

Multiple Machines

- Model Multiple Available resources
 - people
 - time slots
 - queues
 - networks of computers
- Now concerned with both allocation to a machine and ordering on that machine.

$P||C_{\max}$

NP-complete from partition.

Example

j	p_j
1	10
2	8
3	6
4	4
5	2
6	1

- What is the makespan on 2 machines?
- 3 machines ?
- 4 machines ?

Approximation Algorithms

- Cannot come up with an optimal solution in polynomial time
- Will look at **relative error** : $C_{\max}(\text{our algorithm})/C_{\max}(OPT)$
- Challenges:
 - Our algorithm's performance is different on different instances
 - We can't compute $C_{\max}(OPT)$

Approximation Algorithms

- Cannot come up with an optimal solution in polynomial time
- Will look at **relative error** : $C_{\max}(\text{our algorithm})/C_{\max}(OPT)$
- Challenges:
 - Our algorithm's performance is different on different instances
 - We can't compute $C_{\max}(OPT)$

Solution:

- We will use a worst case measure on performance
- We will use a lower bound on $C_{\max}(OPT)$

Approximation Algorithms

An algorithm A is a ρ approximation algorithm for a problem, if for all inputs

$$\frac{C_{\max}(A)}{C_{\max}(OPT)} \leq \rho$$

•

In addition, A must run in polynomial time.

We can't compute $C_{\max}(OPT)$.

Recipe:

- Instead, we compute a lower bound $LB(OPT)$, such that
 - $LB(OPT)$ is easy to compute
 - $LB(OPT) \leq C_{\max}(OPT)$.
- We then show that $C_{\max}(A) \leq \rho LB(OPT)$.

Combining the previous two steps, we have:

$$C_{\max}(A) \leq \rho LB(OPT) \leq \rho C_{\max}(OPT)$$

which can be rewritten as

$$\frac{C_{\max}(A)}{C_{\max}(OPT)} \leq \rho$$

Notes:

- Must come up with a good lower bound
- Can replace C_{\max} with any objective.

Lower Bounds for $P||C_{\max}$

- Average load
- Longest job

Lower Bounds for $P||C_{\max}$

- Average load – $\lceil \sum p_j / m \rceil$
- Longest job – $p_{\max} = \max_j \{p_j\}$

List Scheduling Algorithm

A Greedy Algorithm

1. Make a list of the jobs (in any order)
2. When a machine becomes available, schedule the next job on the list.

List Scheduling Algorithm

A Greedy Algorithm

1. Make a list of the jobs (in any order)
2. When a machine becomes available, schedule the next job on the list.

List Scheduling Algorithm

A Greedy Algorithm

1. Make a list of the jobs (in any order)
2. When a machine becomes available, schedule the next job on the list.

Analysis

- Let t be the last time at which all machines are busy.
- $t \leq \sum_j p_j / m$
- $C_{\max} \leq t + p_{\max} \leq \sum_j p_j / m + p_{\max}$.

Put this together with our lower bound:

$$C_{\max} \leq t + p_{\max} \leq \sum_j p_j / m + p_{\max} \leq 2LB \leq 2OPT$$

Improved Algorithm

- Schedule length is average load plus last job.
- When last job is small, the schedule is shorter.
- Force last job to be small – LPT (Longest Processing Time).

LPT is a 4/3-approximation for $P||C_{\max}$.

Proof Outline

- If last job is small ($\leq 1/3OPT$) then 4/3-approximation
- Otherwise, there are at most 2 jobs per machine and LPT is optimal.

Even better algorithms are possible: . A polynomial-time approximation scheme (PTAS) is an algorithm that, given fixed $\epsilon > 0$, returns at $(1 + \epsilon)$ -approximation in polynomial time. The running time can have a bad dependence on ϵ , such as $n^{O(1/\epsilon)}$.

$P||C_{\max}$ has a PTAS.

Precedence Constraints

- $P_{\infty}|\text{prec}|C_{\max}$ is known as project scheduling.
- $P|\text{prec}|C_{\max}$ has a 2-approximation.


What are good lower bounds for $P|\text{prec}|C_{\max}$?

Precedence Constraints

- $P_{\infty}|\text{prec}|C_{\max}$ is known as project scheduling.
- $P|\text{prec}|C_{\max}$ has a 2-approximation.


What are good lower bounds for $P|\text{prec}|C_{\max}$?

- Average load
- p_{\max}
- any path in the precedence graph
- the **critical path** is the longest path in the precedence graph.




Unit Processing Times

$P|p_j = 1, \text{prec}|C_{\max}$ is NP-hard.




Heuristics

- Critical Path (CP) rule
 - The job at the head of the longest string of jobs in the constraint graph has the highest priority
 - $P|p_j = 1, \text{tree}|C_{\max}$ is solved by CP.
- Largest Number of Successors First (LNS)
 - The job with the largest total number of successors in the constraint graph has highest priority.
 - For in-trees and chains, LNS is identical to CP
 - LNS is also optimal for $P|p_j = 1, \text{outtree}|C_{\max}$
- Generalization to arbitrary processing times is possible



Fixed Number of Processors

- $P2|p_j = 1, \text{prec}|C_{\max}$ is solvable in polynomial time
- $P3|p_j = 1, \text{prec}|C_{\max}$ is a big open question.



2 machines


A	B	E	H	G	I	J	
///	C	D	F	///	///	///	

$$LB = \sum p_i = 10 = 5$$

$$\overline{m} = \frac{10}{2} = 5$$

CP = 8

$C_{max} = 7$

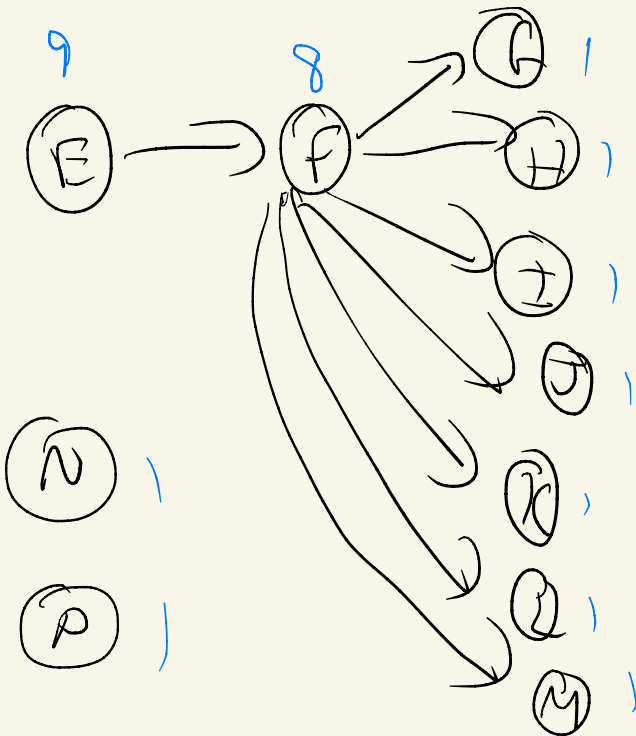



$C_{max} = 6$

Good ideas

- work on jobs on critical paths

- jobs w/ large #s of successors
successors




P3 | prec, $p_0=1$ | C_{max}

Gary + Johnson NP-complete
~ 1978

3 open problems

- Linear Programming (LP)
- Graph Isomorphism



2020 | 1 unresolved problem

Preemptions: $P|pmtn|C_{max}$

- McNaughton's wrap-around rule is optimal.

$$m=3$$

$$LB \quad \frac{\sum p_j}{m} \quad 31$$

Example

j	p_j
A	7
B	10
C	1
D	4
E	9

$$p_{max} = 10$$

$$\max(p_{max}, \frac{\sum p_j}{m}) \text{ is a lower bound}$$

Simple alg. for this problem

Preemptions: $P|pmtn|C_{max}$


- McNaughton's wrap-around rule is optimal.

$m=3$

$LB = \max\left(\frac{31}{3}, 10\right) = \frac{31}{3}$


Example

j	p_j
A	7
B	10
C	1
D	4
E	9



$LB = \max\left(\frac{31}{4}, 10\right) = 10$

time \rightarrow



Preemptions: $P|pmtn|C_{max}$

- McNaughton's wrap-around rule is optimal.

Example

j	p_j
A	7
B	10
C	1
D	4
E	9

Correctness

- Jobs have to fit
- Valid schedule: no job is on ≥ 2 machines at the same time.

time allowed \geq

$$\frac{\sum p_j}{m} \cdot m = \sum p_j$$

$$\text{time} \geq p_{\max}$$

LP for $P|pmtn|C_{max}$

amount of

Variables: x_{ij} is the time that job j runs on machine i . C_{max} is also a variable.

Constraints

- Each job runs for p_j units of time
- Each machine runs for at most C_{max} time.
- C_{max} is more than any processing time.

$$\min C_{max} \quad (1)$$

$$s.t. \quad (2)$$

$$\sum_{i=1}^m x_{ij} = p_j \quad j = 1 \dots n \quad (3)$$

$$\sum_{j=1}^n x_{ij} \leq C_{max} \quad i = 1 \dots m \quad (4)$$

$$\sum_{i=1}^m x_{ij} \leq C_{max} \quad j = 1 \dots n \quad (5)$$

$$x_{ij} \geq 0 \quad \forall i, j \quad (6)$$

Note that LP only assigns pieces of jobs to machines. Need to also assign jobs to times.

corresp. \rightarrow
to
example

$$\begin{array}{cccc} x_{1A} = 7 & x_{2B} = 6\frac{2}{3} & x_{2D} = 2\frac{2}{3} & x_{3E} = 9 \\ x_{1B} = 3\frac{1}{2} & x_{2C} = 1 & x_{3D} = 1\frac{1}{3} & \end{array}$$

Machines with speeds – $Q|\text{pmtn}|C_{\max}$

- Machines M_1, \dots, M_m with speeds v_1, \dots, v_m .
- Assume wlog that $v_1 \geq v_2 \geq v_m$
- Assume wlog that $p_1 \geq p_2 \geq p_n$
- If a job runs for one unit of time on machine M_i , it uses up v_i units of processing.
- If job j runs on machine M_i , then it takes p_j/v_i time units to complete.

Example

j	p_j
A	20
B	16
C	2
D	1

3 machines
 $v_1 = 4$
 $v_2 = 2$
 $v_3 = 1$

Job A can
run
2 units of time on M_1 ,
 $20 - 4(2) = 12$
+ 6 units of time on M_2

What are the lower bounds

Lower bounds for $Q|pmtn|C_{max}$

- What is the analog of p_{max} ?
- What is the analog of average load ?
- Are there others ?

$$P_i / v_i$$

(biggest job on fastest machine)

$$\frac{\sum_j P_j}{\sum_i v_i}$$

(total processing time over sum of speeds)

$$\frac{P_1}{v_1} = \frac{20}{4} = 5$$

$$\frac{\sum P_j}{\sum v_i} = \frac{39}{7} \rightarrow 5.57$$

Example



$$P_1 = 20$$

$$P_2 = 16$$

$$v_1 = 4$$

$$v_2 = 2$$

OPT



Lower bounds for $Q|\text{pmtn}|C_{\max}$

- What is the analog of p_{\max} ? - p_1/v_1
- What is the analog of average load ? - $\sum p_j / \sum v_i$
- Are there others ? - Yes

General Lower Bound

$$C_{\max} \geq \max \left(\frac{p_1}{v_1}, \frac{p_1 + p_2}{v_1 + v_2}, \dots, \frac{\sum_{j=1}^{m-1} p_j}{\sum_{i=1}^{m-1} v_i}, \frac{\sum_{j=1}^n p_j}{\sum_{i=1}^m v_i} \right)$$

Handwritten notes illustrating the derivation of the lower bound:

$\frac{p_1}{v_1} \rightarrow \frac{p_1 + p_2}{v_1 + v_2} \rightarrow \frac{p_1 + p_2 + p_3}{v_1 + v_2 + v_3}, \dots, \frac{p_1 + p_2 + p_3 + \dots + p_{m-1}}{v_1 + v_2 + v_3 + \dots + v_{m-1}}, \frac{p_1 + p_2 + p_3 + \dots + p_n}{v_1 + v_2 + v_3 + \dots + v_m}$

The above sequence shows the progression of the lower bound from individual jobs to the total sum of jobs. The final term represents the average load.

Lower Bound

$$C_{\max} \geq \max \left(\frac{p_1}{v_1}, \frac{p_1 + p_2}{v_1 + v_2}, \dots, \frac{\sum_{j=1}^{m-1} p_j}{\sum_{i=1}^{m-1} v_i}, \frac{\sum_{j=1}^n p_j}{\sum_{i=1}^m v_i} \right)$$

What is the lower bound for our example?

Can we achieve this lower bound?

$$p = (20, 16, 9, 1)$$

$$v = (4, 2, 1)$$

$$C_{\max} \geq \max \left(\frac{20}{4}, \frac{20+16}{4+2}, \frac{20+16+9+1}{4+2+1} \right)$$


$$\max(5, 6, 7)$$

LRPT-FM

Longest Remaining Processing Time on Fastest Machines

Example 1

	j	p_j
4/8	A	20
8/12	B	16
12/16	C	2
16/20	D	1



$v = (4, 2, 1)$
 $A \ 12$
 $B \ 12$
 $A \ 8$
 $B \ 6$
 1 2 3

$C_{max} = 6$

Example 2

j	p_j
A	20
B	16
C	12
D	1

after ϵ time
 $rem(A) = 12 - 4\epsilon$
 $rem(B) = 12 - 2\epsilon$

Notes:

- LRPT-FM is optimal in continuous time
- LRPT-FM is near optimal in discrete time, for small time steps.

LRPT-FM

Longest Remaining Processing Time on Fastest Machines

Example 1

j	p_j
A	20
B	16
C	2
D	1

$$C_{\max} \geq \left(\frac{20}{4}, \frac{20+16}{4+2}, \frac{20+16+12+1}{4+2+1} \right)$$

$$= 5, 6, \frac{49}{7} = 7$$

A	20	16	12	8	7	3	1	2
B	16	14	12	10	6	5	3	
C	12	11	10	9	7	5		
D	1	1	1	1	1	1	1	

$v = (4, 2, 1)$

C	C	C	A	B	A	B
B	B	B	C	C	C	A
A	A	A	B	A	B	C


Example 2

j	p_j
A	20
B	16
C	12
D	1

Notes:

- LRPT-FM is optimal in continuous time
- LRPT-FM is near optimal in discrete time, for small time steps.

Continuous fire



2 \longleftrightarrow 3

switch infinitely often
between the 2 & 3