COMP 322: Fundamentals of Parallel Programming

Lecture 19: Critical Sections and the Isolated Construct

Mack Joyner
mjoyner@rice.edu

http://comp322.rice.edu

Worksheet #18: Cooperative vs Blocking Runtime scheduler

Assume that creating an async causes the task to be pushed into the work queue for execution by any available idle thread.

Fill the following table for the program shown on the right by adding the appropriate number of threads required to execute the program. For the minimum or maximum numbers, your answer must represent a schedule where at some point during the execution all threads are busy executing a task or blocked on some synchronization constraint.

	Minimum number of threads	Maximum number of threads
Cooperative Runtime	1	?
Blocking Runtime	?	?

```
10. finish {
11.   async { S1; }
12.   finish {
13.     async {
14.       finish {
15.         async { S2; }
16.         S3;
17.       }
18.       S4;
19.     }
20.     async {
21.       async { S5; }
22.       S6;
23.     }
24.     S7;
25.   }
26.   S8;
27. }
```

Worksheet #18: Cooperative vs Blocking Runtime scheduler

```
10. finish {
11.   async { S1; }
12.   finish {
13.     async {
14.       finish {
15.         async { S2; }
16.         S3;
17.       }
18.       S4;
19.     }
20.     async {
21.       async { S5; }
22.       S6;
23.     }
24.     S7;
25.   }
26.   S8;
27. }
```

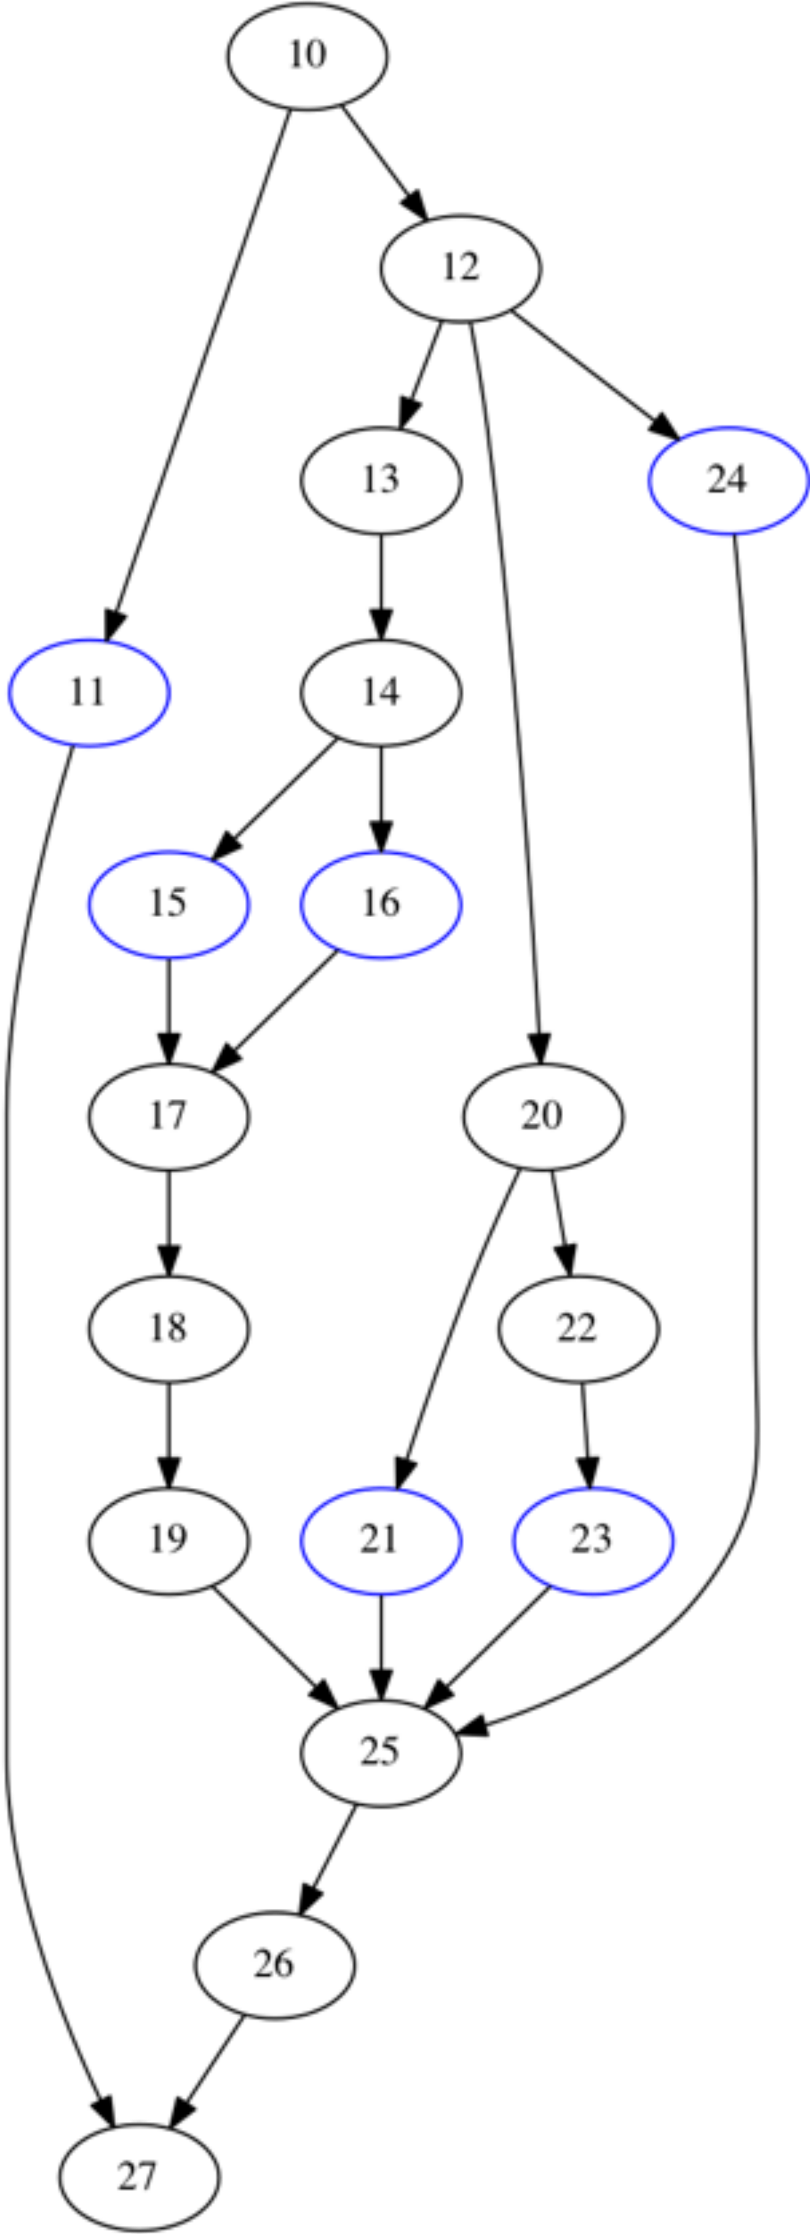

	Maximum number of threads
Cooperative Runtime	6
Blocking Runtime	6

	Minimum number of threads
Cooperative Runtime	1
Blocking Runtime	?

Maximum threads: If we proceed through the graph in top-down manner incrementally, how many maximum leaf nodes can we have?

Worksheet #18: Cooperative vs Blocking Runtime scheduler

```
10. finish {
11.   async { S1; }
12.   finish {
13.     async {
14.       finish {
15.         async { S2; }
16.         S3;
17.       }
18.       S4;
19.     }
20.     async {
21.       async { S5; }
22.       S6;
23.     }
24.     S7;
25.   }
26.   S8;
27. }
```

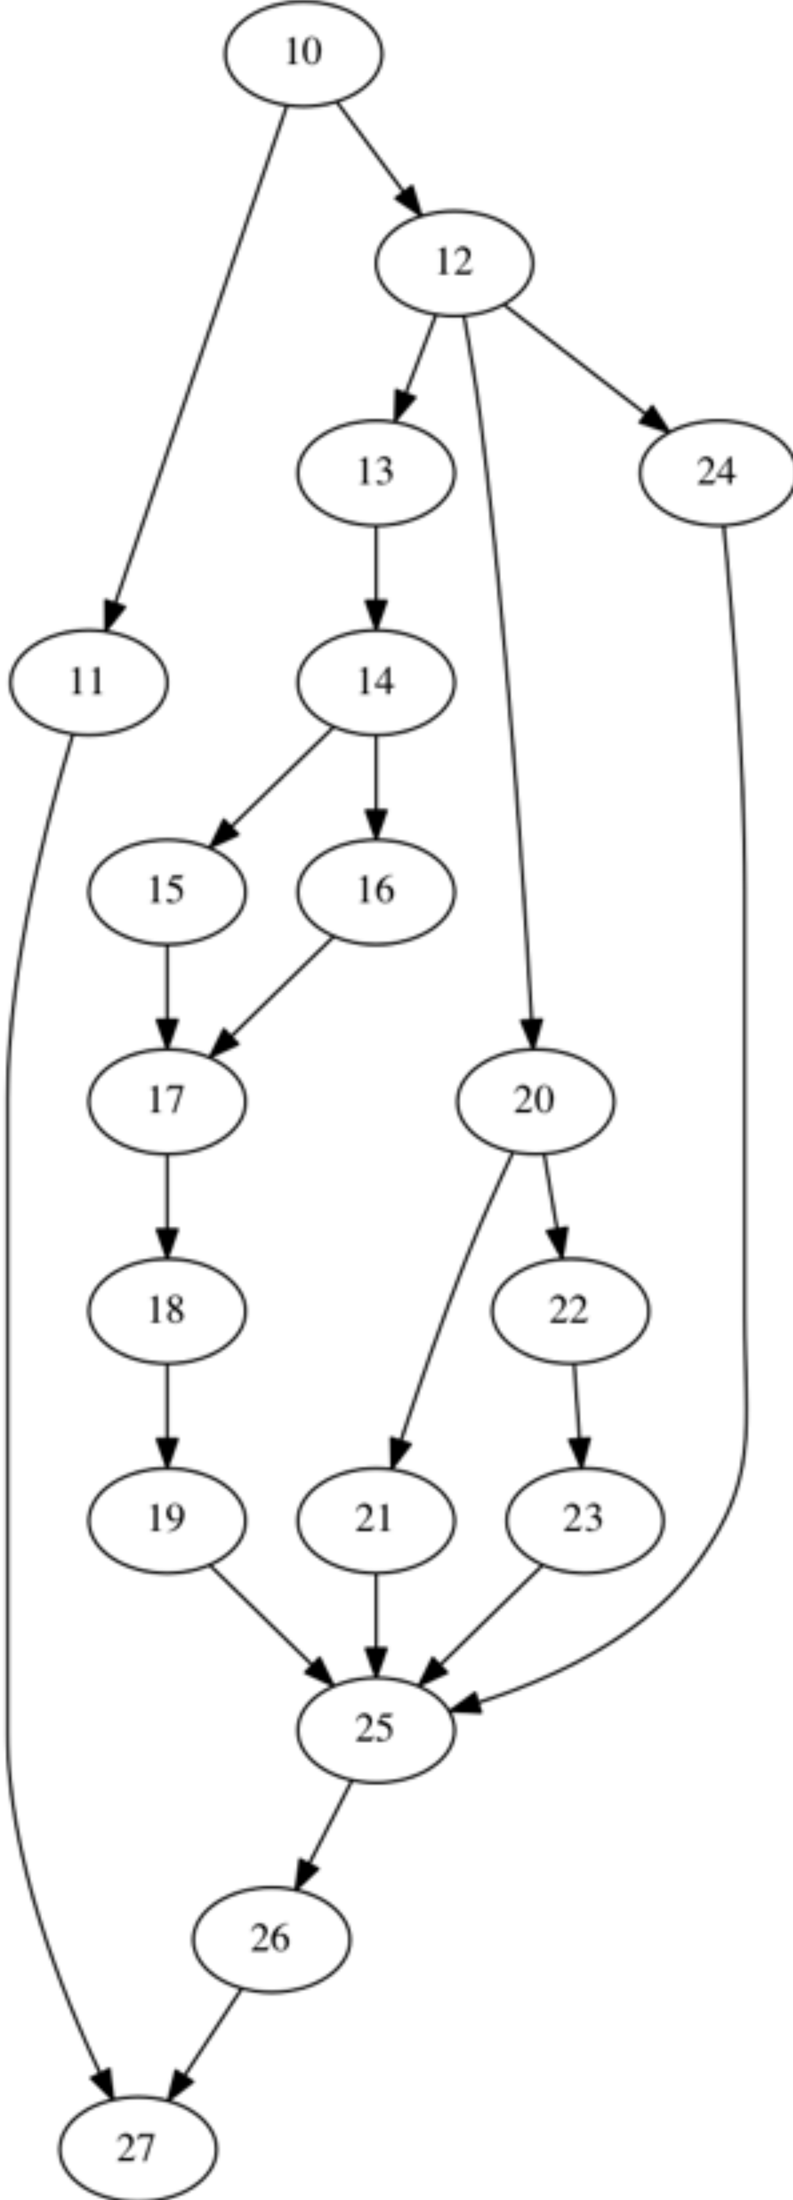

ready queue
11
13
20
15
21

	Minimum number of threads
Blocking Runtime	3

Formal Definition of Data Races (Recap)

Formally, a data race occurs on location L in a program execution with computation graph CG if there exist steps (nodes) S1 and S2 in CG such that:

1. S1 does not depend on S2 and S2 does not depend on S1 i.e., there is no path of dependence edges from S1 to S2 or from S2 to S1 in CG, and
2. Both S1 and S2 read or write L, and at least one of the accesses is a write.

However, there are many cases in practice when two tasks may legitimately need to perform conflicting accesses to shared locations without incurring data races

– How should conflicting accesses be handled in general, when outcome may be nondeterministic?

⇒ Focus of Module 2: “Concurrency” (nondeterministic parallelism)

Example of two tasks performing conflicting accesses --- need for "mutual exclusion"

```
1.  class DoublyLinkedListNode {
2.   DoublyLinkedListNode prev, next;
3.   . . .
4.   void delete() {
5.     { // start of desired mutual exclusion region
6.       this.prev.next = this.next;
7.       this.next.prev = this.prev;
8.     } // end of desired mutual exclusion region
9.     . . . // remaining code in delete() that does not need mutual exclusion
10.   }
11. } // DoublyLinkedListNode
12. . . .
13. static void deleteTwoNodes(final DoublyLinkedListNode L) {
14.  finish(() -> {
15.    DoublyLinkedListNode second = L.next;
16.    DoublyLinkedListNode third = second.next;
17.    async(() -> { second.delete(); });
18.    async(() -> { third.delete(); }); // conflicts with previous async
19.  });
20. }
```

How to enforce mutual exclusion?

- The predominant approach to ensure mutual exclusion proposed many years ago is to enclose the code region in a critical section.
  - "In concurrent programming a critical section is a piece of code that accesses a shared resource (data structure or device) that must not be concurrently accessed by more than one thread of execution. A critical section will usually terminate in fixed time, and a thread, task or process will have to wait a fixed time to enter it (aka bounded waiting). Some synchronization mechanism is required at the entry and exit of the critical section to ensure exclusive use, for example a semaphore."
  - Source: http://en.wikipedia.org/wiki/Critical_section

HJ isolated construct

```
isolated (() -> <body> );
```

- Isolated construct identifies a critical section
- Two tasks executing isolated constructs are guaranteed to perform them in mutual exclusion
  - ➔Isolation guarantee applies to (isolated, isolated) pairs of constructs, not to (isolated, non-isolated) pairs of constructs
- Isolated constructs may be nested
  - — An inner isolated construct is redundant
- Blocking parallel constructs are forbidden inside isolated constructs
  - —Isolated constructs must not contain any parallel construct that performs a blocking operation e.g., finish, future get, next
  - —Non-blocking async operations are permitted, but isolation guarantee only applies to creation of async, not to its execution
- Isolated constructs can never cause a deadlock
  - — Other techniques used to enforce mutual exclusion (e.g., locks — which we will learn later) can lead to a deadlock, if used incorrectly

Use of isolated to fix previous example with conflicting accesses

```
1.  class DoublyLinkedListNode {
2.   DoublyLinkedListNode prev, next;
3.   . . .
4.   void delete() {
5.     isolated(() -> { // start of desired mutual exclusion region
6.       this.prev.next = this.next;
7.       this.next.prev = this.prev;
8.     }); // end of desired mutual exclusion region
9.     . . . // other code in delete() that does not need mutual exclusion
10.  }
11. } // DoublyLinkedListNode
12. . . .
13. static void deleteTwoNodes(final DoublyLinkedListNode L) {
14.  finish(() -> {
15.    DoublyLinkedListNode second = L.next;
16.    DoublyLinkedListNode third = second.next;
17.    async(() -> { second.delete(); });
18.    async(() -> { third.delete(); }); // conflicts with previous async
19.  });
20. }
```

Exercise

Compute the WORK and CPL metrics for this program with a global isolated construct.

```
1.    finish(() -> {
2.        for (int i = 0; i < 5; i++) {
3.          async(() -> {
4.            doWork(2);
5.            isolated(() -> { doWork(1); });
6.            doWork(2);
7.          }); // async
8.        } // for
9.     }); // finish
```

Serialized Computation Graph for Isolated Constructs

- Model each instance of an isolated construct as a distinct step (node) in the CG.
- Need to reason about the *order* in which interfering isolated constructs are executed
  – Complicated because the order of isolated constructs may vary from execution to execution
- Introduce Serialized Computation Graph (SCG) that includes a specific ordering of all interfering isolated constructs.
  – SCG consists of a CG with additional serialization edges.
  – Each time an isolated step, S', is executed, we add a serialization edge from S to S' for each prior "interfering" isolated step, S
    - Two isolated constructs always interfere with each other
    - Interference of "object-based isolated" constructs depends on intersection of object sets
    - Serialization edge is not needed if S and S' are already ordered in CG
  – An SCG represents a set of schedules in which all interfering isolated constructs execute in the same order.

Example of Serialized Computation Graph with Serialization Edges for v10-v16-v11 order

Data race definition can be applied to Serialized Computation Graphs (SCGs) just like regular CGs

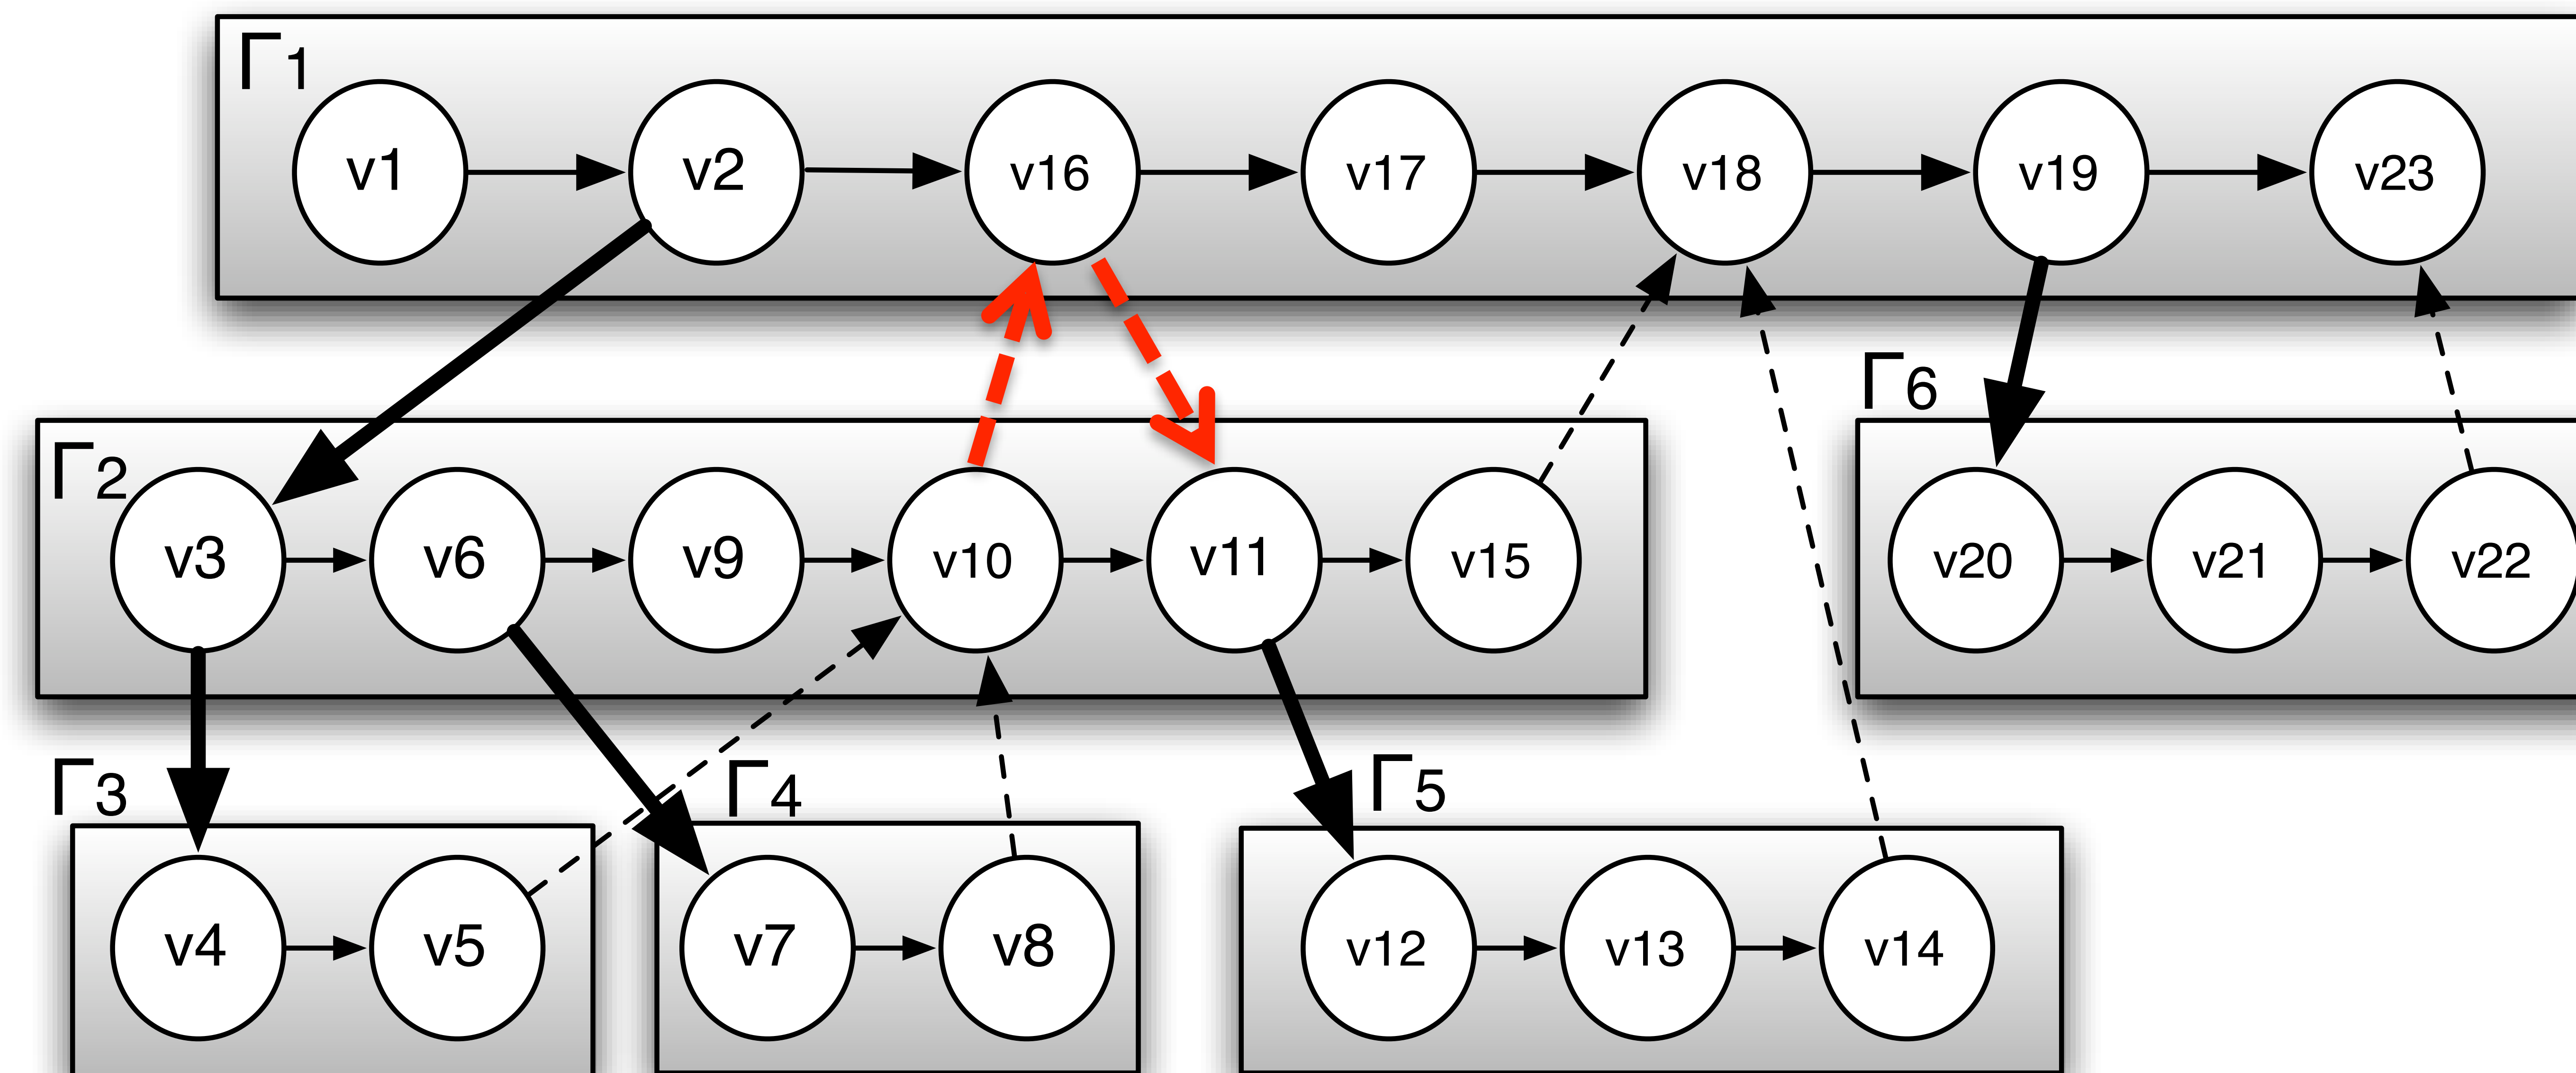

v10: isolated { x ++; y = 10; }
v11: isolated { x++; y = 11; }
v16: isolated { x++; y = 16; }

Need to consider all possible orderings of interfering isolated constructs to establish data race freedom

Object-based isolation

```
isolated(obj1, obj2, …, () -> <body>)
```

- In this case, programmer specifies list of objects for which isolation is required
- Mutual exclusion is only guaranteed for instances of isolated constructs that have a common object in their object lists
  - Serialization edges are only added between isolated steps with at least one common object (non-empty intersection of object lists)
  - Standard isolated is equivalent to "isolated(*)" by default i.e., isolation across all objects
- Inner isolated constructs are redundant — they are not allowed to "add" new objects

DoublyLinkedListNode Example revisited with Object-Based Isolation

```
1.  class DoublyLinkedListNode {
2.   DoublyLinkedListNode prev, next;
3.   . . .
4.   void delete() {
5.     isolated(?, ?,…, () -> { // object-based isolation
6.       this.prev.next = this.next;
7.       this.next.prev = this.prev;
8.     });
9.     . . .
10.  }
11. } // DoublyLinkedListNode
12. . . .
13. static void deleteTwoNodes(final DoublyLinkedListNode L) {
14.  finish(() -> {
15.    DoublyLinkedListNode second = L.next;
16.    DoublyLinkedListNode third = second.next;
17.    async(() -> { second.delete(); });
18.    async(() -> { third.delete(); });
19.  });
20. }
```

Pros and Cons of Object-Based Isolation

- Pros
  - Increases parallelism relative to critical section approach
  - Simpler approach than “locks” (which we will learn later)
  - Deadlock-freedom property is still guaranteed

- Cons
  - Programmer needs to worry about getting the object list right
  - Objects in object list can only be specified at start of the isolated construct
  - Large object lists can contribute to large overheads

Announcements & Reminders

- HW3 CP1 is now due Saturday, Feb 29th at 11:59pm (24-hour extension)
- Midterm exams should be fully graded by Friday, March 6th
- Reserve the right to curve final grades up!