

A Subexponential Time Algorithm for Makespan Scheduling of Unit Jobs with Precedence Constraints

Jesper Nederlof



UU

Céline Swennenhuis



TU/e

Karol Węgrzycki



Saarland University



List of Open Problems by Garey and Johnson 1979

1. Graph Isomorphism
2. Subgraph Homeomorphism
3. Graph genus
4. Chordal graph completion
5. Chromatic index
6. Spanning tree parity problem
7. Partial order dimension
8. Precedence constrained 3-processor scheduling
9. Linear Programming
10. Total unimodularity
11. Composite number
12. Minimum length triangulation

List of Open Problems by Garey and Johnson 1979

1. Graph Isomorphism
2. Subgraph Homeomorphism
3. Graph genus
4. Chordal graph completion
5. Chromatic index
6. Spanning tree parity problem
7. Partial order dimension
8. Precedence constrained 3-processor scheduling
9. Linear Programming
10. Total unimodularity
11. Composite number
12. Minimum length triangulation

List of Open Problems by Garey and Johnson 1979

1. Graph Isomorphism
2. Subgraph Homeomorphism
3. Graph genus
4. Chordal graph completion
5. Chromatic index
6. Spanning tree parity problem
7. Partial order dimension
8. Precedence constrained 3-processor scheduling
9. Linear Programming
10. Total unimodularity
11. Composite number
12. Minimum length triangulation

List of Open Problems by Garey and Johnson 1979

1. Graph Isomorphism
2. Subgraph Homeomorphism
3. Graph genus
4. Chordal graph
5. Chromatic index
6. Spanning tree parity problem
7. Partial order dimension
8. Precedence constrained 3-processor scheduling
9. Linear Programming
10. Total unimodularity
11. Composite number
12. Minimum length triangulation

$2^{O((\log n)^3)}$ time
[Babai 2017]

List of Open Problems by Garey and Johnson 1979

1. Graph Isomorphism

2. Subgraph Homeomorphism

3. Graph genus

4. Chordal graph

5. Chromatic index

6. Spanning tree parity problem

7. Partial order dimension

8. Precedence constrained 3-processor scheduling

9. Linear Programming

10. Total unimodularity

11. Composite numbers

12. Minimum length triangulation

$2^{O((\log n)^3)}$ time
[Babai 2017]

$2^{O(\sqrt{n} \cdot \log n)}$ time
This talk

Literature overview $Pm|prec, p_j = 1|C_{\max}$

Literature overview $Pm|prec, p_j = 1|C_{\max}$

- NP-complete¹ $m = \#$ machines given *as input*

¹Jeffrey D. Ullman. *NP-complete scheduling problems*. Journal of Computer and System sciences, 10(3):384–393, 1975.

Literature overview $Pm|prec, p_j = 1|C_{\max}$

- NP-complete¹ $m = \#$ machines given *as input*

¹Jeffrey D. Ullman. *NP-complete scheduling problems*. Journal of Computer and System sciences, 10(3):384–393, 1975.

- Poly-time solvable² for $m = 2$

²M. Fujii, T. Kasami, and K. Ninomiya. *Optimal sequencing of two equivalent processors*. SIAM Journal on Applied Mathematics, 17(4):784–789, 1969.

Literature overview $Pm|prec, p_j = 1|C_{\max}$

- **NP-complete**¹ $m = \#$ machines given *as input*

¹Jeffrey D. Ullman. *NP-complete scheduling problems*. Journal of Computer and System sciences, 10(3):384–393, 1975.

- **Poly-time solvable**² for $m = 2$

²M. Fujii, T. Kasami, and K. Ninomiya. *Optimal sequencing of two equivalent processors*. SIAM Journal on Applied Mathematics, 17(4):784–789, 1969.

- **????** for $m \geq 3$ **constant** **OPEN**³

³Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

Literature overview $Pm|prec, p_j = 1|C_{\max}$

- **NP-complete**¹ $m = \#$ machines given *as input*

¹Jeffrey D. Ullman. *NP-complete scheduling problems*. Journal of Computer and System sciences, 10(3):384–393, 1975.

- **Poly-time solvable**² for $m = 2$

²M. Fujii, T. Kasami, and K. Ninomiya. *Optimal sequencing of two equivalent processors*. SIAM Journal on Applied Mathematics, 17(4):784–789, 1969.

- **????** for $m \geq 3$ **constant** **OPEN**³

³Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

Before:

$Pm|prec, p_j = 1|C_{\max}$ can be solved in $O\left(2^n \cdot \binom{n}{m}\right)$ time.

Our Result

Our result:

$Pm|prec, p_j = 1|C_{\max}$ can be solved in $\left(1 + \frac{n}{m}\right)^{O(\sqrt{nm})}$ time.

Our Result

Our result:

$Pm|prec, p_j = 1|C_{\max}$ can be solved in $\left(1 + \frac{n}{m}\right)^{O(\sqrt{nm})}$ time.

Corollary:

$Pm|prec, p_j = 1|C_{\max}$ can be solved in $2^{O(\sqrt{n} \cdot \log n)}$ time.

Our Result

Our result:

$Pm|prec, p_j = 1|C_{\max}$ can be solved in $\left(1 + \frac{n}{m}\right)^{O(\sqrt{nm})}$ time.

Corollary:

$Pm|prec, p_j = 1|C_{\max}$ can be solved in $2^{O(\sqrt{n} \cdot \log n)}$ time.

Key points in algorithm:

1. New **decomposition** of schedules
2. Use of **look-up table**
3. Use of **Dynamic Programming** in combining results

$$Pm|prec, p_j = 1|C_{\max}$$

$$Pm|prec, p_j = 1|C_{\max}$$



m identical
parallel machines

$$m = 2$$

$$Pm|prec, p_j = 1|C_{\max}$$

Given:


- n jobs of length 1

$$m = 2$$

$Pm|prec, p_j = 1|C_{\max}$

Given:

- n jobs of length 1
- A precedence graph G




$Pm|prec, p_j = 1|C_{\max}$

Given:

- n jobs of length 1
- A precedence graph G
- $T \in \mathbb{N}$

Q: Is there a schedule of makespan T ?




$Pm|prec, p_j = 1|C_{\max}$

Given:

- n jobs of length 1
- A precedence graph G
- $T \in \mathbb{N}$

Q: Is there a schedule of makespan T ?



time →

	1	2	3	4
1	a	b	d	f
2		c	e	

$Pm|prec, p_j = 1|C_{\max}$


Given:

- n jobs of length 1
- A precedence graph G
- $T \in \mathbb{N}$

Q: Is there a schedule of makespan T ?

Observation:


Jobs of length one \Rightarrow 'timeslots'



time \rightarrow

	1	2	3	4
1	a	b	d	f
2		c	e	


Definitions



Precedence Constraints Graph G

Definitions


- $G \Rightarrow$ partial order:
- $i < j$ if $(i, j) \in G$



Precedence Constraints Graph G

Definitions

- $G \Rightarrow$ partial order:
- $i < j$ if $(i, j) \in G$




Precedence Constraints Graph G

Definitions: Let A be a set of jobs.

$\text{pred}[A] =$

$\text{succ}[A] =$

Definitions



Precedence Constraints Graph G

$G \Rightarrow$ partial order:


- $i < j$ if $(i, j) \in G$

Definitions: Let A be a set of jobs.

$$\text{pred}[A] = \{x \mid \exists a \in A \text{ s.t. } x \preceq a\}$$

$$\text{succ}[A] = \{x \mid \exists a \in A \text{ s.t. } x \succeq a\}$$

Definitions



Precedence Constraints Graph G

$G \Rightarrow$ partial order:

- $i < j$ if $(i, j) \in G$

Definitions: Let A be a set of jobs.


$$\text{pred}[A] = \{x \mid \exists a \in A \text{ s. t. } x \preceq a\}$$

$$\text{sinks}(A) = \max\{A\}$$

$$\text{succ}[A] = \{x \mid \exists a \in A \text{ s. t. } x \succeq a\}$$

$$\text{sources}(A) = \min\{A\}$$

Definitions



Precedence Constraints Graph G

$G \Rightarrow$ partial order:

- $i < j$ if $(i, j) \in G$

Definitions: Let A be a set of jobs.

$$\text{pred}[A] = \{x \mid \exists a \in A \text{ s. t. } x \preceq a\}$$


$$\text{sinks}(A) = \max\{A\}$$

$$\text{succ}[A] = \{x \mid \exists a \in A \text{ s. t. } x \succeq a\}$$

$$\text{sources}(A) = \min\{A\}$$

Let $A = \{c, d\}$, then:

Definitions



Precedence Constraints Graph G

$G \Rightarrow$ partial order:

- $i < j$ if $(i, j) \in G$

Definitions: Let A be a set of jobs.

$$\text{pred}[A] = \{x \mid \exists a \in A \text{ s. t. } x \preceq a\}$$

$$\text{sinks}(A) = \max\{A\}$$


$$\text{succ}[A] = \{x \mid \exists a \in A \text{ s. t. } x \succeq a\}$$

$$\text{sources}(A) = \min\{A\}$$

Let $A = \{c, d\}$, then:

- $\text{pred}[A] = \{a, c, d\}$
- $\text{sinks}(\{a, c, d\}) = \{c, d\} = A$

Definitions



Precedence Constraints Graph G

$G \Rightarrow$ partial order:

- $i < j$ if $(i, j) \in G$

Definitions: Let A be a set of jobs.

$$\text{pred}[A] = \{x \mid \exists a \in A \text{ s.t. } x \preceq a\}$$

$$\text{sinks}(A) = \max\{A\}$$


$$\text{succ}[A] = \{x \mid \exists a \in A \text{ s.t. } x \succeq a\}$$

$$\text{sources}(A) = \min\{A\}$$

Let $A = \{c, d\}$, then:


- $\text{pred}[A] = \{a, c, d\}$
- $\text{sinks}(\{a, c, d\}) = \{c, d\} = A$
- $\text{succ}[A] = \{c, d, e, f\}$
- $\text{sources}(\{c, d, e, f\}) = \{c, d\} = A$


Zero-Adjusted Schedule (D&W)



Zero-Adjusted Schedule (D&W)

Assumption: $n = m \cdot T$

 = sinks(G)





Let $z \in [1, T]$ be the first moment with a sink.

D&W: W.m.a. Each job x after z is a sink or a successor of a job at time z .

Sink-Adjusted Schedule


Assumption: $n = m \cdot T$

 = $\text{sinks}(G)$




Sink-Adjusted Schedule

Assumption: $n = m \cdot T$

 = $\text{sinks}(G)$


$\text{succ}(H_1) \setminus (\text{succ}[H_2] \cup \text{sinks})$



$z \in [1, T]$ is a **sink moment** if there are both sinks and non-sinks at time z .

Sink-Adjusted Schedule


Assumption: $n = m \cdot T$




Sink-Adjusted Schedule

Assumption: $n = m \cdot T$

$$\text{succ}(H_1) \setminus (\text{succ}[H_2] \cup \text{sinks})$$

 = $\text{sinks}(G)$





$z \in [1, T]$ is a **sink moment** if there are both sinks and non-sinks at time z .

Sink-Adjusted Schedule

Assumption: $n = m \cdot T$

$succ(H_1) \setminus (succ[H_2] \cup sinks)$

 = $sinks(G)$




$z \in [1, T]$ is a **sink moment** if there are both sinks and non-sinks at time z .

No sinks in 


Sink-Adjusted Schedule

Assumption: $n = m \cdot T$



Sink-Adjusted Schedule

Assumption: $n = m \cdot T$



$$\sigma_i = \sigma(S_i, S_{i+1})$$

Summary of Algorithm

Input: X a set of jobs, m

Output: minimal makespan of X

1. For each $S_1, S_2 \in \binom{X}{\leq m}$ do:
 - a) Compute recursively makespan of $X \cap \sigma(S_1, S_2)$
2. Find a minimal combination of these subschedules
3. Return found schedule


Simple analysis: $2^n \text{poly}(n)$ time...?

Only n^{2m} subproblems

Use of DP in $\approx n^{2m}$ time

Describing subschedules

Let σ be a *subschedule*.



time

	1	2	3	4
1	a	b	d	f
2		c	e	

Describing subschedules

Let σ be a *subschedule*.


Jobs $V(\sigma)$ can be described as

$$V(\sigma) = \text{succ}[A] \cap \text{pred}[B]$$

where

A = minimal elements = sources

B = maximal elements = sinks



time

	1	2	3	4
1	a	b	d	f
2		c	e	

Describing subschedules

Let σ be a *subschedule*.



Jobs $V(\sigma)$ can be described as

$$V(\sigma) = \text{succ}[A] \cap \text{pred}[B]$$


where

A = minimal elements = sources


B = maximal elements = sinks



Sink-Adjusted Schedule




Sink-Adjusted Schedule



Lemma:

$$V(\sigma_i) \approx \text{succ}[S_i] \setminus (\text{succ}[S_{i+1}] \cup \text{sinks}(G))$$

Sink-Adjusted Schedule



Lemma:

$$\begin{aligned} V(\sigma_i) &\approx \text{succ}[S_i] \setminus (\text{succ}[S_{i+1}] \cup \text{sinks}(G)) \\ &= \text{succ}[S_i] \cap \text{pred}[B] \end{aligned}$$

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \min \text{makespan of } \text{succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \text{min makespan of succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

Lemma:

$$V(\sigma_i) = \text{succ}[S_i] \cap \text{pred}[B]$$

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \text{min makespan of succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

Lemma:

$$V(\sigma_i) = \text{succ}[S_i] \cap \text{pred}[B]$$

Remark: $|S_i| \leq m$!!

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \text{min makespan of } \text{succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

⇒ at most $\binom{n}{m} \cdot \text{\#different } B$

Lemma:

$$V(\sigma_i) = \text{succ}[S_i] \cap \text{pred}[B]$$

Remark: $|S_i| \leq m$!!

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \text{min makespan of } \text{succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

⇒ at most $\binom{n}{m} \cdot \text{\#different } B$

Lemma:

$$V(\sigma_i) = \text{succ}[S_i] \cap \text{pred}[B]$$

Remark: $|S_i| \leq m$!!

Case $|B| \leq \sqrt{n}$

⇒ only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ **different** B 's

Case $|B| > \sqrt{n}$

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \text{min makespan of } \text{succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

⇒ at most $\binom{n}{m} \cdot \text{\#different } B$

Lemma:

$$V(\sigma_i) = \text{succ}[S_i] \cap \text{pred}[B]$$

Remark: $|S_i| \leq m$!!

Case $|B| \leq \sqrt{n}$

⇒ only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ **different B 's** ✓

Case $|B| > \sqrt{n}$

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \text{min makespan of } \text{succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

⇒ at most $\binom{n}{m} \cdot \text{\#different } B$

Lemma:

$$V(\sigma_i) = \text{succ}[S_i] \cap \text{pred}[B]$$

Remark: $|S_i| \leq m$!!

Case $|B| \leq \sqrt{n}$

⇒ only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ **different B 's** ✓


Case $|B| > \sqrt{n}$

In next step: make \sqrt{n} jobs progress!

Algorithm

No sinks in 

$$V(\sigma) = \text{succ}[A] \cap \text{pred}[B]$$



Case $|B| \leq \sqrt{n}$

\Rightarrow only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B 's ✓


Case $|B| > \sqrt{n}$

In next step: make \sqrt{n} jobs progress!

Algorithm


No sinks in 

$$V(\sigma) = \text{succ}[A] \cap \text{pred}[B]$$



Each σ_i contains at most m sinks (and σ_ℓ is extremely easy)

Case $|B| \leq \sqrt{n}$

\Rightarrow only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B 's 

Case $|B| > \sqrt{n}$


In next step: make \sqrt{n} jobs progress!

Algorithm

No sinks in 

$$V(\sigma) = \text{succ}[A] \cap \text{pred}[B]$$

sinks(σ)



Each σ_i contains at most m sinks (and σ_ℓ is extremely easy)

$$|\sigma_i| \leq n - \sqrt{n} + m$$

Case $|B| \leq \sqrt{n}$

\Rightarrow only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B 's ✓

Case $|B| > \sqrt{n}$


In next step: make \sqrt{n} jobs progress!

Algorithm

No sinks in 

$$V(\sigma) = \text{succ}[A] \cap \text{pred}[B]$$

sinks(σ)



Each σ_i contains at most m sinks (and σ_ℓ is extremely easy)

$$|\sigma_i| \leq n - \sqrt{n} + m$$

Case $|B| \leq \sqrt{n}$


\Rightarrow only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B 's ✓

Case $|B| > \sqrt{n}$

In next step: make \sqrt{n} jobs progress! ✓

Algorithm

$\leq n^{2m} = n^{O(1)}$
children



● = $|B| \leq \sqrt{n}$
● = $|B| > \sqrt{n}$

Case $|B| \leq \sqrt{n}$

\Rightarrow only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B 's ✓

Case $|B| > \sqrt{n}$


In next step: make \sqrt{n} jobs progress! ✓

Algorithm

$\leq n^{2m} = n^{O(1)}$
children

Height $\leq \sqrt{n}$

$\leq (n^{O(1)})^{\sqrt{n}}$
 $= 2^{O(\sqrt{n} \cdot \log n)}$
nodes in tree!



● = $|B| \leq \sqrt{n}$
● = $|B| > \sqrt{n}$

Case $|B| \leq \sqrt{n}$

\Rightarrow only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B ✓

Case $|B| > \sqrt{n}$

In next step: make \sqrt{n} jobs progress! ✓

Algorithm

Store subproblem results in **lookup table**:

$$T(A, B) := \min \text{makespan of } \text{succ}[A] \cap \text{pred}[B]$$

Use of **lookup table**, how many *different* subproblems?

⇒ at most $2^{O(\sqrt{n} \cdot \log n)}$ trees with $2^{O(\sqrt{n} \cdot \log n)}$ nodes

⇒ $2^{O(\sqrt{n} \cdot \log n)}$ *different* subproblem!

Case $ B \leq \sqrt{n}$	Case $ B > \sqrt{n}$
⇒ only $\binom{n}{\sqrt{n}} = 2^{O(\sqrt{n} \cdot \log n)}$ different B ✓	In next step: make \sqrt{n} jobs progress! ✓

Summary of Algorithm

Input: A, B, m

Output: minimal makespan of $\text{succ}[A] \cap \text{pred}[B] = \text{Jobs}$

1. For each $S_1, S_2 \in \binom{\text{Jobs}}{\leq m}$ do:
 - a) Compute recursively makespan of $\text{Jobs} \cap (\text{succ}[S_1] \setminus (\text{succ}[S_2] \cup B))$
2. Find a minimal combination of these subschedules
3. Return found schedule

Conclusion

Conclusion

Main result:

$Pm|prec, p_j = 1|C_{\max}$ in $2^{O(\sqrt{n} \cdot \log n)}$ time.

Conclusion

Main result:

$Pm|prec, p_j = 1|C_{\max}$ in $2^{O(\sqrt{n} \cdot \log n)}$ time.

Key idea's:

- New [decomposition](#) of schedules
- Use of [look-up table](#)
- Use of [Dynamic Programming](#) in combining results

Conclusion

Main result:

$Pm|prec, p_j = 1|C_{\max}$ in $2^{O(\sqrt{n} \cdot \log n)}$ time.

Key idea's:

- New [decomposition](#) of schedules
- Use of [look-up table](#)
- Use of [Dynamic Programming](#) in combining results

Future Research:

$P3|prec, p_j = 1|C_{\max}$ in [quasi-polynomial](#) time?

Conclusion

Main result:

$P_m | prec, p_j = 1 | C_{\max}$ in $2^{O(\sqrt{n} \cdot \log n)}$ time.

Key idea's:

- New decomposition of schedules
- Use of look-up table
- Use of Dynamic Programming in combining results

Future Research:

$P_3 | prec, p_j = 1 | C_{\max}$ in quasi-polynomial time?

Thanks for your
attention!