## Load Balancing : 2-approximation to PTAS ${ }^{1}$

- In this lecture, we consider a basic load balancing problem. In this problem we are given $n$ jobs with processing times $p_{1}, \ldots, p_{n}$. We have $m$ machines and we have to schedule the jobs on these machines. The objective is to minimize the time taken for the last job to complete. This value, in the job-scheduling parlance, is called the makespan of the schedule, and this problem is often called makespan minimization on identical machines.
- Let us make one observation. Since we care only about the completion time of the last job, on any single machine it doesn't really matter in which order they are run. If a subset $S \subseteq[n]$ is allocated to a machine, the completion time of the last job among these will be $\sum_{j \in S} p_{j}$. Once we realize this, we see that the load balancing problem is the following min-max allocation problem

$$
\text { Find partition } S_{1}, \ldots, S_{m} \text { of }[n] \text { to, } \quad \text { minimize } \max _{i=1}^{m} p\left(S_{i}\right):=\sum_{j \in S_{i}} p_{j}
$$

- We begin by first noting a simple 2 -approximation for the problem. This algorithm is called the list scheduling algorithm and is ridiculously simple : consider the jobs in any order and then assign it to the least loaded machine at that point.

```
procedure LIST SCHEDULING( \(n\) jobs with proc times \(p_{1}, \ldots, p_{n}\); positive integer \(m\) ):
    Consider jobs in any order \(\left(p_{1}, \ldots, p_{n}\right)\).
    For each machine \(i \in[m]\) maintain load \(_{i}\) initialized to 0 .
    for \(j=1\) to \(n\) do:
        Find machine \(i\) with least load \(_{i}\).
        Assign \(j\) to \(i\), that is, set \(\sigma(j)=i\).
        \(\operatorname{load}_{i} \leftarrow \operatorname{load}_{i}+p_{j}\)
    return \(\sigma\).
```

Theorem 1. List Scheduling is a 2 -approximation algorithm.

Proof. Let opt denote the makespan of the optimal schedule. We now state two simple lower bounds on opt. We then prove the theorem by comparing the performance of List Scheduling to these lower bounds.
a. opt $\geq L_{1}:=\max _{j} p_{j}$. The largest job must be assigned somewhere.
b. opt $\geq L_{2}:=\frac{1}{m} \sum_{j=1}^{n} p_{j}$. The sum of all jobs is at most $m$ times the maximum.

[^0]Let $L:=\max \left(L_{1}, L_{2}\right)$. We claim that alg, the makespan of $\sigma$ returned by List Scheduling is at most $\left(2-\frac{1}{m}\right) \cdot L$. To see this, consider the job $\ell$ whose completion time is alg. Suppose it was assigned to machine $i$. Let $X$ be the set of other jobs assigned to $i$. Then, alg $=p(X)+p_{\ell}$.
Now, since $\ell$ was assigned to machine $i$, at the time $\ell$ was being considered every other machine $i^{\prime}$ must have had load $\geq p(X)$. If not, then $\ell$ would've been assigned to $i^{\prime}$ instead of $i$. In particular, the load of every machine at the end of the algorithm is $\geq p(X)$. This implies that the sum of processing times of every job other than $\ell$ is at least $m$ times $p(X)$. That is,

$$
\left(\sum_{j=1}^{n} p_{j}-p_{\ell}\right) \geq m \cdot p(X) \Rightarrow p(X) \leq \frac{1}{m} \sum_{j=1}^{n} p_{j}-\frac{p_{\ell}}{m}
$$

Adding $p_{\ell}$ to both sides, we get alg $\leq \underbrace{\frac{1}{m} \sum_{j=1}^{n} p_{j}}_{=L_{2}}+\left(1-\frac{1}{m}\right) \underbrace{p_{\ell}}_{\leq L_{1}} \leq\left(2-\frac{1}{m}\right) L$.

- Greedy List Scheduling. When one stares at the proof above, one can say more to improve the approximation factor. Suppose we could ensure that $p_{\ell} \leq \varepsilon \cdot$ opt for some parameter $\varepsilon$, then in fact the approximation factor of LIST SCHEDULING would be $\leq(1+\varepsilon)$.
This suggests the following modification : instead of considering the jobs in an arbitrary order, perhaps we consider them in decreasing order of $p_{j}$ 's. The hope is that the "last job" $p_{\ell}$ would have "low" value. Indeed, this greedy list-scheduling algorithm is a $\frac{4}{3}$-approximation. Suppose it wasn't. Then the previous algorithm implies $p_{\ell}>$ opt $/ 3$. Now, consider just the jobs $\{1, \ldots, \ell\}$. Note that these jobs can be allocated in a way that every machine gets at most 2 jobs and the load on every machine is $\leq$ opt. The following exercise completes the proof.

Exercise: Suppose there is a schedule which assigns at most two jobs in each machine and has makespan $\leq L$, and furthermore every job has processing time $>L / 3$. Then prove that the greedy list scheduling algorithm's makespan is $\leq L$.

## Theorem 2. Greedy List Scheduling gives a 4/3-approximation.

- Big Jobs and Small Jobs. The above idea can be taken one step ahead to get what is called a polynomial time approximation scheme, or PTAS, for the makespan minimization problem. A PTAS takes in a parameter $0<\varepsilon<1$, and returns a $(1+\varepsilon)$-approximation, but the running time is $n^{f(\varepsilon)}$ for some parameter $\varepsilon$. Informally, for any constant $\varepsilon$, there is a poly time $(1+\varepsilon)$-approximation algorithm.
For the time being, let us assume we know the value of opt. Later we see why this is not an issue. Given this value, we partition the jobs in $[n]$ into two classes : small jobs $S=\left\{j: p_{j} \leq \varepsilon \cdot \mathrm{opt}\right\}$, and big jobs $B=\left\{j: p_{j}>\varepsilon \cdot\right.$ opt $\}$. The proof of Theorem 1 implies the following.

Lemma 1. Suppose we have an assignment $\sigma_{B}: B \rightarrow[m]$ of the big jobs with makespan $\operatorname{alg}_{B}$. Then running the List Scheduling algorithm on the jobs on $S$ after having assigned jobs in
$B$ using $\sigma_{B}$ returns an assignment $\sigma$ with makespan $\leq \max \left(\operatorname{alg}_{B},(1+\varepsilon) \mathrm{opt}\right)$.
Proof. Let $\ell$ be the job to end last in $\sigma$ on some machine $i$. If $\ell \in B$, then $\operatorname{load}_{i} \leq \operatorname{alg}_{B}$, as no small jobs were added to $i$ after $B$ was assigned. Otherwise, $\ell \in S$, and then the same argument as in Theorem 1 holds. If $X$ is the set of remaining jobs assigned to $i$, the fact that $\ell$ was assigned to $i$ implies $p(X) \leq \frac{1}{m}\left(\sum_{j \in B \cup S} p_{j}-p_{\ell}\right)<$ opt , and thus, $p(X)+p_{\ell} \leq(1+\varepsilon)$ opt where the last inequality uses that $\ell \in S$.

What the above lemma implies that we can forget about the small jobs. If we can design a polynomial time algorithm to assign $\sigma: B \rightarrow[m]$ with makespan $\operatorname{alg}_{B} \leq(1+\varepsilon)$ opt, then we would have a $(1+\varepsilon)$-approximation for the whole instance. We now proceed to take care of big jobs.
Before we move on, we make one simple observation. The number of big-jobs $|B| \leq \frac{m}{\varepsilon}$. This is simply because opt $\geq \frac{1}{m} \sum_{j \in B} p_{j} \geq \frac{\varepsilon \text { opt. } \cdot|B|}{m}$. We will also assume $\varepsilon<1 / 2$ and all logarithm bases, unless not mentioned, are base $e$.

- Scaling and Grouping. For every $j \in B$, we scale up its processing time to the nearest power of $(1+\varepsilon)$. That is, we find the smallest integer $i$ such that $p_{j} \leq(1+\varepsilon)^{i}$, and change $j$ 's processing time to $p_{j}^{\prime}:=(1+\varepsilon)^{i}$. Note that $p_{j} \leq p_{j}^{\prime} \leq(1+\varepsilon) p_{j}$. Therefore, the optimum value opt' of the makespan wrt these modified processing times also satisfies opt $\leq \mathrm{opt}^{\prime} \leq(1+\varepsilon)$ opt. We now focus on solving the modified instance optimally, and henceforth abuse notation and use the same $p_{j}$ 's to denote the modified processing times.
What does this scaling buy us? It reduces the number of types of jobs. Since for every $j$ we have $\varepsilon \circ \mathrm{pt} \leq p_{j} \leq \operatorname{opt}(1+\varepsilon)$, and every $p_{j}$ is a power of $(1+\varepsilon)$, the number of different processing times is at most $t:=t_{\varepsilon} \leq \log _{1+\varepsilon}\left(1+\frac{1}{\varepsilon}\right) \leq \frac{2 \log (2 / \varepsilon)}{\varepsilon}$ which $^{2}$ is $\tilde{O}(1 / \varepsilon)$. In particular, if $\varepsilon$ is a constant, this is a constant.
Let the $t$ different processing times be $Q:=\left(q_{1} \leq \cdots \leq q_{t}\right)$ where $q_{1}=(1+\varepsilon)^{a}$ and $q_{t}=$ $(1+\varepsilon)^{a+t-1}$ for some integer $a$. Let $B_{s}:=\left\{j \in B: p_{j}=q_{s}\right\}$ for $1 \leq s \leq t$. We have thus partitioned the big jobs into $\tilde{O}(1 / \varepsilon)$-classes $B_{1}$ to $B_{t}$.
- Feasible Configurations and a (slow) PTAS. Since there are only $t$ kinds of jobs, any assignment of jobs to a single machine can be described by a vector $\mathbf{c}=(\mathbf{c}[1], \mathbf{c}[2], \ldots, \mathbf{c}[t])$ where each $\mathbf{c}[t]$ is a non-negative integer indicating the number of big jobs of type $t$. This vector is a feasible configuration if and only if $\sum_{s=1}^{t} q_{s} \mathbf{c}[s] \leq$ opt. Recall that opt is a guess of the optimum we have made up front. Observe that

$$
\begin{equation*}
\text { For any feasible configuration } \mathbf{c}, \quad \forall 1 \leq s \leq t, \quad \mathbf{c}[s] \leq \frac{1}{\varepsilon} \tag{1}
\end{equation*}
$$

This is simply because the smallest $q_{s} \geq \varepsilon \circ \mathrm{pt}$. We use $\mathcal{C}$ to denote the set of all feasible configurations.
Figuring out the allocation of big jobs to the $m$ machines is therefore equivalent to finding $m$ feasible configurations such that for every type $B_{s}$ where $1 \leq s \leq t$ we have allocated all jobs in $B_{s}$. That is, the problem of assigning the big jobs can be re-cast as

$$
\begin{equation*}
\text { Find } \mathbf{c}_{1}, \ldots, \mathbf{c}_{m} \in \mathcal{C} \text { with possible duplicates such that, } \quad \forall 1 \leq s \leq t, \quad \sum_{i=1}^{m} \mathbf{c}_{i}[s]=\left|B_{s}\right| \tag{2}
\end{equation*}
$$

[^1]At this point, we already have enough ammo for a PTAS. First we observe that the number of feasible configurations can be upper bounded. (1) implies that $|\mathcal{C}| \leq\lceil 1 / \varepsilon\rceil^{t} \leq(2 / \varepsilon)^{(2 / \varepsilon) \log (2 / \varepsilon)}=$ $2^{O\left(\log ^{2}(1 / \varepsilon) / \varepsilon\right)}=: f(\varepsilon)$ which, although large, is a constant whenever $\varepsilon$ is. Next, we note that the number of ways $m$ things can be chosen from $\mathcal{C}$ things with repetition allowed and order doesn't matter is $\binom{m+|\mathcal{C}|}{|\mathcal{C}|} \leq m^{f(\varepsilon)}$. And thus, we could enumerate over all such choices and find one which satisfies the condition of (2). If our guess of opt is correct, one such choice is guaranteed to exist, and then along with Lemma 1, we would get an assignment of all the jobs with makespan $\leq(1+\varepsilon) \cdot$ opt. The above enumeration algorithm is wasteful (as the astute reader has noticed) and one should indeed use dynamic programming. We leave this as an exercise for the reader.

Exercise: Use dynamic programming to solve (2) in time $|B|^{O(t)}=m^{\tilde{O}(1 / \varepsilon)}$.

- Guessing of opt. Till now, we have assumed we know the value of opt. In reality, we work with a guess of opt ${ }_{g}$ and we want to find the smallest such guess for which the (2) can be solved. We do so using binary search between $L$ and $2 L$ with a granuality of $\varepsilon L \leq \varepsilon \circ p t$.

```
procedure PTAS FOR LOAD BALANCING( \(n\) jobs with proc times \(p_{1}, \ldots, p_{n}\); positive
integer \(m\); parameter \(\varepsilon\) ):
    \(L \leftarrow \max \left(\max _{j} p_{j}, \frac{1}{m} \sum_{j=1}^{n} p_{j}\right)\).
    \(\mathrm{lb} \leftarrow L ; \mathrm{ub} \leftarrow 2 L\).
    while ub-lb> \(>L\) do:
        Set guess opt \({ }_{g} \leftarrow \frac{\mathrm{lb}+\mathrm{ub}}{2}\)
        Partition jobs into big and small using opt \({ }_{g}\).
        Scale \(p_{j}\) 's of \(B\) into powers of \((1+\varepsilon)\) and obtain the groups \(B_{1}, \ldots, B_{t}\).
        Form the configurations \(\mathcal{C}\) using the guess opt \({ }_{g}\).
        Solve (2). This either returns an assignment \(\sigma_{B}\) of big items, or it returns infeasible.
        In case it returns an assignment \(\sigma_{B}\), use List Scheduling to allocate the small
jobs, and set ub \(\leftarrow\) opt \(_{g}\).
    Otherwise, our guess is too low, and set \(\mathrm{lb} \leftarrow \mathrm{opt}_{g}\).
```

Steps (6) and (7) takes $O(n)$ time, Step (8) takes $O(|\mathcal{C}|)=2^{\tilde{O}(1 / \varepsilon)}$ time, and Step (9), which is the bottleneck step, takes $m^{\tilde{O}(1 / \varepsilon)}$ time using dynamic programming. The number of iterations of the while loop is $O(\log (1 / \varepsilon))$. The final guess is guaranteed to be opt ${ }_{g} \leq \mathrm{opt}+\varepsilon L \leq(1+\varepsilon)$ opt. The makespan of the big jobs at most $(1+\varepsilon) \cdot$ opt $_{g}$ (because of the scaling of the jobs), and by Lemma 1 one gets that the algorithms is a $(1+\varepsilon)^{2} \leq(1+3 \varepsilon)$-approximation.

Theorem 3. PTAS FOR LOAD BALANCING returns an $(1+3 \varepsilon)$-approximation in time $O(n)+$ $2^{\tilde{O}(1 / \varepsilon)}$ plus the time taken to solve (2).

In the next bullet point, we use the fact that the number of feasible allocations is only a constant (when $\varepsilon$ is a constant) to obtain an even faster PTAS. Indeed, the dependence will be reduced to $f(\varepsilon)+\operatorname{poly}(n)$. Note this is much better than $m^{f(\varepsilon)}$. Such a PTAS is called an efficient PTAS, or simply EPTAS.

- Geometry of Numbers and EPTAS*. We now discuss how to solve (2) faster using hammers from the algorithmic theory of numbers. Notice that the recasting (2) is a special case of the following problem : given a collection $\mathcal{V}=\left(\mathbf{v}_{1}, \ldots, \mathbf{v}_{N}\right)$ of $N$ vectors in $d$ dimensions with non-negative integer entries, and a targer $d$-dimension integer vector $\mathbf{b}$, the problem is to decide whether there exists non-negative integers $\mathbf{x}:=\left(x_{1}, \ldots, x_{N}\right)$ such that the $\mathbf{x}$-combination of $\mathcal{V}$ gives the target $\mathbf{b}$.

$$
\begin{equation*}
\exists ? x_{1}, \ldots, x_{N} \in \mathbb{Z}_{\geq 0}: \sum_{i=1}^{N} x_{i} \mathbf{v}_{i}=\mathbf{b}, \quad \text { and if one exists return } x_{i} \text { 's } \tag{ILP}
\end{equation*}
$$

This question is a question in algorithmic lattice theory which is a rich and deep field. One striking result is that the above question can be solved in time $N^{O(N)} \cdot \operatorname{poly}(s)$ time, where $s$ is the number of bits required to describe the input. For our discussion, let's ignore this dependence on $s$.

Theorem 4 (Kannan [7],improving on [8]). The problem (ILP) can be solved in $2^{O(N \log N)}$ time.
For our problem, we have $\mathcal{V}=\mathcal{C}, N=|\mathcal{C}|=2^{\tilde{O}(1 / \varepsilon)}$, and so the above algorithm implies that we can solve (2) in $2^{2^{\sigma(1 / \varepsilon)}}$-time. Plugging this into the guessing procedure, we get an $g(\varepsilon)+\operatorname{poly}(n)$ time algorithm, where $g(\varepsilon)$ is a doubly-exponential function of $\frac{1}{\varepsilon}$.
One can do even better by using another result from geometry of integer vectors.
Theorem 5 (Eisenbrand and Shmonin [3]). If the largest entry $\max _{1 \leq i \leq N} \max _{1 \leq j \leq d} \mathbf{v}_{i}[j] \leq M$, then if there is a feasible solution to (ILP), then there exists a feasible solution to (ILP) with at most $k=O(d \log (d M))$ entries not zero.

For (2), we have that $d=t=\tilde{O}(1 / \varepsilon)$ and $M=\max _{\mathbf{c} \in \mathcal{C}} \max _{s \in[t]} \mathbf{c}[s]$ which, from (1), we know is at most $\leq \frac{1}{\varepsilon}$. Therefore, Theorem 5 tells us that if there is a solution to (2), then there is one involving $\leq \tilde{O}(1 / \varepsilon)$ configurations. We now simply enumerate over all so many configurations of $|\mathcal{C}|$, and then apply Theorem 4 on that subset. The total time is therefore at most $\left(\begin{array}{c}\tilde{O}(1 / \varepsilon)\end{array}\right) \cdot 2^{\tilde{O}(1 / \varepsilon)} \leq 2^{\tilde{O}\left(1 / \varepsilon^{2}\right)}$.

Theorem 6. For any $0<\varepsilon<1$, there is an algorithm obtaining an $(1+\varepsilon)$-approximation in $O(n)+2^{\tilde{O}\left(1 / \varepsilon^{2}\right)}$ time. In particular, for a constant $\varepsilon$, this is a linear time algorithm.

## Notes

The load balancing problem described here is one of the classic job-scheduling problems. This problem is often denoted as $P \| C_{\max }$ where the $P$ denotes all machines are identical, and $C_{\text {max }}$ denotes we are looking at the makespan (min-max) objective. A whole slew of scheduling problems are out there, and we point the interested reader to a recent survey [2] by Bansal. The 2 -approximate list scheduling algorithm is in the paper [4] by Graham. The first PTAS for this problem can be found in the paper [5], and the first EPTAS using ideas from the geometry of numbers is in the paper [1] by Alon, Azar, Woeginger, and Yadid. The dependence on $\varepsilon$ was doubly exponential. The result described here is from the paper [6] by Jansen, which in fact gives the first EPTAS for the more general problem of related machines.

## References

[1] N. Alon, Y. Azar, G. J. Woeginger, and T. Yadid. Approximation schemes for scheduling on parallel machines. Journal of Scheduling, 1(1):55-66, 1998.
[2] N. Bansal. Scheduling open problems: Old and new. Survey talk in the 13th workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP), 2017. http://www.mapsp2017. ma.tum.de/MAPSP2017-Bansal.pdf (Last accessed : 9th January, 2022).
[3] F. Eisenbrand and G. Shmonin. Carathéodory bounds for integer cones. Operations Research Letters, 34(5):564-568, 2006.
[4] R. L. Graham. Bounds for Certain Multiprocessing Anomalies. Bell System Technical Journal, 45:15631581, 1966.
[5] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. Journal of the ACM, 34(1):144-162, 1987.
[6] K. Jansen. An eptas for scheduling jobs on uniform processors: using an milp relaxation with a constant number of integral variables. SIAM Journal on Discrete Mathematics (SIDMA), 24(2):457-485, 2010.
[7] R. Kannan. Minkowski's convex body theorem and integer programming. Math. Oper. Res., 12(3):415440, 1987.
[8] H. W. Lenstra Jr. Integer programming with a fixed number of variables. Math. Oper. Res., 8(4):538548, 1983.


[^0]:    ${ }^{1}$ Lecture notes by Deeparnab Chakrabarty. Last modified : 18th Jan, 2022
    These have not gone through scrutiny and may contain errors. If you find any, or have any other comments, please email me at deeparnab@dartmouth.edu. Highly appreciated!

[^1]:    ${ }^{2}$ The $\tilde{O}(f(n))$ notation is a short-hand for $O\left(f(n) \cdot \log ^{c} f(n)\right)$ for some unspecified constance $c>0$.

