound on the objective value.

Consider the following algorithm. Starting with all vertices being part of their own singleton component, and having color 1, we grow components by inserting edges, and label the vertices with the component they belong to, and with a color that can assume two values; 1 and 2. A color inversion of a vertex changes the color of the vertex (i.e., if it is colored 1, the color is changed to 2, and vice versa). We consider the edges in order of descending weight. When considering the next edge, say  $(j, k)$ , the following 3 cases can occur.

*Case 1.* Vertices  $j$  and  $k$  have the same color, and are in the same component. We end the algorithm. An optimal partitioning of the job set is given by the two color classes, where jobs that have color 1 (respectively, 2) are assigned to machine 1 (respectively, 2) and the objective is equal to the weight of edge  $(j, k)$ .

*Case 2.* Vertices  $j$  and  $k$  have different colors. If the vertices are in different components, then we update the component label for all nodes of the smaller component (breaking ties arbitrarily), so that all vertices have the same label. We then proceed to the next edge.

*Case 3.* Vertices  $j$  and  $k$  have the same color, and are in different components. In this case we invert the color of all nodes in the smaller component (breaking ties arbitrarily), and then proceed as in Case 2.

By construction, two vertices of the same color in the same component are joined by an even-length path. Therefore, when the algorithm terminates in Case 1, we have found an odd cycle in the graph, of which this last edge has lowest weight. Note that the assignment of jobs with the same color to the same machine implies that the makespan of the scenario is bounded by the weight of the last considered edge. Since the weight of any such edge is a lower bound on the objective value, we have found an optimal solution. Its value is given by the maximum of the weight of a monochromatic edge and  $p_{\max} = \max_{j \in J} p_j$ .

The running time follows from the observation that any time we invert the color and/or update the label of a vertex, it ends up in a component of at least twice the size of the component it belonged to before. Hence, the label of a vertex can be updated at most  $\log |J|$  times. The total time can thus be bounded by  $|\mathcal{S}| \log |\mathcal{S}|$  time for sorting the edges by weight, plus  $|J| \log |J|$  time for updating the vertex colors and labels. Finally, we may assume without loss of generality that each job appears in at least one scenario, so  $|\mathcal{S}| \geq |J|/2$ .  $\square$

Another sharp characterization w.r.t. the number of scenarios, is obtained for the case of a constant number of scenarios. For jobs with unit processing times, the problem can be solved exactly: given that the number of scenarios is constant, there is only a constant number of job types, where the type of a job is the set of scenarios it is in. Then, the number of jobs on machine 1 of each type can be guessed. There are only a polynomial number of choices; an extension can also accommodate a constant number of machines in polynomial time. We notice that this also solves SM2 under the same restrictions in polynomial time.

**Theorem 3.** *MM2 and SM2 having jobs with unitary processing times can be solved in polynomial time if the number of scenarios is constant.*

A similar idea, with guessing the optimal value and rounding, leads to a PTAS in the general case under a constant number of scenarios, but this is also implied by the following result.

We conclude this section by noticing that if we consider any number of machines, the problem of minimizing the maximum makespan reduces to the VECTOR SCHEDULING problem, where each coordinate corresponds to a scenario.

**Definition 1.** *In the VECTOR SCHEDULING problem we are given a set  $V$  of  $n$  rational  $d$ -dimensional vectors  $v_1, \dots, v_n$  from  $[0, \infty)^d$  and a number  $m$ . A valid solution is a partition of  $V$  into  $m$  sets  $A_1, \dots, A_m$ . The objective is to minimize  $\max_{1 \leq i \leq m} \|\sum_{v_j \in A_i} v_j\|_\infty$ .*

This problem is a  $d$ -dimensional generalization of the makespan minimization problem, where each job is a  $d$ -dimensional vector and the machines are  $d$ -dimensional objects as well. In our setting, the dimension  $d$  equals the number of scenarios  $|\mathcal{S}|$ . Each coordinate of job  $j$  equals its processing time in the corresponding scenario (either 0 or  $p_j$ ). Results of Chekuri et al. [6] on VECTOR SCHEDULING can directly be translated into our setting.

**Theorem 4 ([6]).** *For the problem of minimizing the maximum makespan over scenarios  $S_i \in \mathcal{S}$  on  $m$  machines,*

1. *there exists a PTAS for the case that  $k = |\mathcal{S}|$  is constant*
2. *there exists a polynomial-time  $O(\log^2 k)$ -approximation for  $k$  scenarios;*
3. *there exists no  $c$ -approximation algorithm for any  $c > 1$ , when dealing with any number of scenarios.*

### 3 Minimizing Sum of Makespans

We now turn our attention to SM2, the problem of minimizing the sum of the makespans over all scenarios, in the case of 2 machines.

We start this section by noting that SM2 is MAX SNP-hard even with unitary processing times and scenarios containing two jobs each.

**Theorem 5.** *SM2 is NP-hard to approximate to within a factor of 1.0196 and UGC-hard to approximation to within a factor of 1.0404, even if all jobs have length 1, and all scenarios contain two jobs.*

The proof is through a reduction from MAX CUT [9], and the hardness of approximation results shown by Håstad [12] and Khot et al. [16]. The details are given in the full version of this paper [8].

In the remainder of this section, we will give approximation results for SM2. As for MM2 in the previous section, we notice that also for this problem any solution is a trivial 2-approximation. In the remainder of this section, we will first show that the algorithm that randomly assigns the jobs to the two machines independently with equal probability gives a 3/2-approximation. We then show two deterministic approximation algorithms, which give good approximation guarantees if the number of jobs per scenario is small.

### 3.1 A Randomized Approximation Algorithm

**Lemma 1.** *Consider a scenario  $S$ , and let  $A, \bar{A}$  be any partitioning of the jobs in  $S$ . When assigning each job of  $S$  to the two machines independently with equal probability, the expected load of the least loaded machine is at least  $\frac{1}{2} \min\{p(A), p(\bar{A})\}$ .*

*Proof.* An assignment of jobs to the two machines induces a partition of  $A$  into sets  $A', A''$ , and a partition of  $\bar{A}$  into sets  $\bar{A}', \bar{A}''$  where the jobs in the same set of the partition are assigned to the same machine. The sets  $A', A'', \bar{A}', \bar{A}''$  are not necessarily all non-empty. We will prove the lemma by showing that, conditioned on the sets  $A', A'', \bar{A}', \bar{A}''$ , the machine load of the least loaded machine is at least  $\frac{1}{2} \min\{p(A), p(\bar{A})\}$ , which implies that the statement also holds unconditionally.

Conditioned on the sets  $A', A'', \bar{A}', \bar{A}''$ , the least loaded machine has a load of  $\min\{p(A') + p(\bar{A}'), p(A'') + p(\bar{A}'')\}$  with probability  $\frac{1}{2}$  (namely, if  $A', \bar{A}'$  are assigned to one machine, and  $A'', \bar{A}''$  to the other machine), and  $\min\{p(A') + p(\bar{A}''), p(A'') + p(\bar{A}')\}$  with probability  $\frac{1}{2}$  (namely, if  $A', \bar{A}''$  are assigned to one machine, and  $A'', \bar{A}'$  are assigned to the other machine). Hence, conditioned on the partition of  $A$  into  $A', A''$  and of  $\bar{A}$  into  $\bar{A}', \bar{A}''$ , the expected load of the least loaded machine is

$$\frac{1}{2} \min\{p(A') + p(\bar{A}'), p(A'') + p(\bar{A}'')\} + \frac{1}{2} \min\{p(A') + p(\bar{A}''), p(A'') + p(\bar{A}')\}.$$

Note that a simple case analysis shows that the sum of the two terms is either at least  $\frac{1}{2} (p(A') + p(A'')) = \frac{1}{2} p(A)$  or at least  $\frac{1}{2} (p(\bar{A}') + p(\bar{A}'')) = \frac{1}{2} p(\bar{A})$ . So the load of the least loaded machine is at least  $\frac{1}{2} \min\{p(A), p(\bar{A})\}$ .  $\square$

**Theorem 6.** *Randomly assigning each job to the two machines independently with equal probability is a 3/2-approximation for SM2.*

*Proof.* Consider a scenario  $S$ , and let  $A$  be the set of jobs processed on machine 1, and  $\bar{A} = S \setminus A$  the set of jobs processed on machine 2 in a schedule with minimum makespan. Hence, the optimal makespan for  $S$  is  $\max\{p(A), p(\bar{A})\}$ . By Lemma 1,

the load of the least loaded machine in scenario  $S$ , if the jobs are randomly assigned to the machines with equal probability, is at least  $\frac{1}{2} \min\{p(A), p(\bar{A})\}$ . Hence, the load of the machine with the highest load is at most  $p(A)+p(\bar{A})-\frac{1}{2} \min\{p(A), p(\bar{A})\} = \max\{p(A), p(\bar{A})\} + \frac{1}{2} \min\{p(A), p(\bar{A})\} \leq \frac{3}{2} \max\{p(A), p(\bar{A})\}$ .

Hence, the expected makespan for scenario  $S$  is at most  $\frac{3}{2}$  times the optimal makespan for scenario  $S$ , which implies that the sum over all scenarios of the expected makespans is at most  $\frac{3}{2}$  times the optimal summed makespan of all scenarios.  $\square$

We remark that the proof of the previous lemma bounds the objective value by comparing the load on a machine in a given scenario to the load for the optimal schedule *for that scenario*, rather than the optimal schedule for our problem.

It is easy to see that the analysis of the simple randomized algorithm is tight, by considering an instance of two jobs  $\{1, 2\}$  with unitary execution time and one scenario  $S_1 = \{1, 2\}$ . The optimal solution is to assign one job to each machine, whereas the randomized algorithm either assigns both jobs to the same machine with probability  $\frac{1}{2}$ , or one job to each machine with probability  $\frac{1}{2}$ .

### 3.2 Deterministic Approximation Algorithms

To obtain a deterministic approximation algorithm, we show that the SM2 problem can be reduced to the WEIGHTED MAX NOT-ALL-EQUAL SATISFIABILITY problem, that we will abbreviate as MAX-NAE SAT.

**Definition 2.** *In MAX-NAE SAT, a boolean expression is given, and a weight for each clause. A clause in the expression is satisfied if it contains both true and false literals. The problem is to find an assignment of true/false values to the variables, such as to maximize the total weight of the clauses satisfied.*

Note that if  $r$  is such that  $|S_i| \leq r$  for all  $S_i \in \mathcal{S}$ , then by adding dummy jobs of processing time 0, we can assume that every scenario contains exactly the same number of jobs, i.e.,  $|S_i| = r$  for all  $S_i \in \mathcal{S}$ . We will reduce the SM2 problem with scenarios of size at most  $r$  to the MAX-NAE SAT problem with clauses of length  $r$  (MAX-NAE  $r$ -SAT).

**Theorem 7.** *A  $(1 - \gamma_r)$ -approximation for MAX-NAE  $r$ -SAT implies a  $(1 + 2^{r-2}\gamma_r)$ -approximation for the SM2 problem with  $|S| \leq r$  for all scenarios  $S \in \mathcal{S}$ .*

*Proof.* We start by formulating the SM2 problem as a MAX-NAE SAT problem. Each job  $j$  corresponds to a variable  $x_j$  in the MAX-NAE SAT instance. An assignment of the variables in the MAX-NAE SAT instance corresponds to an assignment in SM2 as follows: machine 1 is assigned all jobs for which the corresponding variable is set to true, and machine 2 processes all jobs for which the corresponding variable is set to false.



We now construct a set of weighted clauses for each scenario such that the weight of the satisfied clauses for a given assignment is equal to the load of the least loaded machine in the scenario. Hence, maximizing the weight of the satisfied clauses will maximize the weight of the least loaded machine, and it will thus minimize the weight of the machine with the heaviest load, i.e., the makespan.

For a given scenario  $S$  of SM2 with  $r$  jobs, we construct  $2^{r-1}$  clauses of length  $r$  as follows. For each partitioning of  $S$  into two sets  $A$  and  $\bar{A}$ , we create a clause denoted by  $C_S(\{A, \bar{A}\})$ . In clause  $C_S(\{A, \bar{A}\})$ , all variables corresponding to jobs in one set appear negated, all variables corresponding to the other set appear non-negated. Note that  $C_S(\{A, \bar{A}\})$  has the same truth table as  $C_S(\{\bar{A}, A\})$  (namely, a clause is false if and only if all its literals are false, or all its literals are true). Note that this means that if  $A$  is assigned to the first machine and  $\bar{A}$  is assigned to the second machine, then all clauses *except*  $C_S(\{A, \bar{A}\})$  are satisfied.

Denote by  $w_S(\{A, \bar{A}\})$  the weight on the clause  $C_S(\{A, \bar{A}\})$ . To ensure the weight of the satisfied clauses is equal to the weight of the least loaded machine in SM2, we define weights on the clauses to be so that

$$\sum_{B, \bar{B}: B \cup \bar{B} = S, B \cap \bar{B} = \emptyset} w_S(\{B, \bar{B}\}) - w_S(\{A, \bar{A}\}) = \min\{p(A), p(\bar{A})\}.$$

Let  $N = 2^{r-1}$ , i.e.,  $N$  is the number of clauses corresponding to scenario  $S$ . The solution to this system of equations is to set

$$w_S(\{A, \bar{A}\}) = \frac{1}{N-1} \sum_{B, \bar{B}: B \cup \bar{B} = S, B \cap \bar{B} = \emptyset} \min\{p(B), p(\bar{B})\} - \min\{p(A), p(\bar{A})\}.$$

The weights thus defined are not necessarily non-negative: consider a scenario  $S$  that contains  $r = 4$  jobs of unit length. There are four ways of partitioning  $S$  into one set of size one and one set of size three, and there are  $\binom{4}{2}/2 = 3$  ways of partitioning  $S$  into two sets of size two. Therefore  $\sum_{B, \bar{B}} \min\{p(B), p(\bar{B})\} = 10$ , but that means that for a partitioning into sets  $A, \bar{A}$  of size two  $w_S(\{A, \bar{A}\}) = \frac{1}{7}(10) - 2 < 0$ .

To use approximation algorithms for MAX-NAE SAT, we need to make sure that all weights are non-negative. We accomplish this by adding a constant  $K(S)$  to all weights of clauses corresponding to scenario  $S$ , where we set  $-K(S)$  equal to a lower bound on the weights. We derive a lower bound on the weights by noting that (1)  $\frac{1}{N} \sum_{B, \bar{B}} \min\{p(B), p(\bar{B})\}$  is the expected value of the least loaded machine when all jobs are assigned to a machine with probability  $\frac{1}{2}$  independently, hence, by Lemma 1, its value is lower bounded by  $\frac{1}{2} \max_{B, \bar{B}} \min\{p(B), p(\bar{B})\}$ ; and (2) trivially,  $\max_{B, \bar{B}} \min\{p(B), p(\bar{B})\} \leq \frac{1}{2} p(S)$ . Therefore

$$\begin{aligned}
 w_S(\{A, \bar{A}\}) &= \frac{1}{N-1} \sum_{B, \bar{B}: B \cup \bar{B} = S, B \cap \bar{B} = \emptyset} \min\{p(B), p(\bar{B})\} - \min\{p(A), p(\bar{A})\} \\
 &= \frac{N}{N-1} \frac{1}{N} \sum_{B, \bar{B}: B \cup \bar{B} = S, B \cap \bar{B} = \emptyset} \min\{p(B), p(\bar{B})\} - \min\{p(A), p(\bar{A})\} \\
 &\geq \frac{N}{N-1} \frac{1}{2} \max_{B, \bar{B}} \min\{p(B), p(\bar{B})\} - \min\{p(A), p(\bar{A})\} \\
 &\geq \frac{\frac{1}{2}N-(N-1)}{N-1} \max_{B, \bar{B}} \min\{p(B), p(\bar{B})\} \\
 &= -\frac{1}{2} \frac{N-2}{N-1} \max_{B, \bar{B}} \min\{p(B), p(\bar{B})\} \\
 &\geq -\frac{1}{4} \frac{N-2}{N-1} p(S).
 \end{aligned}$$

Thus, we set  $K(S) = \frac{1}{4} \frac{N-2}{N-1} p(S)$ , such that  $\tilde{w}_S(\{A, \bar{A}\}) = w_S(\{A, \bar{A}\}) + K(S) \geq 0$  for all partitionings  $A, \bar{A}$  of  $S$  into two sets.

A solution to the MAX-NAE SAT instance is now mapped to a solution of SM2, by assigning the jobs for which the variable is set to true to machine 1, and scheduling the other jobs on machine 2. We note that the  $w$ -weights of the clauses corresponding to scenario  $S$  were chosen so that the sum of the weights of the clauses that are satisfied is exactly equal to the load on the least loaded machine in scenario  $S$ . Also,  $N - 1$  clauses of scenario  $S$  are satisfied in any solution to the MAX-NAE SAT instance. Therefore the total  $\tilde{w}$ -weight of the clauses for scenario  $S$  that are satisfied in any MAX-NAE SAT solution is equal to the load on the least loaded machine in scenario  $S$  plus an additional  $(N - 1)K(S)$ .

We let  $L = \sum_S p(S)$ , and denote by  $L_{\min}^*$  the sum over all scenarios of the load of the least loaded machine in an optimal solution, and by  $L_{\max}^*$  the sum over all scenarios of the load of the most loaded machine in an optimal solution, so that  $L_{\min}^* + L_{\max}^* = L$ . Note that the additional term  $K(S)$  in the  $\tilde{w}$ -weights of the MAX-NAE SAT solution causes an increase of the objective value with respect to the  $w$ -weights solution by adding an additional  $\sum_S (N - 1)K(S) = \sum_S \frac{1}{4}(N - 2)p(S) = \frac{1}{4}(N - 2)L$  to each solution.

In particular, an optimal solution to the MAX-NAE SAT instance, has objective value  $L_{\min}^* + \frac{1}{4}(N - 2)L$ , and a  $(1 - \gamma)$ -approximation algorithm for the MAX-NAE SAT instance, therefore, has objective value at least  $(1 - \gamma)(L_{\min}^* + \frac{1}{4}(N - 2)L)$ . Let us denote by  $ALG(L_{\min})$  and  $ALG(L_{\max})$  the sum over all scenarios of the least and most loaded machines in the corresponding job assignment. Note that  $ALG(L_{\min}) \geq (1 - \gamma) ((L_{\min}^* + \frac{1}{4}(N - 2)L) - \frac{1}{4}(N - 2)L) = (1 - \gamma)L_{\min}^* - \frac{1}{4}\gamma(N - 2)L$ . Therefore,

$$\begin{aligned}
 ALG(L_{\max}) &= L - ALG(L_{\min}) \leq L - ((1 - \gamma)L_{\min}^* - \frac{1}{4}\gamma(N - 2)L) \\
 &= (1 - \gamma)(L - L_{\min}^*) + \gamma L + \frac{1}{4}\gamma(N - 2)L \\
 &= (1 - \gamma)L_{\max}^* + \frac{1}{4}\gamma(N + 2)L.
 \end{aligned}$$

Noting that  $L \leq 2L_{\max}^*$  gives  $ALG(L_{\max}) \leq (1 - \gamma)L_{\max}^* + \frac{1}{2}\gamma(N + 2)L_{\max}^* = (1 + \frac{1}{2}\gamma N)L_{\max}^*$  which proves the theorem, since  $N = 2^{r-1}$ .  $\square$

For  $r = 3$ , Zwick [18] gives a 0.90871-approximation for MAX-NAE 3-SAT. By the previous lemma, this gives a 1.18258-approximation for SM2 with scenarios of length at most three. For  $r = 4$ , Karloff et al. [14] give a  $\frac{7}{8}$ -approximation for MAX-NAE 4-SAT. By our lemma, this implies a  $\frac{3}{2}$ -approximation for SM2 with scenarios of size 4. Note that this matches the guarantee we proved for the algorithm that randomly assigns each job to one of the two machines. For general  $r$ , the best approximation factor known for MAX-NAE SAT is 0.74996 due to Zhang, et al. [17], and the implied approximation guarantees for our problem are worse than the guarantee for the random assignment.

If every scenario has exactly two jobs, then we can obtain a better approximation guarantee by reducing SM2 to MAX CUT as follows: we create a vertex for every job, and add an edge between  $i$  and  $j$  of weight  $\min\{p_i, p_j\}$  for every scenario that contains jobs  $i$  and  $j$ . For any cut, the weight of the edges crossing the cut is then exactly the sum over all scenarios of the load of the least loaded machine. Since the makespan for a scenario  $S$  is  $p(S)$  minus the load of the least loaded machine, maximizing the load of the least loaded machine, summed over all scenarios, is equivalent to minimizing the sum of the makespans.

If every scenario has at most three jobs, we can also reduce SM2 to MAX CUT, but the reduction, given in the full version [8], is slightly more involved.

**Theorem 8.** *There exists a  $(1 + \gamma)$ -approximation algorithm for the SM2 problem with scenarios containing at most three jobs, where  $1 - \gamma$  is equal to the approximation ratio for MAX CUT.*

The 0.87856-approximation for MAX CUT of Goemans et al. [10] gives us the following corollary.

**Corollary 1.** *There exists a 1.12144-approximation algorithm for the SM2 problem with scenarios containing at most three jobs.*

## 4 Epilogue

This paper presents some first results on a basic scheduling problem under a set of scenarios. The objective is to find a single solution that is applied to all the scenarios specified. We studied this problem for scheduling with two different objectives: minimizing the maximum objective value over all scenarios, the Min-Max version, and minimizing the sum of the objective values of all scenarios, the MinSum version.

To the best of our knowledge, combinatorial optimization problems under a set of fully explicitly specified scenarios has hardly been studied in the literature. Apart from posing theoretically interesting questions as we hope to have shown with this paper, it enhances our ability to model decisions problems where a learning aspect for performing jobs prohibits that job assignments can be adjusted on a day-by-day basis, but merely require a fixed assignment whose quality then necessarily differs over the various instances.

In relation to the MinMax version of the problem, we also like to mention a version of combinatorial optimization which has become known under the

name universal optimization. E.g., [7] study a universal scheduling problem. In such a problem, the scenarios are not explicitly specified, but can be seen to be chosen by an adversary. The quality of an algorithm is then measured by comparing its solution to the optimal solution when the adversarial choices are known beforehand.

For future research, anyone can choose her or his favorite combinatorial optimization problem and study its multiple-scenario version.

We finish with the some questions emerging from our multiple-scenario scheduling problem. The result in [2] suggests a  $3/2$ -approximation for MM2 with 4 jobs per scenario and unitary jobs. Can this be extended to any job sizes? For the SM2 version the question is to close the gap between the  $3/2$ -approximate randomized algorithm for the general case and the 1.0404 lower bound under the Unique Games Conjecture. It would also be interesting to find out if our randomized algorithm can be derandomized.

## References

1. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Protasi, M.: Complexity and approximation: Combinatorial optimization problems and their approximability properties. Springer (1999)
2. Austrin, P., Guruswami, V., Håstad, J.:  $(2 + \epsilon)$ -SAT is NP-hard. Electronic Colloquium on Computational Complexity, TR13-159 (2013)
3. Ben-Tal, A., Nemirovski, A.: Robust optimization—methodology and applications. *Mathematical Programming* 92(3), 453–480 (2002)
4. Bouyahia, Z., Bellalouna, M., Ghedira, K.: Load balancing a priori strategy for the probabilistic weighted flowtime problem. *Comput. Ind. Eng.* 64(1), 1–10 (2013)
5. Bouyahia, Z., Bellalouna, M., Jaillet, P., Ghedira, K.: A priori parallel machines scheduling. *Comput. Ind. Eng.* 58(3), 488–500 (2010), Supply, Production and Distribution Systems
6. Chekuri, C., Khanna, S.: On multidimensional packing problems. *SIAM Journal on Computing* 33(4), 837–851 (2004)
7. Epstein, L., Levin, A., Marchetti-Spaccamela, A., Megow, N., Mestre, J., Skutella, M., Stougie, L.: Universal sequencing on an unreliable machine. *SIAM Journal on Computing* 41, 565–586 (2012)
8. Feuerstein, E., Marchetti-Spaccamela, A., Schalekamp, F., Sitters, R., van der Ster, S., Stougie, L., van Zuylen, A.: Scheduling over scenarios on two machines. CoRR, abs/1404.4766 (2014)
9. Garey, M.R., Johnson, D.S.: Computers and Intractability; A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., New York (1990)
10. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM* 42(6), 1115–1145 (1995)
11. Gupta, A., Pál, M., Ravi, R., Sinha, A.: Sampling and cost-sharing: Approximation algorithms for stochastic optimization problems. *SIAM Journal on Computing* 40(5), 1361–1401 (2011)
12. Håstad, J.: Some optimal inapproximability results. *J. ACM* 48(4), 798–859 (2001)
13. Jaillet, P.: A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Oper. Res.* 36(6), 929–936 (1988)

14. Karloff, H.J., Zwick, U.: A  $7/8$ -approximation algorithm for MAX 3SAT? In: FOCS, pp. 406–415. IEEE Computer Society (1997)
15. Khot, S.: On the power of unique 2-prover 1-round games. In: Proceedings of 34th Annual ACM Symposium on Theory of Computing, pp. 767–775 (2002)
16. Khot, S., Kindler, G., Mossel, E., O’Donnell, R.: Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.* 37(1), 319–357 (2007)
17. Zhang, J., Ye, Y., Han, Q.: Improved approximations for max set splitting and max NAE SAT. *Discrete Applied Mathematics* 142(1-3), 133–149 (2004)
18. Zwick, U.: Outward rotations: A tool for rounding solutions of semidefinite programming relaxations, with applications to MAX CUT and other problems. In: STOC, pp. 679–687 (1999)