COMP 322: Fundamentals of Parallel Programming

Lecture 20: Critical Sections and the Isolated Construct

Vivek Sarkar, Shams Imam
Department of Computer Science, Rice University

Contact email: vsarkar@rice.edu, shams.imam@twosigma.com

http://comp322.rice.edu/

Summary of Module 1: Parallelism

- **Serial-elision subset of HJ**

  - { async, finish, future, forall, forasync }

  - <u>Serial elision property</u>: any HJ program written using the above constructs can be converted to an equivalent sequential program by "eliding" all parallel constructs i.e., by removing async & finish, and replacing future & forasync by equivalent sequential constructs

- **Deadlock-free subset of HJ**

  - { barriers, phasers, forallPhased, asyncPhased } + Serial-elision subset

  - <u>Deadlock-freedom property</u>: if any HJ program written using the above constructs is guaranteed to be data-race-free for a given input, then the execution of that program is guaranteed to be deadlock-free for that input (data-race-freedom is necessary to ensure absence of deadlock with futures)

- **Deterministic subset of HJ**

  - { data driven futures, async await } + Deadlock-free subset

  - <u>Determinism property</u>: if any HJ program written using the above constructs is guaranteed to be data-race-free for a given input, then it must also be functionally deterministic and structurally deterministic for that input i.e., all executions with the same input must generate the same output AND the same computation graph

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.  }
```

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
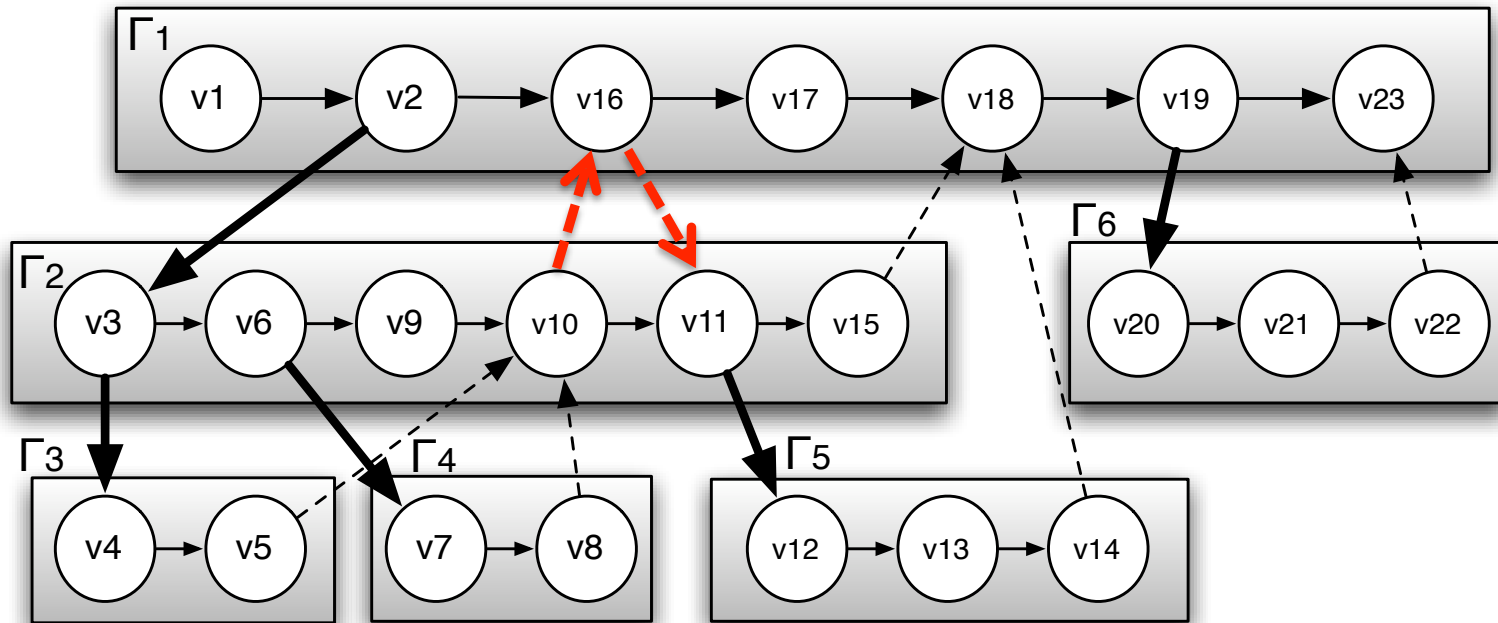
Continue edge   Spawn edge   - - - - ▶ Join edge

- - - ▶ Serialization edge

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 objstec 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

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 (new objects cannot be added later on)

  — Large object lists can contribute to large overheads

DoublyLinkedListNode Example revisited with Object-Based Isolation

```
1.  class DoublyLinkedListNode {
2.   DoublyLinkedListNode prev, next;
3.   . . .
4.   void delete() {
5.     isolated(this.prev, this, this.next, () -> { // 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. }
```

Spanning Tree Definition

- A spanning tree, T, of a connected undirected graph G is

  - rooted at some vertex of G

  - defined by a parent map for each vertex

  - contains all the vertices of G, i.e. spans all vertices

  - contains exactly $|v|$ - 1 edges

    - adding any other edge will create a cycle

  - contains no cycles (a tree!)

    - implies the edges involved in T is a subset of the edges in G

# An Example Graph with 4 possible spanning trees rooted at vertex A

Example Undirected Graph:

Spanning Trees (edges are directed from child to parent):

Vertex	Parent
A	null
B	D
C	A
D	C

Vertex	Parent
A	null
B	A
C	D
D	B

Vertex	Parent
A	null
B	A
C	A
D	B

Vertex	Parent
A	null
B	A
C	A
D	C

Sequential Parallel Spanning Tree Algorithm

```
1.  class V  {
2.    V [] neighbors; // adjacency list for input graph
3.    V parent; // output value of parent in spanning tree

4.    boolean makeParent(V n) {
5.      if (parent == null) { parent = n; return true; }
6.      else return false; // return true if n became parent
7.    } // makeParent

8.    void compute() {
9.      for (int i=0; i<neighbors.length; i++) {
10.       final V child = neighbors[i];
11.       if (child.makeParent(this))
12.         child.compute(); // recursive call
13.     }
14.   } // compute
15. } // class V
16. . . . // main program
17. root.parent = root; // Use self-cycle to identify root
18. root.compute();
19. . . .
```

Announcements

- **Reminder: Checkpoint #2 for Homework 3 is due by this Friday (March 11th), and the entire homework is due by March 18th**

- **The registrar has announced the schedule for the COMP 322 final exam:**

  —**3-MAY-2016**

  —**9:00AM - 12:00PM**

  —**Scheduled Final Exam-OTR Room**

- **Scope of final exam (Exam 2) will be limited to Lectures 20 - 37**