



Chapter 7


Approximation Algorithms

Algorithm Theory
WS 2013/14

Fabian Kuhn

Approximation Algorithms

- Optimization appears everywhere in computer science
- We have seen many examples, e.g.:
 - scheduling jobs
 - traveling salesperson (*minimum*)
 - maximum flow, maximum matching
 - minimum spanning tree
 - minimum vertex cover
 - ...
- Many discrete optimization problems are NP-hard
- They are however still important and we need to solve them
- As algorithm designers, we prefer algorithms that produce solutions which are provably good, even if we can't compute an optimal solution.



Approximation Algorithms: Examples


We have already seen two approximation algorithms

- **Metric TSP:** If distances are positive and satisfy the triangle inequality, the greedy tour is only by a log-factor longer than an optimal tour

- **Maximum Matching and Vertex Cover:** A maximal matching gives solutions that are within a factor of 2 for both problems.

$$\max_{M \in \mathcal{M}} |M| \leq 2 |VC|$$

$$|M| \geq \frac{1}{2} |M^*|$$



Approximation Ratio

An **approximation algorithm** is an algorithm that computes a solution for an optimization with an objective value that is provably within a bounded factor of the optimal objective value.

Formally:

- OPT ≥ 0 : optimal objective value
- ALG ≥ 0 : objective value achieved by the algorithm
- **Approximation Ratio** α

Minimization: $\alpha := \max_{\text{input instances}} \frac{\text{ALG}}{\text{OPT}} = \max \frac{\text{ALG}}{\text{MIN}} \geq \gamma$

Maximization: $\alpha := \max_{\text{input instances}} \frac{\text{OPT}}{\text{ALG}} = \max \frac{\text{MAX}}{\text{ALG}} \geq \gamma$

Example: Load Balancing

We are given:

- m machines M_1, \dots, M_m
- n jobs, processing time of job i is t_i

$n > m$ usually

Goal:

- Assign each job to a machine such that the makespan is **minimized**

makespan: largest total processing time of any machine

"earliest time at which we are done with everything"

The above load balancing problem is **NP-hard** and we therefore want to get a good approximation for the problem.

Greedy Algorithm

There is a simple **greedy algorithm**:

- Go through the jobs in an arbitrary order
- When considering job i , assign the job to the machine that currently has the smallest load.


Example: ^{$m=3$} 3 machines, 12 jobs ^{$n=12$}



Greedy Assignment:



Optimal Assignment:



$$\underline{16} \leq 2 \cdot 13$$

$$\frac{ALG}{OPT} \leq 2$$

Greedy Analysis

- We will show that greedy gives a 2-approximation
- To show this, we need to compare the solution of greedy with an optimal solution (that we can't compute)
- Lower bound on the optimal makespan T^* :

$$T^* \geq \frac{1}{m} \cdot \sum_{i=1}^n t_i$$

← average load on each machine

- Lower bound can be far from T^* :
 - m machines, m jobs of size 1, 1 job of size m

$$T^* = \underline{m}, \quad \frac{1}{m} \cdot \sum_{i=1}^n t_i = \underline{2}$$

$n = m + 1$

Greedy Analysis

- We will show that greedy gives a 2-approximation
- To show this, we need to compare the solution of greedy with an optimal solution (that we can't compute)
- Lower bound on the optimal makespan T^* :

$$T^* \geq \frac{1}{m} \cdot \sum_{i=1}^n t_i =: T_1$$

- Second lower bound on optimal makespan T^* :

$$T^* \geq \max_{1 \leq i \leq n} t_i =: T_2$$

$$T^* \geq \max \{ T_1, T_2 \}$$

algo:


$$T \leq 2T^*$$

Greedy Analysis

Theorem: The greedy algorithm has approximation ratio ≤ 2 , i.e., for the makespan T of the greedy solution, we have $T \leq 2T^*$.

Proof:


- For machine k , let T_k be the time used by machine k
- Consider some machine M_i for which $T_i = T$ ← "worst machine"
- Assume that job j is the last one scheduled on M_i :



$$T = (T - t_j) + t_j$$

- When job j is scheduled, M_i has the minimum load

↳ $\forall R: T_R \geq T - t_j$



Greedy Analysis

Theorem: The greedy algorithm has approximation ratio ≤ 2 , i.e., for the makespan T of the greedy solution, we have $T \leq 2T^*$.

Proof:

- For all machines M_k : load $T_k \geq T - t_j$

$$m \cdot (T - t_j) = \sum_{k=1}^m (T - t_j) \leq \sum_{k=1}^m T_k = \sum_{i=1}^n t_i = m \cdot T_1 \leq m \cdot T^*$$

$T_1 = \frac{1}{m} \cdot \sum_{i=1}^n t_i$

$$\Rightarrow T - t_j \leq T^*$$

$$t_j \leq T_2 \leq T^*$$



$$T = \underbrace{(T - t_j)}_j + \underbrace{t_j}_j \leq 2T^* \quad \square$$

Can We Do Better?

The analysis of the greedy algorithm is almost tight:


- Example with $n = m(m - 1) + 1$ jobs
- Jobs $1, \dots, n - 1 = m(m - 1)$ have $t_i = 1$, job n has $t_n = m$

Greedy Schedule:



Improving Greedy

Bad case for the greedy algorithm:
 One large job in the end can destroy everything



Idea: assign large jobs first

Modified Greedy Algorithm:

- Sort jobs by decreasing length s.t. $t_1 \geq t_2 \geq \dots \geq t_n$
- Apply the greedy algorithm as before (in the sorted order)

Lemma: If $n > m$: $T^* \geq t_m + t_{m+1} \geq 2t_{m+1}$

Proof:

- Two of the first $m + 1$ jobs need to be scheduled on the same machine
- Jobs m and $m + 1$ are the shortest of these jobs

Analysis of the Modified Greedy Alg.

Theorem: The modified algorithm has approximation ratio $\leq 3/2$, i.e., we have $T \leq 3/2 \cdot T^*$.


Proof:

- As before, choose machine M_i with $T_i = T$ *available machine*
- Job t_j is the last one scheduled on machine M_i
- If there is only one job t_j on M_i , we have $T_i = t_j = T^* \geq T_2 = \max_j t_j$
- Otherwise, we have $j \geq m + 1$
 - The first m jobs are assigned to m distinct machines

$$\begin{aligned}
 j &\geq m+1 & t_j &\leq t_{m+1} \leq \frac{T^*}{2} \\
 T &= \underbrace{(T - t_j)}_{\leq T_1 \leq T^*} + t_j \leq T^* + \frac{T^*}{2} \leq \frac{3}{2} T^*
 \end{aligned}$$



Metric TSP



Input:

- Set V of n nodes (points, cities, locations, sites)
- Distance function $d: V \times V \rightarrow \mathbb{R}$, i.e., $d(u, v)$: dist. from u to v
- Distances define a metric on V :

$$d(u, v) = d(v, u) \geq 0, \quad d(u, v) = 0 \iff u = v$$

$$d(u, v) \leq d(u, w) + d(v, w)$$



Solution:

triangle inequality

- Ordering/permutation v_1, v_2, \dots, v_n of vertices
- Length of TSP path: $\sum_{i=1}^{n-1} d(v_i, v_{i+1})$
- Length of TSP tour: $d(v_n, v_1) + \sum_{i=1}^{n-1} d(v_i, v_{i+1})$

Goal:

- Minimize length of TSP path or TSP tour

Metric TSP

- The problem is **NP-hard**
- We have seen that the **greedy** algorithm (always going to the nearest unvisited node) gives an **$O(\log n)$ -approximation**
- Can we get a constant approximation ratio?
- We will see that we can...



TSP and MST

Claim: The length of an optimal TSP path is lower bounded by the weight of a minimum spanning tree

Proof:


- A TSP path is a spanning tree, it's length is the weight of the tree

\Rightarrow "MST" \leq "TSP path" \leq



"TSP tour"

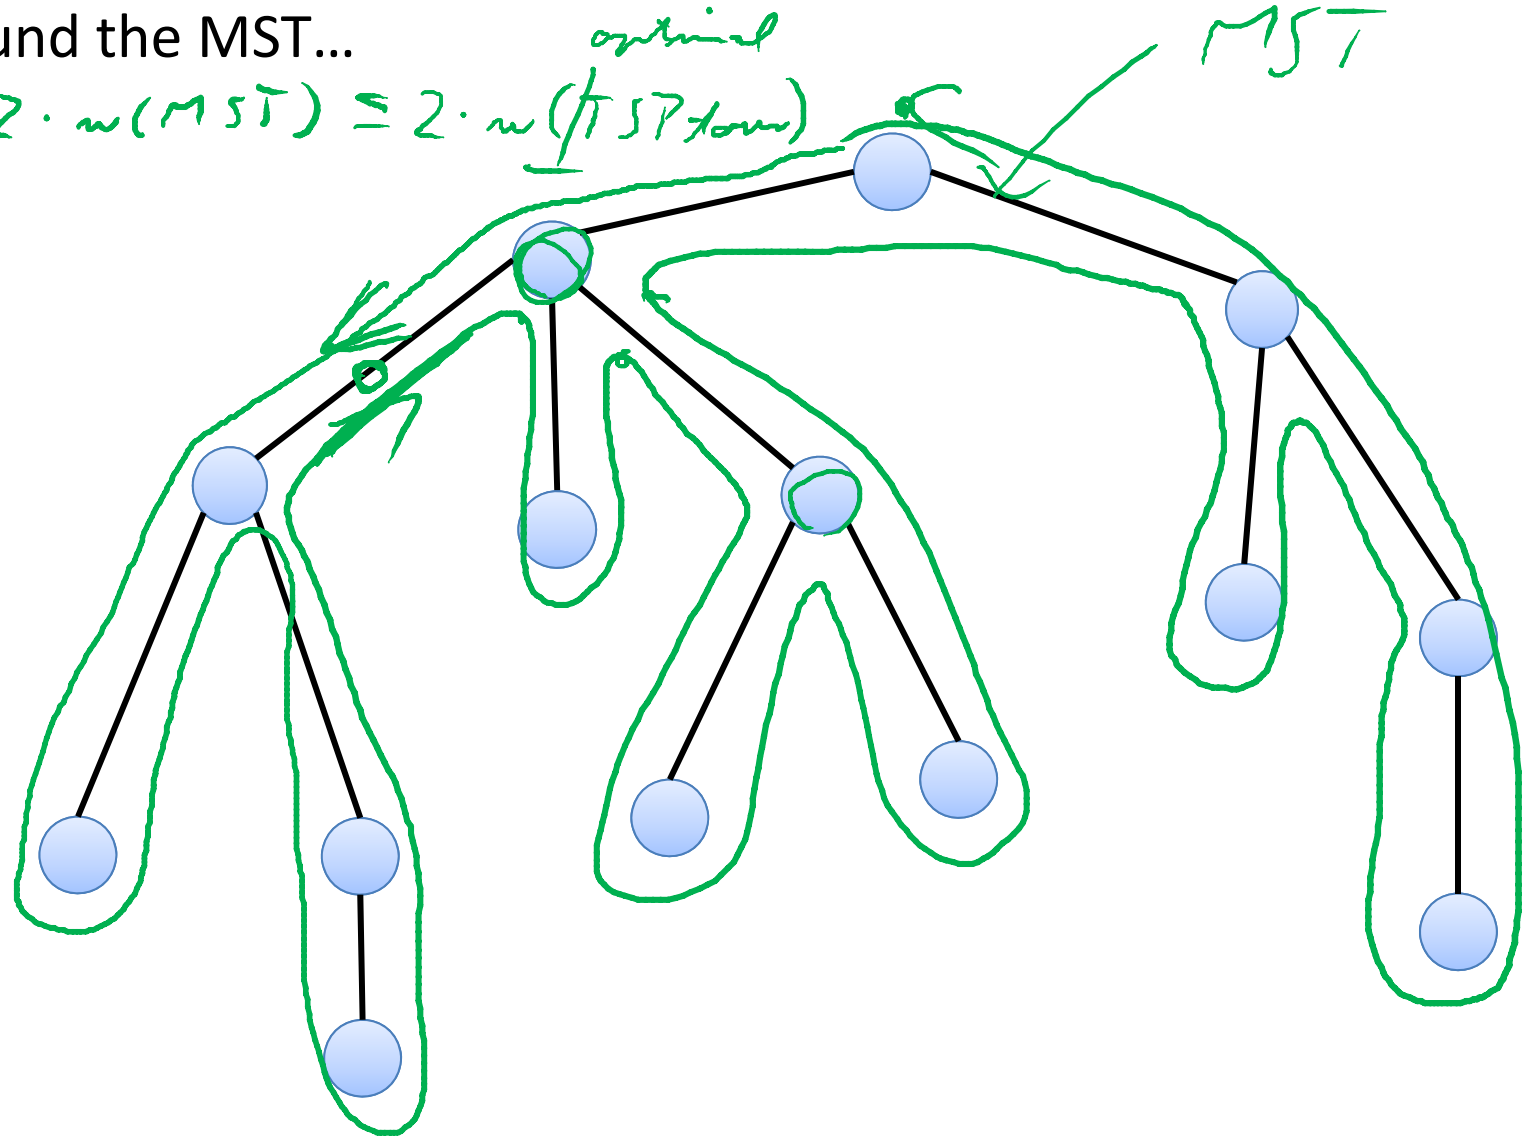
Corollary: Since an optimal TSP tour is longer than an optimal TSP path, the length of an optimal TSP tour is also lower bounded by the weight of a minimum spanning tree.



The MST Tour

Walk around the MST...

walk := $2 \cdot w(\text{MST}) \leq 2 \cdot w(\text{TSP tour})$




The MST Tour

Walk around the MST...

Cost (walk) = 2 · weight(MST)

Cost (tour) < 2 · weight(MST)



Approximation Ratio of MST Tour

Theorem: The MST TSP tour gives a **2-approximation** for the metric TSP problem.

Proof:

- Triangle inequality \rightarrow length of tour is at most $2 \cdot \text{weight}(\text{MST})$
- We have seen that $\text{weight}(\text{MST}) < \text{opt. tour length}$

Can we do even better?

$\textcircled{1}$ $\textcircled{2}$ $\textcircled{3}$ \Rightarrow $\alpha = \frac{3}{2}$

① Metric TSP Subproblems

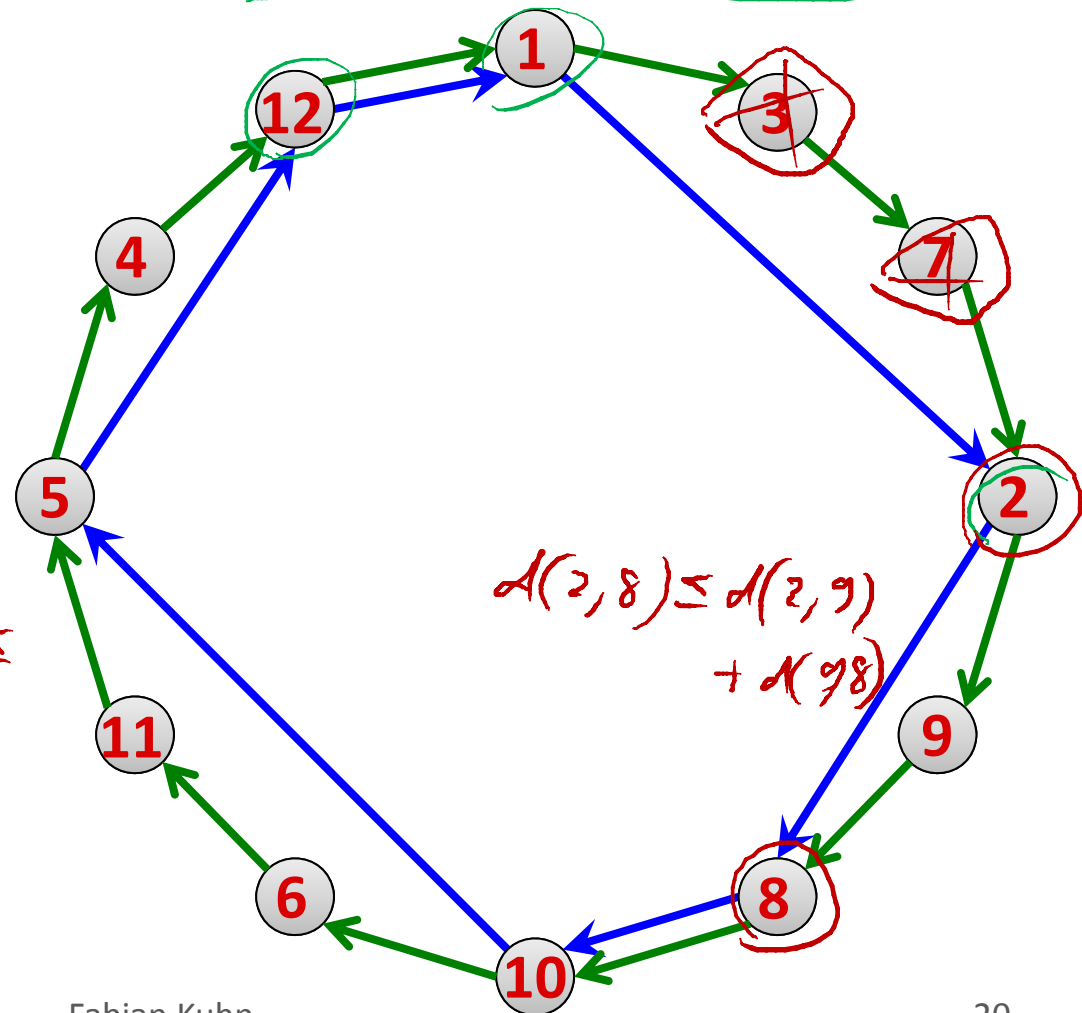
Claim: Given a metric (V, d) and (V', d) for $V' \subseteq V$, the optimal TSP path/tour of (V', d) is at most as large as the optimal TSP path/tour of (V, d) .

Optimal TSP tour of nodes 1, 2, ..., 12

Induced TSP tour for nodes 1, 2, 5, 8, 10, 12

optimal tour on $(\{1, 2, 5, 8, 10, 12\}, d) \leq$
blue tour \leq **green tour**

not necessarily



② TSP and Matching

- Consider a metric TSP instance (V, d) with an even number of nodes $|V|$
- Recall that a perfect matching is a matching $M \subseteq V \times V$ such that every node of V is incident to an edge of M .
- Because $|V|$ is even and because in a metric TSP, there is an edge between any two nodes $u, v \in V$, any partition of V into $|V|/2$ pairs is a perfect matching.
- The weight of a matching M is the sum of the distances represented by all edges in M :

$$w(M) = \sum_{\{u,v\} \in M} d(u, v)$$

TSP and Matching

Lemma: Assume we are given a TSP instance (V, d) with an even number of nodes. The length of an optimal TSP tour of (V, d) is at least twice the weight of a minimum weight perfect matching of (V, d) .

$$\text{opt TSP} \geq 2W \quad \text{where } W = \text{weight of min perfect matching}$$

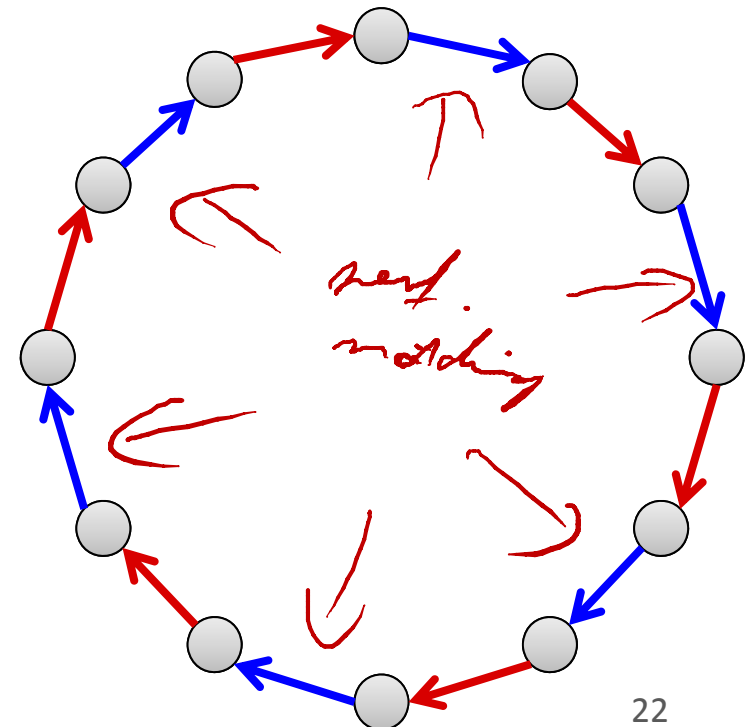
Proof:

- The edges of a TSP tour can be partitioned into 2 perfect matchings

$W \leq \text{blue edges}$

$W \leq \text{red edges}$

$2W \leq \text{red \& blue edges} = \text{opt. TSP tour}$



② Minimum Weight Perfect Matching

Claim: If $|V|$ is even, a minimum weight perfect matching of (V, d) can be computed in polynomial time

Proof Sketch:

- We have seen that a maximum matching in an unweighted graph can be computed in polynomial time
- With a more complicated algorithm, also a maximum weighted matching can be computed in polynomial time *(behrens)*
- In a complete graph, a maximum weighted matching is also a (maximum weight) perfect matching
- Define weight $w(u, v) := D - d(u, v)$ *gives us graph G'* *weights ≥ 0* *(also needed: $|V|$ even)*
 $\leftarrow D = \max_{u,v} d(u, v)$
- A maximum weight perfect matching for (V, w) is a minimum weight perfect matching for (V, d)

$$\boxed{\square} \quad \underline{W_G^{\min}} = \frac{|V|}{2} \cdot D - \underline{W_{G'}^{\max}}$$

Algorithm Outline

Problem of MST algorithm:

- Every edge has to be visited twice *(before taking shortcuts)*

Goal:

- Get a graph G'' on which every edge only has to be visited once (and where still the total edge weight is small compared to an optimal TSP tour)

unlike G' , G'' is not necessarily a tree

Euler Tours:

$\Rightarrow \#edges > n-1$

- A tour that visits each edge of a graph exactly once is called an **Euler tour**
- An Euler tour in a (multi-)graph exists if and only if **every node** of the graph has **even degree**
- That's definitely not true for a tree, but can we modify our MST suitably?



Euler Tour


Theorem: A connected (multi-)graph G has an Euler tour if and only if every node of G has even degree.

Proof:

- If G has an odd degree node, it clearly cannot have an Euler tour
- If G has only even degree nodes, a tour can be found recursively:


"# times entering = # times exiting"

1. Start at some node
2. As long as possible, follow an unvisited edge
 - Gives a partial tour, the remaining graph still has even degree
3. Solve problem on remaining components recursively
4. Merge the obtained tours into one tour that visits all edges



TSP Algorithm


① + ② + ③



1. Compute MST T
2. V_{odd} : nodes that have an odd degree in T ($|V_{\text{odd}}|$ is even)
3. Compute min weight perfect matching M of (V_{odd}, d)
4. $(V, T \cup M)$ is a (multi-)graph with even degrees


$$\sum_{v \in V} d(v) = 2 \cdot |E|$$

red edges: $W \leq \frac{3}{2} \text{TSP}^*$
 black edges: $\text{MST} \leq \text{TSP}^*$
 black & red: $\leq \frac{3}{2} \text{TSP}^*$



TSP Algorithm

5. Compute Euler tour on $(V, T \cup M)$
6. Total length of Euler tour $\leq \frac{3}{2} \cdot \mathbf{TSP_{OPT}}$
7. Get TSP tour by taking shortcuts wherever the Euler tour visits a node twice



TSP Algorithm

- The described algorithm is by Christofides

Theorem: The Christofides algorithm achieves an approximation ratio of at most $3/2$.

Proof:

- The length of the Euler tour is $\leq 3/2 \cdot \text{TSP}_{\text{OPT}}$
- Because of the triangle inequality, taking shortcuts can only make the tour shorter

