

Lecture 4

Friday, March 24, 2017 10:42 AM

Problem

① Makespan Minimization on Identical Machines

Input:

- n jobs with processing times p_1, p_2, \dots, p_n

- m machines.

Identical: Each job j takes same time on each machine.

If a set $S \subseteq \{1, \dots, n\}$ is scheduled on a machine, the total processing time is the sum of processing times.

Output: Allocation of jobs to machines.

Objective: Minimize makespan : max proc. time on any m/c.

Formally, find a partition S_1, S_2, \dots, S_m of $\{1, \dots, n\}$ s.t.

$$\max_{i=1}^m \sum_{j \in S_i} p_j \quad \text{is minimized.}$$

Graham Notation

This problem is $P \parallel C_{\max}$

Identical machines \rightarrow P
 no extra constraints. \rightarrow \parallel
 Makespan \rightarrow C_{\max}

For instance, we could have a machine dependent running time, p_{ij} - job j 's proc. time on machine i , and we could've wanted to minimize the same obj.

That prob is called $R \parallel C_{\max}$.

Today, we stick to $P \parallel C_{\max}$

Algorithm (The "oldest" appx alg) '1966
[Graham]

- Order the jobs in any order.

- Schedule job j on the machine which has the least load so far.

Analysis

- Lower Bounds on OPT

$$\left. \begin{array}{l} P_{\max} := \max_j P_j \\ \bar{P} := \frac{1}{m} \sum_{j=1}^n P_j \end{array} \right\} L := \max(P_{\max}, \bar{P})$$

Obs : $OPT \geq L$

We will use this easier lower bound to compare our algo. It may seem like we're taking a hit, **but we have a better handle on it.**

Back to analysis :

Let i be the most loaded m/c. Consider the last job j being added to i . Let P be the load on m/c i just before j was added.

$$- \text{ALG} = P + P_j \leq P + L$$

- $P \leq$ load on any other m/c
when j was added

$$\Rightarrow P \leq \frac{1}{m} \sum_{j: j \leq i} P_j \leq L$$

$$\therefore ALG \leq 2L$$

Theorem: The above algorithm is a
factor 2 appx. alg. for $P||C_{max}$.

A slightly "simpler" problem: $P2||C_{max}$

So, $m=2$. only 2 m/c.

The problem is still NP-hard.

(Do you see this?)

Classification:

- Fix any $\epsilon > 0$.

- Job j is big if $P_j > \epsilon L$
small, o/w.

Partitioning Output Space:

• B be the set of big jobs.

- Any schedule in particular partitions $B = B_1 \cup B_2$
- Let (B_1^*, B_2^*) be the partition of B in the optimal schedule
- ⊛ We can afford to **guess** this partition.

Why? How many partitions are there?

$$2^{|B|}$$

How large is $|B|$?

$$2L \geq \sum_{j \in B} p_j \geq |B| \cdot \epsilon L \Rightarrow |B| \leq \frac{2}{\epsilon}$$

$$\therefore \text{The \# of partitions of } |B| \leq 2^{2/\epsilon} = O_{\epsilon}(1)$$

\therefore By enumerating, we may assume we have guessed / know (B_1^*, B_2^*)

fancy not" for saying it's a const dependent on ϵ .

Note: in the future, whenever we say

"guess", be aware how much time this guess takes. In this case, the guess is via enumeration and takes $2^{O(1/\epsilon)}$ time.

Algo:

① Guess (B_1^*, B_2^*) // $2^{O(1/\epsilon)}$ time via enum.

② Order small jobs in any order and schedule on the least loaded m/c ... exactly like before.

Analysis

Two cases.

① The max-loaded m/c gets no small jobs. In that case,

$$ALG \leq \max(p(B_1^*), p(B_2^*)) \leq OPT$$

② The max-loaded m/c gets at least one small job j . Now we can repeat the above argument ...

$$ALG \leq P + p_j \leq P + \epsilon L$$

$$\leq (1 + \epsilon) L$$

↑
gain

$$\leq (1 + \epsilon) L$$

$$\leq (1 + \epsilon) \text{OPT}$$

gain

Theorem: For any $\epsilon > 0$, the above algorithm is an $(1 + \epsilon)$ -approximation running in time $\text{poly}(n) \cdot 2^{O(1/\epsilon)}$

Definitions

Polynomial Time Appx Scheme (PTAS):

An alg. which takes an extra parameter $\epsilon > 0$ as input and returns an $(1 + \epsilon)$ -factor appx in time which is a polynomial in n as long as ϵ is a constant.

Example running times: $n^{1/\epsilon}$, $n^{2^{1/\epsilon}}$
 $n^2 \cdot 2^{1/\epsilon}$, n/ϵ

Fully Polynomial Time Appx Scheme (FPTAS)

These are PTAS .. but their

These are PTASes, but their running times are $\text{poly}(n, 1/\epsilon)$.

So, even if $\epsilon \approx \frac{1}{n^{10}}$, the scheme runs in polynomial time.

An FPTAS is the best you can hope for for an NP-hard problem.

So our above theorem can be succinctly stated as

Thm: $P2||C_{\max}$ has a PTAS

$P||C_{\max}$ with many machines

① (Multiplicative Scaling)

"Round up" all job sizes p_j to the smallest power of $(1+\epsilon)$ which is larger than p_j .

$$- \text{OPT}(\mathcal{I}_{\text{old}}) \leq \text{OPT}(\mathcal{I}_{\text{new}}) \leq (1+\epsilon) \cdot \text{OPT}(\mathcal{I}_{\text{old}})$$

- Given any soln to \mathcal{I}_{new} of makespan

$T \Rightarrow$ a soln to Job of $\leq T$ makespan.

② (Guessing OPT) Till now we assumed we didn't know OPT... the number. But we keep working with a 'guess' and bump it by a factor $(1+\epsilon)$ if the final soln. isn't within $\leq (1+\epsilon)$ times guess.

③ (Bucketing Big Jobs)

- Job j is big if $p_j > \epsilon \cdot \text{OPT}$

note this is a power of $(1+\epsilon)$...

guess.

- $B \equiv$ Set of big jobs.

- $B_t \subseteq B$ with the same proc. time
 $1 \leq t \leq N$

- How big is N ? $\log_{1+\epsilon} (1/\epsilon)$

$$\approx \frac{1}{\epsilon} \log \frac{1}{\epsilon}$$

$\lg(1+x) \approx x$
if x is small...

④ (Profile of big jobs)

Every feasible schedule gives each machine a "profile"

$$\vec{v} = (v_1, v_2, \dots, v_N)$$

where $v_t \in \mathbb{Z}_{\geq}$

denotes the # of jobs of B_t is given to this m/c.

$$\text{Each big job } \geq \varepsilon \cdot \text{OPT} \Rightarrow \sum_{t=1}^N v_t \leq \frac{1}{\varepsilon}$$

How many feasible profiles are there?

If \mathcal{V} is collection of all feas. profiles,

$$M = |\mathcal{V}| \leq \left(\frac{1}{\varepsilon}\right)^N \approx 2^{\underbrace{\frac{1}{\varepsilon} \log^2 \frac{1}{\varepsilon}}_{\text{is a constant, believe it or not.}}}$$

⑤ (Histogram of profiles)

Let's order the profiles of \mathbb{V} as
 $\vec{v}^{(1)}, \vec{v}^{(2)}, \dots, \vec{v}^{(M)}$

Any feasible schedule allocates each machine one of these profiles.

Let h_i be the # of m/c's which pick the profile $\vec{v}^{(i)}$.

Here we use the "parallelness" of machines. It doesn't matter who chooses which of these h_i profiles of type $\vec{v}^{(i)}$.

Consistency Conditions:

- $\sum_{i=1}^M h_i = m$

- $\sum_{i=1}^M h_i v_t^{(i)} = |B_t| \quad \forall 1 \leq t \leq N$

of big jobs
of type B_t alloc.
among the m/c's.

⑥ (Guessing OPT's profile)

Let (h_1^*, \dots, h_m^*) be the histogram of OPT. Claim: We can 'guess' h^* by enumeration.

How many consistent histograms are there?

$$\begin{aligned} &\leq m^M \\ &\leq m \cdot 2^{\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon}} \} \text{constant!} \end{aligned}$$

Once we guess OPT's histogram, we give h_i^* arbitrary m/c's the profile $\triangleright^{(i)}$. Since it's consistent, all big jobs are allocated.

The small jobs we allocate GREEDILY. Just as in the $P2 \parallel C_{max}$ case, we get an $(1 + \epsilon) \cdot \text{OPT}$ makespan.