

Near-Optimal Solutions and Large Integrality Gaps for Almost All Instances of Single-Machine Precedence-Constrained Scheduling

Andreas S. Schulz

Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology,
 Cambridge, Massachusetts 02139, schulz@mit.edu

Nelson A. Uhan

School of Industrial Engineering, Purdue University, West Lafayette, Indiana 47907, nuhan@purdue.edu

We consider the problem of minimizing the weighted sum of completion times on a single machine subject to bipartite precedence constraints in which all minimal jobs have unit processing time and zero weight, and all maximal jobs have zero processing time and unit weight. For various probability distributions over these instances—including the uniform distribution—we show several “almost all”-type results. First, we show that almost all instances are prime with respect to a well-studied decomposition for this scheduling problem. Second, we show that for almost all instances, every feasible schedule is arbitrarily close to optimal. Finally, for almost all instances, we give a lower bound on the integrality gap of various linear programming relaxations of this problem.

Key words: scheduling; linear ordering; probabilistic analysis; near-optimal solution; integrality gap

MSC2000 subject classification: Primary: 90B35, 90C27; secondary: 05C80, 06A07, 90C10

OR/MS subject classification: Primary: production/scheduling; sequencing, deterministic, single machine; secondary: mathematics: combinatorics

History: Received May 4, 2009; revised June 15, 2010, and December 6, 2010. Published online in *Articles in Advance* February 2, 2011.

1. Introduction. We consider the following classic scheduling problem. We have a set of jobs $N = \{1, \dots, n\}$ that needs to be scheduled nonpreemptively on a single machine that can process, at most, one job at a time. Each job $i \in N$ has a processing time $p_i \in \mathbb{R}_{\geq 0}$ and weight $w_i \in \mathbb{R}_{\geq 0}$. Precedence constraints are represented by an acyclic, transitively closed directed graph $G = (N, A)$: if $(i, j) \in A$, then job i must be processed before job j . The objective is to schedule these jobs in a way that respects the precedence constraints and minimizes the sum of weighted completion times. In the notation of Graham et al. [10], this problem is denoted as $1 | \text{prec} | \sum w_j C_j$.

The scheduling problem $1 | \text{prec} | \sum w_j C_j$ is strongly NP-hard (Lawler [14], Lenstra and Rinnooy Kan [15]). Currently, the best approximation algorithms all have a performance guarantee of 2 (Hall et al. [11], Chudak and Hochbaum [4], Chekuri and Motwani [3], Margot et al. [16]). On the inapproximability front, Ambühl et al. [1] showed that a PTAS is not possible, assuming NP-complete problems cannot be solved in randomized subexponential time. Bansal and Khot [2] showed that it is NP-hard to compute a $(2 - \varepsilon)$ -approximate schedule for any $\varepsilon > 0$, assuming a stronger version of the Unique Games Conjecture (Khot [13]) holds.

In this work, we focus on 0-1 *bipartite instances*. In a 0-1 bipartite instance (N_1, N_2, A) , the set of jobs is partitioned into $N = N_1 \dot{\cup} N_2$, and precedence constraints take the form of a directed bipartite graph $(N_1 \dot{\cup} N_2, A)$, in which $(i, j) \in A$ implies $i \in N_1$ and $j \in N_2$. The jobs in N_1 have unit processing time and zero weight, and the jobs in N_2 have zero processing time and unit weight. This scheduling problem on 0-1 bipartite instances can equivalently be viewed as a linear ordering problem on a mixed bipartite graph, in which there is an undirected edge between every pair of nodes $i \in N_1, j \in N_2$, for which $(i, j) \notin A$. The goal is to find an orientation B of the undirected edges, such that the resulting directed graph $(N_1 \cup N_2, A \cup B)$ is acyclic and has as few arcs that are directed from N_1 to N_2 as possible.

These 0-1 bipartite instances have further appeal than their simple combinatorial structure: it turns out that these simple instances effectively capture the inherent difficulty of $1 | \text{prec} | \sum w_j C_j$. Chekuri and Motwani [3] used a class of 0-1 bipartite instances to show that the linear programming relaxation in linear ordering variables attributable to Potts [19] has an integrality gap of 2. Moreover, Woeginger [23] showed that a ρ -approximation algorithm for 0-1 bipartite instances of $1 | \text{prec} | \sum w_j C_j$ implies a $(\rho + \varepsilon)$ -approximation algorithm for *arbitrary* instances of $1 | \text{prec} | \sum w_j C_j$; that is, the approximability behavior of 0-1 bipartite instances and arbitrary instances is virtually identical. In fact, the previously mentioned inapproximability result attributable to Bansal and Khot [2] was proved using 0-1 bipartite instances.

We study 0-1 bipartite instances of $1 | \text{prec} | \sum w_j C_j$ with a *probabilistic* lens. One appealing feature of 0-1 bipartite instances is that they are completely defined by their precedence constraints. Because precedence relations in bipartite partial orders are independent, we can apply the model of Erdős and Rényi [7] often used

in random graph theory to define classes of random 0-1 bipartite instances. Our analysis of these random 0-1 bipartite instances yields several “almost all”-type results:

- We show that almost all 0-1 bipartite instances are *non-Sidney-decomposable*. The decomposition technique of Sidney [21] splits an instance of $1 \mid \text{prec} \mid \sum w_j C_j$ into smaller instances so that the concatenation of optimal schedules for the smaller parts yields an optimal schedule for the entire instance. Together with the work of Chekuri and Motwani [3], Margot et al. [16], and Goemans and Williamson [9], our result also implies that for almost all 0-1 bipartite instances, any feasible schedule is a 2-approximation.
- Using two-dimensional Gantt charts, we show that for almost all 0-1 bipartite instances, *all feasible schedules are actually arbitrarily close to optimal*. In particular, we show that for any given $\varepsilon > 0$, any feasible schedule is a $(1 + \varepsilon)$ -approximation with high probability, when the number of jobs is sufficiently large.
- We give a *lower bound on the integrality gap of various linear programming relaxations of $1 \mid \text{prec} \mid \sum w_j C_j$* for almost all 0-1 bipartite instances. For the random models of 0-1 bipartite instances that we study, this lower bound approaches 2 as the precedence constraints become sparser in expectation. This result generalizes a result of Chekuri and Motwani [3].

2. Models for random 0-1 bipartite instances. We form a model for random 0-1 bipartite instances as follows. Let $n \in \mathbb{Z}_{>0}$ and $q \in (0, 1)$. In addition, let $\pi \in \mathbb{R}_{\geq 0}^{n+1}$ be a probability vector; that is, $\sum_{s=0}^n \pi_s = 1$. We define $\mathcal{B}(n, \pi, q)$ as the probability space of 0-1 bipartite instances (N_1, N_2, A) with n jobs such that $\mathbb{P}(|N_1| = s, |N_2| = n - s) = \pi_s$ for $s = 0, \dots, n$ and each arc $(i, j) \in N_1 \times N_2$ appears in A independently with probability q .

In this work, we consider random models of “balanced” 0-1 bipartite instances (N_1, N_2, A) , in the sense that the ratio between the size of N_1 and the size of N_2 is not too far from $\Theta(1)$ with high probability. In particular, we look at models $\mathcal{B}(n, \pi, q)$ with probability vector $\pi \in \mathbb{R}_{\geq 0}^{n+1}$ that satisfy

$$\sum_{s=0}^{\nu(n)-1} \pi_s \leq c_1 n^{c_2 \nu(n)} 2^{-n} \quad \text{and} \quad \sum_{s=n-\nu(n)+1}^n \pi_s \leq c_3 n^{c_4 \nu(n)} 2^{-n} \quad (1)$$

for some function $\nu: \mathbb{Z}_{>0} \rightarrow \mathbb{Z}_{>0}$ such that $\nu(n) \in \Theta(\log^\kappa n)$ for some fixed $\kappa \geq 1$ and for some constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$, when n is sufficiently large.

These conditions on the probability vector π are satisfied for two natural models of random 0-1 bipartite instances in particular. First, consider $\mathcal{B}(n, \bar{\pi}, q)$, in which $\bar{\pi}_s = \binom{n}{s} (1/2)^n$ for $s = 0, \dots, n$: jobs are assigned to N_1 and N_2 with equal probability. Note that $\mathcal{B}(n, \bar{\pi}, 1/2)$ is the uniform distribution over all 0-1 bipartite instances. There exist constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$ so that the probability vector $\bar{\pi}$ satisfies (1) for any $\nu(n) \in \Theta(\log^\kappa n)$ with $\kappa \geq 1$, because

$$\begin{aligned} \sum_{s=0}^{\nu(n)-1} \bar{\pi}_s &= \sum_{s=0}^{\nu(n)-1} \binom{n}{s} 2^{-n} \leq \nu(n) n^{\nu(n)} 2^{-n}, \\ \sum_{s=n-\nu(n)+1}^n \bar{\pi}_s &= \sum_{s=n-\nu(n)+1}^n \binom{n}{s} 2^{-n} \leq \nu(n) n^{\nu(n)} 2^{-n}. \end{aligned}$$

Second, consider $\mathcal{B}(n, \hat{\pi}, q)$, in which $\hat{\pi}_s = 1$ if $s = \alpha n$, and $\hat{\pi}_s = 0$ otherwise, for some fixed $\alpha \in (0, 1)$ such that $\alpha n \in \mathbb{Z}_{>0}$ and $(1 - \alpha)n \in \mathbb{Z}_{>0}$. For any instance (N_1, N_2, A) from $\mathcal{B}(n, \hat{\pi}, q)$, the proportion between the number of jobs in N_1 and the number of jobs in N_2 is always fixed. Clearly, there exist constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$ so that the probability vector $\hat{\pi}$ satisfies (1) for any $\nu(n) \in \Theta(\log^\kappa n)$ with $\kappa \geq 1$, when n is sufficiently large.

3. Sidney-decomposability and 0-1 bipartite instances. Sidney [21] introduced a very useful characterization of optimal schedules to $1 \mid \text{prec} \mid \sum w_j C_j$. We define

$$\rho(S) := \begin{cases} \sum_{j \in S} w_j / \sum_{j \in S} p_j & \text{for any subset of jobs } S \subseteq N \text{ such that } \sum_{j \in S} p_j > 0, \\ +\infty & \text{otherwise.} \end{cases}$$

A set of jobs $I \subseteq N$ is called *initial* if $j \in I$ and $(i, j) \in A$ imply $i \in I$. An initial set I^* is said to be ρ -maximal if $I^* \in \arg \max \{\rho(I) : I \text{ is a nonempty initial set}\}$. Sidney showed that there exists an optimal schedule in which

all jobs in a ρ -maximal initial set S^* are scheduled before those in $N \setminus S^*$. By recursively applying this result, we naturally obtain a partition of jobs (S_1, \dots, S_k) with $\rho(S_1) \geq \dots \geq \rho(S_k)$. Such a partition is called a *Sidney decomposition*. Sidney's decomposition theory can be seen as a generalization of Smith's rule [22] for the problem without precedence constraints. An instance of $1 | \text{prec} | \sum w_j C_j$ is *non-Sidney-decomposable* if the only ρ -maximal initial set is N ; otherwise the instance is called *Sidney-decomposable*. An instance is called *stiff* if $\rho(N) \geq \rho(I)$ for all nonempty initial sets I ; note that stiffness is a necessary condition for an instance to be non-Sidney-decomposable.

A Sidney decomposition can be computed in polynomial time (Lawler [14], Picard and Queyranne [18], Gallo et al. [8], Margot et al. [16]). Independently, Chekuri and Motwani [3] and Margot et al. [16] showed that for stiff instances, every feasible schedule is already a 2-approximation. A geometric proof of this result was subsequently given by Goemans and Williamson [9].

In this section, we show that almost all 0-1 bipartite instances are non-Sidney-decomposable. We begin by giving the following characterization of Sidney-decomposability for 0-1 bipartite instances. For any directed graph (N, A) and any subset of vertices $X \subseteq N$, we define $\Gamma(X) := \{i \in N \setminus X : (i, j) \in A \text{ or } (j, i) \in A \text{ for some } j \in X\}$; in words, $\Gamma(X)$ is the set of neighbors of X .

LEMMA 3.1. *A 0-1 bipartite instance (N_1, N_2, A) of $1 | \text{prec} | \sum w_j C_j$ with $|N_1| = n_1$, $|N_2| = n_2$, and $n_1 + n_2 \geq 2$ is Sidney-decomposable if and only if one of the following three conditions holds: (SD1) $n_1 = 0$; (SD2) $n_2 = 0$; (SD3) (i) there exists a subset $Y \subseteq N_2$ such that $Y \neq \emptyset$, N_2 , and $n_2 |\Gamma(Y)| \leq n_1 |Y|$, or (ii) $|\Gamma(N_2)| \leq n_1 - 1$.*

PROOF. First, note that a 0-1 bipartite instance with $n_1 + n_2 \geq 2$ is Sidney-decomposable when $n_1 = 0$ or $n_2 = 0$, because any nonempty subset of jobs I is initial and satisfies $\rho(I) = \rho(N)$.

Now suppose a 0-1 bipartite instance with $n_1 > 0$ and $n_2 > 0$ is Sidney-decomposable. By definition, this occurs if and only if

$$\text{there exists a } \rho\text{-maximal initial set } I \neq N \text{ such that } \rho(I) \geq n_2/n_1. \quad (2)$$

Recall that by definition, a ρ -maximal initial set is nonempty. Suppose (2) is satisfied with an initial set I such that $I \subseteq N_1 \cup N_2$, but $I \not\subseteq N_1$. Because I is ρ -maximal, $I = \Gamma(Y) \cup Y$ for some $Y \subseteq N_2$ such that $Y \neq \emptyset$. We consider the following cases:

- If $Y \neq N_2$, then (2) holds if and only if $|Y|/|\Gamma(Y)| \geq n_2/n_1$.
- Otherwise, we have $Y = N_2$. In this case, (2) holds if and only if $|\Gamma(N_2)| \leq n_1 - 1$.

Note that (2) cannot be satisfied if $I \subseteq N_1$, because in this case, $\rho(I) = 0 < n_2/n_1 = \rho(N)$. \square

Note that (SD3) implies that a 0-1 bipartite instance (N_1, N_2, A) with $|N_1| = |N_2| \geq 1$ is non-Sidney-decomposable if and only if $|\Gamma(N_2)| = |N_1| = |N_2|$ and $|\Gamma(Y)| > |Y|$ for all $Y \subseteq N_2$ such that $Y \neq \emptyset, N_2$. This is very similar to Hall's marriage theorem [12], which says that an undirected bipartite graph $(N_1 \dot{\cup} N_2, A)$ with $|N_1| = |N_2|$ has a perfect matching if and only if $|\Gamma(Y)| \geq |Y|$ for all $Y \subseteq N_2$.

We now give an analogous characterization of Sidney-decomposable 0-1 bipartite instances that considers subsets of N_1 instead.

LEMMA 3.2. *The condition (SD3) in Lemma 3.1 holds if and only if the following condition holds: (SD3') (i) There exists a subset $X \subseteq N_1$ such that $X \neq \emptyset, N_1$, and $n_1 |\Gamma(X)| \leq n_2 |X|$, or (ii) $|\Gamma(N_1)| \leq n_2 - 1$.*

PROOF. We show that (SD3) implies (SD3'). Suppose that (SD3) holds because there exists a subset $Y \subseteq N_2$ such that $Y \neq \emptyset, N_2$ and $n_2 |\Gamma(Y)| \leq n_1 |Y|$. Let $X = N_1 \setminus \Gamma(Y)$. We consider the following cases:

- $\Gamma(Y) = \emptyset$. In this case, $X = N_1$. Because $Y \neq \emptyset$, this implies that $|\Gamma(N_1)| = |\Gamma(X)| \leq n_2 - 1$.
- $\Gamma(Y) \neq \emptyset, N_1$. In this case, $X \neq \emptyset, N_1$. In addition, we have that $|X| = n_1 - |\Gamma(Y)|$ and $|\Gamma(X)| \leq n_2 - |Y|$. These two observations, in addition to the assumption that $n_2 |\Gamma(Y)| \leq n_1 |Y|$, imply that $n_1 |\Gamma(X)| \leq n_2 |X|$.
- $\Gamma(Y) = N_1$. In this case, because $n_2 |\Gamma(Y)| \leq n_1 |Y|$, we have that $|Y| \geq n_2$, which is a contradiction, because $Y \neq N_2$.

Now suppose that (SD3) holds because $|\Gamma(N_2)| \leq n_1 - 1$. Let $X = N_1 \setminus \Gamma(N_2)$. Note that because $|\Gamma(N_2)| \leq n_1 - 1$, we have that $X \neq \emptyset$. In addition, because $X \cap \Gamma(N_2) = \emptyset$, we have that $\Gamma(X) = \emptyset$. We consider the following cases:

- $\Gamma(N_2) \neq \emptyset$. Then $X \neq \emptyset, N_1$ and $n_1 |\Gamma(X)| = 0 \leq n_2 |X|$.
- $\Gamma(N_2) = \emptyset$. Then $X = N_1$ and $|\Gamma(N_1)| = |\Gamma(X)| = 0 \leq n_2 - 1$.

Showing the reverse direction works in a similar manner. \square

Before we proceed, we need the following lemma.

LEMMA 3.3. For any $a \in (0, 1]$ such that $as \in \mathbb{Z}_{>0}$ and $k = 1, \dots, s$, $\binom{as}{\lfloor ak \rfloor} \leq \binom{s}{k}$.

PROOF. The claim follows directly from the fact that $\binom{n}{x} \geq \binom{n-1}{x-1}$ and $\binom{n}{x} \geq \binom{n-1}{x}$ for any $x = 1, \dots, n$. \square

Using the characterization of Sidney-decomposability in Lemmas 3.1 and 3.2, we can show that almost all 0-1 bipartite instances are non-Sidney-decomposable.

THEOREM 3.1. Fix $q \in (0, 1)$, $\kappa > 1$, and $\nu(n) \in \Theta(\log^\kappa n)$. Let $\pi \in \mathbb{R}_{\geq 0}^{n+1}$ be a probability vector that satisfies (1) for $\nu(n)$ and some constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$, when n is sufficiently large. Then,

$$\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ is non-Sidney-decomposable}) = 1.$$

PROOF. Let $B = (N_1, N_2, A)$ be a random 0-1 bipartite instance from $\mathcal{B}(n, \pi, q)$ with probability vector π that satisfies (1) for $\nu(n)$ and for some constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$, when n is sufficiently large. We show that the probability that B satisfies any of the conditions (SD1)–(SD3) goes to zero as n approaches infinity. For the remainder of this proof, we consider n sufficiently large so that $n \geq 2$ and $\nu(n) \leq \lfloor n/2 \rfloor$.

First, we consider (SD1). We have that

$$\mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD1)}) = \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ has } n_1 = 0) = \pi_0 \leq c_1 n^{c_2 \nu(n)} 2^{-n},$$

and so $\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD1)}) = 0$. Similarly, for (SD2), we have that

$$\mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD2)}) = \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ has } n_2 = 0) = \pi_n \leq c_3 n^{c_4 \nu(n)} 2^{-n},$$

and therefore $\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD2)}) = 0$.

Now we consider (SD3). Observe that any bipartite graph $(N_1 \cup N_2, A)$ with $|N_1| = s$ and $|N_2| = n - s$ with a subset Y of N_2 of size k such that $|\Gamma(Y)| \leq (s/(n-s))k$ can be constructed as follows. Choose a subset Y of N_2 of size k and a subset X of N_1 of size $\lfloor (s/(n-s))k \rfloor$, and forbid all edges between Y and $N_1 \setminus X$. Any bipartite graph $(N_1 \cup N_2, A)$ with $|N_1| = s$ and $|N_2| = n - s$ with a subset X of N_1 of size k such that $|\Gamma(X)| \leq ((n-s)/s)k$ can be constructed similarly. Therefore, by conditioning on the size of N_1 and N_2 and using a union bound, we have

$$\begin{aligned} & \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)}) \\ &= \sum_{s=1}^{n-1} \pi_s \cdot \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)} : |N_1| = s, |N_2| = n-s) \\ &\leq \sum_{s=1}^{\lfloor n/2 \rfloor} \pi_s \cdot \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)} : |N_1| = s, |N_2| = n-s) \\ &\quad + \sum_{s=\lfloor n/2 \rfloor}^{n-1} \pi_s \cdot \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3')} : |N_1| = s, |N_2| = n-s) \\ &\leq \sum_{s=1}^{\lfloor n/2 \rfloor} \pi_s \cdot \left(\sum_{k=1}^{n-s-1} \binom{n-s}{k} \binom{s}{\lfloor (s/(n-s))k \rfloor} (1-q)^{k(s-\lfloor (s/(n-s))k \rfloor)} + s(1-q)^{n-s} \right) \\ &\quad + \sum_{s=\lfloor n/2 \rfloor}^{n-1} \pi_s \cdot \left(\sum_{k=1}^{s-1} \binom{s}{k} \binom{n-s}{\lfloor ((n-s)/s)k \rfloor} (1-q)^{k(n-s-\lfloor ((n-s)/s)k \rfloor)} + (n-s)(1-q)^s \right). \end{aligned}$$

We define

$$\begin{aligned} D_s &:= \pi_s \cdot \sum_{k=1}^{n-s-1} \binom{n-s}{k} \binom{s}{\lfloor (s/(n-s))k \rfloor} (1-q)^{k(s-\lfloor (s/(n-s))k \rfloor)} & \text{for } s = 1, \dots, \lfloor n/2 \rfloor, \\ E_s &:= \pi_s \cdot s(1-q)^{n-s} & \text{for } s = 1, \dots, \lfloor n/2 \rfloor, \\ F_s &:= \pi_s \cdot \sum_{k=1}^{s-1} \binom{s}{k} \binom{n-s}{\lfloor ((n-s)/s)k \rfloor} (1-q)^{k(n-s-\lfloor ((n-s)/s)k \rfloor)} & \text{for } s = \lfloor n/2 \rfloor, \dots, n-1, \\ G_s &:= \pi_s \cdot (n-s)(1-q)^s & \text{for } s = \lfloor n/2 \rfloor, \dots, n-1, \end{aligned}$$

so that

$$\mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)}) \leq \sum_{s=1}^{\lfloor n/2 \rfloor} D_s + \sum_{s=1}^{\lfloor n/2 \rfloor} E_s + \sum_{s=\lceil n/2 \rceil}^{n-1} F_s + \sum_{s=\lceil n/2 \rceil}^{n-1} G_s.$$

For the remainder of this proof, let $r = (1 - q)^{-1}$. Note that $r > 1$. First, we consider the expression F_s in the regime $s = \lceil n/2 \rceil, \dots, n - \nu(n)$. By Lemma 3.3 (letting $a = (n - s)/s$), for all $s = \lceil n/2 \rceil, \dots, n - \nu(n)$, we have that

$$F_s \leq \pi_s \cdot \sum_{k=1}^{s-1} \binom{s}{k}^2 (1 - q)^{k(n-s - \lfloor ((n-s)/s)k \rfloor)} \leq \pi_s \cdot \sum_{k=1}^{s-1} \binom{s}{k}^2 (1 - q)^{((n-s)/s)k(s-k)}.$$

For all $s = \lceil n/2 \rceil, \dots, n - \nu(n)$ and $k = 1, \dots, s - 1$, define

$$H_{s,k} := \binom{s}{k}^2 (1 - q)^{((n-s)/s)k(s-k)},$$

and note that $H_{s,k} = H_{s,s-k}$. We would like to show that $H_{s,k} \geq H_{s,k+1}$ for all $s = \lceil n/2 \rceil, \dots, n - \nu(n)$ and $k = 1, \dots, \lfloor s/2 \rfloor - 1$, or equivalently,

$$2 \log_r \frac{s-k}{k+1} \leq \frac{n-s}{s} (s-2k-1) \quad \text{for } k = 1, \dots, \lfloor s/2 \rfloor - 1 \text{ and } s = \lceil n/2 \rceil, \dots, n - \nu(n). \quad (3)$$

Define

$$\Delta(x) := \frac{n-s}{s} (s-2x-1) - 2 \log_r (s-x) + 2 \log_r (x+1).$$

Taking derivatives, we obtain

$$\frac{\partial \Delta}{\partial x} = -\frac{2(n-s)}{s} + \frac{2}{\log r} \left(\frac{1}{s-x} + \frac{1}{x+1} \right), \quad \frac{\partial^2 \Delta}{\partial x^2} = \frac{2}{\log r} \left(\frac{1}{(s-x)^2} - \frac{1}{(x+1)^2} \right).$$

Note that for $x \in [0, (s-1)/2]$, we have that $\partial^2 \Delta / \partial x^2 \leq 0$, so $\Delta(x)$ is concave on $[0, (s-1)/2]$. We have that $\Delta(0) \geq 0$ for all $s = \lceil n/2 \rceil, \dots, n - \nu(n)$, because

$$\begin{aligned} \Delta(0) &= \frac{n-s}{s} (s-1) - 2 \log_r s + 2 \log_r 1 \\ &= n-s - \frac{n-s}{s} - 2 \log_r s \\ &\geq n-s-1 - 2 \log_r n \quad (\text{because } s \geq n-s \text{ and } s \leq n) \\ &\geq \nu(n) - 1 - 2 \log_r n \quad (\text{because } s \leq n - \nu(n)) \\ &\geq 0 \quad (\text{because } \nu(n) \in \Theta(\log^\kappa n) \text{ and } \kappa > 1). \end{aligned}$$

In addition, we have that $\Delta((s-1)/2) = 0$. Because $\Delta(x)$ is concave on $[0, (s-1)/2]$, it follows that when $s = \lceil n/2 \rceil, \dots, n - \nu(n)$, $\Delta(x) \geq 0$ for all $x \in [0, (s-1)/2]$, which establishes (3). Therefore, $H_{s,k} \geq H_{s,k+1}$ for $s = \lceil n/2 \rceil, \dots, n - \nu(n)$ and $k = 1, \dots, \lfloor s/2 \rfloor - 1$. Because $H_{s,k} = H_{s,s-k}$, it follows that $H_{s,1} \geq H_{s,k}$ for $s = \lceil n/2 \rceil, \dots, n - \nu(n)$ and $k = 1, \dots, s - 1$.

So, for $s = \lceil n/2 \rceil, \dots, n - \nu(n)$, we have that

$$\begin{aligned} F_s &\leq \pi_s \cdot \sum_{k=1}^{s-1} \binom{s}{k}^2 (1 - q)^{((n-s)/s)k(s-k)} \\ &\leq \pi_s \cdot s^3 (1 - q)^{((n-s)/s)(s-1)} \\ &\leq \pi_s \cdot s^3 (1 - q)^{(n-s)/2} \quad \left(\text{because } \frac{s-1}{s} \geq \frac{1}{2} \text{ for } s \geq 2 \right). \end{aligned}$$

Therefore,

$$\sum_{s=\lceil n/2 \rceil}^{n-\nu(n)} F_s \leq \sum_{s=\lceil n/2 \rceil}^{n-\nu(n)} \pi_s \cdot s^3 (1 - q)^{(n-s)/2} \leq n^4 (1 - q)^{\nu(n)/2}.$$

Now we consider F_s in the regime $s = n - \nu(n) + 1, \dots, n - 1$. Note that

$$F_s = \pi_s \cdot \sum_{k=1}^{s-1} \binom{s}{k} \binom{n-s}{\lfloor ((n-s)/s)k \rfloor} (1 - q)^{k(n-s - \lfloor ((n-s)/s)k \rfloor)} \leq \pi_s \cdot 2^{n-s} \sum_{k=1}^{s-1} \binom{s}{k} (1 - q)^k \leq \pi_s \cdot 2^{\nu(n)} (2 - q)^s.$$

It follows that

$$\sum_{s=n-\nu(n)+1}^{n-1} F_s \leq \sum_{s=n-\nu(n)+1}^{n-1} \pi_s \cdot 2^{\nu(n)}(2-q)^s \leq \sum_{s=n-\nu(n)+1}^{n-1} \pi_s \cdot 2^{\nu(n)}(2-q)^n \leq c_3(2n^{c_4})^{\nu(n)} \left(1 - \frac{q}{2}\right)^n.$$

We also have that

$$\sum_{s=\lceil n/2 \rceil}^{n-1} G_s = \sum_{s=\lceil n/2 \rceil}^{n-1} \pi_s \cdot (n-s)(1-q)^s \leq \frac{n^2}{2}(1-q)^{n/2}.$$

Using similar techniques to above, we can also show that

$$\sum_{s=1}^{\nu(n)-1} D_s \leq c_1(2n^{c_2})^{\nu(n)} \left(1 - \frac{q}{2}\right)^n, \quad \sum_{s=\nu(n)}^{\lfloor n/2 \rfloor} D_s \leq n^4(1-q)^{\nu(n)/2}, \quad \sum_{s=1}^{\lfloor n/2 \rfloor} E_s \leq \frac{n^2}{2}(1-q)^{n/2}.$$

Therefore,

$$\begin{aligned} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)}) &\leq \sum_{s=1}^{\lfloor n/2 \rfloor} D_s + \sum_{s=1}^{\lfloor n/2 \rfloor} E_s + \sum_{s=\lceil n/2 \rceil}^{n-1} F_s + \sum_{s=\lceil n/2 \rceil}^{n-1} G_s \\ &\leq c_1(2n^{c_2})^{\nu(n)} \left(1 - \frac{q}{2}\right)^n + c_3(2n^{c_4})^{\nu(n)} \left(1 - \frac{q}{2}\right)^n \\ &\quad + 2n^4(1-q)^{\nu(n)/2} + n^2(1-q)^{n/2}. \end{aligned}$$

Because $\nu(n) \in \Theta(\log^\kappa n)$ for some fixed $\kappa > 1$, it follows that

$$\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)}) = 0.$$

Finally, we put all of the pieces together:

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ is Sidney-decomposable}) &= \lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD1)}) \\ &\quad + \lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD2)}) \\ &\quad + \lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (SD3)}) = 0. \quad \square \end{aligned}$$

In random models of “balanced” 0-1 bipartite instances, the number of jobs in N_1 and the number of jobs in N_2 grow together as the total number of jobs grows. This phenomenon is important for the validity of Theorem 3.1. For example, consider $\mathcal{B}(n, \tilde{\pi}, q)$ with $\tilde{\pi}_s = 1$ if $s = 1$ and $\tilde{\pi}_s = 0$ otherwise: the class of instances in which N_1 consists of one job, and N_2 consists of $n - 1$ jobs. In this case, an instance $B \in \mathcal{B}(n, \tilde{\pi}, q)$ is non-Sidney-decomposable if and only if the job in N_1 must precede all jobs in N_2 . This occurs with probability q^{n-1} , which goes to zero as the total number n of jobs grows.

Finally, we note that Theorem 3.1 still holds for sparser precedence constraints. It is straightforward to show that if the probability $q(n)$ of a precedence constraint appearing is a function of the number n of jobs so that $q(n) \in \omega(1/\log^{\kappa-1} n)$, then the analysis in the proof of Theorem 3.1 holds.

4. Two-dimensional Gantt charts and 0-1 bipartite instances. Two-dimensional (2D) Gantt charts (Eastman et al. [6]) provide an elegant, geometric way of understanding single-machine completion-time-objective scheduling problems. In a traditional Gantt chart, the horizontal axis corresponds to processing time. In a 2D Gantt chart, the horizontal axis corresponds to processing time, and the vertical axis corresponds to weight. Suppose we have an instance $(N, A, (p_i)_{i \in N}, (w_i)_{i \in N})$ of $1|\text{prec}|\sum w_j C_j$. The 2D Gantt chart is constructed for a permutation schedule $(1, \dots, n)$ as follows. Each job $j \in N$ is represented by a rectangle of width p_j and height w_j , whose position in the chart is defined by a startpoint and an endpoint. The startpoint of the first job (job 1) in the schedule is $(0, \sum_{j \in N} w_j)$, and its endpoint is $(p_1, \sum_{j \in N} w_j - w_1)$. For all subsequent jobs in the schedule, the startpoint (t, w) of job j is the endpoint of the previous job $j - 1$, and its endpoint is $(t + p_j, w - w_j)$. The completion time of a job in this schedule is the time component of its endpoint. The *work curve* $W(t)$ formed by the upper side of each rectangle is the total weight of jobs that have not been completed by time t . The area under the work curve is equal to the sum of weighted completion times for the schedule represented by the 2D Gantt chart.




FIGURE 1. An example of a 2D Gantt chart for a 0-1 bipartite instance.

It turns out that the area under the optimal work curve for almost all 0-1 bipartite instances is “large.” We formalize this notion now. Consider the 2D Gantt chart for an optimal schedule of a 0-1 bipartite instance $B = (N_1, N_2, A)$ with $|N_1| = n_1$ and $|N_2| = n_2$. Note that any 2D Gantt chart for such an instance starts at $(0, n_2)$ and ends at $(n_1, 0)$. Also, observe that all jobs in N_1 are represented by a horizontal line segment of length 1, and that all jobs in N_2 are represented by a vertical line segment of length 1. We define R_B to be the region between the optimal work curve and the frontier formed by the lines $\{(t, w) : t = n_1\}$ and $\{(t, w) : w = n_2\}$ (see Figure 1 for an example).

We define the following parametrized condition on a 0-1 bipartite instance B , for any $\alpha \in (0, 1)$:

(R- α) A rectangle of width αn_1 and height αn_2 cannot fit in R_B .

We now show that for any fixed $\alpha \in (0, 1)$, the condition (R- α) is satisfied for almost all 0-1 bipartite instances.

THEOREM 4.1. Fix $q \in (0, 1)$, $\alpha \in (0, 1)$, $\kappa \geq 1$, and $\nu(n) \in \Theta(\log^\kappa n)$. Let $\pi \in \mathbb{R}_{\geq 0}^{n+1}$ be a probability vector that satisfies (1) for $\nu(n)$ and some constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$, when n is sufficiently large. Then,

$$\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies (R-}\alpha) = 1.$$

PROOF. Fix a 0-1 bipartite instance $B = (N_1, N_2, A)$ with $|N_1| = n_1$ and $|N_2| = n_2$. If B does not satisfy (R- α), that is, a rectangle of width αn_1 and height αn_2 can fit in R_B , then there exists a set of $\lceil \alpha n_2 \rceil$ jobs from N_2 that has at most $n_1 - \lceil \alpha n_1 \rceil$ predecessors in N_1 . In other words, if a rectangle of width αn_1 and height αn_2 can fit in R_B , then there exists a set of $\lceil \alpha n_2 \rceil$ jobs from N_2 and a set of $\lceil \alpha n_1 \rceil$ jobs from N_1 with no precedence constraints between them.

Therefore, we have that

$$\begin{aligned} & \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha) \mid |N_1| = s, |N_2| = n - s) \\ & \leq \mathbb{P} \left(\begin{array}{c} |X| = \lceil \alpha s \rceil, \quad |Y| = \lceil \alpha(n - s) \rceil, \\ \exists X \subseteq N_1, Y \subseteq N_2: \quad \text{no precedence constraints} \\ \text{between } X \text{ and } Y \end{array} \mid |N_1| = s, |N_2| = n - s \right) \\ & \leq \binom{s}{\lceil \alpha s \rceil} \binom{n - s}{\lceil \alpha(n - s) \rceil} (1 - q)^{\lceil \alpha s \rceil \lceil \alpha(n - s) \rceil} \leq \binom{s}{\lceil \alpha s \rceil} \binom{n - s}{\lceil \alpha(n - s) \rceil} (1 - q)^{\alpha^2 s(n - s)}. \end{aligned}$$

So, by conditioning on the size of N_1 and N_2 ,

$$\begin{aligned} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha)) &= \sum_{s=1}^{n-1} \pi_s \cdot \mathbb{P} \left(\begin{array}{c} B \in \mathcal{B}(n, \pi, q) \\ \text{does not satisfy (R-}\alpha) \end{array} \mid |N_1| = s, |N_2| = n - s \right) \\ &\leq \sum_{s=1}^{n-1} \pi_s \cdot \binom{s}{\lceil \alpha s \rceil} \binom{n - s}{\lceil \alpha(n - s) \rceil} (1 - q)^{\alpha^2 s(n - s)}. \end{aligned}$$

Let

$$D_s = \pi_s \cdot \binom{s}{\lceil \alpha s \rceil} \binom{n-s}{\lceil \alpha(n-s) \rceil} (1-q)^{\alpha^2 s(n-s)} \quad \text{for } s = 1, \dots, n-1$$

so that

$$\mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha)) \leq \sum_{s=1}^{n-1} D_s.$$

First, for the regime $s = 1, \dots, \nu(n) - 1$, we have that

$$\begin{aligned} \sum_{s=1}^{\nu(n)-1} D_s &= \sum_{s=1}^{\nu(n)-1} \pi_s \cdot \binom{s}{\lceil \alpha s \rceil} \binom{n-s}{\lceil \alpha(n-s) \rceil} (1-q)^{\alpha^2 s(n-s)} \\ &\leq \sum_{s=1}^{\nu(n)-1} \pi_s \cdot 2^n (1-q)^{\alpha^2(n-1)} \leq c_1 n^{c_2 \nu(n)} (1-q)^{\alpha^2(n-1)}. \end{aligned}$$

Similarly, we can show that

$$\sum_{s=n-\nu(n)+1}^{n-1} D_s \leq c_3 n^{c_4 \nu(n)} (1-q)^{\alpha^2(n-1)}.$$

For the regime $s = \nu(n), \dots, n - \nu(n)$, we have that

$$\sum_{s=\nu(n)}^{n-\nu(n)} D_s = \sum_{s=\nu(n)}^{n-\nu(n)} \pi_s \cdot \binom{s}{\lceil \alpha s \rceil} \binom{n-s}{\lceil \alpha(n-s) \rceil} (1-q)^{\alpha^2 s(n-s)} \leq n 2^n (1-q)^{\alpha^2 \nu(n)(n-\nu(n))}.$$

Therefore,

$$\begin{aligned} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha)) \\ \leq \sum_{s=1}^{n-1} D_s \leq c_1 n^{c_2 \nu(n)} (1-q)^{\alpha^2(n-1)} + c_3 n^{c_4 \nu(n)} (1-q)^{\alpha^2(n-1)} + n 2^n (1-q)^{\alpha^2 \nu(n)(n-\nu(n))}. \end{aligned}$$

Because $\nu(n) \in \Theta(\log^\kappa n)$ for a fixed $\kappa \geq 1$, it follows that

$$\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha)) = 0. \quad \square$$

Before we proceed, we need the following version of the Chernoff bound.

LEMMA 4.1 (CHERNOFF BOUNDS; SEE MITZENMACHER AND UPFAL [17]). *Let X_1, \dots, X_m be independent random variables such that for $i = 1, \dots, m$, $\mathbb{P}(X_i = 1) = q$, and $\mathbb{P}(X_i = 0) = 1 - q$ with $q \in (0, 1)$. Then for $S = \sum_{i=1}^m X_i$, $\mu = \mathbb{E}(S) = qm$, and any $\delta \in (0, 1)$, (a) $\mathbb{P}(S \geq (1 + \delta)\mu) \leq e^{-\mu\delta^2/3}$ and (b) $\mathbb{P}(S \leq (1 - \delta)\mu) \leq e^{-\mu\delta^2/2}$.*

As with the non-Sidney-decomposability result in §3, the “balancedness” of the random 0-1 bipartite instances we consider plays a key role in the validity of Theorem 4.1. To illustrate this, as before, fix $q \in (0, 1)$ and consider $\mathcal{B}(n, \tilde{\pi}, q)$ with $\tilde{\pi}_s = 1$ if $s = 1$ and $\tilde{\pi}_s = 0$ otherwise: the class of instances in which N_1 consists of one job, and N_2 consists of $n - 1$ jobs. Take α to be arbitrarily small: in particular, $\alpha < 1 - q$. In this case, an instance $B \in \mathcal{B}(n, \tilde{\pi}, q)$ does not satisfy (R- α) if and only if there exist at least $\lceil \alpha(n - 1) \rceil$ jobs in N_2 that do not have any predecessors in N_1 . Let Z be a binomial random variable with $n - 1$ trials and probability of success $1 - q$. Then, by the lower tail Chernoff bound in Lemma 4.1(b),

$$\begin{aligned} \mathbb{P}(B \in \mathcal{B}(n, \tilde{\pi}, q) \text{ does not satisfy (R-}\alpha)) &= \mathbb{P}(Z \geq \lceil \alpha(n - 1) \rceil) \geq 1 - \mathbb{P}(Z \leq \alpha(n - 1)) \\ &\geq 1 - \exp\left(-\frac{1}{2} \left(1 - \frac{\alpha}{1 - q}\right)^2 (1 - q)(n - 1)\right). \end{aligned}$$

Therefore, $\mathbb{P}(B \in \mathcal{B}(n, \tilde{\pi}, q)$ satisfies (R- α)) goes to zero as the total number n of jobs grows.

With Theorem 4.1 in hand, we can show that for almost all 0-1 bipartite instances, all feasible schedules are arbitrarily close to optimal. Let $\text{opt}(B)$ denote the optimal value of instance B , and let $\text{val}(B, S)$ denote the objective value of (feasible) schedule S for instance B .

THEOREM 4.2. Fix $q \in (0, 1)$, $\alpha \in (0, 1)$, $\kappa \geq 1$, and $\nu(n) \in \Theta(\log^\kappa n)$. Let $\pi \in \mathbb{R}_{\geq 0}^{n+1}$ be a probability vector that satisfies (1) for $\nu(n)$ and some constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$, when n is sufficiently large. Then,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } \frac{\text{val}(B, S)}{\text{opt}(B)} \leq (1 - \alpha)^{-2} \text{ for all feasible schedules } S \right) = 1.$$

PROOF. Consider some 0-1 bipartite instance B with $|N_1| = n_1$ and $|N_2| = n_2$. If (R- α) is satisfied—that is, if a rectangle of width αn_1 and height αn_2 cannot fit in the region R_B —then $\text{opt}(B) > n_1 n_2 (1 - \alpha)^2$. Because the objective value of any feasible schedule of an instance B is at most $n_1 n_2$, this implies that if (R- α) is satisfied, $\text{val}(B, S)/\text{opt}(B) \leq n_1 n_2 / (n_1 n_2 (1 - \alpha)^2) = (1 - \alpha)^{-2}$, which implies the claim. \square

In addition, Theorem 4.1 also implies a nontrivial lower bound on the integrality gap of various linear programming relaxations of $1|\text{prec}|\sum w_j C_j$, for almost all 0-1 bipartite instances. Potts [19] proposed the following integer programming formulation. Define the decision variables $(\delta_{ij})_{i, j \in N: i \neq j}$ as follows: for all $i, j \in N$ such that $i \neq j$, δ_{ij} is equal to 1 if job i is processed before job j , and 0 otherwise. Then $1|\text{prec}|\sum w_j C_j$ can be formulated as

$$[\text{P}] \quad \text{minimize} \quad \sum_{j \in N} p_j w_j + \sum_{i, j \in N: i \neq j} p_i w_j \delta_{ij}, \quad (4a)$$

$$\text{subject to} \quad \delta_{ij} + \delta_{ji} = 1 \quad \text{for all } i, j \in N: i \neq j, \quad (4b)$$

$$\delta_{ij} + \delta_{jk} + \delta_{ki} \leq 2 \quad \text{for all } i, j, k \in N: i \neq j \neq k \neq i, \quad (4c)$$

$$\delta_{ij} = 1 \quad \text{for all } (i, j) \in A, \quad (4d)$$

$$\delta_{ij} \in \{0, 1\} \quad \text{for all } i, j \in N: i \neq j. \quad (4e)$$

It is straightforward to check that [P] is a correct formulation of $1|\text{prec}|\sum w_j C_j$. We denote the LP relaxation of [P] obtained by replacing the binary constraints (4e) with nonnegativity constraints $\delta_{ij} \geq 0$ for all $i, j \in N$ as [P-LP]. Let $\text{lp}(B)$ denote the optimal value of [P-LP].

THEOREM 4.3. Fix $q \in (0, 1)$, $\alpha \in (0, 1)$, $\delta \in (0, 1)$, $\kappa \geq 1$, and $\nu(n) \in \Theta(\log^\kappa n)$. Let $\pi \in \mathbb{R}_{\geq 0}^{n+1}$ be a probability vector that satisfies (1) for $\nu(n)$ and some constants $c_1, c_2, c_3, c_4 \in \mathbb{R}_{>0}$, when n is sufficiently large. Then,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } \frac{\text{opt}(B)}{\text{lp}(B)} > \frac{2(1 - \alpha)^2}{1 + (1 + \delta)q} \right) = 1.$$

PROOF. Consider a 0-1 bipartite instance $B = (N_1, N_2, A)$ with $|N_1| = n_1$ and $|N_2| = n_2$. It is straightforward to show that setting $\delta_{ij} = 1$ if $(i, j) \in A$, and $\delta_{ij} = \frac{1}{2}$ otherwise, is a feasible solution to [P-LP] and that this solution has objective value $\frac{1}{2}(n_1 n_2 + |A|)$. Therefore, $\text{lp}(B) \leq \frac{1}{2}(n_1 n_2 + |A|)$. In the proof of Theorem 4.2, we showed that if B satisfies (R- α), then $\text{opt}(B) > n_1 n_2 (1 - \alpha)^2$. Therefore, if B satisfies (R- α) and $|A| < (1 + \delta)q n_1 n_2$, then

$$\frac{\text{opt}(B)}{\text{lp}(B)} > \frac{n_1 n_2 (1 - \alpha)^2}{\frac{1}{2}(n_1 n_2 + |A|)} > \frac{n_1 n_2 (1 - \alpha)^2}{\frac{1}{2}(n_1 n_2 + (1 + \delta)q n_1 n_2)} = \frac{2(1 - \alpha)^2}{1 + (1 + \delta)q},$$

and so

$$\mathbb{P} \left(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } \frac{\text{opt}(B)}{\text{lp}(B)} \leq \frac{2(1 - \alpha)^2}{1 + (1 + \delta)q} \right) \leq \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha)) \\ + \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } |A| \geq (1 + \delta)q n_1 n_2).$$

By conditioning on the size of N_1 and N_2 , and using the upper tail Chernoff bound from Lemma 4.1(a), we obtain

$$\mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } |A| \geq (1 + \delta)q n_1 n_2) = \sum_{s=1}^{n-1} \pi_s \cdot \mathbb{P} \left(\begin{array}{l} B \in \mathcal{B}(n, \pi, q) \text{ satisfies} \\ |A| \geq (1 + \delta)q n_1 n_2 \end{array} \middle| n_1 = s, n_2 = n - s \right) \\ \leq \sum_{s=1}^{n-1} \pi_s \cdot e^{-qs(n-s)\delta^2/3} \leq ne^{-q(n-1)\delta^2/3}.$$

Therefore, $\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } |A| \geq (1 + \delta)q n_1 n_2) = 0$. By Theorem 4.1, we have that $\lim_{n \rightarrow \infty} \mathbb{P}(B \in \mathcal{B}(n, \pi, q) \text{ does not satisfy (R-}\alpha)) = 0$. It follows that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(B \in \mathcal{B}(n, \pi, q) \text{ satisfies } \frac{\text{opt}(B)}{\text{lp}(B)} \leq \frac{2(1 - \alpha)^2}{1 + (1 + \delta)q} \right) = 0. \quad \square$$

We note that the above result also applies to other formulations of $1|\text{prec}|\sum w_j C_j$, including the further relaxations of [P] attributable to Chudak and Hochbaum [4] and Correa and Schulz [5], and the LP relaxation of $1|\text{prec}|\sum w_j C_j$ based on completion-time variables attributable to Queyranne and Wang [20], because all of these relaxations are no stronger than [P-LP].

Finally, we note that Theorems 4.2 and 4.3 remain valid as long as the probability $q(n)$ of a precedence constraint appearing is a function of the number n of jobs so that $q(n) \in \omega(1/\log^k n)$.

Acknowledgments. The authors thank the associate editor and three anonymous referees, as well as David Gamarnik, Monaldo Mastrolilli, and Maurice Queyranne, for their helpful comments.

References

- [1] Ambühl, C., M. Mastrolilli, O. Svensson. 2007. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. *Proc. 48th Annual IEEE Sympos. Foundations Comput. Sci. (FOCS 2007)*, IEEE Computer Society, Los Alamitos, CA, 329–337.
- [2] Bansal, N., S. Khot. 2009. Optimal long code test with one free bit. *Proc. 50th Annual IEEE Sympos. Foundations Comput. Sci. (FOCS 2009)*, IEEE Computer Society, Los Alamitos, CA, 453–462.
- [3] Chekuri, C., R. Motwani. 1999. Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Discrete Appl. Math.* **98**(1–2) 29–38.
- [4] Chudak, F. A., D. S. Hochbaum. 1999. A half-integral linear programming relaxation for scheduling precedence-constrained jobs on a single machine. *Oper. Res. Lett.* **25**(5) 199–204.
- [5] Correa, J. R., A. S. Schulz. 2005. Single-machine scheduling with precedence constraints. *Math. Oper. Res.* **30**(4) 1005–1021.
- [6] Eastman, W. L., S. Even, I. M. Isaacs. 1964. Bounds for the optimal scheduling of n jobs on m processors. *Management Sci.* **11**(2) 268–279.
- [7] Erdős, P., A. Rényi. 1959. On random graphs I. *Publicationes Mathematicae* **6** 290–297.
- [8] Gallo, G., M. D. Grigoriadis, R. E. Tarjan. 1989. A fast parametric maximum flow algorithm and applications. *SIAM J. Comput.* **18**(1) 30–55.
- [9] Goemans, M. X., D. P. Williamson. 2000. Two-dimensional Gantt charts and a scheduling algorithm of Lawler. *SIAM J. Discrete Math.* **13**(3) 281–294.
- [10] Graham, R. L., E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan. 1979. Optimization and approximation in deterministic sequencing and scheduling: A survey. *Ann. Discrete Math.* **5** 287–326.
- [11] Hall, L. A., A. S. Schulz, D. B. Shmoys, J. Wein. 1997. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Math. Oper. Res.* **22**(3) 513–544.
- [12] Hall, P. 1935. On representatives of subsets. *J. London Math. Soc.* **10**(1) 26–30.
- [13] Khot, S. 2002. On the power of unique 2-prover 1-round games. *Proc. 34th Annual ACM Sympos. Theory Comput. (STOC 2002)*, ACM, New York, 767–775.
- [14] Lawler, E. L. 1978. Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Ann. Discrete Math.* **2** 75–90.
- [15] Lenstra, J. K., A. H. G. Rinnooy Kan. 1978. Complexity of scheduling under precedence constraints. *Oper. Res.* **26**(1) 22–35.
- [16] Margot, F., M. Queyranne, Y. Wang. 2003. Decompositions, network flows, and a precedence constrained single-machine scheduling problem. *Oper. Res.* **51**(6) 981–992.
- [17] Mitzenmacher, M., E. Upfal. 2005. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, Cambridge, UK.
- [18] Picard, J.-C., M. Queyranne. 1980. On the structure of all minimum cuts in a network and applications. *Math. Programming Stud.* **13** 8–16.
- [19] Potts, C. N. 1980. An algorithm for the single machine sequencing problem with precedence constraints. *Math. Programming Stud.* **13** 78–87.
- [20] Queyranne, M., Y. Wang. 1991. Single-machine scheduling polyhedra with precedence constraints. *Math. Oper. Res.* **16**(1) 1–20.
- [21] Sidney, J. B. 1975. Decomposition algorithms for single-machine sequencing with precedence constraints and deferral costs. *Oper. Res.* **23**(2) 283–298.
- [22] Smith, W. E. 1956. Various optimizers for single-stage production. *Naval Res. Logist. Quart.* **3**(1-2) 59–66.
- [23] Woeginger, G. J. 2003. On the approximability of average completion time scheduling under precedence constraints. *Discrete Appl. Math.* **131**(1) 237–252.