

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

A General Framework for Approximating Min Sum Ordering Problems

Felix Happach

Department of Mathematics and TUM School of Management, Technische Universität München, 80333 München, Germany,
felix.happach@tum.de

Lisa Hellerstein

Department of Computer Science and Engineering, New York University Tandon School of Engineering, New York, NY 11201,
lisa.hellerstein@nyu.edu

Thomas Lidbetter

Department of Management Science and Information Systems, Rutgers Business School, Newark, NJ 07102,
tlidbetter@business.rutgers.edu

We consider a large family of problems in which an ordering (or, more precisely, a chain of subsets) of a finite set must be chosen to minimize some weighted sum of costs. This family includes variations of Min Sum Set Cover (MSSC), several scheduling and search problems, and problems in Boolean function evaluation. We define a new problem, called the Min Sum Ordering Problem (MSOP) which generalizes all these problems using a cost and a weight function defined on subsets of a finite set. Assuming a polynomial time α -approximation algorithm for the problem of finding a subset whose ratio of weight to cost is maximal, we show that under very minimal assumptions, there is a polynomial time 4α -approximation algorithm for MSOP. This approximation result generalizes a proof technique used for several distinct problems in the literature. We apply this to obtain a number of new approximation results.

Key words : scheduling, search theory, Boolean function evaluation, Min Sum Set Cover, approximation algorithms

1. Introduction

Many optimization problems require finding a minimum cost (feasible) subset of the elements of a finite set, according to some cost function. For example, in the Set Cover Problem, the given finite set is a collection of subsets of a finite set of “ground” elements X , and the objective is to choose a sub-collection of minimum cardinality whose union is X . In the Minimum Spanning Tree Problem, for a given graph with edge lengths, the

objective is to find a subtree of minimum total length that contains all the vertices of the graph.

This work is concerned with ordering problems in which the objective is not to find a subset of minimum total cost, as in the examples described above, but is instead to find a sequence that minimizes an incremental sum of costs, or equivalently, a weighted average of costs. For example, in many scheduling problems, the objective is to minimize some weighted sum of completion times of a set of jobs. The order in which the jobs are processed may be constrained by some form of precedence constraints. In search problems, one might wish to minimize the expected time or cost incurred in searching for a target or targets that are hidden according to a known probability distribution, and the set of feasible searches may be restricted by some network structure. These problems arise in search and rescue as well as military search operations and can be interpreted as “min sum” versions of such problems as the spanning tree problem. A “min sum” version of Set Cover called *Min Sum Set Cover* was introduced by Feige et al. (2002, 2004). Sequential testing problems also come under this framework: a set of tests (for example, medical tests, database queries or quality tests of computer chip components) must be performed in some order to minimize an expected cost (of forming a diagnosis, of determining whether the query is satisfied or of checking whether the component meets certain quality standards).

We unify problems of this type by introducing a new, very general problem formulation, which we call the *Min Sum Ordering Problem* or *MSOP*. Let V be a finite set of cardinality n and let $f, g : F \rightarrow \mathbb{R}$ be a *cost function* and a *weight function*, respectively, defined on some family of subsets $F \subseteq 2^V$ that contains \emptyset and V (where we use the symbol \subseteq to denote “is a subset of or equal to” and \subset to denote strict inclusion). Further, f and g are non-decreasing with respect to set inclusion and $f(\emptyset) = g(\emptyset) = 0$. In our applications, the set F is typically implicitly defined by the problem setting rather than being part of the input. We assume that f and g are given by value oracles. We define an F -chain to be a sequence of subsets $S = \{S_j\}_{j=0}^k$ for some k such that $\emptyset = S_0 \subset S_1 \subset \dots \subset S_k = V$ and $S_j \in F$ for each j . When there is no ambiguity, we will simply refer to an F -chain as a *chain*.

Then the MSOP is to minimize

$$C_{f,g}(S) \equiv \sum_{j=1}^k f(S_j)(g(S_j) - g(S_{j-1})), \quad (1)$$

over all chains $S = \{S_j\}_{j=0}^k$.

If S minimizes $C_{f,g}(S)$, we say it is *optimal* and if $C_{f,g}(S)$ is at most a factor $\alpha \geq 1$ times the optimal value of the objective, we say S is an α -approximation.

If F contains *all* subsets of V , then the problem is equivalent to minimizing over all permutations of V (see Lemma 1). Here the subsets in a chain correspond to the elements picked “so far” by the permutation. By setting the problem up in the more general way, in terms of maximizing over chains rather than permutations, we ensure that the model is general enough to incorporate the intricacies of precedence constraints or restrictions due to a network structure, for example.

The applications we are interested in include problems of minimizing over a given *subset* of the set of all permutations of V . To be more precise, let $\sigma : V \rightarrow V$ be a permutation of V , and for $j = 1, \dots, n$, let $\$$ be the union of the first j elements of V under the permutation σ . Then for a set Σ of permutations of V and given functions f and g , we wish to solve

$$\min_{\sigma \in \Sigma} \sum_{j=1}^{|X|} f(S_j^\sigma)(g(S_j^\sigma) - g(S_{j-1}^\sigma)). \quad (2)$$

We refer to this problem as the *Min Sum Permutation Problem* or *MSPP*. Neither one of MSOP nor MSPP can be reduced to the other, but we will show in Section 2 that as long as Σ satisfies some technical condition, MSPP can be regarded as a special case of MSOP. Indeed, let F^Σ consist of all *initial* sets of all $\sigma \in \Sigma$ (that is, $S \in F^\Sigma$ if and only if $S = S_j^\sigma$ for some j and some permutation $\sigma \in \Sigma$). Suppose we have an α -approximate solution S to MSOP with $F = F^\Sigma$. Then any permutation σ such that every subset in S is some initial set of σ will be an α -approximate solution to MSPP. This observation is easily verified and we postpone its justification until Section 2.

We choose to write this paper in terms of MSOP rather than MSPP for two main reasons. Firstly, the setting allows us to prove the strongest form of our main results. Secondly, working with chains rather than permutations allows the greedy algorithm we will analyze to consider the “bigger picture” by recursively adding *subsets* of elements rather than adding elements one by one.

The idea of minimizing over chains is novel, but MSPP has been studied previously in some other special cases. In particular, Pisaruk (1992) considered the problem in the case that f is submodular and g is supermodular, giving a 2-approximation algorithm. (The function f is submodular if and only if $f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$ for all S, T ; the

function g is supermodular if $f(S \cup T) + f(S \cap T) \geq f(S) + f(T)$ for all S, T . This special case was also considered more recently in Fokkink et al. (2019), where the problem was introduced as the *submodular search problem*, and previous to that, Fokkink et al. (2017) considered the case of f submodular and g modular.

If g is the cardinality function $g(S) = |S|$ and the sets S_i increase by one element in each step, then the sum in (1) reduces to $\sum_{j=1}^k f(S_j)$. This special case of MSPP was considered in Iwata et al. (2012), for various classes of functions f . In particular, a 4-approximation algorithm was obtained in the case that f is supermodular. We discuss this in more detail in Subsection 3.2, in particular in reference to MSSC and its generalizations. A more general 4-approximation had already been proved by Streeter and Golovin (2008) in the case of f supermodular and g modular.

MSOP is more general than the problems of Iwata et al. (2012) and Pisaruk (1992) in two ways. Firstly, we make weaker assumptions about the form of the cost and weight functions and our main result requires only that the cost function is subadditive. A cost function f is subadditive if and only if $f(S \cup T) \leq f(S) + f(T)$ for all disjoint sets $S, T \in F^1$. Subadditivity is a more general concept than submodularity. Secondly, the aforementioned works take the approach of minimizing over all permutations of V , in contrast to our approach of minimizing over chains.

Since MSOP generalizes several NP-hard problems, MSOP is NP-hard itself, so we consider approximation algorithms for the problem. An important concept in our analysis is that of the *density* $\rho(S)$ of a subset S of V , defined by $\rho(S) = g(S)/f(S)$ (for $f(S) \neq 0$). We will define a simple greedy algorithm for MSOP, which we now briefly describe (see Section 2 for a more precise description). The algorithm constructs a chain by recursively choosing the $(j+1)$ th subset S_{j+1} in the chain in such a way as to maximize the marginal density $(g(S_{j+1}) - g(S_j))/(f(S_{j+1}) - f(S_j))$. We refer to this maximization problem as the *maximum density problem*. If the maximum density problem cannot be solved in polynomial time, but we can approximate it in polynomial time within a factor $\alpha \geq 1$ then we call a chain produced by such an approximation an α -greedy chain. We will prove the following in Section 2.

¹ Some definitions of subadditive set functions in the literature require that $f(S \cup T) \leq f(S) + f(T)$ for non-disjoint S and T as well. When f is non-decreasing with respect to set inclusion, as it is in our case, this is an equivalent definition.

Theorem 1. Suppose f is subadditive and F is closed under union. Then for any $\alpha \geq 1$, an α -greedy chain is a 4α -approximation for an optimal chain for MSOP.

Later, in Section 6, we will derive an analogous result to Theorem 1 for a “backward” greedy algorithm, using the concept of the *dual problem*.

The proof of Theorem 1 is inspired by the elegant proof in Feige et al. (2004) of the 4-approximation algorithm for the problem *Min Sum Set Cover* (MSSC). This is the problem of ordering a ground set V to minimize the sum of “covering times” of a given collection of subsets of V , where the covering time of a subset is the earliest position in the ordering of any element of that subset. The proof uses the idea of representing the cost of the ordering produced by the greedy algorithm and that of an optimal ordering by two histograms, and showing that when the first histogram is shrunk by a factor of two in the horizontal and vertical directions, it fits in the second histogram. The proof idea is generalized in Streeter and Golovin (2008), who proved a 4-approximation result for a class of problems that includes some special cases of MSOP, including MSSC. A different generalization of MSSC is given in Iwata et al. (2012) to prove the 4-approximation for one case of the *Minimum Linear Ordering Problem*. More recently, a similar proof was used in Hermans et al. (2019) to establish an 8-approximation for the *expanding search* problem, and in Happach and Schulz (2020a) to obtain a 4-approximation for *bipartite OR-scheduling*.

While the last three works cited all use a similar proof method, the proof is somewhat different in each case and none of these results directly implies another. The similarity of the proofs strongly suggests that some deeper result is behind all of these problems. We confirm here that this is indeed the case by showing that MSOP generalizes each of them and Theorem 1 generalizes their respective approximation results.

In Section 2 we prove Theorem 1, using a variation of the proof originally devised by Feige et al. (2004). The main difference from the original proof stems from the fact that our algorithm does not greedily pick elements of V one by one, but rather greedily picks subsets in the chain. Also, we do not optimize over permutations but over chains.

In Section 3, we will show that Theorem 1 can be used to recover a number of known results relating to search theory and variants of MSSC. We go on to apply our results to new problems. In particular, we consider scheduling problems with *OR-precedence constraints* in Section 4, where the set of jobs to be processed is represented by the vertices of a directed acyclic graph (DAG), and a job can only be processed after at least one of its

predecessors in the DAG has been processed. We use Theorem 1 to show that there is a polynomial time 4-approximation algorithm for the problem of minimizing the sum of the weighted completion times of a set of jobs that must be scheduled so as to respect some OR-precedence constraints given by a DAG that is in the form of an *inforest* or, more generally, a *multitree* (where inforests and multitrees will be defined in Section 4). We also give a 4-approximation algorithm for a version of MSSC with OR-precedence constraints in the form of an inforest. Finally, in Section 5 we give an 8-approximation algorithm for minimizing the expected cost of non-adaptively evaluating a Boolean read-once formula (AND/OR tree), assuming independent tests. In Section 6 we introduce the dual problem, which leads to further approximation results, and in Section 7 we indicate directions for future work.

2. Approximating MSOP

In this section, we first prove our main result in Subsection 2.1. In Subsection 2.2 we then justify the observation made in the introduction that MSPP is really a special case of MSOP.

2.1. Proof of Main Result

For a set $A \in F$, we write f_A for the function on F given by $f_A(S) = f(S) - f(A)$, and similarly for g_A . For $f_A(S) \neq 0$, let $\rho_A(S) = g_A(S)/f_A(S)$ be the *marginal density* of S (with respect to A); if $f_A(S) = 0$ we set $\rho_A(S) = \infty$. If $A = \emptyset$, we drop the subscript from ρ and simply refer to $\rho(S)$ as the density of S .

We consider a greedy algorithm for MSOP. For $\alpha \geq 1$, we call an F -chain $S = (S_j)_{j=0}^k$ an α -greedy chain if

$$\rho_{S_j}(S_{j+1}) \geq \frac{1}{\alpha} \max_{\{T \in F: S_j \subseteq T\}} \rho_{S_j}(T),$$

for all $j = 0, \dots, k-1$. If $\alpha = 1$, an α -greedy chain is simply one for which S has maximum marginal density with respect to S_j for each $j = 0, \dots, k-1$.

We now prove Theorem 1, which generalizes both the result and the proof in Feige et al. (2004).

Proof of Theorem 1. Let $T = (T_j)_{j=0}^k$ be an optimal chain and let $S = (S_i)_{i=0}^k$ be an α -greedy chain. We first construct a histogram with k columns, the area under which is equal to $C_{f,g}(T)$. The base of the j th column of the histogram is the interval from $g(T_{j-1})$ to

$g(T_j)$ and its height is $f(T_j)$. Thus, the total area under the histogram is equal to $C_{f,g}(T)$. The histogram is depicted in the top left of Figure 1(a).

Next, we construct a second histogram with k columns, the area under which is equal to $C_{f,g}(S)$. Let $\rho_i = \rho_{S_{i-1}}(S_i)$ and let $\phi_i = \rho_i^{-1}(g(V) - g(S_{i-1}))$ (if $\rho_i = \infty$, set $\rho_i^{-1} = 0$). The base of the i th column of this histogram is the interval from $g(S_{i-1})$ to $g(S_i)$ and its height is ϕ_i . Thus the total area A under this histogram is

$$\begin{aligned} A &= \sum_{i=1}^k \phi_i (g(S_i) - g(S_{i-1})) \\ &= \sum_{i=1}^k (g(V) - g(S_{i-1})) (f(S_i) - f(S_{i-1})) \\ &= g(V) \sum_{i=1}^k (f(S_i) - f(S_{i-1})) - \sum_{i=1}^k g(S_{i-1}) (f(S_i) - f(S_{i-1})) \end{aligned} \quad (3)$$

The first sum on the right-hand side above is telescopic and equal to $f(V) - f(\emptyset) = f(V)$. Rearranging the second sum, we obtain

$$\begin{aligned} A &= g(V)f(V) - g(S_{k-1})f(V) + \sum_{i=1}^{k-1} f(S_i)(g(S_i) - g(S_{i-1})) \\ &= \sum_{i=1}^k f(S_i)(g(S_i) - g(S_{i-1})) \\ &= C_{f,g}(S). \end{aligned}$$

The second histogram is depicted in the top right of Figure 1(a). Note that the heights of the columns in the first histogram, from left to right, are non-decreasing.

Now shrink the second histogram by a factor of 2α in the vertical direction, and a factor of 2 in the horizontal direction, and move it to the right so it is flush with the right end $g(V)$, as depicted in Figure 1(b). This results in point (x, y) being mapped to $(\frac{g(V)+x}{2}, \frac{y}{2\alpha})$. The distance of this latter point from the right end is $(g(V) - x)/2$.

We now show that the shrunken (and shifted) histogram is contained in the first histogram, from which it follows that $C_{f,g}(S) \leq 4\alpha C_{f,g}(T)$, proving the theorem. To show that the shrunken histogram is contained in the first histogram, it is sufficient to show that if (a, b) is the top left point of some column i in the shrunken histogram, and (c, d) is the top right point of some column j in the first histogram, then $d < b$ implies that $c < a$. Here $(a, b) = \frac{g(V)+g(S_{i-1})}{2}, \frac{\phi_i}{2\alpha}$ and $(c, d) = (g(T_j), f(T_j))$.

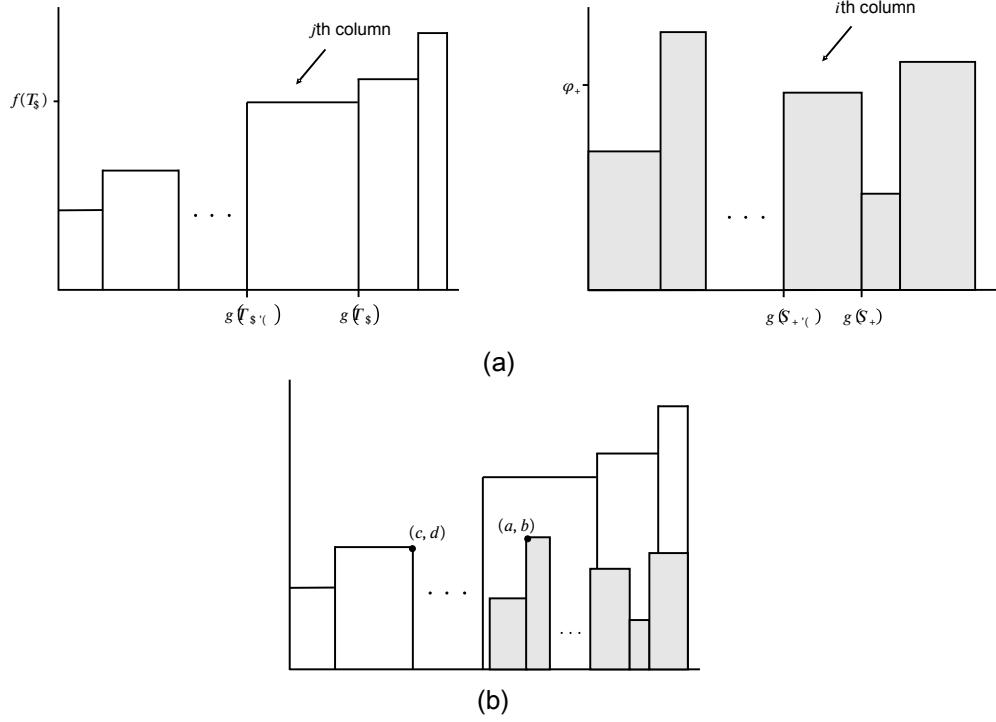


Figure 1 (a) Two histograms with totals $C_g(T)$ (left) and $C_g(S)$ (right); (b) the shrunken version of the second histogram in the first histogram.

So assume $d < b$, or equivalently

$$f(T_j) < \frac{\rho_i^{-1} (g(V) - g(S_{i-1}))}{2\alpha}. \quad (4)$$

Let a^0 be the distance of (a, b) from the right boundary, that is, $a^0 = (g(V) - g(S_{i-1}))/2$ and let c^0 be the distance of (c, d) from the right boundary, that is, $c^0 = g(V) - g(T_j)$. We want to show that $d < b$ implies $c < a$, or equivalently that $a^0 < c^0$. So we want to show that

$$\frac{g(V) - g(S_{i-1})}{2} < g(V) - g(T_j). \quad (5)$$

Since F is closed under union, $S_{i-1} \cup T_j \in F$. We will use the fact that, because S is α -greedy,

$$\rho_i = \rho_{S_{i-1}}(S_i) \geq \frac{1}{\alpha} \rho_{S_{i-1}}(S_{i-1} \cup T_j) = \frac{g(S_{i-1} \cup T_j) - g(S_{i-1})}{\alpha(f(S_{i-1} \cup T_j) - f(S_{i-1}))}. \quad (6)$$

Because f is subadditive and non-decreasing,

$$f(S_{i-1} \cup T_j) - f(S_{i-1}) \leq f(S_{i-1} \cup T_j) - f(S_{i-1} \setminus T_j) \leq f(T_j).$$

Combining that fact with (6) yields

$$\begin{aligned} af(T_j) &\geq \rho_i^{-1}(g(S_{i-1} \cup T_j) - g(S_{i-1})) \\ &= \rho_i^{-1}((g(V) - g(S_{i-1})) - (g(V) - g(S_{i-1} \cup T_j))) \end{aligned}$$

Using our assumption in (4), we thus get

$$\begin{aligned} \frac{\rho_i^{-1}(g(V) - g(S_{i-1}))}{2} &> \rho_i^{-1}((g(V) - g(S_{i-1})) - (g(V) - g(S_{i-1} \cup T_j))) \\ &\geq \rho_i^{-1}((g(V) - g(S_{i-1})) - (g(V) - g(T_j))), \end{aligned}$$

where the second inequality follows from the fact g is non-decreasing. Rearranging gives (5).

Observe that if $F = 2^V$ and f is supermodular and g is submodular, then for $S_j \subseteq T$,

$$\rho_{S_j}(T) \leq \frac{\max_{v \in T \setminus S_j} g(S_j \cup \{v\}) - g(S_j)}{\max_{v \in T \setminus S_j} f(S_j \cup \{v\}) - f(S_j)} \leq \max_{v \in T \setminus S_j} \rho_{S_j}(S_j \cup \{v\}).$$

Hence, a 1-greedy chain can be obtained in polynomial time by adding singletons one-by-one. Suppose f is not just supermodular but also modular. Then f is also subadditive so, by Theorem 1, there is a polynomial time 4-approximation algorithm. We summarize this observation below.

Corollary 1. *Suppose $F = 2^V$. If f is modular and g is submodular then a 1-greedy chain can be constructed in polynomial time and there exists a polynomial time 4-approximation algorithm for MSOP.*

As discussed in Subsection 3.2, the problem MSSC and, more generally, pipelined set cover, are special cases of MSOP where g is submodular and f is modular, so the 4-approximation algorithms for these problems follow from Corollary 1.

2.2. Minimizing Over Permutations

We now turn to the problem MSPP, given in (2), where we wish to minimize a weighted sum over a set of permutations Σ . We will show that provided Σ satisfies a certain technical condition, solving MSPP for Σ is equivalent to solving the corresponding MSOP problem (with the same f and g) for $F = F^\Sigma$.

Suppose F is some family of subsets of V , and suppose $S \equiv (S_j)_{j=0}^k$ is an F -chain. Then if $0 = j_0 \leq j_1 \leq \dots \leq j_k = k$, we say $S^0 \equiv (S_{j_i})_{i=0}^k$ is a *subchain* of S .

Lemma 1. Suppose $S \equiv (S_i)_{i=0}^k$ is a subchain of the F -chain $S^0 \equiv (S_j)_{j=0}^k$. Then

(i) $C_{f,g}(S) \geq C_{f,g}(S^0)$ and

(ii) if S approximates MSOP by a factor of $\alpha \geq 1$ then so does S^0 .

Proof. We perform the following calculation.

$$\begin{aligned}
 C_{f,g}(S) &\equiv \sum_{i=1}^k f(S_{j_i})(g(S_{j_i}) - g(S_{j_{i-1}})) \\
 &= \sum_{i=1}^k f(S_{j_i}) \sum_{j=j_{i-1}+1}^{j_i} (g(S_j) - g(S_{j-1})) \\
 &\geq \sum_{i=1}^k \sum_{j=j_{i-1}+1}^{j_i} f(S_j)(g(S_j) - g(S_{j-1})) \\
 &= \sum_{j=1}^k f(S_j)(g(S_j) - g(S_{j-1})) \equiv C_{f,g}(S^0),
 \end{aligned}$$

where the inequality above follows from the monotonicity of f and g . Part (ii) of the lemma follows immediately.

Suppose now that Σ is a set of permutations of V . Recall that F^Σ consists of all subsets of V that are initial sets of some permutation $\sigma \in \Sigma$. If S is an F^Σ -chain and $\sigma \in \Sigma$ is a permutation, such that each element of S is an initial set of σ , then we say σ is *consistent* with S . If every F^Σ -chain is consistent with some permutation in Σ then we say Σ is *well-founded*.

Lemma 2. Suppose Σ is a set of permutations of V and that Σ is well-founded. If there exists a polynomial time α -approximation algorithm for MSOP with $F = F^\Sigma$ for some $\alpha \geq 1$ then there exists a polynomial time α -approximation algorithm for MSPP with Σ .

Proof. This follows immediately from Lemma 1, part (ii). Indeed, first note that every F^Σ -chain is consistent with a permutation whose objective value is no greater and every permutation is consistent with an F^Σ -chain with the same objective value. It follows that the objective value of any optimal solution to MSOP is equal to the objective value of any optimal solution to MSPP.

Now suppose that S is an α -approximate F^Σ -chain and that σ is consistent with S . Let S^0 be the chain consisting of all the initial sets of σ . Then S is a subchain of S^0 , so S^0 is an α -approximation for MSOP. Equivalently, σ is an α -approximation for MSPP.

Note that Lemma 2 trivially holds in the opposite direction. Indeed, if there exists a polynomial time α -approximation for MSPP, then this also provides a polynomial time α -approximation for MSOP with $F = F^\Sigma$, since every permutation σ defines a feasible F^Σ -chain with the same objective function value.

It is easy to think of examples of Σ that are not well-founded. For example, if $V = \{1, 2, 3\}$ and Σ contains only the permutations $(1, 2, 3)$ and $(3, 1, 2)$, then the F^Σ -chain $\{\{1\}, \{1, 3\}, \{1, 2, 3\}\}$ is not consistent with either of the two permutations, so Σ is not well-founded.

However, for all the examples we consider in this paper, the set of permutations is well-founded. This is easy to check by using the following sufficient condition.

For two permutations σ and τ of V , let $\pi_j(\sigma, \tau)$ be the permutation that follows σ for the first j elements then chooses the remaining elements of V in the order specified by τ , for each $1 \leq j \leq n$. For example, if $V = \{1, 2, 3, 4, 5\}$, and σ and τ are given by $(3, 1, 5, 2, 4)$ and $(4, 5, 1, 2, 3)$, respectively then₂(σ, τ) is given by $(3, 1, 4, 5, 2)$ and₃(σ, τ) is given by $(3, 1, 5, 4, 2)$.

If Σ is a set of permutations for which $\pi_j(\sigma, \tau) \in \Sigma$ for any $\sigma, \tau \in \Sigma$ and $1 \leq j \leq n$, then we say Σ is *closed*.

Lemma 3. *Let Σ be a set of permutations of V . If Σ is closed then it is well-founded and F^Σ is closed under union.*

Proof. Suppose Σ is closed. Let $S = (S_j)_{j=1}^k$ be a F^Σ -chain. We will show that there is some permutation contained in Σ that is consistent with S . Let σ_j be a permutation in Σ that is consistent with S_j for $j = 1, \dots, k$. We set $\tau_1 = \sigma_1$ and for $j = 2, 3, \dots, k$, we recursively define $\tau_j = \pi_{|S_{j-1}|}(\tau_{j-1}, \sigma_j)$, which is contained in Σ , by induction on j and since Σ is closed. Also, S_1 is an initial set of τ_1 , and, by definition of $\pi_{|S_{j-1}|}(\tau_{j-1}, \sigma_j)$ and by induction on j , each of S_1, \dots, S_j are initial sets of τ_j for $j \geq 2$. Therefore, τ_k is consistent with S , so Σ is well-founded.

To see that F^Σ is closed under union, let S and T be elements of F^Σ . Then they are initial sets of some permutations σ and τ in Σ , so $S \cup T$ is an initial set of $\pi_{|S \cup T|}(\sigma, \tau)$, which lies in Σ , since Σ is closed. Therefore $S \cup T \in F^\Sigma$.

We observe that for the expanding search problem, if we take Σ to be the set of expanding searches, then it is easy to check that Σ is closed and therefore well-founded, by Lemma 3.

Furthermore, F^Σ is closed under union. Similarly, for both AND-precedence constraints and OR-precedence constraints, the set of feasible orderings is closed and therefore well founded. If Σ consists of all permutations of V , as in Boolean function evaluation, then Σ is trivially well-founded. It follows from Lemma 2 that for these problems, if we can find a solution (or approximate solution) S to MSOP with $F = F^\Sigma$, then we can recover a solution (or approximate solution) to the original problem by taking any permutation that is consistent with S . Since in each case F^Σ is closed under union, we only need f to be subadditive to apply Theorem 1.

3. Special cases of MSOP

In this section we describe some special cases of MSOP, including those for which our results imply the existence of approximation algorithms that are already known in the literature. We will begin in Subsection 3.1 by discussing problems in search theory and scheduling, including the recent 8-approximation result of Hermans et al. (2019) for expanding search. Then, in Subsection 3.2, we will describe how a number of approximation results for Min Sum Set Cover and its generalizations follow from our results.

3.1. Search Theory and AND-Scheduling

The expanding search problem was introduced in Alpern and Lidbetter (2013), and independently in Averbakh and Pereira (2012) under different nomenclature. A connected graph $G = (V, E)$ is given, and each edge $e \in E$ has a cost c_e . A target is hidden on one of the vertices of the graph according to a known probability distribution, so that the probability it is on vertex $v \in V$ is p_v . An *expanding search*, starting at some distinguished root r , is a sequence of edges e_1, e_2, \dots, e_i chosen so that r is incident to e_1 and every edge e_i ($i > 1$) is incident to some previously chosen edge. For a given expanding search, the *expected search cost* of the target is the expected value of the sum of the costs of all the edges chosen up to and including the first edge that contains the target. The problem is to find an expanding search with minimal expected search cost. The problem was shown to be NP-hard in Averbakh and Pereira (2012), and Hermans et al. (2019) recently gave an 8-approximation algorithm.

To express the problem in the form of MSOP, let Σ be the set of expanding searches and take $F = F^\Sigma$. Then F is closed under union, since its elements consist of connected subgraphs of G containing r . For $S \in F$, let $f(S) = \sum_{e \in S} c_e$ and let $g(S)$ be the sum of

p_v over all vertices v contained in some edge of S . Then f is modular, and Theorem 1 along with the results of Subsection 2.2 on MSPP imply that the greedy algorithm is 4α -approximate, where α is the approximation ratio of the maximum density problem, generalizing the analogous Theorem 2 of Hermans et al. (2019).

Alpern and Lidbetter (2013) gave a solution to the expanding search problem in the case that the graph is a tree. In this case, the problem is equivalent to a special case of the single machine, precedence constrained scheduling problem usually denoted $1|prec| \sum P w_j C_j$, of minimizing the sum of the weighted completion times of a set of jobs. A partial order is given on the jobs, and a job becomes available for processing only after all of its predecessors have been processed. We refer to this type of precedence constraints and precedence constrained scheduling as *AND-precedence constraints* and *AND-scheduling*. The jobs have weights w_j and processing times p , and for a given ordering, the completion time C_j of a job j is the sum of its processing time and all the processing times of the jobs preceding it. The problem is to find a feasible ordering that minimizes the weighted sum $\sum P w_j C_j$ of the completion times. Comparing the weights and the processing times to the probabilities and the edge costs in the expanding search problem, it is easy to see that if the partial order has a tree-like structure, then the scheduling problem and the search problem are equivalent, as pointed out in Fokkink et al. (2019).

A polynomial time algorithm for the scheduling problem $1|prec| \sum P w_j C_j$ on trees was given by Horn (1972). Sidney (1975) proved that any optimal schedule (for general precedence constraints) must respect what is now known as a *Sidney decomposition* obtained by recursively taking subsets of jobs of maximum density. Lawler (1978) gave a polynomial time algorithm for the problem on series-parallel graphs, which was generalized to two-dimensional partial orders in a series of papers of Correa and Schulz (2005) and Ambühl and Mastrolilli (2009). Chekuri and Motwani (1999) and Margot et al. (2003) independently showed that any schedule consistent with a Sidney decomposition is a 2-approximation. Earlier 2-approximation algorithms were derived by Schulz (1996) and Chudak and Hochbaum (1999). Correa and Schulz (2005) showed that all known 2-approximations are consistent with a Sidney decomposition. Sidney's decomposition theorem and the resulting 2-approximation algorithm was generalized to the case of MSPP with f submodular and g supermodular in Fokkink et al. (2019), where further applications to scheduling and search problems were given.

3.2. Min Sum Set Cover and its Generalizations

Min Sum Set Cover was first introduced by Feige et al. (2002). An instance of MSSC consists of a finite ground set V and a collection of subsets (or *hyperedges*) E of V . For a given linear ordering (or permutation) $\pi : V \rightarrow [n] := \{1, \dots, n\}$ of the elements of V , the *covering time* of set $e \in E$ is the first point in time that an element contained in e appears in the linear ordering, i.e., $\pi(e) := \min\{\pi(v) \mid v \in e\}$. The objective is to find a linear ordering π that minimizes the total sum of covering times, $\sum_{e \in E} \pi(e)$.²

Iwata et al. (2012) introduced a generalization of MSSC called the Minimum Linear Ordering Problem, which can be regarded as the special case of MSPP where g is the cardinality function. (In fact, Iwata et al. (2012) perform the summation in the opposite order from (1), but of course the objective is the same.) We give here a slightly different reduction of MSSC to MSPP. Taking F to be 2^V , for a subset $S \in F$, we define $g(S)$ to be the number of hyperedges that contain some element of S and $f(S)$ to be the cardinality of S . Then the total sum of covering times is given by (1). The dual of this problem (see Section 6) corresponds to the reduction of Iwata et al. (2012).

MSSC is closely related to Minimum Color Sum (MCS), which was introduced by Kubicka and Schwenk (1989), and can be shown to be a special case of MSSC (though the reduction is not of polynomial size – see Feige et al. (2002)). MSSC and MCS are min sum variants of the well-known Set Cover and Graph Coloring problems, respectively. Kubicka and Schwenk (1989) observed that MCS can be solved in linear time for trees, and Bar-Noy and Kortsarz (1998) proved that it is APX-hard already for bipartite graphs. For general graphs, Bar-Noy et al. (1998) showed that a greedy algorithm is 4-approximate for MCS. Feige et al. (2002) observed that the greedy algorithm of Bar-Noy et al. (1998) applied to MSSC, which is to choose the element that is contained in the most uncovered sets next, yields a 4-approximation algorithm for MSSC. They simplified the proof by analyzing the performance ratio via a time-indexed linear program instead of comparing the greedy solution directly to the optimum. In the journal version of their paper, Feige et al. (2004) further simplified the proof to an elegant histogram framework, which inspired the results of this paper. They additionally proved that one cannot approximate MSSC strictly better than 4, unless $P = NP$.

² We note that this definition of Min Sum Set Cover, given in Feige et al. (2002), uses a “hitting set” formulation of the problem, in which vertices cover (hit) hyperedges, rather than vice versa.

Munagala et al. (2005) generalized MSSC by introducing non-negative costs c on the elements of V and non-negative weights w_e on the sets in E . The task is to find a linear ordering π of V that minimizes the sum of weighted covering costs of the sets $\sum_{e \in E} w_e C(e)$. The covering cost of $e \in E$ is defined as $C(e) := \min \sum_{u \in V : \pi(u) \leq \pi(v)} c_u \mid v \in e \in E \}$ (that is, the sum of all the costs of all the elements of V chosen up to and including the element that covers e). This problem is known as *pipelined set cover*. A natural extension of the greedy algorithm is to pick the element v that maximizes the ratio of the sum of the weights of the sets covered by v and the cost of v . In fact, Munagala et al. (2005) showed that this greedy algorithm is 4-approximate for pipelined set cover.

Pipelined set cover can be expressed in the form of MSOP by taking $f(S) = \sum_{v \in S} c_v$ for a subset $S \subseteq V$ and $g(S)$ to be the sum of the weights of all the subsets in E that contain at least one element of S . In this case, g is submodular and f is modular, and the fact that the greedy algorithm is 4-approximate follows from Theorem 1 of this paper (or more specifically, from Corollary 1).

Yet another generalization of MSSC is precedence-constrained MSSC. Here, the sets are subject to AND-precedence constraints and the task is to find a linear extension of the partial order on the sets. This problem was studied by McClintock et al. (2017), who proposed a $4\sqrt{n}$ -approximation algorithm for precedence-constrained MSSC using a similar approach to ours: first, apply a \sqrt{n} -approximation for finding a subset of V of maximum density, and then use a histogram-type argument, which yields an additional factor of 4. This result also follows from Theorem 1 of this paper (assuming the approximation result for the maximum density problem).

4. OR-precedence constraints

We now show how our results can be applied to give new approximation algorithms for problems involving OR-precedence constraints. In Subsection 4.1 we will define OR-scheduling and provide 4-approximation algorithms for OR-scheduling on inforests and, more generally, on *multitrees*. We then consider a new OR-precedence constrained version of MSSC in Subsection 4.2 and show that there is also a 4-approximation algorithm for this.

4.1. OR-Scheduling

One can interpret pipelined set cover (described in Subsection 3.2) as a single-machine scheduling problem in the following way. There is a job j_v for every element $v \in V$ with

processing time $p_v = c_v$ and weight $w_{j_v} = 0$, and a job j_e for every $e \in E$ with processing time $p_{j_e} = 0$ and weight $w_{j_e} = w_e$. Further, there are *OR-precedence constraints* between job j_v and all jobs j_e with $v \in e$. That is, job j_e becomes available for processing after at least one of its predecessors in $\{j|v \in e \in E\}$ is completed. Then, finding a linear ordering of V that minimizes the sum of weighted covering costs is equivalent to finding a feasible single-machine schedule that minimizes the sum of weighted completion times.

Formally, OR-scheduling is defined as follows: Let N be a set of jobs that are subject to precedence constraints given by a DAG $G = (N, E)$. An arc $(i, j) \in E$ indicates that job i is an OR-predecessor of j . Any job j with $\{i \in N \mid (i, j) \in E\} \neq \emptyset$ requires that at least one of its predecessors is completed before it can start. A job without predecessors may be scheduled at any point in time. The task is to find a feasible schedule, i.e., each job is processed non-preemptively for p units of time, and at each point in time at most one job is processed, that minimizes the sum of weighted completion times.

To see that OR-scheduling is indeed a special case of MSOP, let Σ be the set of feasible schedules, and let $F = F^\Sigma$. In other words, $S \in F$ if and only if for any job in S with predecessors, at least one of its predecessors is contained in S as well. Clearly, F is closed under union. Further, we set $f(S) = \sum_{j \in S} p_j$ and $g(S) = \sum_{j \in S} w_j$ for every set of jobs $S \subseteq N$. With these modular functions, it is not hard to see that the sum of weighted completion times of a schedule is equal to (2).

Note that, for the above reduction from pipelined set cover, the set of jobs can be partitioned into $N = A \cup B$ such that all arcs in the precedence graph go from A to B . We call such a precedence graph *bipartite*. In a recent paper, Happach and Schulz (2020a) presented a 4-approximation algorithm for scheduling with bipartite OR-precedence constraints using an approach similar to ours. For bipartite OR-scheduling, the maximum density sets can be computed in polynomial time, so a histogram argument yields a 4-approximation algorithm.

Scheduling with OR-precedence constraints was previously considered in the context of AND/OR-networks, see, e.g., Gillies and Liu (1995), Erlebach et al. (2003). In this case, Erlebach et al. (2003) presented the best-known approximation factor, which is linear in the number of jobs, and showed that obtaining a polynomial time constant-factor approximation algorithm is NP-hard. For the case where the AND/OR-constraints are of a similar bipartite structure as above, and no AND-constraints are within $B \times A$, Happach

and Schulz (2020b) obtained a 2Δ -approximation algorithm with Δ being the maximum number of OR-predecessors of any job in B . Johannes (2005) proved that minimizing the sum of weighted completion times with OR-constraints is already NP-hard for unit-processing time jobs. Happach and Schulz (2020a) strengthened this result and showed that the problem remains NP-hard even for bipartite OR-precedence constraints with unit processing times and 0/1 weights, or 0/1 processing times and unit weights.

We now consider a special case of MSOP that can be stated in terms of MSPP. Suppose the elements of V are vertices of a DAG $G = (V, E)$ which represents some OR-precedence constraints. That is, a permutation σ of V is feasible if each element v with a non-empty set of predecessors appears in σ later than at least one of its direct predecessors $P(v)$ (where $P(v)$ is the set of all u such that $(u, v) \in E$). Recall that a DAG $G = (V, E)$ is an *intree* if every vertex has at most one successor. A DAG whose connected components are intrees is an *inforest*.

Theorem 2. *Consider an instance of MSPP for which the set of feasible permutations is derived from some OR-precedence constraints given by an inforest. Then if f is modular and g is submodular, there is a polynomial time 4-approximation algorithm for the problem.*

Proof. Note f is modular and hence subadditive. Also, as pointed out in Subsection 2.2, the set of feasible permutations Σ is closed and therefore, by Lemma 3, the set Σ is well-founded and F^Σ is closed under union. Hence, by Theorem 1 and Lemma 2, it suffices to construct in polynomial time a 1-greedy F^Σ -chain.

We characterize inclusion-minimal sets of maximum density, using the concept of a *stem*. We define a stem in G to be a sequence of vertices v_1, \dots, v_k in V such that v_1 has no predecessors and $v_i \in P(v_{i+1})$ for all $i = 1, \dots, k-1$. We show that any inclusion-minimal subset S_{j+1} of V that maximizes the density $\rho_{S_j}(S_{j+1})$ is a stem and that we can enumerate all stems in polynomial time. Observe that, if we remove a stem S from the instance along with all edges incident to vertices in the stem, the graph decomposes into intrees again; also f_{S_j} is modular and g_{S_j} is submodular. So it suffices to consider only $S = \emptyset$.

Since G is an inforest, the number of paths starting at any vertex is bounded by the total number of vertices. Therefore, the total number of stems is $O(n^2)$. So we can enumerate all stems S that start at a job without a predecessor, and pick the one of maximum density $\rho(S) = g(S)/f(S)$. It remains to show that a stem of maximum density is indeed an OR-initial set of maximum density.

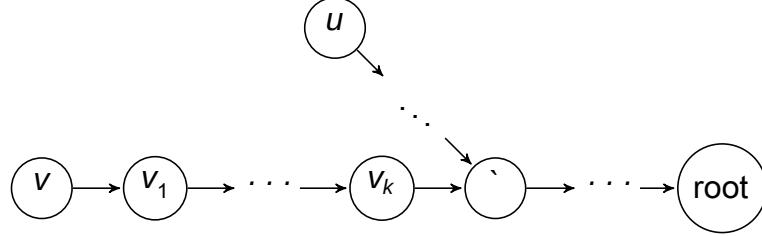


Figure 2 Paths starting at v and u to the root, and their intersection point \cdot .

Let $S \in F$ be an inclusion-minimal set of maximum density and suppose that S is not a stem. Since G is an inforest, S must induce an inforest that contains at least two vertices without a predecessor. Since every vertex has at most one successor, any vertex without a predecessor induces a unique stem to the root of its connected component in S (the root of a component being the unique vertex contained in that component that has no successors). For two such vertices, v and u , let \cdot be the vertex where the unique stems starting at v and u meet, see Figure 2. (Note that \cdot does not exist if v and u are in different connected components.) Now, let $T_v := \{v, v_1, \dots, \cdot\}$ be the set of vertices on the stem from v to \cdot such that $v_k \in P(\cdot)$. If \cdot does not exist, then let T_v be the stem starting at v to its root in S . Clearly, $T_v \in F$, and also $\bar{T} := S \setminus T_v$ is an OR-initial set.

By the submodularity of g and modularity of f , the density $\rho(S)$ satisfies

$$\rho(S) \leq \frac{g(T_v) + g(\bar{T})}{f(T_v) + f(\bar{T})} = \theta\rho(T_v) + (1 - \theta)\rho(\bar{T}), \quad (7)$$

where $\theta = f(T_v)/(f(T_v) + f(\bar{T})) \in [0, 1]$. Hence, by the maximality of $\rho(S)$, both T_v and \bar{T} must have maximum density, contradicting the assumption that S was an inclusion-minimal subset of maximum density.

Consider the OR-scheduling problem on a DAG G in the case that G is an inforest. Since the cost function and the weight function are both modular, the next theorem follows immediately from Theorem 2.

Theorem 3. *There is a polynomial time 4-approximation algorithm for OR-scheduling of inforests.*

In fact, we derive a more general result for OR-scheduling of *multitrees*, introduced in Furnas and Zacks (1994). A DAG is called a multitree if, for every vertex, its successors form an outtree (where an outtree is a DAG such that each vertex has at most one predecessor). Equivalently, there is at most one directed path between any two vertices. Inforests are examples of multitrees.

Theorem 4. *There is a polynomial time 4-approximation algorithm for OR-scheduling of multitrees.*

Proof. By Theorem 1 and Lemma 2, it is sufficient to find a 1-greedy F -chain, where $F = F^\Sigma$ and Σ is the set of feasible schedules.

We will show that the inclusion-minimal sets of maximum density are outtrees. This means that we can find a maximum density subset of F by considering each vertex v with no predecessors and finding a maximum density subtree $v \in F$ of the outtree formed by v and its successors. This can be done in polynomial time using the dynamic programming algorithm of Horn (1972), for example. We then choose a subtree with maximum density over all vertices v with no predecessors.

The proof that the inclusion-minimal sets of maximum density are outtrees is similar to the proof that the inclusion-minimal subsets of Theorem 2 were stems, so we do not go into detail. It can be shown that if S is an inclusion-minimal set of maximum density that is not an outtree, then it can be expressed as the disjoint union of an outtree T and another set in F . By an identical calculation as in (7), the set T must have maximum density, contradicting S being inclusion-minimal. This completes the proof.

It is worth pointing out that approximating the problem of minimizing the sum of weighted completion times for OR-scheduling appears to be harder than the analogous problem for AND-scheduling in the following sense. As discussed in Subsection 3.1, for AND-scheduling there is a polynomial time algorithm for series-parallel DAGs and polynomial time 2-approximation algorithms for arbitrary DAGs, whereas for OR-scheduling, we have given polynomial time 4-approximation algorithms for inforests and, more generally, multitrees. It is not possible that better approximations exist for OR-scheduling on multitrees (or even for bipartite graphs) unless $P = NP$, since the same is true of MSSC (Feige et al. 2004), which is a special case of OR-scheduling on bipartite graphs. Of course, for outtrees, OR-scheduling and AND-scheduling are equivalent, but there are no polynomial time algorithms known for OR-scheduling on any other classes of DAGs.

4.2. OR-Precedence Constrained MSSC and Pipelined Set Cover

Consider a variation of MSSC in which the order that the elements of V are chosen must be consistent with some OR-precedence constraints given by a DAG G . As in pipelined set cover, we additionally assume that there is a non-negative cost c_v for each vertex v and

a non-negative weight w_e for each hyperedge $e \in E$, and the objective is to minimize the weighted sum of covering times of the edges. As for OR-scheduling, we take F to be the collection of OR-initial sets of G and as for pipelined set cover, we take $f(S) = \sum_{v \in S} c_v$ and $g(S)$ to be the sum of the weights of all hyperedges in E that contain at least one element of S , for $S \in F$. Then f is modular and g is submodular, so we can again apply Theorem 2.

Theorem 5. *There is a polynomial time 4-approximation algorithm for pipelined set cover with OR-precedence constraints that take the form of an inforest.*

To see that Theorem 5 is a generalization of Theorem 3, simply observe that if the set of hyperedges E consists of all the singletons of V , then pipelined set cover with OR-precedence constraints is equivalent to OR-scheduling.

5. Evaluation of Read-Once Formulas

In this section we give an 8-approximation algorithm for a non-adaptive version of a Boolean function evaluation problem, involving the evaluation of a read-once formula on an initially unknown input in a stochastic setting. We call this the *na-ROF evaluation problem*.

A read-once formula, also called an AND/OR tree, is a rooted tree with the following properties. Each internal node of the tree is labeled either OR or AND (corresponding to OR or AND gates). The leaves of the tree are labeled with Boolean variables x_1, \dots, x_n , where n is the number of leaves, with each x_i appearing in exactly one leaf. Given a Boolean assignment to the variables in the leaves, the value of the formula on that assignment is defined recursively in the usual way: the value of a leaf labeled x_i is the assignment to x_i , and the value of a tree whose root is labeled OR (respectively, AND) is the Boolean OR (respectively, AND) of the values of the subtrees of that root. Read-once formulas are equivalent to series-parallel systems (cf. Ünlüyurt (2004)).

An example read-once formula, corresponding to the expression $\varphi(x_1, x_2, x_3, x_4, x_5) = x_1 \wedge x_2 \wedge ((x_3 \wedge x_4) \vee x_5)$, is shown in Figure 3.

In the na-ROF evaluation problem, we are given a read-once formula $\varphi(x_1, \dots, x_n)$ that must be evaluated on an initially unknown random assignment to its input variables. For each of the n Boolean variables x_i , we are given a probability p_i , where $0 < p_i < 1$, and a positive integer cost c_i . The random assignment on which we must evaluate φ is assumed

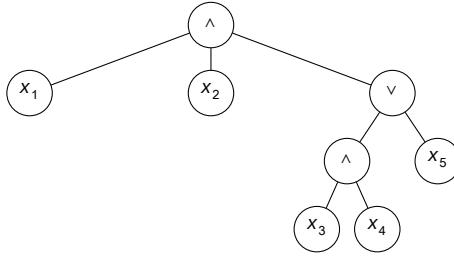


Figure 3 Read-once formula.

to be drawn from the product distribution defined by the p 's, that is, the joint distribution in which $p_i = P[x_i = 1]$ and the x_i are independent. The value of an x_i in the random assignment can only be ascertained by performing a test, which we call *test* i . Performing test i incurs cost c_i , and its outcome is the value of x_i . Tests are performed sequentially until there is enough information to determine the value of φ . The problem is to develop a linear ordering (permutation) of the tests, such that performing the tests in that order, until the function value is determined, minimizes the expected cost incurred in testing.

For example, consider evaluating the formula in Figure 3 in the order given by the permutation $(3, 4, 5, 2, 1)$. Suppose the first test reveals that $x_3 \neq 0$, the second that $x_4 = 1$, and the third that $x_5 = 0$. Then testing will stop after that third test, when it can be determined that the value of φ is 0. The probability of this happening is $(1 - p_3)p_4(1 - p_5)$, and the incurred cost in this case is $c_3 + c_4 + c_5$. More generally, for each prefix R of $(3, 4, 5, 2, 1)$ let P_R denote the probability that testing stops at the end of that prefix. Using q_i to denote $1 - p_i$, we have e.g., $P_{(3)} = P_{(3,4)} = 0$, $P_{(3,4,5)} = (1 - p_3)p_4q_5$ and $P_{(3,4,5,2)} = (1 - P_{(3,4,5)})q_2$. The expected cost of evaluating the formula according to the permutation $(3, 4, 5, 2, 1)$ is $c_3 + c_4 + c_5 + c_2(1 - P_{(3,4,5)}) + c_1(1 - P_{(3,4,5,2)})$. More generally, the expected cost associated with an arbitrary permutation of the five tests is equal to $\sum_{j=1}^5 f(S_j)(g(S_j) - g(S_{j-1}))$, where S_j is the set consisting of the first j elements of the permutation $(3, 4, 5, 2, 1)$, $f(S) = \sum_{i \in S} c_i$, and $g(S_j)$ is the probability that the value of φ can be determined by performing just the tests in S_j . Thus $g(S_j) - g(S_{j-1}) = P_{R^j}$, where R^j denotes the prefix consisting of the first j elements of the permutation.

The na-ROF evaluation problem corresponds to the MSOP with $F = 2^V$, where $V = \{1, \dots, n\}$ is the set of tests, $f(S) = \sum_{i \in S} c_i$, and $g(S)$ is equal to the probability that the value of φ can be determined from the outcomes of the tests in S . (We note that the same correspondence holds for analogous evaluation problems involving other types of

formulas and functions. Gkenosis et al. (2018) studied the analogous evaluation problem for symmetric Boolean functions.)

In Section 5 we show that there is an 8-approximation algorithm for the na-ROF evaluation problem that runs in pseudo-polynomial time for general costs, and polynomial time for unit costs. We do this by giving a dynamic programming algorithm producing a 2-approximate solution for the associated maximum density problem.

To our knowledge, the na-ROF evaluation problem has not been previously studied. However, an “adaptive” version of this evaluation problem has been studied since the 1970’s, under a variety of names, including “satisficing strategies for AND/OR trees” (Greiner et al. (2006)), “sequential testing of series-parallel systems” (Ünlüyurt (2004)), and “the Stochastic Boolean Function Evaluation (SBFE) problem for read-once formulas” (Deshpande et al. (2016)). This version seeks an optimal *adaptive* strategy for ordering the tests used to evaluate the read-once formula. In an adaptive strategy, the choice of the next test can depend on the outcomes of the previous tests, and thus the testing order corresponds to a decision tree, rather than to a permutation. It is still an open question whether the adaptive version of the problem is NP-hard, or whether it can be solved by a polynomial-time algorithm, even in the unit-cost case. It is also unknown, even in the unit-cost case, whether it can be solved by a polynomial-time algorithm achieving a sublinear approximation factor. In contrast, we establish here that there is a polynomial-time constant-factor approximation algorithm for the na-ROF evaluation problem in the unit-cost case. However, as with the adaptive version of the problem, it remains open whether the na-ROF evaluation problem is, in fact, NP-hard.

An easy case of the na-ROF evaluation problem is where φ is the Boolean OR function, $\varphi(x_1, \dots, x_n) = x_1 \vee \dots \vee x_n$. In this case, the optimal solution is to perform the tests in decreasing order of the ratio p/c_i until the value of f can be determined (which occurs as soon as an $x_i = 1$ is found, or after all x_i have been found to equal 0). This optimal solution has been rediscovered many times (cf Ünlüyurt (2004)). It is also optimal for the adaptive version of the problem. However, in contrast to the case of the OR function, an optimal linear order for evaluating a read-once formula will generally incur higher expected cost than an optimal adaptive strategy. This is because, for example, learning that a variable $x_i = 0$ when its parent node is labeled AND allows us to “prune” all other subtrees of that

AND node, making it unnecessary to perform tests on any of the variables that were in the pruned subtrees.

Boros and Ünlüyurt (2000) and Işık and Ünlüyurt (2013) considered what might be called a “partially-adaptive” version of the na-ROF evaluation problem, where the strategy is specified by a permutation, but tests in the permutation are skipped if the results of previous tests have rendered them irrelevant. They presented results on evaluating read-once formulas of small depth. There does not appear to be a way to express this partially-adaptive version of the na-ROF evaluation problem in MSOP form.

Charikar et al. (2000) studied the evaluation problem for read-once formulas in the worst-case online (non-stochastic) setting, where the goal is to minimize the so-called “competitive ratio”. They gave an interesting exact algorithm for minimizing the competitive ratio, with pseudo-polynomial running time. It is not applicable to our problem, where the goal is to minimize expected cost.

5.1. Preliminaries

Recall the inputs to the na-ROF evaluation problem: (1) a read-once formula $\varphi(x_1, \dots, x_n)$, (2) for each x_i , the value $p := P[x_i = 1]$ where $0 < p < 1$, and (3) for each x_i , the associated integer test cost c_i , which is greater than 0. We may assume without loss of generality that each AND and OR gate of φ has exactly two inputs (since, e.g., $\wedge x_2 \wedge x_3 = (x_1 \wedge x_2) \wedge x_3$). We consider each input x_i in φ to also be a gate (an *input gate*) of φ . The set of tests is $V = \{1, \dots, n\}$.

We use partial assignments to represent the outcomes of a subset of the tests. In a partial assignment $b \in \{0, 1, *\}^n$, $b_i = *$ means that test i has not been performed and the value of x_i is unknown, otherwise b is the outcome of test i . For a (full) assignment $a \in \{0, 1\}^n$ and S a subset of $V = \{1, \dots, n\}$, $a|_S$ is the partial assignment $b \in \{0, 1, *\}^n$ where $b = a_i$ for $i \in S$, and $b_i = *$ otherwise. Given partial assignment $b \in \{0, 1, *\}^n$, an *extension* of b is a (full) assignment $a \in \{0, 1\}^n$ where $a = b_i$ for all i such that $b_i \neq *$. If for all extensions a of b , $\varphi(a)$ has the same value $\varphi(b)$, the value of φ is determined by b and we write $\varphi(b) = \varphi$. Otherwise, we write $\varphi(b) = *$. Intuitively, for $S \subset V$ and $a \in \{0, 1\}^n$, $\varphi(a|_S) = *$ means that the outcomes of the tests in S , as specified by a , are not sufficient to determine the value of φ .

Let A_1, \dots, A_n be independent Bernoulli random variables, where $P[A_i = 1] = p_i$. Thus $A = [A_1, \dots, A_n]$ is a random variable which takes on a value $a \in \{0, 1\}^n$, corresponding to the outcome of the n tests.

In the MSOP formulation of the na-ROF evaluation problem, given above, we defined $g(S)$ to be equal to the probability that the value of φ can be determined from the outcomes of the tests in S . Thus, $g(S) = P[\varphi(A|_S) \neq *]$.

We will obtain an 8-approximate solution to the problem by constructing a 2-greedy chain. In particular, we will give a (pseudo) polynomial time 2-approximation algorithm solving the associated density problem.

5.2. Background for the Dynamic Programming Algorithm

The dynamic programming algorithm relies on the following definitions and observations.

For $S \subseteq T \subseteq V$, let

$$\rho_S(T) = \frac{P[\varphi(A|_T) \neq *] - P[\varphi(A|_S) \neq *]}{\sum_{i \in T \setminus S} c_i} \quad (8)$$

For $S \subseteq V$ and $\alpha > 0$, call $R \subseteq V \setminus S$ an α -approximate max-density *supplement* for S if

$$\rho_S(S \cup R) \geq \frac{1}{\alpha} \max_{\{T \subseteq V : S \subseteq T\}} \rho_S(T).$$

Our dynamic programming algorithm computes a 2-approximate max-density supplement for an input subset S . It does so in a bottom-up fashion, calculating values at each of the gates G of φ . For gate G of φ , define $\text{tests}(G)$ to be the set of $i \in V$ such that x_i is a descendant of G in the tree φ . We consider a gate to be its own descendant, so if G is an input gate x_i , then $i \in \text{tests}(G)$.

Each gate G of φ is the root of a subtree of φ . Define φ_G to be the subformula corresponding to the subtree of φ that is rooted at G . Thus φ_G is a read-once formula over the variable set $\{x_i \mid i \in \text{tests}(G)\}$. We treat φ_G as computing a function over $\{0, 1\}^n$, whose output depends only on the values of the variables in $\{x_i \mid i \in \text{tests}(G)\}$. For $b \in \{0, 1, *\}^n$, we refer to $\varphi_G(b)$ as the output of gate G on partial assignment b , which may be either 0, 1, or *.

We note that given a subset $S \subseteq V$, and $\cdot \in \{0, 1\}$, the value of $P[\varphi(A|_S) = \cdot]$ for each gate G of φ can be computed in time linear in n by processing the gates of φ in bottom-up order. Consider the case where $\cdot = 1$ and let $p_G = P[\varphi_G(A|_S) = 1]$. If G is an input gate x_i , then $p_G = p_i$ if $x_i \in S$, otherwise $p_G = 0$. If G is an AND gate with children G^0 and G^{00} , then because φ is read-once and the x_i are independent, $p_G = p_{G^0} \cdot p_{G^{00}}$. If G is an OR gate, then $p_G = p_{G^0} + p_{G^{00}} - p_{G^0} \cdot p_{G^{00}}$.

Dually, consider the case where $\cdot = 0$ and let $q_G = P[\varphi_G(A|_S) = 0]$. If G is an input gate x_i , then $q_G = 1 - p_i$ if $x_i \in S$, otherwise $q_G = 0$. If G is an OR gate with children G^0 and G^{00} , then $q_G = q_{G^0} q_{G^{00}}$. If G is an AND gate, then $q_G = q_{G^0} + q_{G^{00}} - q_{G^0} q_{G^{00}}$.

5.3. The na-ROF Evaluation Algorithm

Our algorithm for na-ROF evaluation relies on the dynamic programming algorithm presented in the proof of the following lemma.

Lemma 4. *Given $S \subseteq V$, a 2-approximate max-density supplement R for S can be computed in time polynomial in n and $\sum_{i=1}^n c_i$.*

Proof. We prove the lemma for the case of unit costs, where all the c_i 's are equal to 1. We then explain how to extend the proof to handle arbitrary costs.

Assume the c_i 's are all equal to 1. We describe an algorithm that we call *FindSupp* that finds a max-density subset R for a given input subset S .

Fix S . For $R \subseteq V \setminus S$, let $\sigma(R) = \rho|_S(S \cup R)$. Since we assumed the c_i 's are equal to 1,

$$\sigma(R) = \frac{P[\varphi(A|_{S \cup R}) = 1] - P[\varphi(A|_S) = 1]}{|R|}.$$

Clearly, $P[\varphi(A|_S) = 1] = P[\varphi(A|_S) = 0] + P[\varphi(A|_S) = 1]$ and similarly for $A|_{S \cup R}$.

For $\cdot \in \{0, 1\}$, define

$$\sigma_\cdot(R) = \frac{P[\varphi(A|_{S \cup R}) = \cdot] - P[\varphi(A|_S) = \cdot]}{|R|}.$$

Thus

$$\sigma(R) = \sigma_1(R) + \sigma_0(R) \tag{9}$$

The idea behind *FindSupp* is to compute two subsets R and R^0 , maximizing σ_1 and σ_0 respectively. By (9), the R with the larger value of $\sigma(R)$ is a 2-approximate max-density supplement for S .

For gate G of φ , let $T(G) = \text{tests}(G) \setminus S$. For $t \in \{0, \dots, |T(G)|\}$, and $\cdot \in \{0, 1\}$, let $R_{G,t,\cdot}$ be a subset R that maximizes the value of $P[\varphi(A|_{S \cup R}) = \cdot]$ subject to the constraints that $R \subseteq T(G)$ and $|R| = t$.

Let $p_{G,t,\cdot}$ be the value of $P[\varphi_G(A|_{S \cup R}) = \cdot]$ for $R = R_{G,t,\cdot}$. Let \tilde{G} denote the root gate of φ . Thus for $t \in \{1, \dots, |V \setminus S|\}$ and $\cdot \in \{0, 1\}$, setting $R = R_{\tilde{G},t,\cdot}$ maximizes the value of $\sigma(R)$, over all $R \subseteq V \setminus S$ of size t .

Algorithm FindSupp:

FindSupp first runs a procedure ComputeRp that computes $R_{G,t,`}$ and $p_{G,t,`}$ for all $t \in \{1, \dots, |V|S\}$ and $` \in \{0, 1\}$. We describe the details of ComputeRp below. After running ComputeRp, FindSupp uses it to obtain the two subsets, R^0 and R^1 , maximizing σ_1 and σ_0 respectively. It does this as follows. First, using the linear-time procedure described above, it computes the value of $P[\varphi(A|S) = `]$, for $` \in \{0, 1\}$.

Then, for each $t \in \{1, \dots, |V|S\}$, for $` \in \{0, 1\}$, FindSupp computes the value of

$$\sigma(`(R_{G,t,`})) = \frac{p_{G,t,`} - P[\varphi(A|S) = `]}{t}$$

For each $` \in \{0, 1\}$, the algorithm then finds the value of t which yielded the highest value for $\sigma(`(R_{G,t,`}))$. Let t^* denote that value. Let R^* be the value of $R_{G,t,`}$ for $t = t^*$.

Because $R_{G,t,`}$ maximizes σ among candidate subsets of size t , setting $R = R^*$ maximizes $\sigma(R)$ among candidate subsets of all possible sizes. The algorithm returns R^0 if $\sigma_0(R^0) > \sigma_1(R^1)$, and returns R^1 otherwise.

Procedure ComputeRp:

For all gates G of φ , ComputeRp computes the values of $p_{G,t,`}$ and $R_{G,t,`}$ for all $t \in \{0, \dots, |T(G)|\}$, and $` \in \{0, 1\}$. It processes the gates G of φ in bottom-up order, from the leaves to the root.

We begin by describing how ComputeRp computes $p_{G,t,`}$ and $R_{G,t,`}$ when G is an AND gate, $t \in \{0, \dots, |T(G)|\}$, and $` = 1$. Suppose that G^0 and G^{00} are the children of AND gate G , and that $R_{G^0,t^0,1}$, $p_{G^0,t^0,1}$, $R_{G^{00},t^{00},1}$, and $p_{G^{00},t^{00},1}$ have already been computed, for all $t^0 \in \{0, \dots, |T(G^0)|\}$ and $t^{00} \in \{0, \dots, |T(G^{00})|\}$. ComputeRp first computes the product $p_{G^0,t^0,1} \cdot p_{G^{00},t^{00},1}$ for all j such that $j \in \{0, \dots, |T(G^0)|\}$ and $t - j \in \{0, \dots, |T(G^{00})|\}$. It then sets j^* to be the value of j maximizing that product, and sets $R_{G,t,1} = R_{G^0,j^*,1} \cup R_{G^{00},t-j^*,1}$ and $p_{G,t,1} = p_{G^0,j^*,1} \cdot p_{G^{00},t-j^*,1}$.

The correctness of these settings follows from the fact that $R_{G,t,1}$ must consist of a subset R^0 of $T(G^0)$ of some size j^* , and a subset R^{00} of $T(G^{00})$ of size $t - j^*$. $R_{G,t,1}$ maximizes the probability that G outputs 1 (among subsets of $T(G)$ of size t , when added to S). Since φ is a read-once formula, $\text{tests}(G^0)$ and $\text{tests}(G^{00})$ are disjoint. R^0 and R^{00} are thus sets that maximize the probability that G^0 and G^{00} output 1 (when added to S , among subsets of $T(G^0)$ and $T(G^{00})$ of sizes j^* and $t - j^*$ respectively). Thus, given j^* , $R_{G,t,1}$ can be set to

$R^0 \cup R^{00}$ where $R^0 = R_{G^0 j^*, 1}$, $R^{00} = R_{G^{00} t-j^*, 1}$, and $p_{G, t, 1}$ can be set to $p_{G^0 j^*, 1} \cdot p_{G^{00} t-j^*, 1}$. Since ComputeRp is not given the value of j^* , it must try all possible j .

Similarly, suppose G is an AND gate, $t \in \{0, \dots, |T(G)|\}$, and $\hat{\cdot} = 0$. In this case, ComputeRp computes $p_{G, 0} + p_{G^{00} t-j, 0} - p_{G^0 j, 0} \cdot p_{G^{00} t-j, 0}$ for all possible j , and then sets j^* to be the value that maximized the expression. It then sets $R_{G, t, 0} = R_{G^0 j^*, 0} \cup R_{G^{00} t-j^*, 0}$ and $p_{G, t, 0} = p_{G^0 j^*, 0} + p_{G^{00} t-j^*, 0} - p_{G^0 j^*, 0} \cdot p_{G^{00} t-j^*, 0}$. The correctness in this case follows from the fact that the output of G is 0 if either of its child gates outputs 0. Thus to maximize the probability that AND gate G outputs 0, one needs to maximize the probability that each of its child gates outputs 0.

The case where G is an OR gate is dual and we omit the details.

The remaining case is where G is an input gate x_i , $t \in \{0, \dots, |T(G)|\}$, and $\hat{\cdot} \in \{0, 1\}$. Note that since G is an input gate, if $i \in S$, then $|T(G)| = 0$. If $i \notin S$, then $|T(G)| = 1$. If $t = 0$, then for $\hat{\cdot} \in \{0, 1\}$, ComputeRp sets $R_{G, t, \hat{\cdot}} = \emptyset$. Then, if $i \in S$ it sets $p_{G, t, 1} = p_i$ and $p_{G, t, 0} = 1 - p_i$. If $i \notin S$ it sets both $p_{G, t, 0} = 0$ and $p_{G, t, 1} = 0$. If $t = 1$ (and therefore $i \notin S$), it sets $R_{G, t, \hat{\cdot}} = x_i$, $p_{G, t, 1} = p_i$ and $p_{G, t, 0} = 1 - p_i$. The correctness of these settings is straightforward.

To compute the running time of ComputeRp, note that because φ is a read-once formula with n input variables, it has $O(n)$ gates. For each gate G , $|T(G)|$ is $O(n)$, and thus there are $O(n^2)$ values computed by ComputeRp. For each $G, t, \hat{\cdot}$ where G is an AND or OR gate, ComputeRp spends time linear in $|T(G)|$ to find j^* , which yields the values for $p_{G, \hat{\cdot}}$ and $R_{G, t, \hat{\cdot}}$. Thus the running time of ComputeRp is $O(n^3)$.

Generalization to arbitrary costs:

The algorithm FindSupp can easily be modified to handle arbitrary non-negative integer costs. The main difference is that t is used to represent the total cost of a set R of tests, rather than just the size of the set.

Consider a gate G of φ . If there is at least one subset $R \subseteq T(G)$ such that $t = \sum_{i \in R} c_i$, call t a feasible value for G .

For t a feasible value for G , let $R_{t, \hat{\cdot}}$ be the subset maximizing $\rho(S)$ (whose denominator is now $\sum_{i \in R} c_i$) subject to the constraints that $R \subseteq T(G)$ and $\sum_{i \in R} c_i = t$. Because of our assumption that the costs c_i are integers, the feasible values of t are all in the set $\{1, \dots, \sum_{i \in V} c_i\}$.

ComputeRp computes $R_{G,t}$ for all feasible values t for G , as follows. Suppose G is an AND or OR gate. Let G^0 and G^{00} be its children. ComputeRp identifies the feasible values t for G , which are the values $t = t^0 + t^{00}$ such that t^0 is feasible for G^0 and t^{00} is feasible for G^{00} . To compute $R_{G,t}$ for a feasible value t for G , ComputeRp tries all j such that j is feasible for G^0 and $t - j$ is feasible for G^{00} . The other modifications in the algorithm are straightforward.

Because there are at most $\sum_{i \in V} c_i$ feasible values t for gates G , the running time of FindSupp is $O(n(\sum_{i \in V} c_i)^2)$.

The algorithm described in Lemma 4 can be used to form a 2-greedy chain $S_0 \subset S_1 \subset \dots \subset S_k$, where each S_{j+1} is generated from S_j by running FindSupp with $S = S_j$ to produce R , and then setting $S_{j+1} = S_j \cup R$.

The theorem now follows immediately from Lemma 4 and Theorem 1.

Theorem 6. *There is a pseudo-polynomialtime 8-approximation algorithm solving the na-ROF evaluation problem. The algorithm runs in polynomial time in the unit cost case.*

6. Backward Greedy Algorithms

The greedy algorithm we presented for MSOP works by starting with the empty set of elements, and greedily adding subsets in each greedy step. An alternative greedy approach starts with the set V , and greedily removes subsets from V in each greedy step. A backward greedy approach was used by Iwata et al. (2012) in their work on some special cases of MSPP.

In this section, we describe a backward analog to our greedy algorithm for MSOP and its relationship to a *dual* MSOP problem, analogous to the dual problem introduced in Fokkink et al. (2019).

We call an F -chain $S = (S_j)_{j=0}^k$ a *backward α -greedy chain* if

$$\rho_{S_j}(S_{j-1}) \leq \alpha \min_{\{T \in F : T \subseteq S_j\}} \rho_{S_j}(T),$$

for all $j = 1, \dots, k$.

A backward α -greedy chain can be seen to be equivalent to an α -greedy chain for the dual problem. To describe the dual problem, we write the cost function $C(f, g)$ in another, equivalent way. Fixing $F \subseteq 2^V$, we first define $F^\#$ as the family of complements of sets in F . That is, $F^\# = \{S \subseteq V : V \setminus S \in F\}$. Note that F is closed under intersection if

and only if $F^\#$ is closed under union. There is a one-to-one correspondence between F -chains and $F^\#$ -chains, obtained by mapping an element S of an F -chain to $V \setminus S$ and reversing the order. We refer to the corresponding $F^\#$ -chain of a F -chain S as its *dual chain*, which we denote by $S^\#$. We also denote the *dual function* of f by $f^\# : F^\# \rightarrow \mathbb{R}$, given by $f^\#(S) = f(V) - f(V \setminus S)$, and similarly for g . Note that a set function is the dual of its dual, as is an F -chain, and that $f^\#$ and $g^\#$ are non-decreasing. Also, f is submodular if and only if $f^\#$ is supermodular.

Given an MSOP with inputs f , g and F , the dual problem is an MSOP with inputs $g^\#$, $f^\#$ and $F^\#$. In other words, the dual problem is to minimize

$$C_{g^\#, f^\#}(T) = \sum_{j=1}^{|T|} g^\#(T_j)(f^\#(T_j) - f^\#(T_{j-1})),$$

over all $F^\#$ -chains $T = (T_j)_{j=0}^k$. Observe that an MSOP is the dual of its dual.

It is now easy to see that an F -chain $S = (S_j)_{j=0}^n$ is a backward α -greedy chain if and only if its dual chain is an α -greedy $F^\#$ -chain for the dual problem.

The following is immediate and generalizes a similar observation from Fokkink et al. (2019).

Lemma 5. *If S is an F -chain, then $C_{f,g}(S) = C_{g^\#, f^\#}(S^\#)$ and S is an α -approximation for an instance of MSOP if and only if $S^\#$ is an α -approximation for its dual.*

Proof. To prove the first statement, we point out that the area A under the second histogram in the proof of Theorem 1, given by the sum in (3), is equal to $C_{g^\#, f^\#}(S^\#)$. This area is also shown to be equal to $C_g(S)$ later in the same proof. The second statement in the lemma follows directly from the first.

Applying Theorem 1 to the dual problem, we obtain the following theorem as a corollary.

Theorem 7. *Suppose $\#$ is subadditive and F is closed under intersection. Then for any $\alpha \geq 1$, a backward α -greedy chain is a 4α -approximation for an optimal chain for MSOP.*

Note that if g is supermodular, then $g^\#$ is submodular and hence subadditive, and therefore satisfies the hypothesis of Theorem 7.

We may also apply Corollary 1 to the dual problem to obtain an additional corollary.

Corollary 2. *Suppose $F = 2^V$. If f is supermodular and g is modular then a backward 1-greedy chain can be found in polynomial time and there exists a polynomial time 4-approximation algorithm for MSOP.*

7. Future Work

We have created a general framework for min-sum ordering problems, and while Theorem 1 relies on very modest assumptions, it is only useful if the maximum density problem can be efficiently approximated. More work is needed in this area in order to further exploit our approximation result.

Particular problems of interest include Generalized Min Sum Set Cover (GMSSC), introduced in Azar et al. (2009). Unlike MSSC, where a hyperedge is “covered” the first time any of its vertices are chosen, in GMSSC, each hyperedge has its own “covering requirement”, which specifies how many of its vertices must be chosen before it is “covered”. The objective is to minimize the sum of covering times, as in MSSC. If the associated maximum density problem could be approximated within a factor of α , this would yield a 4α -approximation algorithm for GMSSC. The best approximation factor obtained to date for GMSSC is 4.642, due to Bansal et al. (2020). Their algorithm solves an LP relaxation and then applies a kernel transformation and randomized rounding. By adding costs to vertices and weights to hyperedges, one could further generalize GMSSC, giving rise to a more general maximum density problem of interest.

Another problem that comes under our framework is the Unreliable Job Scheduling Problem (UJP), introduced in Agnetis et al. (2009). In the basic setting, a set of jobs with given rewards must be scheduled by a single machine to maximize the total expected reward. There is a probability of failure associated with each job when it is scheduled, and if failure occurs the machine cannot schedule any further jobs. The problem has a neat “index” solution. Natural generalizations of the problem would consider the possibility of AND- or OR-precedence constraints, and therefore their associated maximum density problems.

Acknowledgments

The authors would like to thank Andreas S. Schulz for helpful comments in improving the presentation of the paper. This material is based upon work supported by the National Science Foundation under Grant Numbers IIS-1909446 and IIS-1909335 and by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF).

References

Agnetis A, Detti P, Pranzo M, Sodhi MS (2009) Sequencing unreliable jobs on parallel machines. *Journal of Scheduling* 12(1):45.

Alpern S, Lidbetter T (2013) Mining coal or finding terrorists: The expanding search paradigm. *Operations Research* 61(2):265–279.

Ambühl C, Mastrolilli M (2009) Single machine precedence constrained scheduling is a vertex cover problem. *Algorithmica* 53(4):488–503.

Averbakh I, Pereira J (2012) The flowtime network construction problem. *IIE Transactions* 44(8):681–694.

Azar Y, Gamzu I, Yin X (2009) Multiple intents re-ranking. *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, 669–678.

Bansal N, Batra J, Farhadi M, Tetali P (2020) Improved approximations for min sum vertex cover and generalized min sum set cover. *CoRR* abs/2007.09172, URL <https://arxiv.org/abs/2007.09172>.

Bar-Noy A, Bellare M, Halldórsson MM, Shachnai H, Tamir T (1998) On chromatic sums and distributed resource allocation. *Information and Computation* 140(2):183–202.

Bar-Noy A, Kortsarz G (1998) Minimum color sum of bipartite graphs. *Journal of Algorithms* 28(2):339–365.

Boros E, Ünlüyurt T (2000) Sequential testing of series-parallel systems of small depth. Laguna M, Velarde JLG, eds., *Computing Tools for Modeling, Optimization and Simulation: Interfaces in Computer Science and Operations Research*, 39–73 (Boston, MA: Springer US).

Charikar M, Fagin R, Guruswami V, Kleinberg J, Raghavan P, Sahai A (2000) Query strategies for priced information (extended abstract). *Proceedings of the 32nd annual ACM Symposium on Theory of Computing*, 582–591, STOC '00 (New York, NY, USA: ACM), ISBN 1-58113-184-4, URL <http://dx.doi.org/10.1145/335305.335382>.

Chekuri C, Motwani R (1999) Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Disc. Appl. Math.* 98(1):29–38.

Chudak FA, Hochbaum DS (1999) A half-integral linear programming relaxation for scheduling precedence-constrained jobs on a single machine. *Operations Research Letters* 25(5):199–204.

Correa JR, Schulz AS (2005) Single-machine scheduling with precedence constraints. *Mathematics of Operations Research* 30(4):1005–1021.

Deshpande A, Hellerstein L, Kletenik D (2016) Approximation algorithms for stochastic submodular set cover with applications to boolean function evaluation and min-knapsack. *ACM Trans. Algorithms* 12(3):42:1–42:28.

Erlebach T, Käab V, Möhring RH (2003) Scheduling AND/OR-networks on identical parallel machines. *International Workshop on Approximation and Online Algorithms*, volume 2909 of *LNCS*, 123–136 (Springer).

Feige U, Lov'asz L, Tetali P (2002) Approximating min-sum set cover. *International Workshop on Approximation Algorithms for Combinatorial Optimization*, volume 2462 of *LNCS*, 94–107 (Springer).

Feige U, Lov'asz L, Tetali P (2004) Approximating min sum set cover. *Algorithmica* 40(4):219–234.

Fokkink R, Kikuta K, Ramsey D (2017) The search value of a set. *Annals of Oper. Res.* 256(1):63–73.

Fokkink R, Lidbetter T, V'egh L (2019) On submodular search and machine scheduling. *Mathematics of Operations Research* 44(4):1431–1449.

Furnas GW, Zacks J (1994) Multitrees: enriching and reusing hierarchical structure. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 330–336.

Gillies DW, Liu JWS (1995) Scheduling tasks with and/or precedence constraints. *SIAM Journal on Computing* 24(4):797–810.

Gkenosis D, Grammel N, Hellerstein L, Kletenik D (2018) The stochastic score classification problem. *Proceedings of the 26th Annual European Symposium on Algorithms, ESA*, 36:1–36:14.

Greiner R, Hayward R, Jankowska M, Molloy M (2006) Finding optimal satisficing strategies for and-or trees. *Artificial Intelligence* 170(1):19–58.

Happach F, Schulz AS (2020a) Approximation algorithms and LP relaxations for scheduling problems related to min-sum set cover. *arXiv:2001.07011*.

Happach F, Schulz AS (2020b) Precedence-constrained scheduling and min-sum set cover. *Approximation and Online Algorithms*, volume 11926 of *LNCS*, 170–187 (Springer).

Hermans B, Leus R, Matuschke J (2019) Exact and approximation algorithms for the expanding search problem. *arXiv:1911.08959*.

Horn W (1972) Single-machine job sequencing with treelike precedence ordering and linear delay penalties. *SIAM Journal on Applied Mathematics* 23(2):189–202.

İşik G, Ünlüyurt T (2013) Sequential testing of 3-level deep series-parallel systems. *IEEE International Conference on Industrial Engineering and Engineering Management*, URL <http://dx.doi.org/10.1109/IEEM.2013.6962515>.

Iwata S, Tetali P, Tripathi P (2012) Approximating minimum linear ordering problems. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 206–217 (Springer).

Johannes B (2005) On the complexity of scheduling unit-time jobs with OR-precedence constraints. *Operations Research Letters* 33(6):587–596.

Kubicka E, Schwenk AJ (1989) An introduction to chromatic sums. *Proceedings of the 17th conference on ACM Annual Computer Science Conference*, 39–45 (ACM).

Lawler EL (1978) Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Annals of Discrete Mathematics*, volume 2, 75–90 (Elsevier).

Margot F, Queyranne M, Wang Y (2003) Decompositions, network flows and a precedence constrained single machine scheduling problem. *Oper. Res.* 51(6):981–992.

McClintock J, Mestre J, Wirth A (2017) Precedence-constrained min sum set cover. *28th International Symposium on Algorithms and Computation*, volume 92 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 55:1–55:12.

Munagala K, Babu S, Motwani R, Widom J (2005) The pipelined set cover problem. *International Conference on Database Theory*, volume 3363 of *LNCS*, 83–98 (Springer).

Pisaruk N (1992) The boundaries of submodular functions. *Comp. Math. Math. Phys.* 32(12):1769–1783.

Schulz AS (1996) Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. *Proceedings of the 5th International Conference on Integer Programming and Combinatorial Optimization*, volume 1084 of *LNCS*, 301–315 (Springer).

Sidney J (1975) Decomposition algorithms for single-machine sequencing with precedence relations and deferral costs. *Oper. Res.* 23(2):283–298.

Streeter MJ, Golovin D (2008) An online algorithm for maximizing submodular functions. Koller D, Schuurmans D, Bengio Y, Bottou L, eds., *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, 1577–1584 (Curran Associates, Inc.).

Ünlüyurt T (2004) Sequential testing of complex systems: a review. *Discrete Applied Mathematics* 142(1-3):189–205.