Improved Results for Data Migration and Open Shop Scheduling

RAJIV GANDHI
Department of Computer Science
Rutgers University - Camden
and
MAGNÚS M. HALLDÓRSSON
Department of Computer Science
University of Iceland
and
GUY KORTSARZ
Department of Computer Science
Rutgers University - Camden
and
HADAS SHACHNAI
Department of Computer Science
The Technion

The *data migration* problem is to compute an efficient plan for moving data stored on devices in a network from one configuration to another. We consider this problem with the objective of minimizing the sum of completion times of all storage devices. It is modeled by a transfer graph, where vertices represent the storage devices, and the edges indicate the data transfers required between pairs of devices. Each vertex has a non-negative weight, and each edge has a release time and a processing time. A vertex completes when all the edges incident on it complete; the constraint is that two edges incident on the same vertex cannot be processed simultaneously. The objective is to minimize the sum of weighted completion times of all vertices. Kim (*Journal of Algorithms, 55:42-57, 2005*) gave a 9-approximation algorithm for the problem when edges have arbitrary processing times and are released at time zero. We improve Kim's result by giving a 5.06-approximation algorithm. We also address the *open shop scheduling* problem, $O|r_j|\sum w_j C_j$, and show that it is a special case of the data migration problem. Queyranne and Sviridenko (*Journal of Scheduling, 5:287-305, 2002*) gave a 5.83-approximation algorithm for the non-preemptive version of the open shop problem. They state as an obvious open question whether there exists an algorithm for open shop scheduling that gives a performance guarantee better than 5.83. Our 5.06 algorithm for data migration proves the existence of such an algorithm. Crucial to our improved result is a property of the linear programming relaxation for the problem. Similar linear programs have been used for various other scheduling problems. Our technique may be useful in obtaining improved results for these problems as well.

1. INTRODUCTION

The *data migration* problem arises in large storage systems, such as *Storage Area Networks* [Khuller et al. 2004], where a dedicated network of disks is used to store multimedia data. As the data access pattern changes over time, the load across the disks needs to be rebalanced so as to continue providing efficient service. This is done by computing a new data layout and then "migrating" data to convert the initial data layout to the target data layout. While migration is being performed, the storage system is running suboptimally, therefore it is important to compute a data migration schedule that converts the initial layout to the target layout quickly.

This problem can be modeled as a *transfer graph* [Kim 2005], in which the vertices represent the storage disks and an edge between two vertices u and v corresponds to a data object that must be transferred from u to v, or vice-versa. Each edge has a processing time (or length) that represents the transfer time of a data object between the disks corresponding to the end points of the edge. An important constraint is that any disk can be involved in at most one transfer at any time.

Several variations of the data migration problem have been studied. These variations arise either due to different objective functions or due to additional constraints. One common objective function is to minimize the *makespan* of the migration schedule, i.e., the time by which all migrations complete. Coffman et al. [1985] introduced this problem. They showed that when edges may have arbitrary lengths, a class of greedy algorithms yields a 2-approximation to the minimum makespan. In the special case where the edges have equal (unit) lengths, the prob-

lem reduces to edge coloring of the transfer (multi)graph of the system for which an asymptotic approximation scheme is now known [Sanders and Steurer 2005].

Hall et al. [2001] studied the data migration problem with unit edge lengths and *capacity constraints*; that is, the migration schedule must respect the storage constraints of the disks. The paper gives a simple 3/2-approximation algorithm for the problem. Hall et al. [2001] and Anderson et al. [2001] also present approximation algorithms for the makespan minimization problem with the following constraints: (*i*) data can only be moved, i.e, no new copies of a data object can be created, (*ii*) additional nodes can assist in data transfers, and (*iii*) each disk has a unit of spare storage. Khuller et al. [2004] solved a more general problem, where each data object can also be copied. They gave a constant factor approximation algorithm for the problem.

Another objective function is to minimize the *average completion time* over all data migrations. This corresponds to minimizing the average edge completion time in the transfer graph. For the case of unit edges lengths, Bar-Noy et al. [1998] showed that the problem is NP-hard and gave a simple 2-approximation algorithm. For arbitrary edge lengths, Halldórsson et al. [2003] gave a 12-approximation algorithm for the problem. This was improved to 10 by Kim [2005].

In this paper, we study the data migration problem with the objective of minimizing the average completion time over all storage disks. Indeed, this objective favors the individual storage devices, which are often geographically distributed over a large network. It is therefore natural to try and minimize the average amount of time that each of these (independent) devices is involved in the migration process. For the case where vertices have arbitrary weights, and the edges have unit length, Kim [2005] proved that the problem is NP-hard and showed that the list scheduling algorithm [Graham 1966], when guided by an optimal solution to a linear programming relaxation, gives an approximation ratio of 3. Kim also gave a 9-approximation algorithm for the case where edges have arbitrary lengths. We show that the analysis of the 3-approximation algorithm is tight, and for the case where edges have release times and arbitrary lengths, we give a 5.06-approximation algorithm.

A problem related to the data migration problem is non-preemptive *open shop scheduling*, denoted by $O|r_j|\sum w_j C_j$ in the standard three-field notation [Lawler et al. 1993]. In this problem, we have a set of jobs, $\mathcal{J}$, and a set of machines $M_1, \ldots, M_m$. Each job $J_j \in \mathcal{J}$ consists of a set of m operations: $o_{j,i}$ has the processing time $p_{j,i}$ and must be processed on M_i, $1 \le i \le m$. Each machine can process a single operation at any time, and two operations that belong to the same job cannot be processed simultaneously. Also, each job J_j has a positive weight, w_j, and a release time, r_j, which means that no operation of J_j can start before r_j. The objective is to minimize the sum of weighted completion times of all jobs. This problem is MAX-SNP hard [Hoogeveen et al. 1998]. Chakrabarti et al. [1996] gave a $(5.78 + \epsilon)$-approximation algorithm for the case where the number of machines, m, is some fixed constant. They also gave a $(2.89 + \epsilon)$-approximation algorithm for the preemptive version of the problem and fixed number of machines. For arbitrary number of machines, Queyranne and Sviridenko [2002b] presented algorithms that yield approximation factors of 5.83 and 3 for the non-preemptive and preemptive

versions of the problems, respectively. The approximation factor for the preemptive version was subsequently improved to $(2 + \epsilon)$ [Queyranne and Sviridenko 2002a].

Our Contribution: Since the open shop scheduling problem is a special case of the data migration problem all of our positive results for data migration apply to open shop scheduling. Note that the MAX-SNP hardness of the data migration problem follows from the MAX-SNP hardness of open shop scheduling [Hoogeveen et al. 1998]. Our main result is a 5.06-approximation algorithm for the data migration problem with arbitrary edge lengths. Our algorithm is based on rounding a solution of a linear programming (LP) relaxation of the problem. The general idea of our algorithm is that edges have to wait before they are actually processed (i.e., data transfer begins). This idea has been used for other scheduling problems as well [Chekuri et al. 2000; Halldórsson et al. 2003; Skutella and Uetz 2005]. Even though the high-level idea is similar to the one used by Halldórsson et al. [2003], there are subtle differences that are crucial to the improved results that we present here. Our method combines solutions obtained by using two different wait functions. It is interesting to note that each solution (when all edges are released at time zero) is a 5.83-approximate solution, which is the approximation ratio obtained by Queyranne and Sviridenko [2002b]. In their algorithm artificial precedence constraints are defined and after that a schedule is obtained by a greedy algorithm whereas in our work artificial wait function for the edges (edges have to wait for some time before being processed) is defined after which a schedule is obtained using a greedy approach. To obtain an approximation ratio better than 5.83, we crucially use a property of the LP relaxation that we prove in Lemma 3.1. Although the LP relaxation has been used earlier [Wolsey 1985; Queyranne 1993; Schulz 1996; Hall et al. 1997; Queyranne and Sviridenko 2002b; Kim 2005], we are not aware of any previous work that uses such a property of the LP. Our technique may be useful for deriving improved results for other shop scheduling problems.

For the case where edges have unit lengths, we show, by giving a tight example, that the list scheduling analysis of Kim [2005] is tight. This illustrates the limitations of the LP relaxation. Finally, we study the open shop problem under *operations completion time* criteria (cf. [Queyranne and Sviridenko 2002b]); that is, we sum the completion times of all operations for every job. For the special case of unit length operations with arbitrary non-negative weights, we show that an algorithm of Halldórsson et al. [2003] yields a 1.796 approximation algorithm for the problem.

2. RELATION OF DATA MIGRATION AND OPEN SHOP SCHEDULING

In this section, we formally state the data migration and open shop scheduling problems and show that the latter is a special case of the former.

Data Migration Problem: We are given a graph $G = (V, E)$. The vertices represent storage devices, and the edges correspond to data transmissions among the devices. We denote by $E(u)$ the set of edges incident on a vertex u. Each vertex v has weight w_v and processing time 0. Each edge e has a length, or processing time, p_e. Moreover, each edge e can be processed only after its release time r_e. All release times and processing times are non-negative integers. The completion time

of an edge is simply the time at which its processing is completed. Each vertex v can complete only after all the edges in $E(v)$ are completed. Since each vertex v has the processing time 0, the completion time, C_v, of v is the latest completion time of any edge in $E(v)$. The crucial constraint is that two edges incident on the same vertex cannot be processed at the same time. The objective is to minimize $\sum_{v \in V} w_v C_v$.

Open Shop Scheduling Problem: We are given a set of jobs $\mathcal{J} = \{J_1, \ldots, J_n\}$, to be scheduled on a set of machines $\mathcal{M} = \{M_1, \ldots, M_m\}$. Each job J_j has a non-negative weight w_j; also, J_j consists of a set of m operations $o_{j,1}, \ldots, o_{j,m}$, with the corresponding processing times $p_{j,i}$, $1 \leq i \leq m$; the operation $o_{j,i}$ must be processed on the machine M_i. Each machine can process at most one operation at any time, and no two operations belonging to the same job can be processed simultaneously. The completion time C_j of each job J_j is the latest completion time of any of its operations. The objective is to minimize $\sum_{J_j \in \mathcal{J}} w_j C_j$.

The open shop scheduling problem is a special case of the data migration problem, as shown by the following reduction. Given an instance of the open shop scheduling problem, construct a bipartite graph $B = (J, M, F)$ as follows. Each vertex $j_\ell \in J$ represents a job $J_\ell \in \mathcal{J}$, and each vertex $m_i \in M$ represents a machine $M_i \in \mathcal{M}$. The edge $(j_\ell, m_i) \in F$ with processing time $p_{\ell,i}$ corresponds to the operation $o_{\ell,i}$, $1 \leq i \leq m$. Assign $w_{m_i} = 0$ to each vertex $m_i \in M$, and $w_{j_\ell} = w_\ell$ (i.e., the weight of the job J_ℓ) to each vertex $j_\ell \in J$. It is now easy to verify that any data migration schedule for B is a valid solution for the corresponding open shop problem.

In the remainder of the paper, we consider only the data migration problem, with the understanding that all of our results apply to open shop scheduling.

3.   A LINEAR PROGRAMMING RELAXATION

The linear programming relaxation for the data migration problem (without release times) was given by Kim [2005]. Such relaxations have been proposed earlier by Wolsey [1985] and Queyranne [1993] for single machine scheduling problems and by Schulz [1996] and Hall et al. [1997] for parallel machines and flow shop problems. For the sake of completeness, we state below the LP relaxation for the data migration problem.

For an edge e (vertex v) let C_e (C_v) be the variable that represents the completion time of e (resp., v) in the LP relaxation. For any set of edges $S \subseteq E$, let $p(S) = \sum_{e \in S} p_e$ and $p(S^2) = \sum_{e \in S} p_e^2$.

$$(LP) \qquad \text{minimize} \quad \sum_{v \in V} w_v C_v \qquad (1)$$

$$\text{subject to:} \qquad C_v \geq r_e + p_e, \qquad \forall v \in V, e \in E(v) \qquad (2)$$

$$C_v \geq p(E(v)), \qquad \forall v \in V \qquad (3)$$

$$C_v \geq C_e, \qquad \forall v \in V, e \in E(v) \qquad (4)$$

$$\sum_{e \in S(v)} p_e C_e \geq \frac{p(S(v))^2 + p(S(v)^2)}{2}, \qquad \forall v \in V, S(v) \subseteq E(v) \qquad (5)$$

$$C_e \geq 0, \qquad \forall e \in E \qquad (6)$$

$$C_v \geq 0, \qquad \forall v \in V \qquad . \qquad (7)$$

The set of constraints represented by (2), (3), and (4) are due to the different lower bounds on the completion times of a vertex. The justification for constraints (5) is as follows. By the problem definition, no two edges incident on the same vertex can be scheduled at the same time. Consider any ordering of the edges in $S(v) \subseteq E(v)$. If an edge $e \in S(v)$ is the jth edge to be scheduled among the edges in $S(v)$ then, setting $C_j = C_e$ and $p_j = p_e$, we get

$$\sum_{j=1}^{|S(v)|} p_j C_j \geq \sum_{j=1}^{|S(v)|} p_j \sum_{k=1}^{j} p_k = \sum_{j=1}^{|S(v)|} \sum_{k=1}^{j} p_j p_k = \frac{p(S(v))^2 + p(S(v)^2)}{2}.$$

Although there are exponentially many constraints, the above LP can be solved in polynomial time via the ellipsoid algorithm [Queyranne 1993].

3.1 A Property of the LP

In this section, we state and prove a property of the LP that plays a crucial role in the analysis of our algorithm. Let $X(v, t_1, t_2) \subseteq E(v)$ denote the set of edges that complete in the time interval $(t_1, t_2]$ in the LP solution (namely, their fractional value belongs to this interval). Hall et al. [1997] showed that $p(X(v, 0, t)) \leq 2t$. In Lemma 3.1 we prove a stronger property of a solution given by the above LP. Intuitively, our property states that if too many edges complete early, then the completion times of other edges must be delayed. For example, as a consequence of our property, for any $t > 0$ if $p(X(v, 0, t/2)) = t$ then $p(X(v, t/2, t)) = 0$, which means that no edges in $E(v) \setminus X(v, 0, t/2)$ complete before t in the LP solution. We now formally state and prove the lemma.

LEMMA 3.1. *Consider a vertex v and times $t_1 > 0$ and $t_2 \geq t_1$. If $p(X(v, 0, t_1)) = \lambda_1$ and $p(X(v, t_1, t_2)) = \lambda_2$, then λ_1 and λ_2 are related by*

$$\lambda_2 \leq t_2 - \lambda_1 + \sqrt{t_2^2 - 2\lambda_1(t_2 - t_1)}$$

PROOF. Using the constraint (5) of the LP relaxation for vertex v, we get

$$\frac{p(X(v,0,t_2))^2}{2} \leq \sum_{e \in X(v,0,t_1)} p_e C_e + \sum_{e \in X(v,t_1,t_2)} p_e C_e$$

$$\leq p(X(v,0,t_1))t_1 + p(X(v,t_1,t_2))t_2$$

$$\therefore (\lambda_1 + \lambda_2)^2 \leq 2\lambda_1 t_1 + 2\lambda_2 t_2$$

$$\therefore (\lambda_1 + \lambda_2 - t_2)^2 \leq t_2^2 - 2\lambda_1(t_2 - t_1)$$

$$\therefore \lambda_2 \leq t_2 - \lambda_1 + \sqrt{t_2^2 - 2\lambda_1(t_2 - t_1)}$$

□

The following result [Hall et al. 1997] follows from Lemma 3.1 by substituting $t_1 = 0$, $\lambda_1 = 0$, and $t_2 = t$.

COROLLARY 3.2. *For any vertex v and time $t \geq 0$, $p(X(v,0,t)) \leq 2t$.*

4. ALGORITHM

Note that if an edge has processing time zero, it can be processed as soon as it is released, without consuming any time steps. Hence, without loss of generality, we assume that the processing time of each edge is a positive integer.

The algorithm is parameterized by a *wait function* $W : E \to \mathcal{R}^+$. The idea is that each edge e must *wait* for W_e ($W_e \geq r_e$) time steps before it can actually start processing. The algorithm processes the edges (that have waited enough time) in non-decreasing order of their completion times in the LP solution. When e is being processed, we say that e is *active*. Once it becomes active, it remains active for p_e time steps, after which it is *finished*. A not-yet-active edge can be *waiting* only if none of its neighboring edges are active; otherwise, it is said to be *delayed*. Thus, at any time, an edge is in one of four modes: *delayed, waiting, active,* or *finished*. When adding new active edges, among those that have done their waiting duty, the algorithm uses the LP completion time as priority.

The precise rules are given in the pseudocode in Fig. 1. Let $wait(e,t)$ denote the number of time steps that e has waited until the end of time step t. Let $Active(t)$ be the set of active edges during time step t. Let $\widetilde{C_e}$ ($\widetilde{C_u}$) be the completion time of edge e (vertex u) in our algorithm. The algorithm in Fig. 1, implemented as is, would run in pseudo-polynomial time; however, it is easy to implement the algorithm in strongly polynomial time, by increasing t in each iteration by the smallest remaining processing time of an active edge.

One property of our processing rules, that distinguishes it from the wait functions used by Halldórsson et al. [2003] for the sum of edge completion times, is that multiple edges touching the same vertex can wait at the same time.

We run the algorithm for two different wait functions W and choose the better of the two solutions. For any vertex v (edge e), let C_v^* (C_e^*) be its completion time in the optimal LP solution. In the first wait function, for each edge e we choose $W_e = \lfloor \beta_1 C_e^* \rfloor$, $\beta_1 \geq 1$ and in the second one, we choose $W_e = \lfloor \beta_2 \max\{r_e, p(S_e(u)), p(S_e(v))\} \rfloor$, where $S_e(u) = \{f | f \in E(u), C_f^* \leq C_e^*\}$ and $\beta_2 \geq 1$. Note that the choice of wait functions ensures that the edges become

SCHEDULE($G = (V, E), W$)

```
1     Solve the LP relaxation for the given instance optimally.
2     t ← 0
3     Finished ← Active(t) ← ∅
4     for each e ∈ E do
5          wait(e, t) ← 0
6     while (Finished ≠ E) do
7          t ← t + 1
8          Active(t) ← {e | e ∈ Active(t − 1) and e ∉ Active(t − pₑ)}
9          for each edge e ∈ Active(t − 1) \ Active(t) do
10              C̃ₑ ← t − 1
11              Finished ← Finished ∪ {e}              // e is finished
12          for each edge e = (u, v) ∈ E \ (Active(t) ∪ Finished)
13              in non-decreasing order of LP completion time do
14              if (Active(t) ∩ (E(u) ∪ E(v)) = ∅) and (wait(e, t − 1) = Wₑ) then
15                  Active(t) ← Active(t) ∪ {e}          // e is active
16          for each edge e = (u, v) ∈ E \ (Active(t) ∪ Finished) do
17              if (Active(t) ∩ (E(u) ∪ E(v)) = ∅) then
18                  wait(e, t) ← wait(e, t − 1) + 1      // e is waiting
19              else wait(e, t) ← wait(e, t − 1)          // e is delayed
20    return C̃
```

Fig. 1. Algorithm for Data Migration

active only after they are released. When all release times are 0, we can choose β_1 and β_2 such that $\beta_1 > 0$ and $\beta_2 > 0$.

5. ANALYSIS

Consider a vertex x and an edge $e = (x, y)$, and recall that C_x^* and C_e^* are their completion times in the optimal LP solution. Let $B_e(x) = \{f | f \in E(x), C_f^* > C_e^*, \widetilde{C_f} < \widetilde{C_e}\}$, i.e., edges in $E(x)$ that finish after e in the LP solution, but finish before e in our algorithm. Recall that $S_e(x) = \{f | f \in E(x), C_f^* \leq C_e^*\}$. Note that $e \in S_e(x)$. Let $\overline{S_e(x)} = S_e(x) \backslash \{e\}$. By constraint (3), we have $p(S_e(x)) + p(B_e(x)) \leq C_x^*$.

We analyze our algorithm separately for the two wait functions defined in Section 4. In each case, we analyze the completion time of an arbitrary but fixed vertex $u \in V$. Without loss of generality, let $e_u = (u, v)$ be the edge that finishes last among the edges in $E(u)$. By constraint (4), we have $C_{e_u}^* \leq C_u^*$, since the LP solution we are considering is optimal, we can assume that $C_{e_u}^* = C_u^*$. We analyze our algorithm for the case where all edges in $\overline{S_{e_u}(u)} \cup \overline{S_{e_u}(v)}$ finish before e_u in our algorithm. If this is not true then our results can only improve. Let

$$p(S_{e_u}(v)) = \lambda_{e_u} C_{e_u}^* = \lambda_{e_u} C_u^*, \qquad 0 \leq \lambda_{e_u} \leq 2 \qquad (8)$$

The upper bound on λ_{e_u} follows from Corollary 3.2.

For ease of reference, in Table I we summarize the key notation used in the rest of this section. Lemma 5.1 gives an upper bound on the completion time of any vertex in our algorithm for any wait function. Lemma 5.2 and Lemma 5.3 use Lemma 5.1 to give upper bounds for the two wait functions.

LEMMA 5.1. $\widetilde{C_u} \leq W_{e_u} + C_u^* + p(S_{e_u}(v)) + p(B_{e_u}(v))$

Symbol	Description
C_e^* (C_v^*)	completion time of edge e (vertex v) in the LP solution.
$\widetilde{C_e}$ ($\widetilde{C_u}$)	completion time of edge e (vertex v) in our algorithm.
$p(Z)$	$\sum_{e \in Z} p_e$.
r_e	release time of edge e.
$E(u)$	set of edges incident on vertex u.
$S_e(u)$	$\{f \mid f \in E(u) \text{ and } C_f^* \leq C_e^*\}$
$\overline{S_e(u)}$	$S_e(u) \setminus \{e\}$.
$B_e(u)$	$\{f \mid f \in E(u) \text{ and } C_f^* > C_e^* \text{ and } \widetilde{C_f} < \widetilde{C_e}\}$, i.e., edges in $E(u)$ that finish after e in the LP solution, but finish before e in our algorithm.
W_e	waiting time of edge e.
e_u	edge that finishes last in our algorithm among all edges in $E(u)$.
λ_{e_u}	a constant in $[0, 2]$ used to define $p(S_{e_u}(v))$ where $e_u = (u, v)$.

Table I.   Summary of notation

PROOF. Observe that when e_u is in delayed mode it must be that some edge in $\overline{S_{e_u}(u)} \cup B_{e_u}(u) \cup \overline{S_{e_u}(v)} \cup B_{e_u}(v)$ must be active. Hence, we have

$$\widetilde{C_{e_u}} \leq W_{e_u} + p(S_{e_u}(u)) + p(B_{e_u}(u)) + p(S_{e_u}(v)) + p(B_{e_u}(v))$$
$$\therefore \widetilde{C_u} \leq W_{e_u} + C_u^* + p(S_{e_u}(v)) + p(B_{e_u}(v))$$

□

Define $f(\beta_1, \lambda) = \beta_1 + 1 + \frac{\beta_1 + 1}{\beta_1} + \sqrt{\left(\frac{\beta_1 + 1}{\beta_1}\right)^2 - \frac{2\lambda}{\beta_1}}$.

LEMMA 5.2. *If* $W_{e_u} = \lfloor \beta_1 C_{e_u}^* \rfloor$ *then* $\widetilde{C_u} \leq f(\beta_1, \lambda_{e_u}) C_u^*$

PROOF. Let $e_b \in B_{e_u}(v)$ be the edge with the largest completion time in the LP solution among all the edges in $B_{e_u}(v)$. Note that when e_b is in waiting mode it must be that either e_u is waiting or an edge in $\overline{S_{e_u}(u)} \bigcup B_{e_u}(u)$ is active (i.e., e_u is delayed). Thus, we get

$$W_{e_b} \leq W_{e_u} + p(\overline{S_{e_u}(u)}) + p(B_{e_u}(u)).$$

Hence, we have that

$$\lfloor \beta_1 C_{e_b}^* \rfloor \leq \lfloor \beta_1 C_{e_u}^* \rfloor + p(S_{e_u}(u)) - p_{e_u} + p(B_{e_u}(u)).$$

Since $p_{e_u} \geq 1$, it follows that $\beta_1 C_{e_b}^* - 1 \leq \beta_1 C_u^* + C_u^* - 1$, and $C_{e_b}^* \leq \frac{\beta_1 + 1}{\beta_1} C_u^*$.

Note that in the LP solution, each edge in $S_{e_u}(v)$ finishes at or before C_u^* and each edge in $B_{e_u}(v)$ finishes after C_u^* and at or before $C_{e_b}^*$. Thus, by substituting $t_1 = C_u^*$, $t_2 = \frac{\beta_1 + 1}{\beta_1} C_u^*$, $\lambda_1 = p(S_{e_u}(v)) = \lambda_{e_u} C_u^*$, and $\lambda_2 = p(B_{e_u}(v))$ in Lemma 3.1 we get

$$p(B_{e_u}(v)) \leq \left(\frac{\beta_1 + 1}{\beta_1} - \lambda_{e_u} + \sqrt{\left(\frac{\beta_1 + 1}{\beta_1}\right)^2 - 2\lambda_{e_u}\left(\frac{\beta_1 + 1}{\beta_1} - 1\right)}\right) C_u^*$$

$$= \left(\frac{\beta_1 + 1}{\beta_1} - \lambda_{e_u} + \sqrt{\left(\frac{\beta_1 + 1}{\beta_1}\right)^2 - \frac{2\lambda_{e_u}}{\beta_1}}\right) C_u^*$$

Then, using (8), we have that

$$p(S_{e_u}(v)) + p(B_{e_u}(v)) \leq \left(\frac{\beta_1 + 1}{\beta_1} + \sqrt{\left(\frac{\beta_1+1}{\beta_1}\right)^2 - \frac{2\lambda_{e_u}}{\beta_1}} \right) C_u^*$$

The lemma now follows from Lemma 5.1 and the fact that $C_{e_u}^* = C_u^*$. □

Define $h(\beta_2, \lambda) = (\beta_2 + 1)\max\{1, \lambda\} + \frac{\beta_2+1}{\beta_2}$.

LEMMA 5.3. *If* $W_{e_u} = \lfloor \beta_2 \max\{r_{e_u}, p(S_{e_u}(u)), p(S_{e_u}(v))\} \rfloor$ *then* $\widetilde{C_u} \leq h(\beta_2, \lambda_{e_u})C_u^*$

PROOF. By constraints (2) and (4) $r_{e_u} \leq C_{e_u}^* = C_u^*$, and from constraint (3), $p(S_{e_u}(u)) \leq C_u^*$. Also, recall that $p(S_{e_u}(v)) = \lambda_{e_u}C_u^*$, $0 \leq \lambda_{e_u} \leq 2$. Hence,

$$W_{e_u} \leq \lfloor \beta_2 \max\{C_u^*, \lambda_{e_u} C_u^*\} \rfloor. \tag{9}$$

In the following we upper bound $p(S_{e_u}(v)) + p(B_{e_u}(v))$. Let $z \in \overline{S_{e_u}(v)} \cup B_{e_u}(v)$ be the edge with the *largest waiting time*, i.e., $W_z = \max_{f \in \overline{S_{e_u}(v)} \cup B_{e_u}(v)}\{W_f\}$. When z is in waiting mode it must be that either e_u is waiting or an edge in $\overline{S_{e_u}(u)} \bigcup B_{e_u}(u)$ is active. Thus, using (9), we get

$$\begin{aligned} W_z &\leq W_{e_u} + p(\overline{S_{e_u}(u)}) + p(B_{e_u}(u)) \\ &\leq \lfloor \beta_2 \max\{C_u^*, \lambda_{e_u}C_u^*\}\rfloor + p(S_{e_u}(u)) - p_{e_u} + p(B_{e_u}(u)) \\ &\leq \beta_2 \max\{C_u^*, \lambda_{e_u}C_u^*\} + C_u^* - 1. \end{aligned} \tag{10}$$

Let l be the edge with the *largest completion time* in the LP solution among the edges in $\overline{S_{e_u}(v)} \cup B_{e_u}(v)$, i.e., $C_l^* = \max_{f \in \overline{S_{e_u}(v)} \cup B_{e_u}(v)}\{C_f^*\}$. Since $W_l \leq W_z$, we have

$$\lfloor \beta_2(p(S_{e_u}(v)) + p(B_{e_u}(v)))\rfloor = \lfloor \beta_2 \cdot p(S_l(v))\rfloor \leq W_l \leq W_z. \tag{11}$$

Combining (10) and (11) we get

$$\lfloor \beta_2(p(S_{e_u}(v)) + p(B_{e_u}(v)))\rfloor \leq \beta_2 \max\{C_u^*, \lambda_{e_u}C_u^*\} + C_u^* - 1.$$

Hence,

$$p(S_{e_u}(v)) + p(B_{e_u}(v)) \leq \left(\max\{1, \lambda_{e_u}\} + \frac{1}{\beta_2}\right)C_u^*.$$

Now, from Lemma 5.1 we have that

$$\widetilde{C_u} \leq W_{e_u} + C_u^* + (\max\{1, \lambda_{e_u}\} + \frac{1}{\beta_2})C_u^*,$$

and from (9),

$$\widetilde{C_u} \leq \beta_2 C_u^* \cdot \max\{1, \lambda_{e_u}\} + C_u^* + C_u^*(\max\{1, \lambda_{e_u}\} + \frac{1}{\beta_2}).$$

This gives the statement of the lemma. □

Combining the two solutions: To obtain an improved performance ratio, we run our algorithm with each of the two wait functions and return the better result.

The cost of that solution is, by Lemmas 5.2 and 5.3, bounded by

$$ALG \leq \min \left(\sum_{v \in V} f(\beta_1, \lambda_v) C_v^*, \sum_{v \in V} h(\beta_2, \lambda_v) C_v^* \right).$$

This expression is unwieldy for performance analysis, but can fortunately be simplified. The following lemma allows us to assume without loss of generality that there are only two different λ-values: 0, or a constant $\hat{\lambda}$ in the range $[1, 2]$. To analyze the performance ratio, we then numerically optimize a function over three variables: β_1, β_2, and $\hat{\lambda}$.

We first introduce some notation. Let $OPT^* = \sum_v C_v^*$ be the cost of the optimal LP solution. Partition V into $V_0 = \{v \in V : \lambda_v \in [0, 1]\}$ and $V_1 = \{v \in V : \lambda_v \in (1, 2]\}$. Let $w = (\sum_{v \in V_0} C_v^*)/OPT^*$ be the contribution of V_0 to the LP solution. Given an instance with values λ_v, LP solution vector C_v^*, and the notation above, we define two functions: Let $F(t) = [w \cdot f(\beta_1, 0) + (1 - w) \cdot f(\beta_1, t)]$ and $H(t) = [w \cdot h(\beta_2, 0) + (1 - w) \cdot h(\beta_2, t)]$.

LEMMA 5.4. *There is a $\hat{\lambda} \in [1, 2]$ such that $ALG \leq \min(H(\hat{\lambda}), F(\hat{\lambda})) OPT^*$.*

PROOF. Since h is constant on $[0, 1]$ and f is decreasing, $h(\beta_2, \lambda_v) = h(\beta_2, 0)$ and $f(\beta_1, \lambda_v) \leq f(\beta_1, 0)$, for $v \in V_0$. Let $\hat{\lambda} = \left(\sum_{v \in V_1} \lambda_v C_v^* \right)/|V_1|$ be the average λ-value of nodes in V_1 weighted by their LP value. Since h is linear on $[1, 2]$, we have that $\sum_{v \in V_1} h(\beta_2, \lambda_v) C_v^* = \sum_{v \in V_1} h(\beta_2, \hat{\lambda}) C_v^*$, and since f is convex, $\sum_{v \in V_1} f(\beta_1, \lambda_v) C_v^* \leq \sum_{v \in V_1} f(\beta_1, \hat{\lambda}) C_v^*$. It follows that

$$ALG \leq \sum_{v \in V} f(\beta_1, \lambda_v) C_v^* \leq \sum_{v \in V_0} f(\beta_1, 0) C_v^* + \sum_{v \in V_1} f(\beta_1, \hat{\lambda}) C_v^* = F(\hat{\lambda}) OPT^*.$$

The same type of bound holds for H, establishing the lemma. □
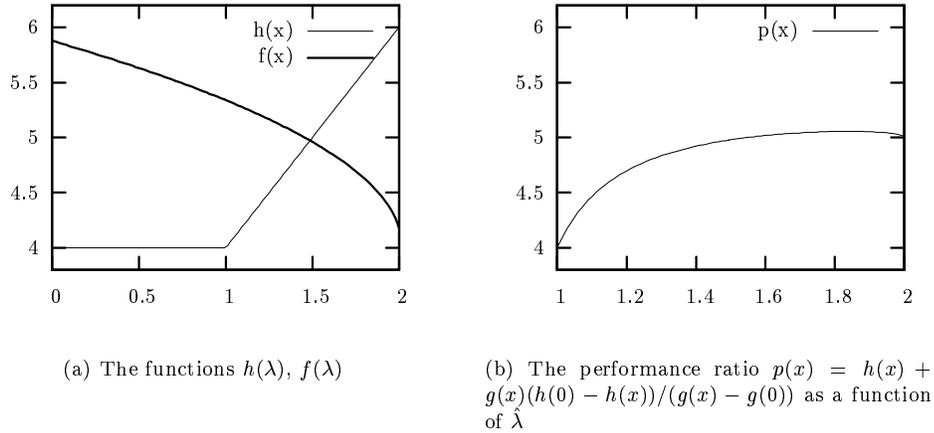
The worst case occurs when $h'(t) = f'(t)$; let $\hat{\lambda}$ denote that value of t. We can then determine the value of w from the other variables. Namely, defining $g(x) = g_{\beta_1, \beta_2}(x) = h(\beta_2, x) - f(\beta_1, x)$, we have $w = g(\hat{\lambda})/(g(\hat{\lambda}) - g(0))$. The performance ratio of the algorithm is then bounded by

$$\rho \leq \max_{\hat{\lambda} \in [1, 2]} \left(h(\hat{\lambda}) + g(\hat{\lambda}) \frac{h(0) - h(\hat{\lambda})}{g(\hat{\lambda}) - g(0)} \right).$$

We find the best choice of parameters β_1 and β_2 numerically. When the release times can be non-zero, we need to ensure that each operation e does not begin executing before its release time. This is satisfied if the β-values are at least 1, ensuring that $W_e \geq r_e$. Setting $\beta_1 = 1.177$ and $\beta_2 = 1.0$, the worst-case is then achieved at about $\hat{\lambda} = 1.838$, giving a ratio of $\rho \leq 5.0553$.

THEOREM 5.5. *There exists a 5.06-approximation algorithm for the data migration problem, as well as for the open shop scheduling problem.*

When all release times are zero, we can widen the search to all non-zero values. We then obtain a slightly improved ratio of 5.03, when choosing $\beta_1 = 1.125$ and $\beta_2 = 0.8$.

(a) The functions $h(\lambda)$, $f(\lambda)$

(b) The performance ratio $p(x) = h(x) + g(x)(h(0) - h(x))/(g(x) - g(0))$ as a function of $\hat{\lambda}$

Fig. 2. The performance functions when $\beta_1 = 1.177$ and $\beta_2 = 1.0$

6. UNIT PROCESSING TIMES

It is known that the open shop scheduling problem is polynomially solvable when all operations are of unit length (c.f. [Brucker 2001]). On the other hand, Kim [2005] showed that the more general data migration problem is NP-hard in the case of unit edge lengths. For such instances, when all edges are released at time zero, Kim [2005] showed that the list scheduling algorithm [Graham 1966], guided by an optimal solution to the LP relaxation[1] gives a 3-approximate solution; the algorithm is called *Ordered List Scheduling (OLS)*. The problem of obtaining a better than 3-approximate solution remained open. In Section 6.1, we show by giving a tight example that OLS cannot achieve a ratio better than 3. The tight example also illustrates the limitations of the LP solution. For the sake of completeness, we describe the OLS algorithm and its analysis here. The edges are sorted in non-decreasing order of their completion times in the LP solution. At any time, an edge $e = (u, v)$ is scheduled iff no edge in $S_e(u) \cup S_e(v)$ is scheduled at that time. (Recall that $S_e(u) = \{f | f \in E(u), C_f^* \leq C_e^*\}$.) For any vertex u, if e_u is the edge that finishes last among the edges in $E(u)$, and if $\widetilde{C_u}$ is the completion time of u in OLS, then $\widetilde{C_u} \leq p(S_{e_u}(u)) + p(S_{e_u}(v))$. Combining the fact that $p(S_{e_u}(u)) \leq C_u^*$ along with $p(S_{e_u}(v)) \leq 2C_{e_u}^* = 2C_u^*$ (Corollary 3.2), we get $\widetilde{C_u} \leq 3C_u^*$ and hence a 3-approximation ratio.

6.1 A tight example

Consider a tree rooted at vertex r. Let $S = \{s_1, s_2, \ldots, s_k\}$ be the children of r. For each vertex s_i, let $L_i = \{l_1^i, l_2^i, \ldots, l_h^i\}$ be the children of s_i. Let $L = \cup_{i=1}^{k} L_i$. Let $k = (n+1)/2$ and $h = n - 1$. For each vertex $u \in S$, let $w_u = \epsilon$ and for each vertex $v \in L$, let $w_v = 0$. Let $w_r = M$. For each edge e, let $C_e^* = (n+1)/2$ be its

---

[1]See Section 3.

completion time in the LP solution. For each vertex $v \in L \cup \{r\}$, $C_v^* = (n+1)/2$ and for each vertex $v \in S$, $C_v^* = n$. The completion times of vertices in L do not matter as the weights of all those vertices are zero. It is easy to verify that this is an optimal LP solution. The cost of the LP solution equals

$$w_r \left(\frac{n+1}{2} \right) + \sum_{i=1}^{k} w_{s_i} n \;=\; M \left(\frac{n+1}{2} \right) + k\epsilon n = M \left(\frac{n+1}{2} \right) + \left(\frac{n(n+1)}{2} \right) \epsilon \;.$$

OLS could process the edges in the following order. At any time t, $1 \leq t \leq n-1$, OLS processes all edges in $\{(s_1, l_t^1), (s_2, l_t^2), \ldots, (s_k, l_t^k)\}$. At time $t = n + z$, $0 \leq z < (n+1)/2$, OLS processes edge (r, s_{z+1}). The cost of the solution in OLS is at least

$$w_r \left(n - 1 + \frac{n+1}{2} \right) + \sum_{i=1}^{k} w_{s_i} n = M \left(\frac{3n-1}{2} \right) + \left(\frac{n(n+1)}{2} \right) \epsilon \;.$$

For large n, if $M \gg \epsilon$, the ratio of the cost of the OLS solution to the cost of the LP solution approaches 3.

6.2   Open shop and sum of operation completion times

Consider now the open shop problem, where each operation has unit processing time and a non-negative weight, and the objective is to minimize the weighted sum of completion times of all operations. Even though the problem of minimizing the sum of job completion times is polynomially solvable for open shop scheduling, minimizing the weighted sum of completion times of unit operations is NP-hard. This holds even for uniform weights as recently proved by Marx [2004]. We relate this problem to a result of Halldórsson et al. [2003] for the *sum coloring* problem. The input to sum coloring is a graph G, where each vertex corresponds to a unit length job. We need to assign a positive integer (color) to each vertex (job) so as to minimize the sum of the colors over all vertices. The constraint is that adjacent vertices receive distinct colors. In the weighted case, each vertex (job) is associated with a non-negative weight, and the goal is to minimize the weighted sum of the vertex colors.

In the *maximum k-colorable subgraph* problem, we are given an undirected graph G and a positive integer k; we need to find a maximum size subset $U \subseteq V$ such that $G[U]$, the graph induced by U, is k-colorable. In the weighted version, each vertex has a non-negative weight and we seek a maximum weight k-colorable subgraph. The following theorem is proved by Halldórsson et al. [2003].

THEOREM 6.1. *The weighted sum coloring problem admits a 1.796 ratio approximation algorithm on graphs for which the maximum weight k-colorable subgraph problem is polynomially solvable.*

We can relate this theorem to the above variant of the open shop problem, by defining the bipartite graph $B = (J, M, F)$ (see in Section 2) and setting $G = L(B)$, i.e., G is the line graph of B. Recall that in $L(B)$ the vertices are the edges of B; two vertices are neighbors if the corresponding edges in B share a vertex.

In order to apply Theorem 6.1, we need to show that the maximum weight k-colorable subgraph problem is polynomial on $L(B)$. Note that this is the problem

of finding a maximum weight collection of edges in B that is k-colorable (i.e., can be decomposed into k disjoint matchings in B). Observe that, on bipartite graphs, this problem is equivalent to the well-known weighted *b-matching* problem. In weighted b-matching, we seek a maximum weight set of edges that induces a subgraph of maximum degree at most k. Recall that a bipartite graph always admits a matching touching every vertex of maximum degree (c.f. [Gabow and Kariv 1982]). It follows, that the chromatic index of a bipartite graph is equal to its maximum degree. Since weighted b-matching is solvable in polynomial time (c.f. [Cook et al. 1998]), the same holds for the weighted k-colorable subgraph problem on $L(B)$. Hence, we have shown

THEOREM 6.2. *Open shop scheduling of unit jobs, under weighted sum of operation completion time criteria, admits a* 1.796 *ratio approximation.*

REFERENCES

ANDERSON, E., HALL, J., HARTLINE, J., HOBBES, M., KARLIN, A., SAIA, J., SWAMINATHAN, R., AND WILKES, J. 2001. An experimental study of data migration algorithms. In *Proceedings of the Workshop on Algorithm Engineering*. Springer-Verlag, 145–158.

BAR-NOY, A., BELLARE, M., HALLDÓRSSON, M. M., SHACHNAI, H., AND TAMIR, T. 1998. On chromatic sums and distributed resource allocation. *Information and Computation 140*, 183–202.

BRUCKER, P. 2001. *Scheduling Algorithms*. Springer Verlag.

CHAKRABARTI, S., PHILLIPS, C. A., SCHULZ, A. S., SHMOYS, D. B., STEIN, C., AND WEIN, J. 1996. Improved scheduling problems for minsum criteria. In *Proceedings of the Twenty-Third International Colloquium on Automata, Languages, and Programming*. Springer, 646–657.

CHEKURI, C., MOTWANI, R., NATARAJAN, B., AND STEIN, C. 2000. Approximation techniques for average completion time scheduling. *SIAM Journal on Computing 31*, 1, 146–166.

COFFMAN, E. G., GAREY, M. R., JOHNSON, D. S., AND LAPAUGH, A. S. 1985. Scheduling file transfers. *SIAM Journal on Computing 14*, 3, 744–780.

COOK, W. J., CUNNINGHAM, W. H., PULLEYBLANK, W. R., AND SCHRIJVER, A. 1998. *Combinatorial Optimization*. John Wiley and Sons.

GABOW, H. AND KARIV, O. 1982. Algorithms for edge coloring bipartite graphs and multigraphs. *SIAM Journal on Computing 11*, 1, 117–129.

GRAHAM, R. 1966. Bounds for certain multiprocessing anomalies. *Bell System Technical Journal 45*, 1563–1581.

HALL, J., HARTLINE, J., KARLIN, A., SAIA, J., AND WILKES, J. 2001. On algorithms for efficient data migration. In *Proceedings of the Twelfth ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 620–629.

HALL, L., SCHULZ, A. S., SHMOYS, D. B., AND WEIN, J. 1997. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Mathematics of Operations Research 22*, 513–544.

HALLDÓRSSON, M. M., KORTSARZ, G., AND SHACHNAI, H. 2003. Sum coloring interval graphs and k-claw free graphs with applications for scheduling dependent jobs. *Algorithmica 37*, 187–209.

HOOGEVEEN, H., SCHUURMAN, P., AND WOEGINGER, G. 1998. Non-approximability results for scheduling problems with minsum criteria. In *Proceedings of the Sixth International Conference on Integer Programming and Combinatorial Optimization*. Springer, 353–366.

KHULLER, S., KIM, Y., AND WAN, Y. C. 2004. Algorithms for data migration with cloning. *SIAM Journal on Computing 33*, 2, 448–461.

KIM, Y. 2005. Data migration to minimize the average completion time. *Journal of Algorithms 55*, 42–57.

LAWLER, E. L., LENSTRA, J. K., RINNOOY-KAN, A. H. G., AND SHMOYS, D. B. 1993. Sequencing and scheduling: Algorithms and complexity. *In S. C. Graves et al. Handbooks in Operations Research and Management Science, Vol. 4: Logistics of Production and Recovery*, 445–522.

MARX, D. 2004. Complexity results for minimum sum edge coloring. Unpublished manuscript.

QUEYRANNE, M. 1993. Structure of a simple scheduling polyhedron. *Mathematical Programming 58*, 263–285.

QUEYRANNE, M. AND SVIRIDENKO, M. 2002a. A $(2 + \epsilon)$-approximation algorithm for generalized preemptive open shop problem with minsum objective. *Journal of Algorithms 45*, 202–212.

QUEYRANNE, M. AND SVIRIDENKO, M. 2002b. Approximation algorithms for shop scheduling problems with minsum objective. *Journal of Scheduling 5*, 287–305.

SANDERS, P. AND STEURER, D. 2005. An asymptotic approximation scheme for multigraph edge coloring. In *Proceedings of the Sixteenth ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 897–906.

SCHULZ, A. S. 1996. Scheduling to minimize total weighted completion time: Performance guarantees of lp-based heuristics and lower bounds. In *Proceedings of the Fifth International Conference on Integer Programming and Combinatorial Optimization*. Springer, 301–315.

SKUTELLA, M. AND UETZ, M. 2005. Stochastic machine scheduling with precedence constraints. *SIAM Journal on Computing 34*, 4, 788–802.

WOLSEY, L. 1985. Mixed integer programming formulations for production planning and scheduling problems. Invited talk at the Twelfth International Symposium on Mathematical Programming, MIT, Cambridge.