
COMP 322: Fundamentals of Parallel Programming

Lecture 14: Data-Driven Tasks and Data-Driven Futures


Vivek Sarkar
Department of Computer Science, Rice University
vsarkar@rice.edu

<https://wiki.rice.edu/confluence/display/PARPROG/COMP322>




Dataflow Computing

- Original idea: replace machine instructions by a small set of dataflow operators



Example instruction sequence and its dataflow graph

```
x = a + b;  
y = b * 7;  
z = (x-y) * (x+y);
```




An operator executes when all its input values are present; copies of the result value are distributed to the destination operators.

No separate control flow



Macro-Dataflow Programming



- “Macro-dataflow” = extension of dataflow model from instruction-level to task-level operations
- General idea: build an arbitrary task graph, but restrict all inter-task communications to single-assignment variables
 - Static dataflow ==> graph fixed when program execution starts
 - Dynamic dataflow ==> graph can grow dynamically
- Semantic guarantees: race-freedom, determinism
 - Deadlocks are possible due to unavailable inputs (but they are deterministic)



Extending HJ Futures for Macro-Dataflow: Data-Driven Futures (DDFs) and Data-Driven Tasks (DDTs)

```
HjDataDrivenFuture<T1> ddfA = newDataDrivenFuture();
```

- Allocate an instance of a data-driven-future object (container)
- Object in container must be of type T1
- HjDataDrivenFuture interface extends the HjFuture interface

```
asyncAwait(ddfA, ddfB, ..., () -> Stmt);
```

- Create a new data-driven-task to start executing **Stmt** after all of **ddfA**, **ddfB**, ... become available (i.e., after task becomes “enabled”)

```
ddfA.put(V) ;
```

- Store object V (of type T1) in **ddfA**, thereby making **ddfA** available
- Single-assignment rule: at most one put is permitted on a given DDF

```
ddfA.get()
```

- Return value (of type T1) stored in **ddfA**
- Throws an exception if put() has not been performed
 - Should be performed by **async**'s that contain **ddfA** in their await clause, or if there's some other synchronization to guarantee that the put() was performed



Implementing Future Tasks using DDFs

- **Future version**

```
1. final HjFuture<T> f = future(() -> { return g(); });
2. S1
3. async(() -> {
4.     ... = f.get();
5.     S2;
6.     S3;
7. });
```


- **DDF version**

```
1. HjDataDrivenFuture<T> f = newDataDrivenFuture();
2. async(() -> { f.put(g()) });
3. S1
4. asyncAwait(f, () -> {
5.     ... = f.get();
6.     S2;
7.     S3;
8. });
```



Use of DDFs with dummy objects (like future<Void>)

```
1. finish(() -> {
2.     HjDataDrivenFuture<Void> ddfA = newDataDrivenFuture();
3.     HjDataDrivenFuture<Void> ddfB = newDataDrivenFuture();
4.     HjDataDrivenFuture<Void> ddfC = newDataDrivenFuture();
5.     HjDataDrivenFuture<Void> ddfD = newDataDrivenFuture();
6.     HjDataDrivenFuture<Void> ddfE = newDataDrivenFuture();
7.     async(() -> { ... ; ddfA.put(null); }); // Task A
8.     asyncAwait(ddfA, () -> { ... ; ddfB.put(null); }); // Task B
9.     asyncAwait(ddfA, () -> { ... ; ddfC.put(null); }); // Task C
10.    asyncAwait(ddfB, ddfC, ()->{ ... ; ddfD.put(null); }); // Task D
11.    asyncAwait(ddfC, () -> { ... ; ddfE.put(null); }); // Task E
12.    asyncAwait(ddfD, ddfE, () -> { ... }); // Task F
13. }); // finish
```



Differences between Futures and DDFs/DDTs

- Consumer task blocks on `get()` for each future that it reads, whereas `async-await` does not start execution till all DDFs are available
- Future tasks cannot deadlock, but it is possible for a DDT to block indefinitely (“deadlock”) if one of its input DDFs never becomes available
- DDTs and DDFs are more general than futures
 - Producer task can only write to a single future object, where as a DDT can write to multiple DDF objects
 - The choice of which future object to write to is tied to a future task at creation time, where as the choice of output DDF can be deferred to any point with a DDT
- DDTs and DDFs can be more implemented more efficiently than futures
 - An “`asyncAwait`” statement does not block the worker, unlike a `future.get()`
 - You will never see the following message with “`asyncAwait`”
 - “ERROR: Maximum number of hj threads per place reached”



Two Exception (error) cases for DDFs that do not occur in futures

- **Case 1:** If two put's are attempted on the same DDF, an exception is thrown because of the violation of the single-assignment rule
 - There can be at most one value provided for a future object (since it comes from the producer task's return statement)
- **Case 2:** If a get is attempted by a task on a DDF that was not in the task's await list, then an exception is thrown because DDF's do not support blocking gets
 - Futures support blocking gets



Deadlock example with DDTs

```
1. HjDataDrivenFuture left = newDataDrivenFuture();
2. HjDataDrivenFuture right = newDataDrivenFuture();
3. finish(() -> {
4.     asyncAwait(left, () -> {
5.         right.put(rightWriter()); });
6.     asyncAwait(right, () -> {
7.         left.put(leftWriter()); });
8. });
```

- HJ-Lib has deadlock detection mode

- Enabled using:

- `System.setProperty(HjSystemProperty.trackDeadlocks.propertyKey(), "true");`
- Reports an `edu.rice.hj.runtime.util.DeadlockException` when deadlock detected




“Adatum Dashboard” Example: Sequential Version

```

1. public MarketRecommendation DoAnalysisSequential() {
2.   StockDataCollection nyseData = LoadNyseData();
3.   StockDataCollection nasdaqData = LoadNasdaqData();
4.   StockDataCollection mergedMarketData =
5.     MergeMarketData(new[] {nyseData, nasdaqData});
6.   StockDataCollection normalizedMarketData =
7.     NormalizeData(mergedMarketData);
8.   StockDataCollection fedHistoricalData =
9.     LoadFedHistoricalData();
10.  StockDataCollection normalizedHistoricalData =
11.    NormalizeData(fedHistoricalData);
12.  StockAnalysisCollection analyzedStockData =
13.    AnalyzeData(normalizedMarketData);
14.  MarketModel modeledMarketData = RunModel(analyzedStockData);
15.  StockAnalysisCollection analyzedHistoricalData =
16.    AnalyzeData(normalizedHistoricalData);
17.  MarketModel modeledHistoricalData = RunModel(analyzedHistoricalData);
18.  MarketRecommendation recommendation =
19.    CompareModels(new[] {modeledMarketData, modeledHistoricalData});
20.  return recommendation;
21. }

```



Source: <http://programming4.us/enterprise/3004.aspx>



“Adatum Dashboard” Example (Pseudocode): Parallel Version using DDTs and DDFs

```
1. public MarketRecommendation DoAnalysisParallelDDT() {
2.   async nyseData.put(LoadNyseData());
3.   async nasdaqData.put(LoadNasdaqData());
4.   async await(nyseData, nasdaqData)
5.     mergedMarketData.put(MergeMarketData(new[] {nyseData.get(), nasdaqData.get()}));
6.   async await(mergedMarketData) normalizedMarketData.put(NormalizeData(mergedMarketData.get()));
7.   async fedHistoricalData.put(LoadFedHistoricalData());
8.   async await(fedHistoricalData) normalizedHistoricalData.put(NormalizeData(fedHistoricalData.get()));
9.   async await(normalizedMarketData) analyzedStockData.put(AnalyzeData(normalizedMarketData.get()));
10.  async await(analyzedStockData) modeledMarketData.put(RunModel(analyzedStockData.get()));
11.  async await(normalizedHistoricalData) analyzedHistoricalData.put(AnalyzeData(normalizedHistoricalData.get()));
12.  async await(analyzedHistoricalData) modeledHistoricalData.put(RunModel(analyzedHistoricalData.get()));
13.  MarketRecommendation recommendation =
14.    CompareModels(new[] {modeledMarketData.get(), modeledHistoricalData.get()});
15.  return recommendation;
16. }
```

Note that the put, await, and get clauses follow directly from the data flow structure of the program!



Worksheet 14: Data Driven Futures

Name: _____

Netid: _____

For the example below, will reordering the five `async` statements change the meaning of the program? If so, show two orderings that exhibit different behaviors. If not, explain why not. (You can use the space below this slide for your answer.)

```
1. DataDrivenFuture left = new DataDrivenFuture();
2. DataDrivenFuture right = new DataDrivenFuture();
3. finish {
4.     async await(left) leftReader(left); // Task3
5.     async await(right) rightReader(right); // Task5
6.     async await(left, right)
7.         bothReader(left, right); // Task4
8.     async left.put(leftWriter()); // Task1
9.     async right.put(rightWriter()); // Task2
10. }
```

