

Fibonacci Heaps

Outline for Today

- ***Review from Last Time***
 - Quick refresher on binomial heaps and lazy binomial heaps.
- ***The Need for decrease-key***
 - An important operation in many graph algorithms.
- ***Fibonacci Heaps***
 - A data structure efficiently supporting ***decrease-key***.
- ***Representational Issues***
 - Some of the challenges in Fibonacci heaps.

Review: (Lazy) Binomial Heaps

Building a Priority Queue


- Group nodes into “packets” with the following properties:
 - Size must be a power of two.
 - Can efficiently fuse packets of the same size.
 - Can efficiently find the minimum element of each packet.
 - Can efficiently “fracture” a packet of 2^k nodes into packets of 1, 2, 4, 8, ..., 2^{k-1} nodes.

Binomial Trees

- A **binomial tree of order k** is a type of tree recursively defined as follows:


A binomial tree of order k is a single node whose children are binomial trees of order $0, 1, 2, \dots, k - 1$.

- Here are the first few binomial trees:



Binomial Trees

- A **heap-ordered binomial tree** is a binomial tree whose nodes obey the heap property: all nodes are less than or equal to their descendants.
- We will use heap-ordered binomial trees to implement our “packets.”



The Binomial Heap

- A **binomial heap** is a collection of heap-ordered binomial trees stored in ascending order of size.
- Operations defined as follows:
 - **meld**(pq_1, pq_2): Use addition to combine all the trees.
 - Fuses $O(\log n)$ trees. Total time: $O(\log n)$.
 - **$pq.enqueue(v, k)$** : Meld pq and a singleton heap of (v, k) .
 - Total time: $O(\log n)$.
 - **$pq.find-min()$** : Find the minimum of all tree roots.
 - Total time: $O(\log n)$.
 - **$pq.extract-min()$** : Find the min, delete the tree root, then meld together the queue and the exposed children.
 - Total time: $O(\log n)$.

Lazy Binomial Heaps

- A ***lazy binomial heap*** is a variation on a standard binomial heap in which ***melds*** are done lazily by concatenating tree lists together.
- Tree roots are stored in a doubly-linked list.
- An extra pointer is required that points to the minimum element.
- ***extract-min*** eagerly coalesces binomial trees together and runs in amortized time $O(\log n)$.

The Overall Analysis

- Set $\Phi(D)$ to be the number of trees in D .
- The *amortized* costs of the operations on a lazy binomial heap are as follows:
 - ***enqueue***: $O(1)$
 - ***meld***: $O(1)$
 - ***find-min***: $O(1)$
 - ***extract-min***: $O(\log n)$
- Details are in the previous lecture.
- Let's quickly review ***extract-min***'s analysis.

Analyzing Extract-Min

- Initially, we expose the children of the minimum element. This takes time $O(\log n)$.
- Suppose that at this point there are T trees. The runtime for the coalesce is $\Theta(T)$.
- When we're done merging, there will be $O(\log n)$ trees remaining, so $\Delta\Phi = -T + O(\log n)$.
- Amortized cost is

$$\begin{aligned} & O(\log n) + \Theta(T) + O(1) \cdot (-T + O(\log n)) \\ &= O(\log n) + \Theta(T) - O(1) \cdot T + O(1) \cdot O(\log n) \\ &= \mathbf{O(\log n)}. \end{aligned}$$

A Detail in the Analysis

- The amortized cost of an *extract-min* is


$$O(\log n) + \Theta(T) + O(1) \cdot (-T + O(\log n))$$

- Where do these $O(\log n)$ terms come from?
 - First $O(\log n)$: Removing the minimum element might expose $O(\log n)$ children, since the maximum order of a tree is $O(\log n)$.
 - Second $O(\log n)$: Maximum number of trees after a coalesce is $O(\log n)$.
- **A different intuition:** Let $M(n)$ be the maximum possible order of a tree in a lazy binomial heap.
- Amortized runtime is $O(M(n))$.

The Need for *decrease-key*

Review: Dijkstra's Algorithm

- Dijkstra's algorithm solves the single-source shortest paths (SSSP) problem in graphs with nonnegative edge weights.




Dijkstra and Priority Queues

- At each step of Dijkstra's algorithm, we need to do the following:
 - Find the node at v minimum distance from s .
 - Update the candidate distances of all the nodes connected to v . (Distances only decrease in this step.)
- This first step sounds like an ***extract-min*** on a priority queue.
- How would we implement the second step?


Review: Prim's Algorithm

- Prim's algorithm solves the minimum spanning tree (MST) problem in undirected graphs.



Review: Prim's Algorithm

- Prim's algorithm solves the minimum spanning tree (MST) problem in undirected graphs.



Prim and Priority Queues

- At each step of Prim's algorithm, we need to do the following:
 - Find the node v outside of the spanning tree with the lowest-cost connection to the tree.
 - Update the candidate distances from v to nodes outside the set S .
- This first step sounds like an ***extract-min*** on a priority queue.
- How would we implement the second step?

The *decrease-key* Operation

- Some priority queues support the operation $pq.\mathbf{decrease-key}(v, k)$, which works as follows:

Given a pointer to an element v in pq , lower its key (priority) to k . It is assumed that k is less than the current priority of v .

- This operation is crucial in efficient implementations of Dijkstra's algorithm and Prim's MST algorithm.

Dijkstra and *decrease-key*

- Dijkstra's algorithm can be implemented with a priority queue using
 - $O(n)$ total *enqueues*,
 - $O(n)$ total *extract-mins*, and
 - $O(m)$ total *decrease-keys*.
- Dijkstra's algorithm runtime is

$$O(n T_{\text{enq}} + n T_{\text{ext}} + m T_{\text{dec}})$$

Prim and *decrease-key*

- Prim's algorithm can be implemented with a priority queue using
 - $O(n)$ total *enqueues*,
 - $O(n)$ total *extract-mins*, and
 - $O(m)$ total *decrease-keys*.
- Prim's algorithm runtime is

$$O(n T_{\text{enq}} + n T_{\text{ext}} + m T_{\text{dec}})$$

Standard Approaches

- In a binary heap, *enqueue*, *extract-min*, and *decrease-key* can be made to work in time $O(\log n)$ time each.
- Cost of Dijkstra's / Prim's algorithm:
$$O(n T_{\text{enq}} + n T_{\text{ext}} + m T_{\text{dec}})$$
$$= O(n \log n + n \log n + m \log n)$$
$$= \mathbf{O(m \log n)}$$

Standard Approaches

- In a binomial heap, n **enqueues** takes time $O(n)$, each **extract-min** takes time $O(\log n)$, and each **decrease-key** takes time $O(\log n)$.
- Cost of Dijkstra's / Prim's algorithm:
$$O(n T_{\text{enq}} + n T_{\text{ext}} + m T_{\text{dec}})$$
$$= O(n + n \log n + m \log n)$$
$$= \mathbf{O(m \log n)}$$

Where We're Going

- The **Fibonacci heap** has these runtimes:
 - *enqueue*: $O(1)$
 - *meld*: $O(1)$
 - *find-min*: $O(1)$
 - *extract-min*: $O(\log n)$, amortized.
 - *decrease-key*: $O(1)$, amortized.

- Cost of Prim's or Dijkstra's algorithm:


$$\begin{aligned} & O(n T_{\text{enq}} + n T_{\text{ext}} + m T_{\text{dec}}) \\ &= O(n + n \log n + m) \\ &= \mathbf{O(m + n \log n)} \end{aligned}$$

- This is theoretically optimal for a comparison-based priority queue in Dijkstra's or Prim's algorithms.

The Challenge of *decrease-key*


A Simple Implementation

- It is possible to implement *decrease-key* in time $O(\log n)$ using lazy binomial heaps.
- **Idea:** “Bubble” the element up toward the root of the binomial tree containing it and (potentially) update the *min* pointer.




A Simple Implementation

- It is possible to implement *decrease-key* in time $O(\log n)$ using lazy binomial heaps.
- **Idea:** “Bubble” the element up toward the root of the binomial tree containing it and (potentially) update the *min* pointer.



A Simple Implementation


- It is possible to implement *decrease-key* in time $O(\log n)$ using lazy binomial heaps.
- **Idea:** “Bubble” the element up toward the root of the binomial tree containing it and (potentially) update the *min* pointer.




The Challenge

- **Goal:** Implement *decrease-key* in amortized time $O(1)$.
- Why is this hard?
 - Lowering a node's priority might break the heap property.
 - Correcting the imbalance $O(\log n)$ layers deep in a tree might take time $O(\log n)$.
- We will need to change our approach.


A Crazy Idea



A Crazy Idea



A Crazy Idea




A Crazy Idea

- To implement *decrease-key* efficiently:
 - Lower the key of the specified node.
 - If its key is greater than or equal to its parent's key, we're done.
 - Otherwise, cut that node from its parent and hoist it up to the root list, optionally updating the min pointer.
- Time required: $O(1)$.
 - This requires some changes to the tree representation; more details later.


Tree Sizes and Orders

- **Recall:** A binomial tree of order k has 2^k nodes and the root has k children.
- Going forward, we'll say that the **order** of a node is the number of children it has.
- **Concern:** If trees can be cut, a tree of order k might have many fewer than 2^k nodes.

The Problem



The Problem



Number of nodes: $\Theta(k^2)$

Number of trees: $\Theta(n^{1/2})$

The Problem

- **Recall:** The amortized cost of an **extract-min** is $O(M(n))$, where $M(n)$ is the maximum order of a tree in the heap.
- With true binomial trees, this is $O(\log n)$.
- With our “damaged” binomial trees, this can be $\Theta(n^{1/2})$.
- We've lost our runtime bounds!

Time-Out for Announcements!

Problem Sets

- Problem Set Three was due at the start of class today.
 - Want to use late days? Feel free to submit it by Saturday at 3:00PM.
- The next problem set goes out on Tuesday. Enjoy a little break from the problem sets!



HUNTINGTON'S OUTREACH PROJECT FOR EDUCATION, AT STANFORD

The Huntington's Outreach Project for Education, at Stanford (HOPES) is an educational service project working to build a web resource on Huntington's disease (HD). Our mission is to make scientific information about HD more readily accessible to patients, their families, and the general public. We are currently hiring **student researchers** (writers), **graphic designers**, and **web developers** for the **2016-2017 school year**.

Student researcher positions:

As a researcher, you will be responsible for researching a specific HD-related topic, writing articles based on your research, and planning the graphics to go along with your article. You are also expected to play a big role in the editorial process for both your own work and the work of other group members. Applicants should have a strong background in biology, human biology, or anthropological sciences, including a good working knowledge of genetics. Strong writing, editing, and communication skills are also necessary.

Graphic designer positions:

Graphic designers work on the most popular parts of the site including illustrated books, articles and interactive tutorials. Responsibilities include collaborating with researchers to visually enhance the educational text, creating interactive tutorials, and brainstorming new projects for the HOPES website. Although not required, college-level biology background is a strong asset. Other skills we are looking for include: experience with Adobe Photoshop or another graphics editing program, Illustrator, Flash, vector or 3D graphics, digital video editing, and web design.

Web developer positions:

We are hiring a web designer to improve our Wordpress site. This is a great opportunity for someone with an interest in web design to learn and improve as a web designer, and develop a competitive portfolio. Experience preferred, but not necessary.

Compensation:

Units or pay (units through Anthropology or HumBio; starting salary of **\$16** per hour)

Commitment:

Full time or part-time during the summer and part-time throughout the school year. During the school year, weekly hours are flexible and most work is independent, but you must be able to average 6-10 hours of work per week. There will also be group workshops, outreach events, and weekly meetings, all of which will be scheduled according to the availability of the team.

Faculty Coordinator:

Prof. Bill Durham

HOW TO APPLY:

Applications for all positions are due on **Sunday, May 1 at 11:59pm**

Please send a **current resume, letter of application, and unofficial transcript** to HOPES Project Leader Kristen Powers (kapowers@stanford.edu) with the subject line "YOUR LAST NAME - HOPES Application". The letter should include a candid discussion of your qualifications for the position, your other time commitments, your leadership skills, and your reasons for interest in the position.

Student researchers: Please attach two writing samples, science-related and/or research-based in nature.

Graphic designers: Please send in 3 recent design samples with a brief description about each (tools used, time spent, purpose/client, etc).

Web developers: Please send links to any web-design work you may have done.

For more information, please visit our website at <http://hopes.stanford.edu> and/or email the Project Leader.

Back to CS166!

The Problem


- This problem arises because we have lost one of the guarantees of binomial trees:
 A binomial tree of order k has 2^k nodes.
- When we cut low-hanging trees, the root node won't learn that these trees are missing.
- However, communicating this information up from the leaves to the root might take time $O(\log n)$!

The Tradeoff

- If we don't impose any structural constraints on our trees, then trees of large order may have too few nodes.
 - Leads to $M(n)$ getting too high, wrecking our runtime bounds for *extract-min*.
- If we impose too many structural constraints on our trees, then we have to spend too much time fixing up trees.
 - Leads to *decrease-key* taking too long.
- How can we strike a balance?


The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.



The Compromise

- Every non-root node is allowed to lose at most one child.
- If a non-root node loses two children, we cut it from its parent. (This might trigger more cuts.)
- We will **mark** nodes in the heap that have lost children to keep track of this fact.




The Compromise

- To cut node v from its parent p :
 - Unmark v .
 - Cut v from p .
 - If p is not already marked and is not the root of a tree, mark it.
 - If p was already marked, recursively cut p from its parent.

The Compromise

- If we do a few **decrease-keys**, then the tree won't lose “too many” nodes.
- If we do many **decrease-keys**, the information slowly propagates to the root.





Assessing the Impact

- The amortized cost of an *extract-min* is $O(M(n))$, where $M(n)$ is the maximum possible order of a tree.
- This used to be $O(\log n)$ because our trees had exponentially many nodes in them.
- What is it now?

Two Extremes

- If we never do any *decrease-keys*, then the trees in our data structure are all binomial trees.
- Each tree of order k has 2^k nodes in it, the maximum possible order is $O(\log n)$.
- On the other hand, suppose that all trees in the binomial heap have lost the maximum possible number of nodes.
- In that case, how many nodes will each tree have?


Maximally-Damaged Trees



A *maximally-damaged tree of order k* is a node whose children are maximally-damaged trees of orders

$0, 0, 1, 2, 3, \dots, k - 2.$

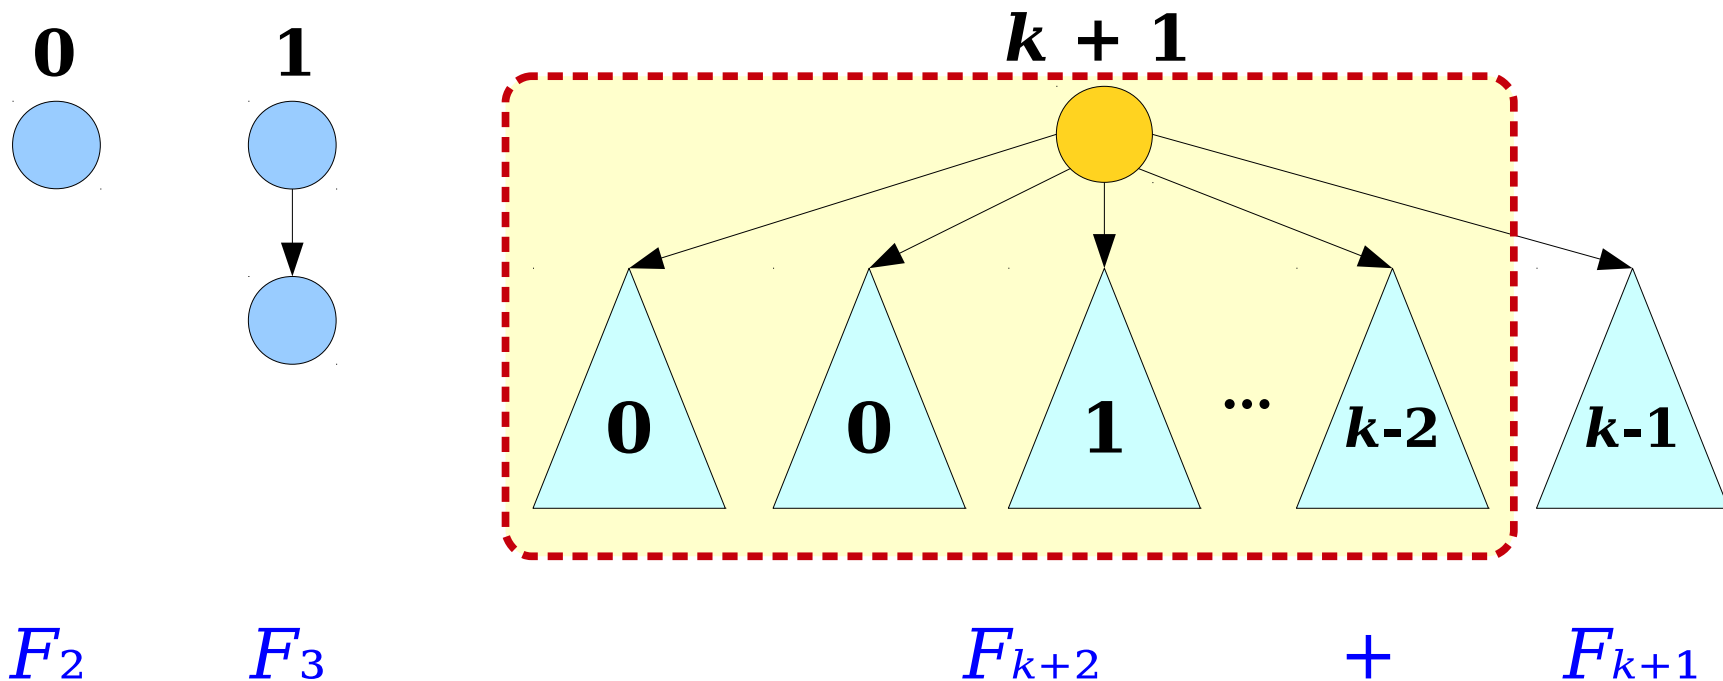
Maximally-Damaged Trees



Claim: The minimum number of nodes in a tree of order k is F_{k+2}

Maximally-Damaged Trees

- **Theorem:** The number of nodes in a maximally-damaged tree of order k is F_{k+2} .
- **Proof:** Induction.



φ -bonacci Numbers

- **Fact:** For $n \geq 2$, we have $F_n \geq \varphi^{n-2}$, where φ is the golden ratio:

$$\varphi \approx 1.61803398875\dots$$

- **Claim:** In our modified data structure, we have $M(n) = O(\log n)$.
- **Proof:** In a tree of order k , there are at least $F_{k+2} \geq \varphi^k$ nodes. Therefore, the maximum order of a tree in our data structure is $\log_{\varphi} n = O(\log n)$. ■

Fibonacci Heaps

- A ***Fibonacci heap*** is a lazy binomial heap where ***decrease-key*** is implemented using the earlier cutting-and-marking scheme.
- Operation runtimes:
 - ***enqueue***: $O(1)$
 - ***meld***: $O(1)$
 - ***find-min***: $O(1)$
 - ***extract-min***: $O(\log n)$ amortized
 - ***decrease-key***: Up next!

Analyzing *decrease-key*

- In the best case, *decrease-key* takes time $O(1)$ when no cuts are made.
- In the worst case, *decrease-key* takes time $O(C)$, where C is the number of cuts made.
- What is the *amortized* cost of a *decrease-key*?

Refresher: Our Choice of Φ

- In our amortized analysis of lazy binomial heaps, we set Φ to be the number of trees in the heap.
- With this choice of Φ , we obtained these amortized time bounds:
 - *enqueue*: $O(1)$
 - *meld*: $O(1)$
 - *find-min*: $O(1)$
 - *extract-min*: $O(\log n)$

Rethinking our Potential

- Intuitively, a cascading cut only occurs if we have a long chain of marked nodes.
- Those nodes were only marked because of previous *decrease-key* operations.
- **Idea:** Backcharge the work required to do the cascading cut to each preceding *decrease-key* that contributed to it.
- Specifically, change Φ as follows:
$$\Phi = \#trees + \#marked$$
- **Note:** Since only *decrease-key* interacts with marked nodes, our amortized analysis of all previous operations is still the same.

The (New) Amortized Cost

- Using our new Φ , a *decrease-key* makes C cuts, it
 - Marks one new node (+1),
 - Unmarks C nodes ($-C$), and
 - Adds C trees to the root list ($+C$).

- Amortized cost is

$$\begin{aligned} & \Theta(C) + O(1) \cdot \Delta\Phi \\ = & \Theta(C) + O(1) \cdot (1 - C + C) \\ = & \Theta(C) + O(1) \cdot 1 \\ = & \Theta(C) + O(1) \\ = & \Theta(C) \end{aligned}$$

- Hmm... that didn't work.

The Trick

- Each *decrease-key* makes extra work for *two* future operations:
 - Future *extract-mins* that now have more trees to coalesce, and
 - Future *decrease-keys* that might have to do cascading cuts.
- We can make this explicit in our potential function:

$$\Phi = \#trees + 2 \cdot \#marked$$

The (Final) Amortized Cost

- Using our new Φ , a *decrease-key* makes C cuts, it
 - Marks one new node (+2),
 - Unmarks C nodes ($-2C$), and
 - Adds C trees to the root list ($+C$).

- Amortized cost is

$$\begin{aligned} & \Theta(C) + O(1) \cdot \Delta\Phi \\ &= \Theta(C) + O(1) \cdot (2 - 2C + C) \\ &= \Theta(C) + O(1) \cdot (2 - C) \\ &= \Theta(C) - O(C) + O(1) \\ &= \Theta(1) \end{aligned}$$

- We now have amortized $O(1)$ *decrease-key*!

The Story So Far



- The Fibonacci heap has the following amortized time bounds:
 - ***enqueue***: $O(1)$
 - ***find-min***: $O(1)$
 - ***meld***: $O(1)$
 - ***decrease-key***: $O(1)$ amortized
 - ***extract-min***: $O(\log n)$ amortized
- This is about as good as it gets!

The Catch: Representation Issues



Representing Trees

- The trees in a Fibonacci heap must be able to do the following:
 - During a merge: Add one tree as a child of the root of another tree.
 - During a cut: Cut a node from its parent in time $O(1)$.
- ***Claim:*** This is trickier than it looks.

Representing Trees

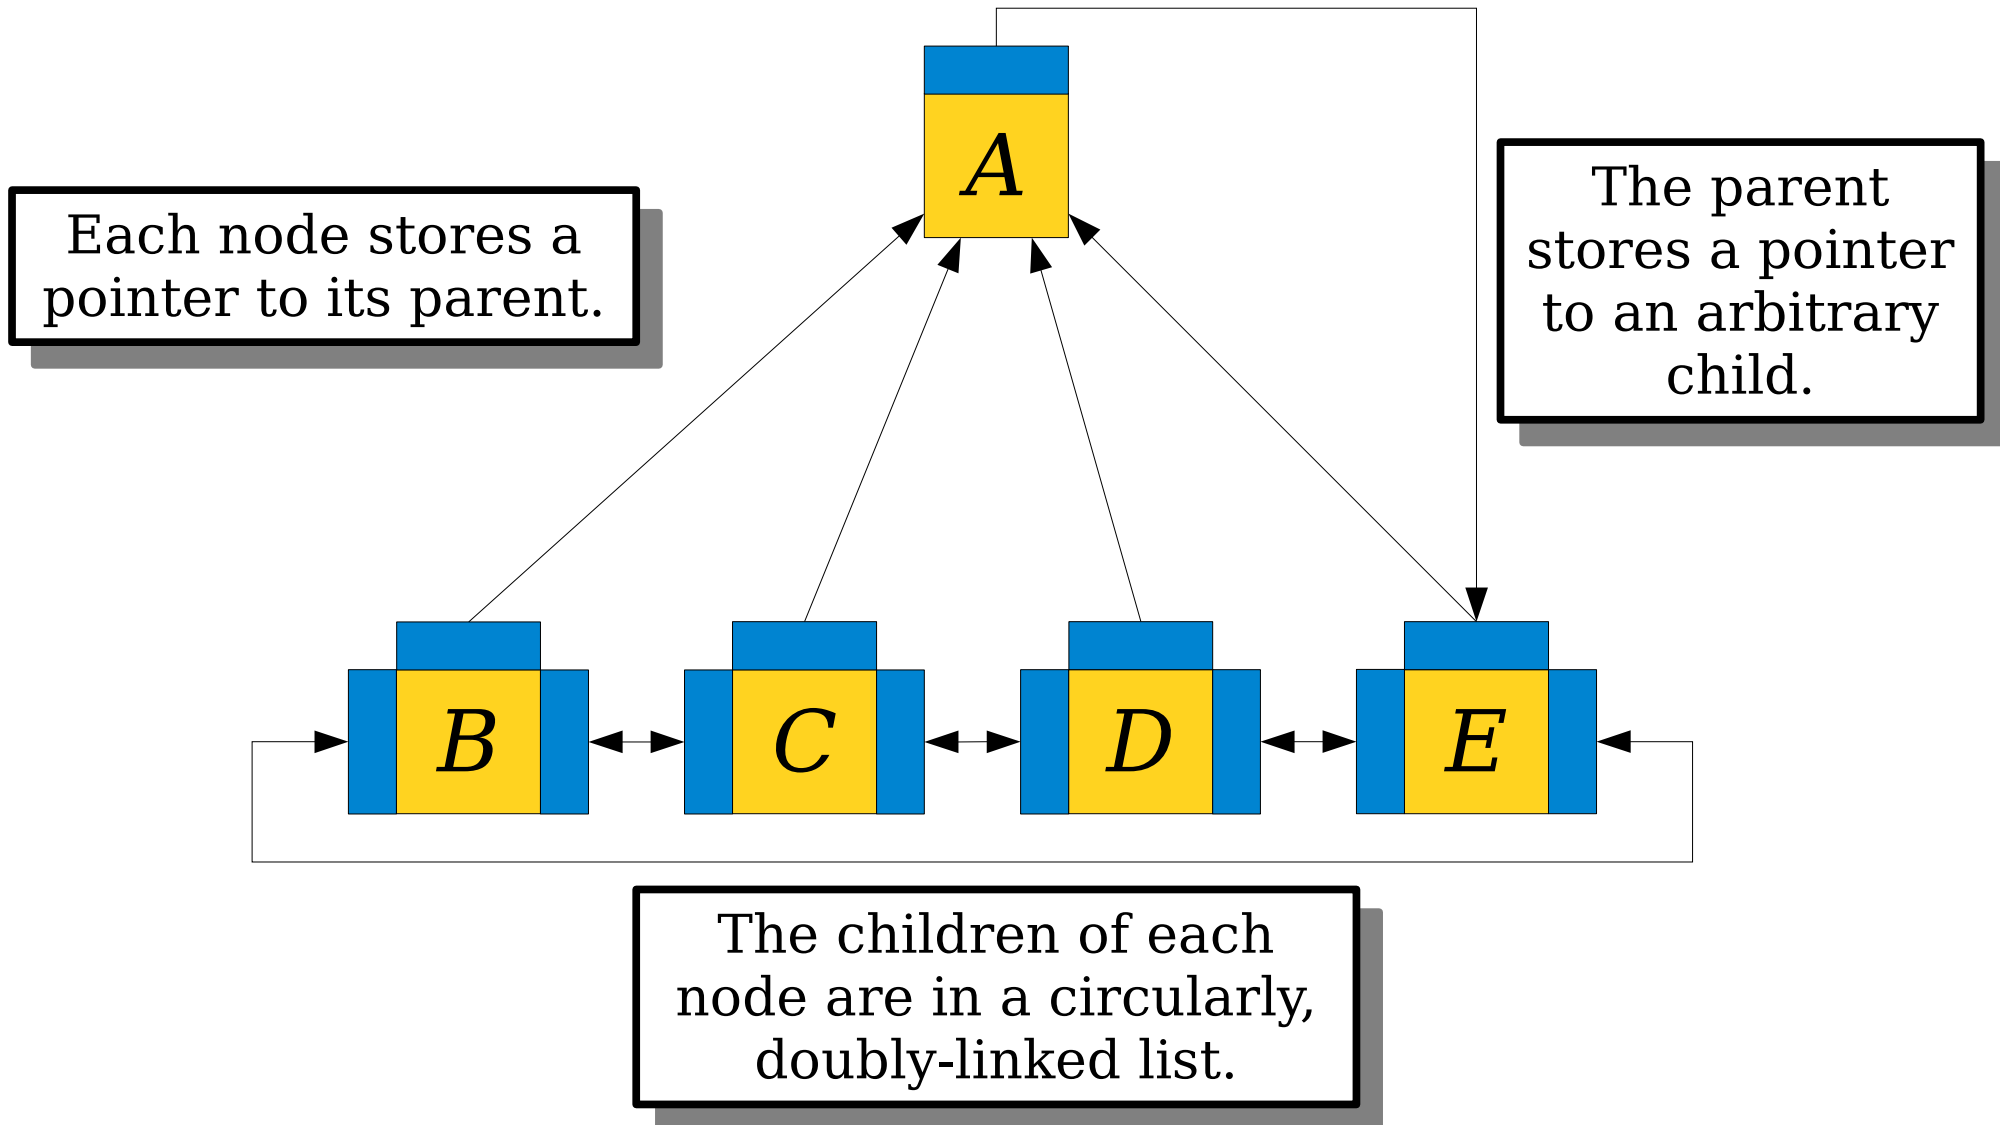


Representing Trees



Finding this pointer might take time $\Theta(\log n)$!

The Solution



Awful Linked Lists

- Trees are stored as follows:
 - Each node stores a pointer to *some* child.
 - Each node stores a pointer to its parent.
 - Each node is in a circularly-linked list of its siblings.
- Awful, but the following possible are now possible in time $O(1)$:
 - Cut a node from its parent.
 - Add another child node to a node.
- This is the main reason Fibonacci heaps are so complex.

Fibonacci Heap Nodes

- Each node in a Fibonacci heap stores
 - A pointer to its parent.
 - A pointer to the next sibling.
 - A pointer to the previous sibling.
 - A pointer to an arbitrary child.
 - A bit for whether it's marked.
 - Its order.
 - Its key.
 - Its element.

In Practice

- In practice, Fibonacci heaps are slower than other heaps with worse asymptotic performance.
- Why?
 - Huge memory requirements per node.
 - High constant factors on all operations.
 - Poor locality of reference and caching.

In Theory

- That said, Fibonacci heaps are worth knowing about for several reasons:
 - Clever use of a two-tiered potential function shows up in lots of data structures.
 - Implementation of *decrease-key* forms the basis for many other advanced priority queues.
 - Gives the theoretically optimal comparison-based implementation of Prim's and Dijkstra's algorithms.

Summary

- *decrease-key* is a useful operation in many graph algorithms.
- Implement *decrease-key* by cutting a node from its parent and hoisting it up to the root list.
- To make sure trees of high order have lots of nodes, add a marking scheme and cut nodes that lose two or more children.
- Represent the data structure using Awful Linked Lists.
- Can prove that the number of trees is $O(\log n)$ by most maximally damaged trees in the heap.

Next Time

- ***Splay Trees***
 - Amortized-efficient balanced trees.
- ***Static Optimality***
 - Is there a single best BST for a set of data?
- ***Dynamic Optimality***
 - Is there a single best BST for a set of data if that BST can change over time?