

Deadlock Avoidance in Parallel Programs with Futures

Why Parallel Tasks Should Not Wait for Strangers

Tiago Cogumbreiro, Rishi Surendran, Francisco Martins,
Vivek Sarkar, Vasco Vasconcelos, Max Grossman

OOPSLA, Vancouver, 2017

Futures

Fork-join model + Data

Widespread use of futures

1. Asynchronous programming

- Language support (async, await): Python, Javascript, Rust

2. Task parallel programming

- Language support: Java, C#, C++, Kotlin
- Library support: C++ (TBB, Kokkos, Charm++), Java (HJ-Lib, Quasar)

Uses of futures with shared memory

Task-DAG parallelism

- **Data-flow** parallelism
- **Shared** collections of futures (matrices)

Uses of futures with shared memory

Task-DAG parallelism

- **Data-flow** parallelism
- **Shared** collections of futures (matrices)

Problem: Cyclic data-dependencies cause *deadlocks!*

Off-by-one errors cause deadlocks

Shared memory and futures

Intuition:

■ Root-cause of future-deadlocks are data races.

1. Is it true?
2. Why?
3. How can we use this property for verification?

Outline

1. Futures and its deadlocks
2. **Known Joins** \Rightarrow **DF**: Deadlock avoidance with futures & benchmarks
3. **DRF** \Rightarrow **Known Joins**: How DRF enjoys Deadlock-Freedom
4. Conclusion and future work

1. Futures and its deadlocks

Futures: Tasks that "return" values

async: $(\text{unit} \rightarrow T) \rightarrow \text{Future}\langle T \rangle$

- **Control:** Forks a task A
- **Data:** Returns the future value of type $\text{Future}\langle T \rangle$

get: $\text{Future}\langle T \rangle \rightarrow T$

- **Control:** Joins with task A
- **Data:** Returns the value of type T "produced" by task A

Deadlocked example

```
// Task P  
1 shared Future<Integer> x, y;  
2 x = async(() -> y.get()); // Task A  
3 y = async(() -> x.get()); // Task B
```

1. P forks A writes to x
 - A waits for the task in y
2. P forks B writes to y
 - B waits for the task in x

Data-race causes 2 traces

Trace 1 (no deadlock)

```
1 shared Future<Integer> x, y;  
2 x = async(() -> y.get()); // y = null  
3 y = async(() -> x.get());
```

<i>Task P</i>	<i>Task A</i>	<i>Task B</i>
fork A	read y null	read x A
write x A		get A
fork B		
write y B		

Trace 2 (deadlock)

```
1 shared Future<Integer> x, y;  
2 x = async(() -> y.get()); // y = B  
3 y = async(() -> x.get());
```

<i>Task P</i>	<i>Task A</i>	<i>Task B</i>
fork A	read y B	read x A
write x A	get B	get A
fork B		
write y B		

Proving DRF \Rightarrow DF

DRF \Rightarrow \$POLICY \Rightarrow DF

Deadlock-freedom policy valid in all DRF programs

2. Known Joins \Rightarrow DF

Deadlock avoidance with futures & benchmarks

Known-Joins implementation overview

Program start (empty-known set)

async

1. Before: parent copies known-set to child
2. After: parent extends known-set with new task

get

1. Before: membership-test fail ⇒ **POLICY ABORT**
2. After: merge known-set of task

Running example knowledge

Knowledge: {}

```
1 shared Future<Integer> x, y;  
2 x = async(() -> y.get());
```

Knowledge: {A}

```
3 y = async(() -> x.get());
```

Knowledge: {A, B}

Know-Joins in practice

Habanero-Java: A Java 8 parallel programming library

Extends the deadlock-free API subset with futures!

- `isolated`: mutual-exclusion
- `phaser`: barrier and producer-consumer
- `finish`: descendant task termination
- **`future` (with the known-joins policy)**

Evaluation

- 2,300 assignments checked (1 unknown join, deadlocked example)
- 5 benchmarks

<i>Benchmark</i>	<i># of async</i>	<i># of get</i>
Jacobi	15,872	37,696
Smith-Waterman	21,000	4,641
Crypt	16,384	16,384
Strassen	30,811	44,816
Series	999,999	999,999

Evaluation: time overhead


<i>Benchmark</i>	<i>Snapshot-sets</i>
Jacobi	0.99×
Smith-Waterman	0.96×
Crypt	1.04×
Strassen	1.07×
Series	1.06×

Evaluation: memory overhead

<i>Benchmark</i>	<i>Snapshot-sets</i>
Jacobi	1.00×
Smith-Waterman	1.00×
Crypt	1.08×
Strassen	1.28×
Series	2.34×


3. DRF \Rightarrow Known Joins

Computation Graphs



- Nodes: instruction instances
- Edges: happens-before dependencies (async, get, and sequential)
- Node annotations: **known** tasks and local **memory**


Knowledge flows with reachability



v_4 knows g
 v_4 happens-before w_2

 w_2 knows g

Knowledge must contain tasks in memory



In DRF graphs

Main results

Mechanized and proved in Coq

1. Known-Joins \Rightarrow Deadlock Freedom
2. Data-Race Freedom \Rightarrow Known-Joins
3. Know-Joins interpretation as a causality query

```
CoqIDE
File Edit View Navigation Try Tactics Templates Queries Tools Compile Windows Help
SafeJoins.v
1028 exists (eval (J x y) k).
1029 apply safe_cons; auto.
1030 apply can_check_join; auto.
1031 eauto using edge_to_f_edge.
1032 Qed.
1033
1034 Lemma safe_to_trace:
1035   forall l k,
1036   Safe l k ->
1037   Trace l.
1038 Proof.
1039   intros.
1040   induction H. {
1041     auto using trace_nil.
1042   }
1043   destruct o as ([], a, b). {
1044     inversion H; subst; clear H.
1045     apply trace_fork; auto.
1046     unfold not; intros N.
1047     contradiction H4.
1048     destruct N as (e, (Hi, He)).
1049     unfold In.
1050     apply edge_to_f_edge with (k:=k) in Hi; eauto.
1051   }
1052   inversion H; subst.
1053   eauto using f_edge_to_edge, trace_join.
1054 Qed.
1055
1056 End Trace.
1057
1058 Section Examples.
1059
1060 Let check_fork_nil:
1061   forall x y,
1062   x <> y ->
1063   CheckOp nil f l op t := FORK; op src := x; op dst := ...
1 subgoal
o : op
l : trace
k : list (tid * tid)
H : CanCheckOp k o
H0 : Safe l k
IHsafe : Trace l
Trace (o :: l)
Messages Errors
```

Conclusion

- Introduced a theory of futures and shared memory (CG)
- Showed that data-races are the root cause of deadlocks
- Talked about a deadlock avoidance tool (1.06× time-overhead for 1 million tasks)

Future work

- Promises lack runtime-information to derive deadlock detection
- Extend the theoretical framework for nondeterminism

Deadlock Avoidance in Parallel Programs with Futures

Why Parallel Tasks Should Not Wait for Strangers

Tiago Cogumbreiro, Rishi Surendran, Francisco Martins,
Vivek Sarkar, Vasco Vasconcelos, Max Grossman

OOPSLA, Vancouver, 2017