

Deadlocks

CSE 332 Summer 2021

Instructor: Kristofer Wong

Teaching Assistants:

Alena Dickmann

Arya GJ

Finn Johnson

Joon Chong

Kimi Locke

Peyton Rapo

Rahul Misal

Winston Jodjana

Announcements

- ❖ Reminder: Ex 11, 12 submit on Gradescope
- ❖ EC exercise due on Gradescope **separately**
- ❖ P3 Fix pushed yesterday
- ❖ Final Exam

Oral Final Exam: 8/16 – 8/18

- ❖ 5 – 10 min
 - Link to sign up on Google Calendar posted today, due 8/8 @11:59
 - If you cannot make any of the available timeslots, email staff before 8/8 @11:59
- ❖ Conversational style; we will ask 2 questions from **any topics covered in lecture**
 - What can you tell us about X?
 - How would changing Y in Z affect Z?
 - Explain why A is good/better/useful.
 - Please give us an example of B.
- ❖ Grading:
 - 40 pts - Checkboxes (did you talk about x?)
 - 10 pts - Accuracy (do you know what you're talking about)
 - 10 pts - Clarity & Communication
 - 10 pts - Post-exam Reflection (Opportunity to get half credit back on any missed checkboxes)
- ❖ Optionally recorded
 - Accuracy, Clarity, and Communication can sometimes be subjective. **You may request a regrade iff you opt in to recording.** We will only watch recordings if you request that we review your exam.

Lecture Outline

- ❖ **Five Guidelines to Avoid Race Conditions**
- ❖ Deadlocks
- ❖ Graphs!!


Avoiding Race Conditions

- ❖ Avoiding *race conditions* on shared resources is difficult
 - What ‘seems fine’ in a sequential world got us into trouble when we introduced concurrency

- ❖ Decades of bugs have led to some techniques known to work
 - More info:
 - “Java Concurrency in Practice”, ch 2
 - Grossman Notes, section 8
 - None of these techniques are specific to Java or a particular book!
 - Hard to appreciate right now, but refer back to these!


3 Choices: Categorizing Memory Locations

- ❖ Every *memory location* (e.g., object field) in your program must obey at least one of the following:
 1. *Thread-local*: Do not use the location in > 1 thread
 2. *Immutable*: Do not write to the memory location
 3. *Shared-and-mutable*: Use synchronization to control access



Category #1: Thread-local

- ❖ Whenever possible, do not share resources
 - Easier for each thread to maintain its own thread-local *copy* of a resource than to have a global resource with shared updates
 - Correct only if threads don't need to communicate via the resource
 - Remember: the stack is thread-local, so never need to synchronize on local variables




Guideline #(-2): Minimize shared memory

The vast majority of objects should be thread-local

Category #2: Immutable

- ❖ Sometimes, objects don't need to be mutated at all!
 - In practice, programmers over-use mutation
- ❖ When possible, don't mutate objects; make new ones instead!
 - A key tenet of functional programming (see CSE 341); functional programming concepts helpful in a concurrent setting!
 - Yes, uses more space
 - Sometimes is a tradeoff for better concurrency




Guideline #(-1): Read-only locations do not require synchronization
Simultaneous reads are not races and not a problem

Category #3: Keep the Rest Synchronized

❖ After minimizing the amount of memory that is ...

1. thread-shared
2. mutable



❖ ... we need guidelines for keeping other data consistent

Guideline #0: No *data races*

- Never allow two threads to read/write or write/write the same location at the same time
- Use locks! Even if it “seems safe”

❖ *Necessary:*

- a Java or C program with a *data race* is almost always wrong

❖ *But Not Sufficient:*

- Our peek() example had no *data races*, and was still wrong ...

Guideline #1: Use Consistent Locking (1 of 2)


Guideline #1: Use consistent locking

For each location needing synchronization, have a lock that is always held when reading *or* writing the location

- ❖ We say the lock *guards* the location
 - Clearly document the guard for each location
- ❖ In Java, the guard is often the object containing the location
 - E.g.: `this` when inside the object's methods
- ❖ The same lock can (and often should) guard multiple locations
 - E.g.: multiple fields in a class
 - But also often guard a larger structure with one lock to ensure mutual exclusion on the entire structure

Guideline #1: Use Consistent Locking (2 of 2)

- ❖ The mapping from locations to locks is conceptual
 - Must be enforced by you as the programmer!
 - Partitions the shared-and-mutable locations by “which lock”




- ❖ Consistent locking is:
 - *Not sufficient*: Prevents **data races** but still allows **bad interleavings**
 - *(Aside) Not Necessary*: You could have different locking protocols for different phases of your program as long as all threads are coordinated when moving from one phase to next
 - eg. at start of program, data structure is being updated (needs locks); later it is not modified so can be read simultaneous (no locks)

Guideline #2: Lock Granularity (1 of 2)

❖ Lock *granularity* is a continuum. The two ends are:

■ **Coarse-grained:** Fewer locks, i.e., more objects per lock

- E.g.: One lock for entire data structure (e.g., array)
- E.g.: One lock for all bank accounts



■ **Fine-grained:** More locks, i.e., fewer objects per lock

- E.g.: One lock per data element (e.g., array index)
- E.g.: One lock per bank account



Guideline #2: Lock Granularity (2 of 2)

- ❖ There are tradeoffs at either end of the continuum:
 - **Coarse-grained** advantages:
 - Simpler to implement, especially implementing operations that access multiple locations (because all guarded by the same lock)
 - Much easier for operations that modify data-structure shape
 - **Fine-grained** advantages:
 - Enables more simultaneous access; coarse-grained locking may lead to unnecessary blocking)
 - Can make multi-node operations more difficult: say, rotations in an AVL tree

Guideline #2: Start with coarse-grained

Optimize for implementation simplicity, and move to fine-grained only if contention on the coarser locks becomes an issue

Lock Granularity Example: Separate Chaining Hashtable

- ❖ Continuum:
 - Coarse-grained: One lock for entire hashtable
 - Fine-grained: One lock for each bucket
- ❖ Which supports more concurrency for insert and lookup?
 - *Fine-grained allows simultaneous access to different buckets*
- ❖ Which makes `resize()`'s implementation easier? How would you do it?
 - *Coarse-grained; just grab one lock and proceed*
- ❖ If there is a `numElements` field, maintaining it will destroy the benefits of using separate locks for each bucket. Why?
 - *Updating it on each mutation without a lock creates a data race*
 - *Updating it on each mutation with a lock is coarse-grained locking*

Guideline #3: Critical Section Granularity

- ❖ A second, orthogonal granularity issue is critical-section size; i.e. “how much work should I do while holding lock(s)?”
- ❖ What happens if critical sections are too long?
 - *Performance loss because other threads are blocked*
- ❖ What happens if critical sections are too short?
 - *Bugs! You broke up something that shouldn't have been broken up; other threads can see intermediate state*

Guideline #3: Keep critical sections as small as possible while still being correct

Don't do expensive computations or I/O in critical sections, but also don't introduce race conditions

Critical Section Granularity: Example #1

- ❖ Change a key's value within a hashtable without removing it from the table
 - Assume `lock` guards the whole table
 - `expensive()` takes in the old value, and computes a new one, but takes a long time

Too long!

*(entire table locked
during expensive call)*

```
synchronized(lock) {  
    v1 = table.lookup(k);  
    v2 = expensive(v1);  
    table.remove(k);  
    table.insert(k, v2);  
}
```

Critical Section Granularity: Example #2

- ❖ Change a key's value within a hashtable without removing it from the table
 - Assume `lock` guards the whole table
 - `expensive()` takes in the old value, and computes a new one, but takes a long time

Too short!

(if another thread updated k 's value, we will lose their update)

```
synchronized(lock) {  
    v1 = table.lookup(k);  
}  
v2 = expensive(v1);  
synchronized(lock) {  
    table.remove(k);  
    table.insert(k, v2);  
}
```

Critical Section Granularity: Example #3

- ❖ Change a key's value within a hashtable without removing it from the table
 - Assume `lock` guards the whole table
 - `expensive()` takes in the old value, and computes a new one, but takes a long time

Just right

(if another update occurred, retry update again)

```
done = false;
while (!done) {
    synchronized(lock) {
        v1 = table.lookup(k);
    }
    v2 = expensive(v1);
    synchronized(lock) {
        if(table.lookup(k) == v1) {
            done = true;
            table.remove(k);
            table.insert(k, v2);
        }
    }
}
```

Guideline #4: Atomicity

- ❖ An operation is *atomic* if no other thread can see it partly executed
 - “Atomic”, as in “appears indivisible”
 - Typically want ADT operations atomic, even to other threads running operations on the same ADT

Guideline #4: Think about atomicity first, and locks second

Think in terms of what operations need to be atomic, and make critical sections just long enough to preserve atomicity. Only then should you design the locking protocol to implement the critical sections

Guideline #5: Don't Roll Your Own

- ❖ In “real life”, writing a data structure from scratch is ... rare
 - Standard libraries provide most of what you need
 - Team/Department/Company libraries usually provide the rest
 - CSE332 teaches key trade-offs, abstractions, and analysis of such implementations
- ❖ This is especially true for concurrent data structures!
 - Hard to write *correct* and *performant* on the first try; you're much more likely to write code with ***race conditions***

Guideline #5: Use libraries whenever they meet your needs

Standard libraries like ConcurrentHashMap were written by world experts. Do you really want to spend your time chasing down your bugs?

Lecture Outline

- ❖ Five Guidelines to Avoid Race Conditions
- ❖ **Deadlocks**
- ❖ Graphs!!

The Problem (1 of 2)

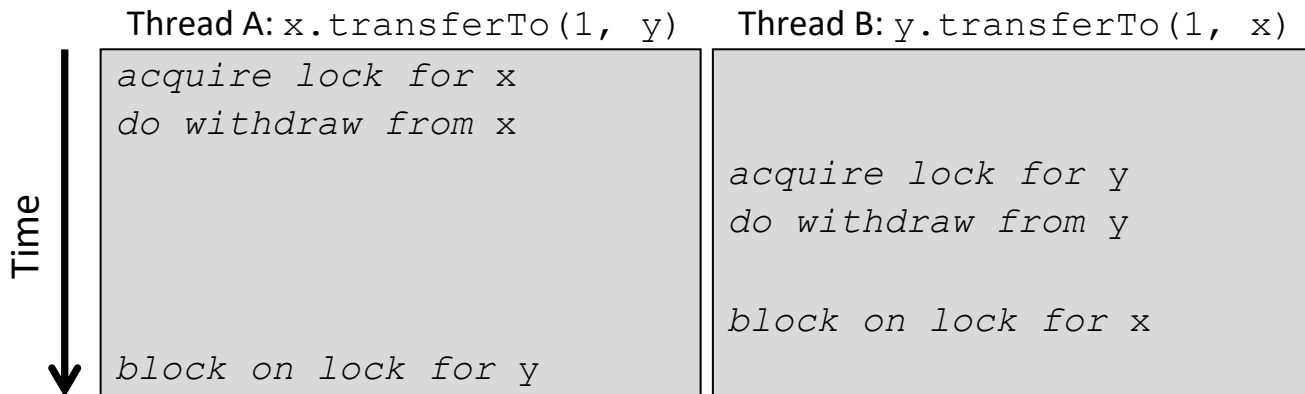
- ❖ Consider a method to transfer money between bank accounts

```
class BankAccount {  
    ...  
    synchronized public void withdraw(int amt) {...}  
    synchronized public void deposit(int amt) {...}  
    synchronized public void transferTo(int amt,  
                                         BankAccount a) {  
        this.withdraw(amt);  
        a.deposit(amt);  
    }  
}
```

- ❖ Potential problems?
 - During call to a.deposit(), thread holds two locks
 - Need to investigate whether (when?) this may be a problem


The Problem (2 of 2)

- ❖ Suppose x and y are different accounts



The Dining Philosophers

- ❖ Classic formulation of a computer science problem!
 - 5 philosophers go to dinner at an Italian restaurant
 - They sit at a round table with *one fork per setting*
 - When the spaghetti arrives, each philosopher first attempts to grab their right fork, then their left fork
 - If they successfully grab two forks, they can eat
- ❖ 'Locking' for fork results in a **deadlock**



Deadlock

- ❖ A **deadlock** occurs when there are threads T_1, \dots, T_n such that:
 - For $i=1, \dots, n-1$, T_i is waiting for a resource held by $T_{(i+1)}$
 - T_n is waiting for a resource held by T_1

- ❖ In other words, there is a *cycle of waiting*
 - If we model the waiting as a graph of dependencies, cycles are bad!
 - Deadlock avoidance is basically ensuring a cycle can never arise

Back to Our BankAccount Example: Transfers

```
1  class BankAccount {
2      ...
3      private int acctNumber; // must be unique
4      public void transferTo(int amt, BankAccount a) {
5          synchronized(this) {
6              synchronized(a) {
7                  this.withdraw(amt);
8                  a.deposit(amt);
9              }
10         }
11     }
```

- ❖ Deadlock occurs when two threads attempt to transfer to each other. How might this occur? Refer to the line numbers.

How do we fix Deadlock?

❖ Options for deadlock-proof transfer:

1. Make a smaller critical section: “unsynchronize” transferTo()
 - Exposes intermediate state after withdraw and before deposit
 - Might be okay here, but bank will have wrong total amount (transiently)
2. Coarsen lock granularity: one lock for all accounts
 - Allows transfers, but sacrifices concurrent deposits/withdrawals
3. Assign each account an ordering; consistently acquire locks in order
 - If entire program obeys this ordering, can avoid cycles
 - Code that acquires only one lock can ignore the order

Consistently Ordering Locks

```
class BankAccount {
    ...
    private int acctNumber; // must be unique
    public void transferTo(int amt, BankAccount a) {
        BankAccount first = a;
        BankAccount second = this;
        if (this.acctNumber < a.acctNumber){
            first = this;
            second = a;
        }
        synchronized(first) {
            synchronized(second) {
                this.withdraw(amt);
                a.deposit(amt);
            }
        }
    }
}
```

Aside: Another Example

- ❖ The Java standard library's `StringBuffer`

```
class StringBuffer {
    private int count;
    private char[] value;
    ...
    synchronized append(StringBuffer sb) {
        int len = sb.length();
        if(this.count + len > this.value.length)
            this.expand(...);
        sb.getChars(0, len, this.value, this.count);
    }

    synchronized getChars(int x, int, y,
                           char[] a, int z) {
        "copy this.value[x..y] into a starting at z"
    }
}
```

Aside: Two problems with `StringBuffer`

- ❖ **Problem #1:** `sb`'s lock not held between `sb.length` and `sb.getChars`
 - `sb` could get longer, causing `append()` to throw `ArrayBoundsException`
- ❖ **Problem #2:** Deadlock potential if two threads try to append in opposite directions, just like in the bank-account first example
- ❖ Not easy to fix both problems without extra copying 😞
 - Do not want unique ids on every `StringBuffer`
 - Do not want one lock for all `StringBuffer` objects
- ❖ Actual Java library: fixed neither
 - Left code as is and changed javadoc
 - Up to clients to avoid such situations with own protocols

Summary: Deadlocks

- ❖ Code that modifies multiple objects, like account-transfer *and* *string-buffer append*, may introduce deadlock
- ❖ **Easier case:** objects have different (logical) types
 - Define a fixed order among types
 - E.g.: “When moving an item from the hashtable to the work queue, never acquire the queue lock while holding the hashtable lock”
- ❖ **Easier case:** objects are in an acyclic structure
 - Use the structure to determine a fixed order
 - E.g.: “If holding a tree node’s lock, do not acquire other nodes’ locks unless they are children in the tree”
- ❖ Many of these techniques depend on developer discipline and documentation 😞

Summary: Concurrency (1 of 2)

- ❖ **Concurrent programming** allows multiple threads to access shared resources, possibly increasing throughput
 - e.g. hash table, work queue
- ❖ It also introduces new sources of 🐛 bugs 🐛:
 - **Race conditions**: **data races** and **bad interleavings**
 - Critical sections too small or use wrong locks
 - Deadlocks

Summary: Concurrency (2 of 2)

- ❖ Concurrency requires *synchronization*
 - *Locks*, to ensure for *mutual exclusion*
 - Other synchronization primitives (see Grossman notes):
 - Reader/Writer Locks
 - Condition variables for signaling others

- ❖ Guidelines for correct use can help avoid common pitfalls

- ❖ Shared memory model is not the only approach, but other approaches (e.g., message passing, streams) are not painless

Lecture Outline

- ❖ Five Guidelines to Avoid Race Conditions
- ❖ Deadlocks
- ❖ **Graphs!!**

Graphs

CSE 332 Summer 2021

Instructor: Kristofer Wong

Teaching Assistants:

Alena Dickmann	Arya GJ	Finn Johnson
Joon Chong	Kimi Locke	Peyton Rapo
Rahul Misal	Winston Jodjana	

Graphs

- ❖ A **graph** represents relationships among items
 - Very general definition because very general concept

- ❖ A **graph** is a pair: $G = (V, E)$


- A set of **vertices**, also known as **nodes**

$$V = \{v_1, v_2, \dots, v_n\}$$

- A set of **edges**, possibly **directed**

$$E = \{e_1, e_2, \dots, e_m\}$$

- Each edge e_i is a pair of vertices (v_j, v_k)
- An edge “connects” the vertices



$$V = \{\text{Han}, \text{Leia}, \text{Luke}\}$$

$$E = \{(\text{Luke}, \text{Leia}), (\text{Han}, \text{Leia}), (\text{Leia}, \text{Han})\}$$

- ❖ We've seen these before: DAGs

Is a Graph an ADT?

- ❖ Kind of, but not in the way we're used to
 - They have operations like `hasEdge ((vj, vk))`
 - But it is unclear what the “standard operations” are

- ❖ Instead:
 - We tend to develop algorithms over graphs and then use whatever data structure is efficient for that algorithms

- ❖ Many important problems can be solved by:
 1. Formulating them in terms of graphs
 2. Applying a standard graph algorithm


Some Graph Examples

- ❖ For each of the following, what are the vertices and the edges?
 - Web pages with links
 - Facebook friends
 - Methods in a program that call each other
 - Road maps (e.g., Google maps)
 - Airline routes
 - Family trees
 - Course pre-requisites

- ❖ **Wow!** Using the same algorithms for problems across so many domain... Almost as if... It's a core concept...?

Undirected Graphs


- ❖ In *undirected graphs*, edges have no specific direction
 - Edges are always “two-way”




- ❖ Thus, $(u, v) \in E$ implies $(v, u) \in E$
 - Only one of these edges needs to be in the set; the other is implicit
- ❖ *Degree* of a vertex: number of edges containing that vertex
 - i.e.: the number of adjacent vertices

Directed Graphs

- ❖ In *directed graphs* (aka *digraphs*), edges have a *direction*



or




- ❖ Thus, $(u, v) \in E$ does not imply $(v, u) \in E$
 - $(u, v) \in E$ means $u \rightarrow v$; u is the *source* and v the *destination*
- ❖ *In-Degree* of a vertex: number of in-bound edges
 - i.e.: edges where the vertex is the destination
- ❖ *Out-Degree* of a vertex: number of out-bound edges
 - i.e.: edges where the vertex is the source

Self-edges

- ❖ A *self-edge* (aka a *loop*) is an edge of the form (u, u)
- ❖ Depending on the use/algorithm, a graph may have:
 - No self edges
 - Some self edges
 - All self edges (therefore often implicit, but we will be explicit)
- ❖ A node can have a degree / in-degree / out-degree of zero

Adjacency (1 of 2)

- ❖ If $(u, v) \in E$
 - Then v is a *neighbor* of u , i.e., v is *adjacent* to u
 - For directed edges, order matters
 - u is not adjacent to v unless $(v, u) \in E$



$V = \{A, B, C, D\}$

$E = \{(C, B), (A, B), (B, A), (C, D)\}$

Adjacency (2 of 2)

❖ For a graph $G = (V, E)$:

- $|V|$ is the number of vertices
- $|E|$ is the number of edges

- Minimum size?

- 0


- Maximum size for an undirected graph with no self-edges?

- $|V|(|V-1|)/2 \in O(|V|^2)$

- Maximum for a directed graph with no self-edges?

- $|V|(|V-1|) \in O(|V|^2)$

- If self-edges are allowed, add $|V|$ to the answers above (applies to both undirected and directed graphs)




Graph Examples, Again

- ❖ For each of the following, which would use *directed edges*? Which would have *self-edges*? Which might have *0-degree nodes*?
- (A) Web pages with links
 - (B) Facebook friends
 - (C) Methods in a program that call each other
 - (D) Road maps (e.g., Google maps)
 - (E) Airline routes
 - (F) Family trees
 - (G) Course pre-requisites

Weighted Graphs

- ❖ In a weighed graph, each edge has a **weight** a.k.a. **cost**
 - Typically numeric (most examples will use ints)
 - Some graphs allow *negative weights*; many don't




Graph Examples, Once More Unto the Breach

- ❖ Do *weights* make sense for each of the following graphs? What would they represent, and could those weights be *negative*?
 - Web pages with links
 - Facebook friends
 - Methods in a program that call each other
 - Road maps (e.g., Google maps)
 - Airline routes
 - Family trees
 - Course pre-requisites

Paths and Cycles (1 of 2)

- ❖ A **path** is a list of vertices $[v_0, v_1, \dots, v_n]$ such that $(v_i, v_{i+1}) \in E$ for all $0 \leq i < n$
 - You'd call it a path from v_0 to v_n



- ❖ A **cycle** is a path that begins and ends at the same node
 - i.e., $v_0 == v_n$



- ❖ Example path:
 - [Seattle, SLC, Chicago, Dallas, SF, Seattle]
 - Also happens to be a cycle!


Paths and Cycles (2 of 2)

- ❖ A graph that does not contain any cycles is *acyclic*



Path Length and Cost


- ❖ **Path length:** Number of edges in a path
 - Also called “unweighted cost”
- ❖ **Path cost:** Sum of the weights of each edge in a path
- ❖ Example: $P = [\text{Seattle, SLC, Chicago, Dallas, SF}]$



length(P) = 4
cost(P) = 9.5


❖ Is there a path from A to D? Does the graph contain any cycles?

- A. Yes / Yes
- B. Yes / No
- C. No / Yes
- D. No / No
- E. I'm not sure ...




Connectivity: Undirected Graphs

- ❖ An undirected graph is **connected** if for all pairs of vertices u, v , there exists a *path* from u to v




Connected graph



Disconnected graph


- ❖ An undirected graph is **complete** (aka **fully connected**) if for all pairs of vertices u, v , there exists an *edge* from u to v




(not pictured: self edges)

Connectivity: Directed Graphs


- ❖ A directed graph is **strongly connected** if for all pairs of vertices u, v , there exists a *path* from u to v



- ❖ A directed graph is **weakly connected** if for all pairs of vertices u, v , there exists a path from u to v *ignoring direction of edges*



- ❖ A directed graph is **complete** (aka **fully connected**) if for all pairs of vertices u, v , there exists an *edge* from u to v



(not pictured: self edges)


Much Example. Very Graph. Wow.

- ❖ For undirected graphs: *connected?*
- ❖ For directed graphs: *strongly connected?* *weakly connected?*
 - (A) Web pages with links
 - (B) Facebook friends
 - (C) Methods in a program that call each other
 - (D) Road maps (e.g., Google maps)
 - (E) Airline routes
 - (F) Family trees
 - (G) Course pre-requisites

Trees as Graphs

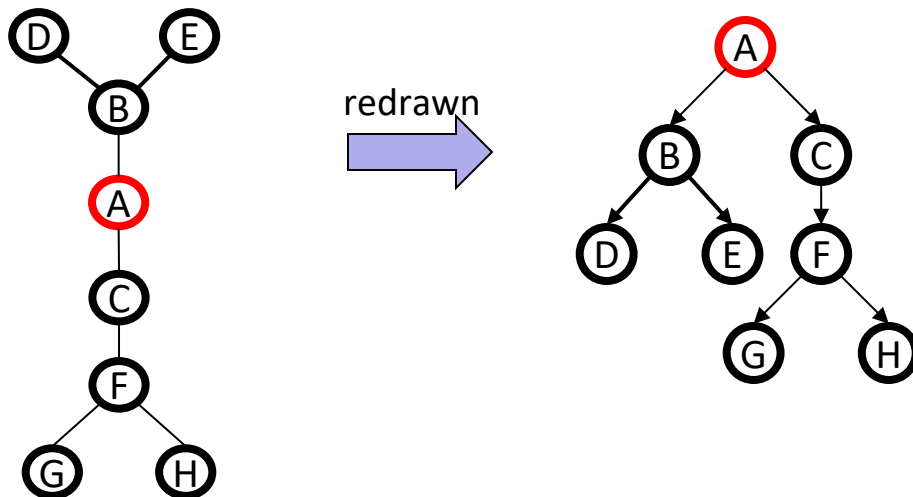
- ❖ A **tree** is a graph that is:
 - undirected
 - acyclic
 - connected
- ❖ So all trees are graphs, but not all graphs are trees
- ❖ How does this relate to the trees we know and love?...

Example:




Rooted Trees (1 of 2)

- ❖ We've previously worked with *rooted trees*, where:
 - We identify a unique ("special") vertex: the root
 - We think of edges as **directed**: parent to children
- ❖ The same tree can be redrawn as multiple rooted trees depending on which node you pick as the root




Rooted Trees (2 of 2)

- ❖ We've previously worked with *rooted trees*, where:
 - We identify a unique ("special") vertex: the root
 - We think of edges as **directed**: parent to children
- ❖ The same tree can be redrawn as multiple rooted trees depending on which *node you pick as the root*




Directed Acyclic Graphs (aka DAGs)

- ❖ A **DAG** is a directed graph with no directed cycles
- ❖ Every rooted directed tree is a DAG
 - But not every DAG is a rooted directed tree:



← Not a rooted directed tree;
has an undirected cycle

- ❖ Every DAG is a directed graph (**D**AG: It's in the name)
 - But not every directed graph is a DAG:



← Not a DAG; has a
directed cycle

Graph Examples: One more time

- ❖ Which of our *directed*-graph examples do you expect to be a **DAG**?
 - Web pages with links
 - Facebook friends
 - Methods in a program that call each other
 - Road maps (e.g., Google maps)
 - Airline routes
 - Family trees
 - Course pre-requisites

Density / Sparsity (1 of 2)

❖ Recall:

- In an undirected graph, $0 \leq |E| < |V|^2$
- In a directed graph: $0 \leq |E| \leq |V|^2$

So for any graph,
 $|E| \in O(|V|^2)$

❖ One more fact:

- In a *connected* undirected graph, $|E| \geq |V|-1$
- In a *weakly connected* directed graph, $|E| \geq |V|-1$
- In a *strongly connected* directed graph, $|E| \geq |V|$

So for any
connected graph,
 $|E| \in \Omega(|V|^2)$

Density / Sparsity (2 of 2)

- ❖ We do not always approximate as $|E|$ as $O(|V|^2)$
 - This is a *correct* bound, it's just oftentimes not *tight*
- ❖ If it is tight, i.e. $|E| \in \Theta(|V|^2)$, we say the graph is **dense**
 - Intuitively: “lots of edges”
- ❖ If $|E| \in O(|V|)$ we say the graph is **sparse**
 - Sparse: “most (of the possible) edges missing”

