## CSE 421: Introduction to Algorithms

## Greedy Algorithms

Paul Beame

## Greedy Algorithms

- Hard to define exactly but can give general properties
- Solution is built in small steps
- Decisions on how to build the solution are made to maximize some criterion without looking to the future
- Want the 'best' current partial solution as if the current step were the last step
- May be more than one greedy algorithm using different criteria to solve a given problem


## Greedy Algorithms

- Greedy algorithms
- Easy to produce
- Fast running times
- Work only on certain classes of problems
- Hard part is showing that they are correct
- Two methods for proving that greedy algorithms do work
- Greedy algorithm stays ahead
- At each step any other algorithm will have a worse value for some criterion that eventually implies optimality
- Exchange Argument
- Can transform any other solution to the greedy solution at no loss in quality


## Interval Scheduling

- Interval Scheduling
- Single resource
- Reservation requests
- Of form "Can I reserve it from start time $s$ to finish time f?"
$-\mathbf{S}<\mathbf{f}$


## Interval Scheduling

- Interval scheduling.
- Job $\mathbf{j}$ starts at $\mathbf{s}_{\mathrm{j}}$ and finishes at $\mathrm{f}_{\mathrm{j}}>\mathbf{s}_{\mathrm{j}}$.
- Two jobs $i$ and $j$ compatible if they don't overlap: $f_{i} \leq s_{j}$ or $f_{j} \leq s_{i}$
- Goal: find maximum size subset of mutually compatible jobs.



## Greedy Algorithms for Interval Scheduling

- What criterion should we try?


## Greedy Algorithms for Interval Scheduling

- What criterion should we try?
- Earliest start time si
- Shortest request time $\mathrm{f}_{\mathrm{i}}-\mathbf{s}_{\mathrm{i}}$
- Earliest finish fime $f_{i}$


## Greedy Algorithms for Interval Scheduling

- What criterion should we try?
- Earliest start time $\mathbf{s}_{\mathbf{i}}$
- Doesn't work
- Shortest request time $\mathrm{f}_{\mathrm{i}}-\mathbf{S}_{\mathrm{i}}$
- Doesn't work

- Even fewest conflicts doesn't work

- Earliest finish fime $\mathrm{f}_{\mathrm{i}}$
- Works


## Greedy Algorithm for Interval Scheduling

$\mathbf{R} \leftarrow$ set of all requests
$\mathrm{A} \leftarrow \varnothing$
While $\mathbf{R} \neq \varnothing$ do
Choose request $\mathbf{i} \in \mathbf{R}$ with smallest finishing time $f_{i}$
Add request it to $\mathbf{A}$
Delete all requests in $\mathbf{R}$ that are not compatible with request i
Return A

## Greedy Algorithm for Interval Scheduling

- Claim: $\mathbf{A}$ is a compatible set of requests and these are added to $\mathbf{A}$ in order of finish time


## Greedy Algorithm for Interval Scheduling

- Claim: A is a compatible set of requests and these are added to $\mathbf{A}$ in order of finish time
- When we add a request to A we delete all incompatible ones from $\mathbf{R}$


## Greedy Algorithm for Interval Scheduling

- Claim: A is a compatible set of requests and these are added to $\mathbf{A}$ in order of finish time
- When we add a request to A we delete all incompatible ones from $\mathbf{R}$
- Claim: For any other set $\mathbf{O} \subseteq$ R of compatible requests then if we order requests in $\mathbf{A}$ and $\mathbf{O}$ by finish time then for each $\mathbf{k}$ :


## Greedy Algorithm for Interval Scheduling

- Claim: A is a compatible set of requests and these are added to $\mathbf{A}$ in order of finish time
- When we add a request to A we delete all incompatible ones from $\mathbf{R}$
- Claim: For any other set $\mathbf{O} \subseteq$ R of compatible requests then if we order requests in $\mathbf{A}$ and $\mathbf{O}$ by finish time then for each $\mathbf{k}$ :
- If $\mathbf{O}$ contains a $\mathbf{k}^{\text {th }}$ request then so does $\mathbf{A}$ and


## Greedy Algorithm for Interval Scheduling

- Claim: A is a compatible set of requests and these are added to $\mathbf{A}$ in order of finish time
- When we add a request to A we delete all incompatible ones from R
- Claim: For any other set $\mathbf{O} \subseteq$ R of compatible requests then if we order requests in $\triangle$ and 0 by finish time then for each $k$. Enough to prove that $A$ is optimal
- If O contains a $\mathbf{k}^{\text {th }}$ request then so does $\mathbf{A}$ and


## Greedy Algorithm for Interval Scheduling

- Claim: A is a compatible set of requests and these are added to $\mathbf{A}$ in order of finish time
- When we add a request to A we delete all incompatible ones from R
- Claim: For any other set $\mathbf{O} \subseteq$ R of compatible requests then if we order requests in $\triangle$ and 0 by finish time then for each $k$. Enough to prove that $A$ is optimal
- If $\mathbf{O}$ contains a $\mathbf{k}^{\text {th }}$ request then so does A and
- the finish time of the $\mathbf{k}^{\text {th }}$ request in $\mathbf{A}$, is $\leq$ the finish time of the $\mathrm{k}^{\text {th }}$ request in $\mathbf{O}$, i.e. " $\mathrm{a}_{\mathrm{k}} \leq \mathbf{o}_{\mathrm{k}}$ " where $\mathbf{a}_{\mathrm{k}}$ and $\mathbf{o}_{\mathrm{k}}$ are the respective finish times
- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Inductive Step: Suppose $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$
- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Inductive Step: Suppose $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$
- By definition of compatibility
- If $\mathbf{O}$ contains a $\mathbf{k + 1} \mathbf{1}^{\text {st }}$ request $\mathbf{r}$ then the start time of that request must be after $\mathbf{o}_{\mathrm{k}}$ and thus after $\mathrm{a}_{\mathrm{k}}$
- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Inductive Step: Suppose $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$
- By definition of compatibility
- If O contains a $\mathbf{k}+1^{\text {st }}$ request $\mathbf{r}$ then the start time of that request must be after $\mathbf{o}_{\mathbf{k}}$ and thus after $\mathrm{a}_{\mathrm{k}}$
- Thus $r$ is compatible with the first $k$ requests in $\mathbf{A}$
- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Inductive Step: Suppose $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$
- By definition of compatibility
- If $\mathbf{O}$ contains a $\mathbf{k + 1} \mathbf{1}^{\text {st }}$ request $r$ then the start time of that request must be after $\mathbf{o}_{\mathbf{k}}$ and thus after $\mathrm{a}_{\mathrm{k}}$
- Thus $r$ is compatible with the first $k$ requests in $\mathbf{A}$
- Therefore
- A has at least k+1 requests since a compatible one is available after the first $k$ are chosen


## Inductive Proof of Claim: $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$

- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Inductive Step: Suppose $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$
- By definition of compatibility
- If $\mathbf{O}$ contains a $k+1^{\text {st }}$ request $r$ then the start time of that request must be after $0_{k}$ and thus after $a_{k}$
- Thus $\mathbf{r}$ is compatible with the first $\mathbf{k}$ requests in $\mathbf{A}$
- Therefore
- A has at least $\mathbf{k + 1}$ requests since a compatible one is available after the first $k$ are chosen
- r was among those considered by the greedy algorithm for that $\mathbf{k + 1} \mathbf{1 s t}^{\text {st }}$ request in $\mathbf{A}$


## Inductive Proof of Claim: $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$

- Base Case: This is true for the first request in A since that is the one with the smallest finish time
- Inductive Step: Suppose $\mathbf{a}_{\mathbf{k}} \leq \mathbf{o}_{\mathbf{k}}$
- By definition of compatibility
- If $\mathbf{O}$ contains a $\mathbf{k + 1} \mathbf{1}^{\text {st }}$ request $\mathbf{r}$ then the start time of that request must be after $0_{k}$ and thus after $a_{k}$
- Thus $\mathbf{r}$ is compatible with the first $\mathbf{k}$ requests in $\mathbf{A}$
- Therefore
- A has at least $\mathbf{k + 1}$ requests since a compatible one is available after the first $k$ are chosen
- r was among those considered by the greedy algorithm for that $\mathbf{k}+\mathbf{1}^{\text {st }}$ request in $\mathbf{A}$
- Therefore by the greedy choice the finish time of $r$ which is $\mathbf{O}_{\mathbf{k}+\mathbf{1}}$ is at least the finish time of that $\mathbf{k + 1}{ }^{\text {st }}$ request in $\mathbf{A}$ which is $a_{k+1}$


## Interval Scheduling: Analysis

Therefore we have:

- Theorem. Greedy algorithm is optimal.
- Alternative Proof. (by contradiction)
- Assume greedy is not optimal, and let's see what happens.
- Let $a_{1}, a_{2}, \ldots a_{t}$ denote set of jobs selected by greedy.
- Let $\mathbf{o}_{1}, \mathbf{o}_{2}, \ldots \mathbf{o}_{\mathrm{m}}$ denote set of jobs in the optimal solution with $a_{1}=0_{1}, a_{2}=0_{2}, \ldots, a_{k}=\mathbf{o}_{k}$ for the largest possible value of $k$.



## Interval Scheduling: Greedy Algorithm Implementation

$O(n \log n)$ Sort jobs by finish times so that $0 \leq f_{1} \leq f_{2} \leq \ldots \leq f_{n}$.

## Scheduling All Intervals: Interval Partitioning

- Interval partitioning.
- Lecture j starts at $\mathrm{s}_{\mathrm{j}}$ and finishes at $\mathrm{f}_{\mathrm{j}}$.
- Goal: find minimum number of classrooms to schedule all lectures so that no two occur at the same time in the same room.
- Example: This schedule uses 4 classrooms to schedule 10 lectures.



## Interval Partitioning

- Interval partitioning.
- Lecture j starts at $\mathrm{s}_{\mathrm{j}}$ and finishes at $\mathrm{f}_{\mathrm{j}}$.
- Goal: find minimum number of classrooms to schedule all lectures so that no two occur at the same time in the same room.
- Example: This schedule uses only 3 classrooms



## Interval Partitioning: Lower Bound on Optimal Solution

- Definition. The depth of a set of open intervals is the maximum number that contain any given time.
- Key observation. Number of classrooms needed $\geq$ depth.
- Ex: Depth of schedule below $=3 \Rightarrow$ schedule below is optimal.

- Q. Does there always exist a schedule equal to depth of intervals?


## A simple greedy algorithm

Sort requests in increasing order of start times $\left(\mathbf{s}_{1}, f_{1}\right), \ldots,\left(\mathbf{s}_{n}, \mathbf{f}_{n}\right)$
For $\mathbf{i}=1$ to $\mathbf{n}$
$\mathrm{j} \leftarrow 1$
While (request i not scheduled)
last $_{\mathrm{j}} \leftarrow$ finish time of the last request currently scheduled on resource j
if $\mathbf{s}_{i} \geq$ last $\boldsymbol{t}_{j}$ then schedule request $i$ on resource ${ }^{j}$
$\mathrm{j} \leftarrow \mathrm{j}+1$
End While
End For

## Interval Partitioning: Greedy Analysis

- Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.


## Interval Partitioning: Greedy Analysis

- Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.
- Theorem. Greedy algorithm is optimal.


## Interval Partitioning: Greedy Analysis

- Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.
- Theorem. Greedy algorithm is optimal.
- Proof.
- Let $\mathbf{d}=$ number of classrooms that the greedy algorithm allocates.


## Interval Partitioning: Greedy Analysis

- Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.
- Theorem. Greedy algorithm is optimal.
- Proof.
- Let $\mathrm{d}=$ number of classrooms that the greedy algorithm allocates.
- Classroom d is opened because we needed to schedule a job, say j , that is incompatible with all d -1 other classrooms.


## Interval Partitioning: Greedy Analysis

- Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.
- Theorem. Greedy algorithm is optimal.
- Proof.
- Let $\mathbf{d}=$ number of classrooms that the greedy algorithm allocates.
- Classroom d is opened because we needed to schedule a job, say j , that is incompatible with all d -1 other classrooms.
- Since we sorted by start time, all these incompatibilities are caused by lectures that start no later than $\mathrm{s}_{\mathrm{j}}$.
- Thus, we have d lectures overlapping at time $s_{j}+\varepsilon$.


## Interval Partitioning: Greedy Analysis

- Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.
- Theorem. Greedy algorithm is optimal.
- Proof.
- Let $\mathbf{d}=$ number of classrooms that the greedy algorithm allocates.
- Classroom d is opened because we needed to schedule a job, say j , that is incompatible with all d-1 other classrooms.
- Since we sorted by start time, all these incompatibilities are caused by lectures that start no later than $\mathrm{s}_{\mathrm{j}}$.
- Thus, we have d lectures overlapping at time $\mathbf{s}_{\mathrm{j}}+\varepsilon$.
- Key observation $\Rightarrow$ all schedules use $\geq$ d classrooms. -


## A simple greedy algorithm

Sort requests in increasing order of start times $\left(\mathbf{s}_{1}, \mathbf{f}_{1}\right), \ldots,\left(\mathbf{s}_{n}, \mathbf{f}_{\mathrm{n}}\right)$
For $\mathrm{i}=1$ to n
$\mathrm{j} \leftarrow 1$
While (request i not scheduled)
last $_{j} \leftarrow$ finish time of the last request currently scheduled on resource j
if $\mathbf{s}_{i} \geq$ last $_{j}$ then schedule request $i$ on resource j
$\mathrm{j} \leftarrow \mathrm{j}+1$
End While
End For

May be slow
O(nd)
which may be $\Omega\left(n^{2}\right)$

## A more efficient implementation

Sort requests in increasing order of start times $\left(\mathbf{s}_{1}, f_{1}\right), \ldots,\left(\mathbf{s}_{n}, \mathbf{f}_{n}\right)$
$O(n \log n)$
$\mathrm{d} \leftarrow 1$
Schedule request 1 on resource 1
last $_{1} \leftarrow \mathbf{f}_{1}$
Insert 1 into priority queue $\mathbf{Q}$ with key = last $_{1}$
For $\mathrm{i}=2$ to n
$\mathrm{j} \leftarrow$ findmin $(\mathbf{Q})$
if $\mathrm{s}_{\mathrm{i}} \geq$ last $_{\mathrm{j}}$ then
schedule request $i$ on resource $j$
last $_{j} \leftarrow \mathrm{f}_{\mathrm{i}}$
Increasekey( $\mathbf{j}, \mathbf{Q}$ ) to last ${ }_{\mathrm{j}}$
else
$\mathrm{d} \leftarrow \mathrm{d}+1$
schedule request $i$ on resource $d$ last $_{\mathrm{d}} \leftarrow \mathrm{f}_{\mathrm{i}}$
Insert d into priority queue $\mathbf{Q}$ with key = last $_{\mathrm{d}}$
End For
$O(n \log n)$ time

## Greedy Analysis Strategies

- Greedy algorithm stays ahead. Show that after each step of the greedy algorithm, its solution is at least as good as any other algorithm's.
- Structural. Discover a simple "structural" bound asserting that every possible solution must have a certain value. Then show that your algorithm always achieves this bound.


## Greedy Analysis Strategies

- Greedy algorithm stays ahead. Show that after each step of the greedy algorithm, its solution is at least as good as any other algorithm's.
- Structural. Discover a simple "structural" bound asserting that every possible solution must have a certain value. Then show that your algorithm always achieves this bound.
- Exchange argument. Gradually transform any solution to the one found by the greedy algorithm without hurting its quality.


## Scheduling to Minimize Lateness

- Scheduling to minimize lateness
- Single resource as in interval scheduling but instead of start and finish times request i has
- Time requirement $t_{i}$ which must be scheduled in a contiguous block
- Target deadline $d_{i}$ by which time the request would like to be finished
- Overall start time s


## Scheduling to Minimize Lateness

- Scheduling to minimize lateness
- Single resource as in interval scheduling but instead of start and finish times request i has
- Time requirement $t_{i}$ which must be scheduled in a contiguous block
- Target deadline $d_{i}$ by which time the request would like to be finished
- Overall start time s
- Requests are scheduled by the algorithm into time intervals $\left[\mathbf{s}_{i}, f_{i}\right.$ ] such that $\mathrm{t}_{\mathrm{i}}=\mathrm{f}_{\mathrm{i}}-\mathrm{s}_{\mathrm{i}}$


## Scheduling to Minimize Lateness

- Scheduling to minimize lateness
- Single resource as in interval scheduling but instead of start and finish times request i has
- Time requirement $t_{i}$ which must be scheduled in a contiguous block
- Target deadline $d_{i}$ by which time the request would like to be finished
- Overall start time s
- Requests are scheduled by the algorithm into time intervals [ $\mathrm{s}_{\mathrm{i}}, \mathrm{f}_{\mathrm{i}}$ ] such that $\mathrm{t}_{\mathrm{i}}=\mathrm{f}_{\mathrm{i}}-\mathrm{s}_{\mathrm{i}}$
- Lateness of schedule for request $i$ is
- If $d_{i}<f_{i}$ then request $i$ is late by $L_{i}=f_{i}-d_{i}$ otherwise its lateness $\mathrm{L}_{\mathrm{i}}=0$


## Scheduling to Minimize Lateness

- Scheduling to minimize lateness
- Single resource as in interval scheduling but instead of start and finish times request i has
- Time requirement $t_{i}$ which must be scheduled in a contiguous block
- Target deadline $d_{i}$ by which time the request would like to be finished
- Overall start time s
- Requests are scheduled by the algorithm into time intervals [ $\mathbf{s}_{i}, \mathrm{f}_{\mathrm{i}}$ ] such that $\mathrm{t}_{\mathrm{i}}=\mathrm{f}_{\mathrm{i}}-\mathrm{s}_{\mathrm{i}}$
- Lateness of schedule for request $i$ is
- If $d_{i}<f_{i}$ then request $i$ is late by $L_{i}=f_{i}-d_{i}$ otherwise its lateness $L_{i}=0$
- Maximum lateness $L=$ max $_{i} L_{i}$


## Scheduling to Minimize Lateness

- Scheduling to minimize lateness
- Single resource as in interval scheduling but instead of start and finish times request i has
- Time requirement $t_{i}$ which must be scheduled in a contiguous block
- Target deadline $d_{i}$ by which time the request would like to be finished
- Overall start time s
- Requests are scheduled by the algorithm into time intervals [ $\mathrm{s}_{\mathrm{i}}, \mathrm{f}_{\mathrm{i}}$ ] such that $\mathrm{t}_{\mathrm{i}}=\mathrm{f}_{\mathrm{i}}-\mathrm{s}_{\mathrm{i}}$
- Lateness of schedule for request $i$ is
- If $d_{i}<f_{i}$ then request $i$ is late by $L_{i}=f_{i}-d_{i}$ otherwise its lateness $L_{i}=0$
- Maximum lateness $L=$ max $_{i} L_{i}$
- Goal: Find a schedule for all requests (values of $s_{i}$ and $f_{i}$ for each request i) to minimize the maximum lateness, $L$


## Scheduling to Minimizing Lateness

- Example:

|  | 1 | 2 | 3 | 4 | 5 | 6 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $t_{j}$ | 3 | 2 | 1 | 4 | 3 | 2 |
| $d_{j}$ | 6 | 8 | 9 | 9 | 14 | 15 |



## Minimizing Lateness: Greedy Algorithms

- Greedy template. Consider jobs in some order.
- [Shortest processing time first] Consider jobs in ascending order of processing time $\mathbf{t}_{\mathrm{j}}$.
- [Earliest deadline first] Consider jobs in ascending order of deadline $\mathrm{d}_{\mathrm{j}}$.
- [Smallest slack] Consider jobs in ascending order of slack $\mathrm{d}_{\mathrm{j}}-\mathrm{t}_{\mathrm{j}}$.


## Minimizing Lateness: Greedy Algorithms

- Greedy template. Consider jobs in some order.
- [Shortest processing time first] Consider jobs in ascending order of processing time $\mathbf{t}_{\mathrm{j}}$.
- [Smallest slack] Consider jobs in ascending order of slack $\mathrm{d}_{\mathrm{j}}-\mathrm{t}_{\mathrm{j}}$.


## Minimizing Lateness: Greedy Algorithms

- Greedy template. Consider jobs in some order.
- [Shortest processing time first] Consider jobs in ascending order of processing time $\mathbf{t}_{\mathrm{j}}$.

- [Smallest slack] Consider jobs in ascending order of slack $\mathbf{d}_{\mathrm{j}}-\mathrm{t}_{\mathrm{j}}$.



## Greedy Algorithm: Earliest Deadline First

- Order requests in increasing order of deadlines
- Schedule the request with the earliest deadline as soon as the resource becomes available


## Minimizing Lateness: Greedy EDF Algorithm

- Greedy Earliest Deadline First (EDF).

```
Sort deadlines in increasing order (d
f}\leftarrow\mathbf{s
for i\leftarrow1 to n to
    si}\leftarrow
    fi}\leftarrow\mp@subsup{\mathbf{S}}{\mathbf{i}}{\mathbf{+}
    f}\leftarrow\mp@subsup{f}{i}{
end for
```


## Scheduling to Minimizing Lateness

- Example:

|  | 1 | 2 | 3 | 4 | 5 | 6 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $t_{j}$ | 3 | 2 | 1 | 4 | 3 | 2 |
| $d_{j}$ | 6 | 8 | 9 | 9 | 14 | 15 |

Original Schedule


## Scheduling to Minimizing Lateness

- Example:


EDF Schedule
$\max$ lateness $=1$


## Proof for Greedy EDF Algorithm: Exchange Argument

- We will show that if there is another schedule O (think optimal schedule) then we can gradually change $\mathbf{O}$ so that
- at each step the maximum lateness in $\mathbf{O}$ never gets worse
- it eventually becomes the same cost as $\mathbf{A}$


## Minimizing Lateness: No Idle Time

- Observation. There exists an optimal schedule with no idle time.

- Observation. The greedy EDF schedule has no idle time.


## Minimizing Lateness: Inversions

- Definition. An inversion in schedule $\mathbf{S}$ is a pair of jobs $i$ and $j$ such that $d_{i}<d_{j}$ but $j$ scheduled before $i$.



## Minimizing Lateness: Inversions

- Definition. An inversion in schedule $\mathbf{S}$ is a pair of jobs $i$ and $j$ such that $d_{i}<d_{j}$ but $j$ scheduled before $i$.

- Observation. Greedy EDF schedule has no inversions.


## Minimizing Lateness: Inversions

- Definition. An inversion in schedule $\mathbf{S}$ is a pair of jobs $i$ and $j$ such that $d_{i}<d_{j}$ but $j$ scheduled before $i$.

- Observation. Greedy EDF schedule has no inversions.
- Observation. If a schedule (with no idle time) has an inversion, it has one with a pair of inverted jobs scheduled consecutively (by transitivity of <).


## Minimizing Lateness: Inversions

- Definition. An inversion in schedule $\mathbf{S}$ is a pair of jobs $i$ and $j$ such that $d_{i}<d_{j}$ but $j$ scheduled before $i$.

- Claim. Swapping two adjacent, inverted jobs reduces the number of inversions by one and does not increase the max lateness.


## Minimizing Lateness: Inversions

- If $d_{j}>d_{i}$ but $j$ is scheduled in $\mathbf{O}$ immediately before $i$ then swapping requests $i$ and $j$ to get schedule $\mathbf{O}^{\prime}$ does not increase the maximum lateness

- Lateness $L_{i}{ }^{\prime} \leq L_{i}$ since $f_{i}>f_{i}$
- Requests i and j together occupy same total time slot
- All other requests $k \neq i$, $j$ have $L_{k}{ }^{\prime}=L_{k}$
- $f_{j}{ }^{\prime}=f_{i}$ so $L_{j}{ }^{\prime}=f_{j}{ }_{j}-d_{j}=f_{i}-d_{j}<f_{i}-d_{i}=L_{i}$
- Maximum lateness has not increased!


## Optimal schedules and inversions

- Claim: There is an optimal schedule with no idle time and no inversions
- Proof:
- By previous argument there is an optimal schedule $\mathbf{O}$ with no idle time
- If O has an inversion then it has a consecutive pair of requests in its schedule that are inverted and can be swapped without increasing lateness


## Optimal schedules and inversions

- Eventually these swaps will produce an optimal schedule with no inversions
- Each swap decreases the number of inversions by 1
- There are a bounded number of (at most $\mathrm{n}(\mathrm{n}-1) / \mathbf{2}$ ) inversions (we only care that this is finite.)

QED

## Idleness and Inversions are the only issue

- Claim: All schedules with no inversions and no idle time have the same maximum lateness
- Proof
- Schedules can differ only in how they order requests with equal deadlines
- Consider all requests having some common deadline d
- Maximum lateness of these jobs is based only on the finish time of the last of these jobs but the set of these requests occupies the same time segment in both schedules
- Last of these requests finishes at the same time in any such schedule.


## Earliest Deadline First is optimal

- We know that
- There is an optimal schedule with no idle time or inversions
- All schedules with no idle time or inversions have the same maximum lateness
- EDF produces a schedule with no idle time or inversions
- Therefore
- EDF produces an optimal schedule


## Greedy Analysis Strategies

- Greedy algorithm stays ahead. Show that after each step of the greedy algorithm, its solution is at least as good as any other algorithm's.
- Structural. Discover a simple "structural" bound asserting that every possible solution must have a certain value. Then show that your algorithm always achieves this bound.
- Exchange argument. Gradually transform any solution to the one found by the greedy algorithm without hurting its quality.


## Single-source shortest paths

- Given an (un)directed graph G=(V,E) with each edge e having a non-negative weight $w(e)$ and a vertex $v$
- Find length of shortest paths from $\mathrm{s}=\mathrm{v}$ to each vertex in $G$


## A greedy algorithm

- Dijkstra's Algorithm:
- Maintain a set S of vertices whose shortest paths are known
- initially $\mathrm{S}=\{\mathbf{s}\}$
- Maintaining current best lengths of paths that only go through $S$ to each of the vertices in $G$
- path-lengths to elements of $S$ will be right, to V-S they might not be right
- Repeatedly add vertex $\mathbf{v}$ to S that has the shortest path-length of any vertex in V-S
- update path lengths based on new paths through v


## Dijsktra's Algorithm

Dijkstra(G,w,s)
$\mathbf{S} \leftarrow\{\mathbf{s}\}$
$\mathrm{d}[\mathrm{s}] \leftarrow 0$
while $\mathbf{S}=\mathbf{V}$ do
of all edges $\mathbf{e}=(\mathbf{u}, \mathbf{v})$ s.t. $\mathbf{v} \notin \mathbf{S}$ and $\mathbf{u} \in \mathbf{S}$ select* one with the minimum value of $\mathbf{d}[\mathbf{u}]+\mathbf{w}(\mathbf{e})$

$$
\begin{aligned}
& \mathbf{S} \leftarrow \mathbf{S} \cup\{\mathbf{v}\} \\
& \mathbf{d}[\mathbf{v}] \leftarrow \mathbf{d}[\mathbf{u}]+\mathbf{w}(\mathbf{e}) \\
& \operatorname{pred}[\mathbf{v}] \leftarrow \mathbf{u}
\end{aligned}
$$

*For each $v \notin S$ maintain $d^{\prime}[v]=$ minimum value of $d[u]+w(e)$ over all vertices $u \in S$ s.t. $e=(u, v)$ is in of $G$

## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm

Add to $S$


## Dijkstra's Algorithm

Update distances


## Dijkstra's Algorithm



## Dijkstra's Algorithm

Update distances


## Dijkstra's Algorithm



## Dijkstra's Algorithm

Update distances


## Dijkstra's Algorithm



## Dijkstra's Algorithm

Update distances


## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm



## Dijkstra's Algorithm

Update distances


## Dijkstra's Algorithm



## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in V-S
$\therefore$ distance value of vertex in $V$-S=length of shortest path from s with only last edge leaving $S$


## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in V-S
$\therefore$ distance value of vertex in V -S=length of shortest path from s with only last edge leaving $S$


Suppose some other path to $v$ and $x=$ first vertex on this path not in $S$

## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in V-S
$\therefore$ distance value of vertex in V -S=length of shortest path from s with only last edge leaving $S$


Suppose some other path to $v$ and $x=$ first vertex on this path not in $S$

## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in $V$-S
$\therefore$ distance value of vertex in V -S $=$ length of shortest path from s with only last edge leaving $S$


Suppose some other path to $v$ and $x=$ first vertex on this path not in $S$

$$
d^{\prime}(v) \leq d^{\prime}(x)
$$

## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in V -S
$\therefore$ distance value of vertex in V -S $=$ length of shortest path from s
 with only last edge leaving $S$

Suppose some other path to $v$ and $x=$ first vertex on this path not in $S$

$$
\begin{aligned}
& d^{\prime}(v) \leq d^{\prime}(x) \\
& x-v \text { path length } \geq 0
\end{aligned}
$$

## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in V-S
$\therefore$ distance value of vertex in V -S $=$ length of shortest path from s


Suppose some other path to $v$ and $x=$ first vertex on this path not in $S$

$$
\begin{aligned}
& d^{\prime}(v) \leq d^{\prime}(x) \\
& x-v \text { path length } \geq 0
\end{aligned}
$$

$\therefore$ other path is longer

## Dijkstra's Algorithm Correctness

Suppose all distances to vertices in $S$ are correct and $u$ has smallest current value in $V$-S
$\therefore$ distance value of vertex in V -S $=$ length of shortest path from s


Suppose some other path to $v$ and $x=$ first vertex on this path not in $S$

$$
\begin{aligned}
& d^{\prime}(v) \leq d^{\prime}(x) \\
& x-v \text { path length } \geq 0
\end{aligned}
$$

$\therefore$ other path is longer
Therefore adding v to S keeps correct distances

## Dijkstra's Algorithm

- Algorithm also produces a tree of shortest paths to $v$ following pred links
- From w follow its ancestors in the tree back to v
- If all you care about is the shortest path from $v$ to $w$ simply stop the algorithm when $w$ is added to $S$


## Dijsktra's Algorithm

Dijkstra(G,w,s)
$\mathbf{S} \leftarrow\{\mathbf{s}\}$
$\mathrm{d}[\mathrm{s}] \leftarrow 0$
while $\mathbf{S}=\mathbf{V}$ do
of all edges $\mathbf{e}=(\mathbf{u}, \mathbf{v})$ s.t. $\mathbf{v} \notin \mathbf{S}$ and $\mathbf{u} \in \mathbf{S}$ select* one with the minimum value of $\mathbf{d}[\mathbf{u}]+\mathbf{w}(\mathbf{e})$

$$
\begin{aligned}
& \mathbf{S} \leftarrow \mathbf{S} \cup\{\mathbf{v}\} \\
& \mathbf{d}[\mathbf{v}] \leftarrow \mathbf{d}[\mathbf{u}]+\mathbf{w}(\mathbf{e}) \\
& \operatorname{pred}[\mathbf{v}] \leftarrow \mathbf{u}
\end{aligned}
$$

*For each $v \notin S$ maintain $d^{\prime}[v]=$ minimum value of $d[u]+w(e)$ over all vertices $u \in S$ s.t. $e=(u, v)$ is in of $G$

## Implementing Dijkstra's Algorithm

- Need to
- keep current distance values for nodes in V-S
- find minimum current distance value
- reduce distances when vertex moved to $S$


## Data Structure Review

- Priority Queue:
- Elements each with an associated key
- Operations
- Insert
- Find-min
- Return the element with the smallest key
- Delete-min
- Return the element with the smallest key and delete it from the data structure
- Decrease-key
- Decrease the key value of some element
- Implementations
- Arrays: $\mathbf{O}(\mathrm{n})$ time find/delete-min, $\mathbf{O}(1)$ time insert/ decrease-key
- Heaps: O(log n) time insert/decrease-key/delete-min, O(1) time find-min


## Dijkstra's Algorithm with Priority Queues

- For each vertex u not in tree maintain cost of current cheapest path through tree to u
- Store u in priority queue with key = length of this path
- Operations:
- $\mathrm{n}-1$ insertions (each vertex added once)
- n-1 delete-mins (each vertex deleted once)
- pick the vertex of smallest key, remove it from the priority queue and add its edge to the graph
- <m decrease-keys (each edge updates one vertex)


## Dijskstra's Algorithm with Priority Queues

- Priority queue implementations
- Array
- insert $O(1)$, delete-min $O(n)$, decrease-key $O(1)$
- total $O\left(n+n^{2}+m\right)=O\left(n^{2}\right)$
- Heap
- insert, delete-min, decrease-key all O(log n)
- total $O(m \log n)$
- d-Heap ( $d=m / n$ )
- insert, decrease-key $O\left(\log _{m / n} n\right)$
- delete-min $O\left((m / n) \log _{m / n} n\right)$
- total $O\left(m \log _{m / n} n\right)$


## Minimum Spanning Trees (Forests)

- Given an undirected graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$ with each edge e having a weight $w(e)$
- Find a subgraph T of G of minimum total weight s.t. every pair of vertices connected in G are also connected in T
- if $G$ is connected then $T$ is a tree otherwise it is a forest


## Weighted Undirected Graph



## Greedy Algorithm

- Prim's Algorithm:
- start at a vertex s
- add the cheapest edge adjacent to s
- repeatedly add the cheapest edge that joins the vertices explored so far to the rest of the graph
- Exactly like Dijsktra's Algorithm but with a different metric


## Dijsktra's Algorithm

Dijkstra(G,w,s)
$\mathbf{S} \leftarrow\{\mathbf{s}\}$
$\mathrm{d}[\mathrm{s}] \leftarrow 0$
while $\mathbf{S}=\mathbf{V}$ do
of all edges $\mathbf{e}=(\mathbf{u}, \mathbf{v})$ s.t. $\mathbf{v} \notin \mathbf{S}$ and $\mathbf{u} \in \mathbf{S}$ select* one with the minimum value of $\mathbf{d}[\mathbf{u}]+\mathbf{w}(\mathbf{e})$

$$
\begin{aligned}
& \mathbf{S} \leftarrow \mathbf{S} \cup\{\mathbf{v}\} \\
& \mathbf{d}[\mathbf{v}] \leftarrow \mathbf{d}[\mathbf{u}]+\mathbf{w}(\mathbf{e}) \\
& \operatorname{pred}[\mathbf{v}] \leftarrow \mathbf{u}
\end{aligned}
$$

*For each $v \notin S$ maintain $d^{\prime}[v]=$ minimum value of $d[u]+w(e)$ over all vertices $u \in S$ s.t. $e=(u, v)$ is in of $G$

## Prim's Algorithm

$\operatorname{Prim}(\mathbf{G}, \mathbf{w}, \mathbf{s})$
$\mathbf{S} \leftarrow\{\mathbf{s}\}$
while $\mathbf{S} \neq \mathbf{V}$ do
of all edges $\mathbf{e}=(\mathbf{u}, \mathbf{v})$ s.t. $\mathbf{v} \notin \mathbf{S}$ and $\mathbf{u} \in \mathbf{S}$ select* one with the minimum value of $\mathbf{w}(\mathbf{e})$

$$
\mathbf{S} \leftarrow \mathbf{S} \cup\{\mathbf{v}\}
$$

## $\operatorname{pred}[\mathbf{v}] \leftarrow \mathbf{u}$

*For each $v \notin S$ maintain small[v]=minimum value of $w(e)$ over all vertices $u \in S$ s.t. $e=(u, v)$ is in of $G$

## Second Greedy Algorithm

- Kruskal's Algorithm
- Start with the vertices and no edges
- Repeatedly add the cheapest edge that joins two different components. i.e. that doesn't create a cycle


## Why greed is good

- Definition: Given a graph $G=(V, E)$, a cut of G is a partition of V into two non-empty pieces, $S$ and V-S
- Lemma: For every cut (S,V-S) of G, there is a minimum spanning tree (or forest) containing any cheapest edge crossing the cut, i.e. connecting some node in $S$ with some node in V-S.
- call such an edge safe


## Cuts and Spanning Trees



## The greedy algorithms always choose safe edges

- Prim's Algorithm
- Always chooses cheapest edge from current tree to rest of the graph
- This is cheapest edge across a cut which has the vertices of that tree on one side.


## Prim's Algorithm



## The greedy algorithms always choose safe edges

- Kruskal's Algorithm
- Always chooses cheapest edge connecting two pieces of the graph that aren't yet connected
- This is the cheapest edge across any cut which has those two pieces on different sides and doesn't split any current pieces.


## Kruskal's Algorithm



## Kruskal's Algorithm



## Why greed is good

- Definition: Given a graph $G=(V, E)$, a cut of G is a partition of V into two non-empty pieces, $S$ and V-S
- Lemma: For every cut (S,V-S) of G, there is a minimum spanning tree (or forest) containing any cheapest edge crossing the cut, i.e. connecting some node in $S$ with some node in V-S.
- call such an edge safe


## Proof of Lemma: An Exchange Argument

Suppose you have an MST T not using cheapest edge e


Endpoints of e, u and v must be connected in T

## Proof of Lemma

Suppose you have an MST T not using cheapest edge e


Endpoints of $e, u$ and $v$ must be connected in $T$

## Proof of Lemma

Suppose you have an MST T not using cheapest edge e


Endpoints of $e, u$ and $v$ must be connected in $T$

## Proof of Lemma

Suppose you have an MST T not using cheapest edge e

$$
w(e) \leq w(h)
$$



Endpoints of e, u and v must be connected in T

## Proof of Lemma

Replacing h by e does not increase weight of T

$$
w(e) \leq w(h)
$$



All the same points are connected by the new tree

## The general exchange argument

So far we have shown that any minimum spanning tree can be modified to include any safe edge (minimum edge crossing some cut)...

This just says that the first edge is OK.
We need something stronger (which has the same proof):

- Lemma: For every subset F of edges of some minimum spanning tree/forest and any cut (S,V-S) of G that does not cut any edge of $F$, there is a minimum spanning tree/forest containing both F and any safe edge crossing this cut.


## Kruskal's Algorithm Implementation \& Analysis

- First sort the edges by weight $\mathrm{O}(\mathrm{m} \log \mathrm{m})$
- Go through edges from smallest to largest
- if endpoints of edge e are currently in different components
- then add to the graph
- else skip
- Union-find data structure handles last part
- Total cost of last part: $\mathrm{O}(\mathrm{m} \alpha(\mathrm{n}))$ where $\alpha(\mathrm{n}) \ll \log \mathrm{m}$
- Overall O(m log $n$ )


## Union-find disjoint sets data structure

- Maintaining components
- start with $n$ different components
- one per vertex
- find components of the two endpoints of e
- 2 m finds
- union two components when edge connecting them is added
- n -1 unions


## Prim's Algorithm with Priority Queues

- For each vertex u not in tree maintain current cheapest edge from tree to $u$
- Store u in priority queue with key = weight of this edge
- Operations:
- n -1 insertions (each vertex added once)
- n -1 delete-mins (each vertex deleted once)
- pick the vertex of smallest key, remove it from the p.q. and add its edge to the graph
- <m decrease-keys (each edge updates one vertex)


## Prim's Algorithm with Priority Queues

- Priority queue implementations
- Array
- insert $O(1)$, delete-min $O(n)$, decrease-key $O(1)$
- total $O\left(n+n^{2}+m\right)=O\left(n^{2}\right)$
- Heap
- insert, delete-min, decrease-key all O(log n)
- total $O(m \log n)$
- d-Heap ( $d=m / n$ )
- insert, decrease-key $O\left(\log _{m / n} n\right)$
- delete-min $O\left((m / n) \log _{m / n} n\right)$
- total $O\left(m \log _{m / n} n\right)$


## Boruvka's Algorithm (1927)

- A bit like Kruskal's Algorithm
- Start with n components consisting of a single vertex each
- At each step, each component chooses its cheapest outgoing edge to add to the spanning forest
- Two components may choose to add the same edge
- Useful for parallel algorithms since components may be processed (almost) independently


## Many other minimum spanning tree algorithms, most of them greedy

- Cheriton \& Tarjan
- O(m loglog n) time using a queue of components
- Chazelle
- O(m $\alpha(\mathrm{m}) \log \alpha(\mathrm{m}))$ time
- Incredibly hairy algorithm
- Karger, Klein \& Tarjan
- $\mathrm{O}(\mathrm{m}+\mathrm{n})$ time randomized algorithm that works most of the time


## Applications of Minimum Spanning Tree Algorithms

- Minimum cost network design:
- Build a network to connect all locations $\left\{\mathbf{v}_{1}, \ldots, \mathbf{v}_{\mathrm{n}}\right\}$
- Cost of connecting $\mathbf{v}_{\mathrm{i}}$ to $\mathbf{v}_{\mathrm{j}}$ is $\mathbf{w}\left(\mathbf{v}_{\mathrm{i}}, \mathbf{v}_{\mathrm{j}}\right)>0$
- Choose a collection of links to create that will be as cheap as possible
- Any minimum cost solution is an MST
- If there is a solution containing a cycle then we can remove any edge and get a cheaper solution


## Applications of Minimum Spanning Tree Algorithms

- Maximum Spacing Clustering
- Given
- a collection U of $n$ objects $\left\{p_{1}, \ldots, p_{n}\right\}$
- Distance measure $d\left(p_{i}, p_{j}\right)$ satisfying
- $\mathrm{d}\left(\mathrm{p}_{\mathrm{i}}, \mathrm{p}_{\mathrm{i}}\right)=0$
- $\mathrm{d}\left(\mathrm{p}_{\mathrm{i}}, \mathrm{p}_{\mathrm{j}}\right)>0$ for $\mathrm{i} \neq \mathrm{j}$
- $\mathrm{d}\left(\mathrm{p}_{\mathrm{i}}, \mathrm{p}_{\mathrm{j}}\right)=\mathrm{d}\left(\mathrm{p}_{\mathrm{j}}, \mathrm{p}_{\mathrm{i}}\right)$
- Positive integer $k \leq n$
- Find a k-clustering, i.e. partition of $U$ into $k$ clusters $\mathrm{C}_{1}, \ldots, \mathrm{C}_{\mathrm{k}}$, such that the spacing between the clusters is as large possible where
spacing $=\min \left\{d\left(p_{i}, p_{j}\right): p_{i}\right.$ and $p_{j}$ in different clusters


## Greedy Algorithm

- Start with $n$ clusters each consisting of a single point
- Repeatedly find the closest pair of points in different clusters under distance d and merge their clusters until only k clusters remain


## Greedy Algorithm

- Start with $n$ clusters each consisting of a single point
- Repeatedly find the closest pair of points in different clusters under distance d and merge their clusters until only k clusters remain
- Gets the same components as Kruskal's Algorithm does!
- The sequence of closest pairs is exactly the MST
- Alternatively we could run Kruskal's algorithm once and for any $k$ we could get the maximum spacing k -clustering by deleting the $\mathrm{k}-1$ most expensive edges in the MST


## Proof that this works

- Removing the k-1 most expensive edges from an MST yields $k$ components $C_{1}, \ldots, C_{k}$ and the spacing for them is precisely the cost $\mathbf{d}^{*}$ of the $\mathbf{k}-\mathbf{1}^{\text {st }}$ most expensive edge in the tree
- Consider any other k-clustering $\mathbf{C}_{1}^{\prime}, \ldots, \mathbf{C}_{\mathrm{k}}$
- Since they are different and cover the same set of points there is some pair of points $p_{i}, p_{j}$ such that $p_{i}, p_{j}$ are in some cluster $\mathrm{C}_{\mathrm{r}}$ but $\mathrm{p}_{\mathrm{i}}, \mathrm{p}_{\mathrm{j}}$ are in different clusters $\mathrm{C}_{\mathrm{s}}{ }^{\prime}$ and $\mathrm{C}_{\mathrm{t}}^{\prime}$
- Since $p_{i}, p_{j} \in C_{r}, p_{i}$ and $p_{j}$ have a path between them all of whose edges have distance at most $\mathrm{d}^{*}$
- This path must cross between clusters in the C' clustering so the spacing in $\mathrm{C}^{\prime}$ is at most $\mathrm{d}^{*}$


## Optimal Caching/Paging

- Memory systems
- many levels of storage with different access times
- smaller storage has shorter access time
- to access an item it must be brought to the lowest level of the memory system
- Consider the management problem between adjacent levels
- Main memory with $\mathbf{n}$ data items from a set $\mathbf{U}$
- Cache can hold $\mathbf{k}<\mathbf{n}$ items
- Simplest version with no direct-mapping or other restrictions about where items can be
- Suppose cache is full initially
- Holds k data items to start with


## Optimal Offline Caching

- Caching.
- Cache with capacity to store $\mathbf{k}$ items.
- Sequence of $m$ item requests $d_{1}, d_{2}, \ldots, d_{m}$.
- Cache hit: item already in cache when requested.
- Cache miss: item not already in cache when requested: must bring requested item into cache, and evict some existing item, if full.
- Goal. Eviction schedule that minimizes number of cache misses (actually, \# of evictions).
- Example: $\mathbf{k}=\mathbf{2}$, initial cache $=\mathbf{a b}$, requests: $\mathbf{a , ~ b , ~ c , ~ b , ~ c , ~ a , ~ a , ~ b . ~}$
- Optimal eviction schedule: 2 cache misses.

| a | a | b |
| :---: | :---: | :---: |
| b | a | b |
| c | c | b |
| b | c | b |
| C | c | b |
| a | a | b |
| a | a | b |
| b | a | b |
| requests | cac | he |

## Optimal Offline Caching: Farthest-InFuture

- Farthest-in-future. Evict item in the cache that is not requested until farthest in the future.

```
current cache: a b b c c d l e f
```



- Theorem. [Bellady, 1960s] FIF is an optimal eviction schedule.
- Proof. Algorithm and theorem are intuitive; proof is subtle.


## Other Algorithms

- Often there is flexibility, e.g.
- $k=3, C=\{a, b, c\}$
- $D=a b c d a d e a d b c$
$\begin{array}{llll}\text { - } S_{\text {FIF }}= & c & b & \text { ed } \\ \text { - } S & = & b & c\end{array} \quad \mathrm{de}$
- Why aren't other algorithms better?
- Least-Frequenty-Used-In-Future?
- Exchange Argument
- We can swap choices to convert other schedules to Farthest-In-Future without losing quality


## Reduced Eviction Schedules

- Definition. A reduced schedule is a schedule that only inserts an item into the cache in a step in which that item is requested.
- Intuition. Can transform an unreduced schedule into a reduced one with no more cache misses.

an unreduced schedule

| $a$ | $a$ | $b$ | $c$ |
| :--- | :--- | :--- | :--- |
| $a$ | $a$ | $b$ | $c$ |
| $c$ | $a$ | $b$ | $c$ |
| $d$ | $a$ | $d$ | $c$ |
| $a$ | $a$ | $d$ | $c$ |
| $b$ | $a$ | $d$ | $b$ |
| $c$ | $a$ | $c$ | $b$ |
| $a$ | $a$ | $c$ | $b$ |
| $a$ | $a$ | $c$ | $b$ |

a reduced schedule

## Reduced Eviction Schedules

- Claim. Given any unreduced schedule S, can transform it into a reduced schedule $S^{\prime}$ with no more cache misses.
- Proof. (by induction on number of unreduced items)
- Suppose $\mathbf{S}$ brings $\mathbf{d}$ into the cache at time $\mathbf{t}$, without a request.
- Let $\mathbf{c}$ be the item $\mathbf{S}$ evicts when it brings $\mathbf{d}$ into the cache.


## Reduced Eviction Schedules

- Claim. Given any unreduced schedule S, can transform it into a reduced schedule $S^{\prime}$ with no more cache misses.
- Proof. (by induction on number of unreduced items)
- Suppose $\mathbf{S}$ brings $d$ into the cache at time $t$, without a request.
- Let $\mathbf{c}$ be the item $\mathbf{S}$ evicts when it brings $\mathbf{d}$ into the cache.
- Case 1: d evicted at time $\mathbf{t}^{\prime}$, before next request for $\mathbf{d}$.


Case 1

## Reduced Eviction Schedules

- Claim. Given any unreduced schedule S, can transform it into a reduced schedule $S^{\prime}$ with no more cache misses.
- Proof. (by induction on number of unreduced items)
- Suppose $S$ brings $d$ into the cache at time $t$, without a request.
- Let $\mathbf{c}$ be the item $\mathbf{S}$ evicts when it brings $\mathbf{d}$ into the cache.
- Case 1: d evicted at time $\mathbf{t}^{\prime}$, before next request for $\mathbf{d}$.


Case 1

## Reduced Eviction Schedules

- Claim. Given any unreduced schedule S, can transform it into a reduced schedule $S^{\prime}$ with no more cache misses.
- Proof. (by induction on number of unreduced items)
- Suppose $\mathbf{S}$ brings $\mathbf{d}$ into the cache at time $\mathbf{t}$, without a request.
- Let $\mathbf{c}$ be the item $\mathbf{S}$ evicts when it brings $\mathbf{d}$ into the cache.
- Case 1: d evicted at time t', before next request for d.
- Case 2: $\mathbf{d}$ requested at time $\mathrm{t}^{\prime}$ before $\mathbf{d}$ is evicted.


Case 1


Case 2

## Reduced Eviction Schedules

- Claim. Given any unreduced schedule S, can transform it into a reduced schedule $S^{\prime}$ with no more cache misses.
- Proof. (by induction on number of unreduced items)
- Suppose $\mathbf{S}$ brings $\mathbf{d}$ into the cache at time $\mathbf{t}$, without a request.
- Let $\mathbf{c}$ be the item $\mathbf{S}$ evicts when it brings $\mathbf{d}$ into the cache.
- Case 1: d evicted at time t', before next request for d.
- Case 2: $\mathbf{d}$ requested at time $\mathrm{t}^{\prime}$ before $\mathbf{d}$ is evicted.


Case 1


Case 2

## Farthest-In-Future: Analysis

- Theorem. FIF is optimal eviction algorithm.
- Proof. (by induction on number or requests j)

Invariant: There exists an optimal reduced schedule S that makes the same eviction schedule as $\mathbf{S}_{\text {FIF }}$ through the first $\mathbf{j}+\mathbf{1}$ requests.

- Let $\mathbf{S}$ be reduced schedule that satisfies invariant through j requests. We produce $\mathrm{S}^{\prime}$ that satisfies invariant after $\mathrm{j}+1$ requests.
- Consider ( $\mathbf{j}+\mathbf{1}$ ) st request $\mathbf{d}=\mathrm{d}_{\mathrm{j}+1}$.
- Since $S$ and $\mathrm{S}_{\text {FIF }}$ have agreed up until now, they have the same cache contents before request $\mathrm{j}+1$.
- Case 1: (d is already in the cache). $\mathbf{S}^{\prime}=\mathbf{S}$ satisfies invariant.
- Case 2: (d is not in the cache and $\mathbf{S}$ and $\mathbf{S}_{\text {FIF }}$ evict the same element). $\mathbf{S}^{\prime}=\mathbf{S}$ satisfies invariant.


## Farthest-In-Future: Analysis

- Proof. (continued)
- Case 3: ( $\mathbf{d}$ is not in the cache; $\mathbf{S}_{\text {FIF }}$ evicts $\mathbf{e}$; $\mathbf{S}$ evicts $\mathbf{f} \neq \mathbf{e}$ ).


## Farthest-In-Future: Analysis

- Proof. (continued)
- Case 3: ( $\mathbf{d}$ is not in the cache; $\mathbf{S}_{\text {FIF }}$ evicts $\mathbf{e}$; $\mathbf{S}$ evicts $\mathbf{f} \neq \mathbf{e}$ ).
- begin construction of $S^{\prime}$ from $S$ by evicting e instead of $f$


## Farthest-In-Future: Analysis

- Proof. (continued)
- Case 3: ( $\mathbf{d}$ is not in the cache; $\mathbf{S}_{\text {FIF }}$ evicts $\mathbf{e}$; $\mathbf{S}$ evicts $\mathbf{f} \neq \mathbf{e}$ ).
- begin construction of $S^{\prime}$ from $S$ by evicting e instead of $f$



## Farthest-In-Future: Analysis

- Proof. (continued)
- Case 3: (d is not in the cache; $\mathbf{S}_{\text {FIF }}$ evicts $\mathbf{e}$; $\mathbf{S}$ evicts $\mathbf{f} \neq \mathbf{e}$ ).
- begin construction of $S^{\prime}$ from $S$ by evicting e instead of $f$

- now S' agrees with $\mathbf{S}_{\text {FIF }}$ on first j+1 requests; we show that having element $\mathbf{f}$ in cache is no worse than having element $\mathbf{e}$


## Farthest-In-Future: Analysis

- Proof. (continued)
- Case 3: (d is not in the cache; $\mathbf{S}_{\text {FIF }}$ evicts $\mathbf{e}$; $\mathbf{S}$ evicts $\mathbf{f} \neq \mathbf{e}$ ).
- begin construction of $S^{\prime}$ from $S$ by evicting e instead of $f$

- now S' agrees with $\mathbf{S}_{\text {FIF }}$ on first j+1 requests; we show that having element $f$ in cache is no worse than having element $\mathbf{e}$
- Continue building S' to be the same as $\mathbf{S}$ until forced to be different


## Farthest-In-Future: Analysis

- Let j' be the first time after j+1 that $\mathbf{S}$ and $\mathbf{S}^{\prime}$ must take a different action, and let $\mathbf{g}$ be item requested at time $j^{\prime}$.
$j^{\prime}$

- Case 3a: $\mathbf{g}=\mathbf{e}$. Can't happen: e was evicted by Farthest-In-Future so there must be a request for $f$ before $e$.
- Case 3b: $\mathbf{g}=\mathbf{f}$. Element $\mathbf{f}$ can't be in cache of $\mathbf{S}$, so let $\mathrm{e}^{\prime}$ be the element that S evicts.
- if $\mathbf{e}^{\prime}=\mathbf{e}, \mathbf{S}$ ' accesses f from cache; now $\mathbf{S}$ and $\mathbf{S}^{\prime}$ have same cache
- if $\mathbf{e}^{\prime} \neq \mathbf{e}, \mathbf{S}^{\prime}$ evicts $\mathbf{e}^{\prime}$ and brings $\mathbf{e}$ into the cache; now $\mathbf{S}$ and $\mathbf{S}^{\prime}$ have the same cache


## Farthest-In-Future: Analysis

- Let j' be the first time after j+1 that $\mathbf{S}$ and $\mathbf{S '}^{\prime}$ must take a different action, and let g be item requested at time $\mathrm{j}^{\prime}$.
$j^{\prime}$
same
e

| same | $f$ |
| :---: | :---: |
| $S^{\prime}$ |  |

## Farthest-In-Future: Analysis

- Let j' be the first time after j+1 that $\mathbf{S}$ and $\mathbf{S '}^{\prime}$ must take a different action, and let $\mathbf{g}$ be item requested at time $j^{\prime}$.
$j^{\prime}$

- Case 3a: $\mathbf{g}=\mathbf{e}$. Can't happen: e was evicted by Farthest-In-Future so there must be a request for $f$ before $e$.


## Farthest-In-Future: Analysis

- Let j' be the first time after j+1 that $\mathbf{S}$ and $\mathrm{S}^{\prime}$ must take a different action, and let g be item requested at time $\mathrm{j}^{\prime}$.
$j^{\prime}$
same
e
S

S'
- Case 3a: $\mathbf{g}=\mathbf{e}$. Can't happen: e was evicted by Farthest-In-Future so there must be a request for $f$ before $e$.
- Case 3b: $\mathbf{g = f}$. Element $\mathbf{f}$ can't be in cache of $\mathbf{S}$, so let $\mathbf{e}$ ' be the element that S evicts.
- if e' = e, S' accesses from cache; now $\mathbf{S}$ and $\mathbf{S}^{\prime}$ have same cache
- if $\mathbf{e}^{\prime} \neq \mathbf{e}, \mathbf{S}^{\prime}$ evicts $\mathbf{e}^{\prime}$ and brings $\mathbf{e}$ into the cache; now $\mathbf{S}$ and $\mathbf{S}^{\prime}$ have the same cache


Note: $\mathbf{S}^{\prime}$ is no longer reduced, but can be transformed into
a reduced schedule that agrees with $\mathrm{S}_{\text {FIF }}$ through step j+1

## Farthest-In-Future: Analysis

- Let j' be the first time after $\mathbf{j}+1$ that $\mathbf{S}$ and $\mathrm{S}^{\prime}$ must take a different action, and let $\mathbf{g}$ be item requested at time $j^{\prime}$.


> otherwise S' would take the same action

- Case Bc: $\mathbf{g} \neq \mathbf{e}$ and $\mathbf{g} \neq \mathbf{f}$. S must evict e.

Make $S^{\prime}$ evict $f$; now $S$ and $S^{\prime}$ have the same cache. •


## Farthest-In-Future: Analysis

- Let j' be the first time after j+1 that S and S' must take a different action, and let $g$ be item requested at time $j$ '.
$j^{\prime}$
same $e$

S
same
S'
otherwise S' would take the same action

- Case $3 \mathrm{c}: \mathbf{g} \neq \mathbf{e}$ and $\mathrm{g} \neq \mathrm{f}$. S must evict $\mathbf{e}$.

Make $\mathbf{S}^{\prime}$ evict f; now $\mathbf{S}$ and $\mathbf{S}$ ' have the same cache. •


In each case can now extend $S^{\prime}$ using rest of $S$ at no extra cost. $\mathrm{S}^{\prime}$ is optimal, reduced, and agrees with $\mathrm{S}_{\text {FIF }}$ for $\mathrm{j}+1$ steps Optimality of $\mathrm{S}_{\text {FIF }}$ follows by induction.

## Caching Perspective

- Online vs. offline algorithms.
- Offline: full sequence of requests is known a priori.
- Online (reality): requests are not known in advance.
- Caching is among most fundamental online problems in CS.
- LIFO. Evict page brought in most recently.
- LRU. Evict page whose most recent access was earliest.

FIF with direction of time reversed!

- Theorem. FIF is optimal offline eviction algorithm.
- Provides basis for understanding and analyzing online algorithms.
- LRU is k-competitive. [Section 13.8]
- LIFO is arbitrarily bad.


## Greedy Analysis Strategies

- Greedy algorithm stays ahead. Show that after each step of the greedy algorithm, its solution is at least as good as any other algorithm's.
- Exchange argument. Gradually transform any solution to the one found by the greedy algorithm without hurting its quality.
- Structural. Discover a simple "structural" bound asserting that every possible solution must have a certain value. Then show that your algorithm always achieves this bound.


## A Note on Optimal Caching

- In real operating conditions one typically needs an on-line algorithm
- make the eviction decisions as each memory request arrives
- To design and analyze these algorithms it is also important to understand how the best possible decisions can be made if one did know the future
- Field of on-line algorithms compares the quality of on-line decisions to that of the optimal off-line schedule
- Bellady's FIF algorithm tells us what an optimal schedule looks like so we have a baseline to compare online algorithms


## Belady's Greedy Algorithm: Farthest-In-Future

- Given sequence $\mathrm{D}=\mathrm{d}_{1}, \mathrm{~d}_{2}, \ldots, \mathrm{~d}_{\mathrm{m}}$
- When $d_{i}$ needs to be brought into the cache evict the item that is needed farthest in the future
- Let $\operatorname{NextAccess} \mathrm{s}_{\mathrm{i}}(\mathrm{d})=\min \left\{\mathrm{j} \geq \mathrm{i}: \mathrm{d}_{\mathrm{j}}=\mathrm{d}\right\}$ be the next point in $\mathbf{D}$ that item $d$ will be requested
- Evict $\mathbf{d}$ such that $\operatorname{NextAccess}_{i}(\mathbf{d})$ is largest


## Optimal Caching/Paging

- Given a memory request d from U
- If $d$ is stored in the cache we can access it quickly
- If not then we call it a cache miss and (since the cache is full)
- we must bring it into cache and evict some other data item from the cache
- which one to evict?
- Given a sequence $\mathrm{D}=\mathrm{d}_{1}, \mathrm{~d}_{2}, \ldots, \mathrm{~d}_{\mathrm{m}}$ of elements from $U$ corresponding to memory requests
- Find a sequence of evictions (an eviction schedule) that has as few cache misses as possible


## Caching Example

- $\mathrm{n}=3, \mathrm{k}=2, \mathrm{U}=\{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$
- Cache initially contains $\{\mathbf{a}, \mathbf{b}\}$
- D=abcbcab
- $\mathrm{S}=\mathrm{a} \quad \mathrm{c}$
$=C=\left[\begin{array}{ll}a & l \\ b & b \\ c\end{array}\right]$
- This is optimal


## Interval scheduling

- Formally
- Requests 1,2,...,n
- request $\mathbf{i}$ has start time $\mathbf{s}_{\mathbf{i}}$ and finish time $\mathbf{f}_{\mathrm{i}}>\mathbf{s}_{\mathrm{i}}$
- Requests i and j are compatible iff either
- request $i$ is for a time entirely before request $j$
$-f_{i} \leq s_{j}$
- or, request $j$ is for a time entirely before request i
$-f_{j} \leq s_{i}$
- Set A of requests is compatible iff every pair of requests $i, j \in \mathbf{A}, i \neq j$ is compatible
- Goal: Find maximum size subset A of compatible requests


## Implementing the Greedy Algorithm

- Sort the requests by finish time
- O(nlog n) time
- Maintain current latest finish time scheduled
- Keep array of start times indexed by request number
- Only eliminate incompatible requests as needed
- Walk along array of requests sorted by finish times skipping those whose start time is before current latest finish time scheduled
- O(n) additional time for greedy algorithm


## Scheduling all intervals

- Interval Partitioning Problem: We have resources to serve more than one request at once and want to schedule all the intervals using as few of our resources as possible
- Obvious requirement: At least the depth of the set of requests


