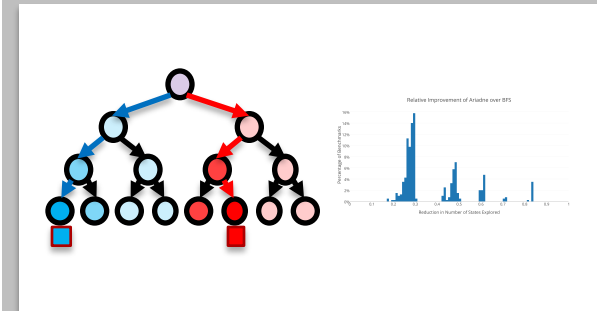




*Exceptional service in the national interest*



# Ariadne: Hybridizing Directed Model

## Checking and Static Analysis

Reed Milewicz and Peter Pirkelbauer

March 2017



# Overview

- Who am I?
- Background; Research Questions
- Related Work
- Introducing Ariadne
- Design and Implementation
- Experimental Results

# Who am I?

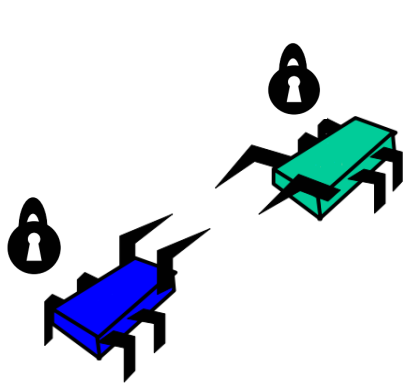
- Postdoctoral researcher at Sandia National Laboratories
- My areas of interest
  - Empirical Software Engineering
  - Static and Dynamic Analysis
  - Source-to-Source Transformation Systems
  - Machine Learning



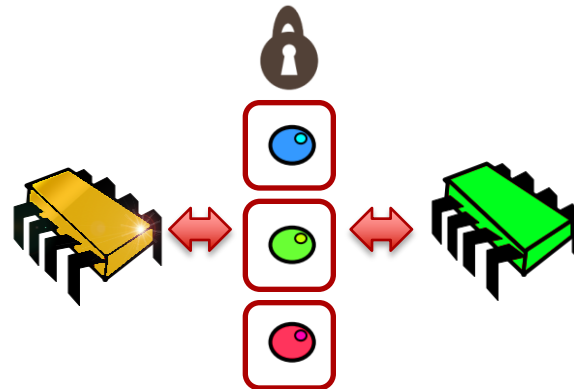
# Background

# Background


- Concurrency bugs are notoriously resistant to testing, highly latent, and can be very dangerous.



Deadlocks



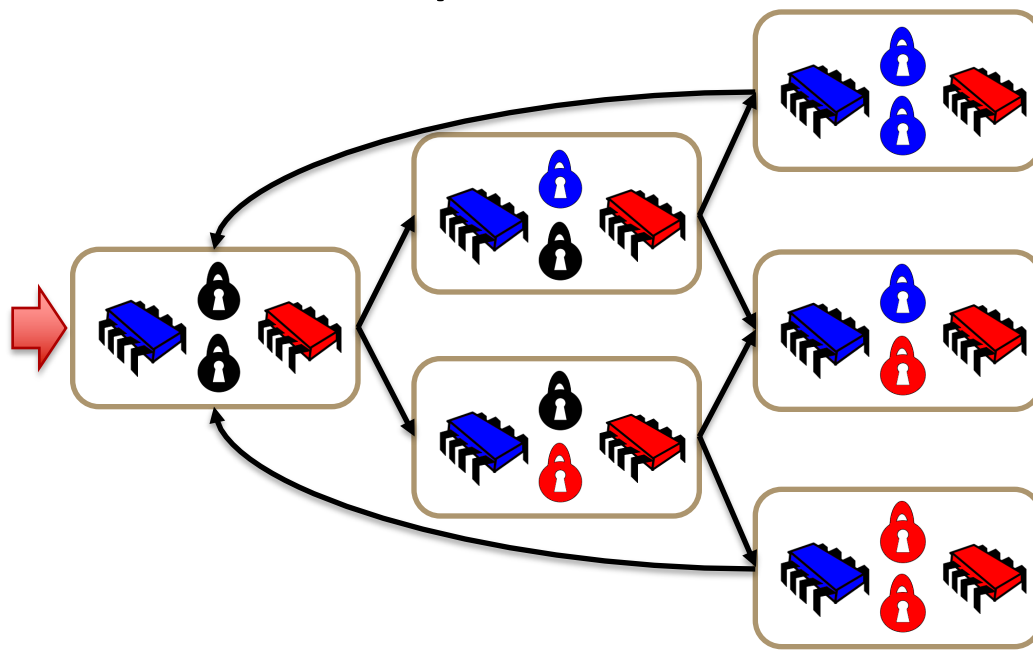
Atomicity Violations



Race Conditions

# Background

- Directed model checking (DMC) is a powerful verification strategy against concurrency bugs.
- Combinatorial explosion is managed through combinations of parallel and heuristic search.



**Two Dining  
Philosophers as  
Kripke Structure**

# Problem Statement

- **Problem:** It does not matter how good a technique is on paper if no one uses it!
- DMC, like most formal verification techniques, is heavy-weight and demands *significant* labor inputs. This is a barrier to adoption. (Engler and Musuvathi 2004)
  - The need for design, selection, and parameterization of heuristics is a barrier to practical use. Research on metaheuristics applicable to this domain has stalled. (Sörensen 2015)
  - Massive parallelism is now inexpensive but that doesn't remove the need for aggressive state space reduction.

# Problem Statement

- Development teams often use multiple tools to test and verify their software. These sources of information can provide insights to benefit the model checking process.
- Model checking as a platform for the synthesis for different analyses.
- Challenges:
  - Need for strategies to manage imperfect and/or conflicting evidence at scale.
  - Need for loose coupling between model checker and external sources.
  - Need for rigorous, empirical methods to validate effectiveness.



- This is not the first work on this subject, nor will it be the last:
  - Refining static analysis through model checking (Post et al. 2008; Darke et al. 2012; Muske and Khedker 2015)
  - Bi-directional collaboration between static analysis and model checking (Beyer et al. 2007; Chen and MacDonald 2008)
  - Analogous work for dynamic analysis and model checking (Groce and Joshi 2008; Milewicz and Pirkelbauer 2016)
- How this work differs:
  - Our emphasis on producing more informed heuristics than state space reduction.
  - Our aim is to elaborate strategies for using model checking as a loosely-coupled platform for the synthesis of analyses.



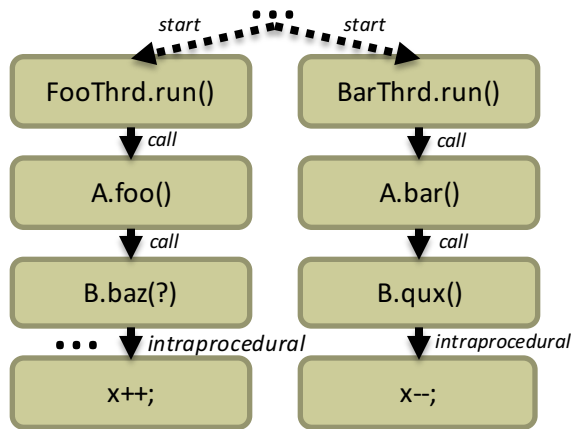
# Ariadne: Design and Implementation

# Introducing Ariadne

- An open-source toolchain and algorithm
  - Translates reports of suspected race conditions of a static analyzer (Petablox)...
  - To dynamic metadata annotations using a source-to-source compiler (ROSE)...
  - Which are exploited by directed model checker (Java Pathfinder).



# Design and Implementation



Call graph,  
Interthread CFG



Abstract Syntax  
Tree



Kripke Structure  
(State Space)

```
1 public static void foo(){B.baz(false)}  
2 public static void bar(){B.qux()}
```

Code 5.1: class A (Client)

```
1 private static int x = 0;  
2 public static synchronized void baz(boolean flag){  
3   if(flag)  
4     x++;  
5   else  
6     B.norf();  
7 }  
8 private static void norf() { B.baz(true) };  
9 public static void qux(){x--;}
```


Code 5.2: class B (Library)

```
1 Possible race condition detected!  
2 Path of abstract thread #1:  
3   A.foo() in Thread.run()  
4   B.baz(false) in A.foo  
5   x++ in B.baz  
6 Path of abstract thread #2:  
7   A.bar() in Thread.run()  
8   B.qux() in A.bar  
9   x-- in B.qux
```

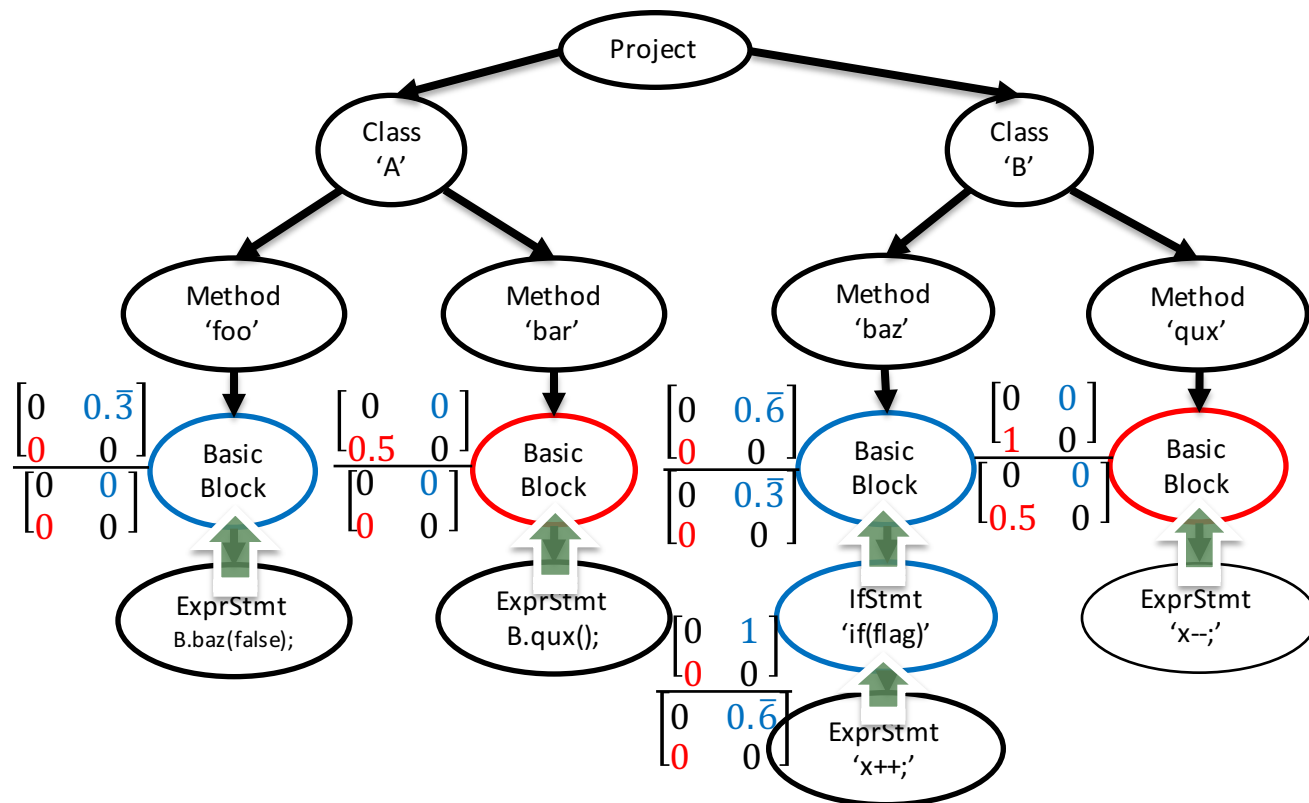
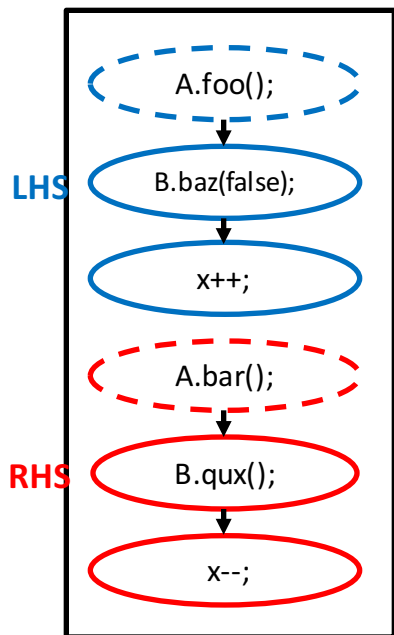
## Petablox Bug Report

# Translating Petablox Reports

- Race conditions can be understood as paths through the call graph, which gives way to a measure of relative distance to a violation.



# Translating Petablox Reports




# Translating Petablox Reports

```
1 public static void foo(){
2   float[] rmatrix_foo_enter = {{0.333},{0.0}};
3   Ariadne.annotate(rmatrix_foo_enter);
4   B.baz(false);
5   float[] rmatrix_foo_exit = {{0.0},{0.0}};
6   Ariadne.annotate(rmatrix_foo_exit);
7 }
8 public static void bar(){
9   float[] rmatrix_bar_enter = {{0.0},{0.5}};
10  Ariadne.annotate(rmatrix_bar_enter);
11  B.qux()
12  float[] rmatrix_bar_exit = {{0.0},{0.0}};
13  Ariadne.annotate(rmatrix_bar_exit);
14 }
```

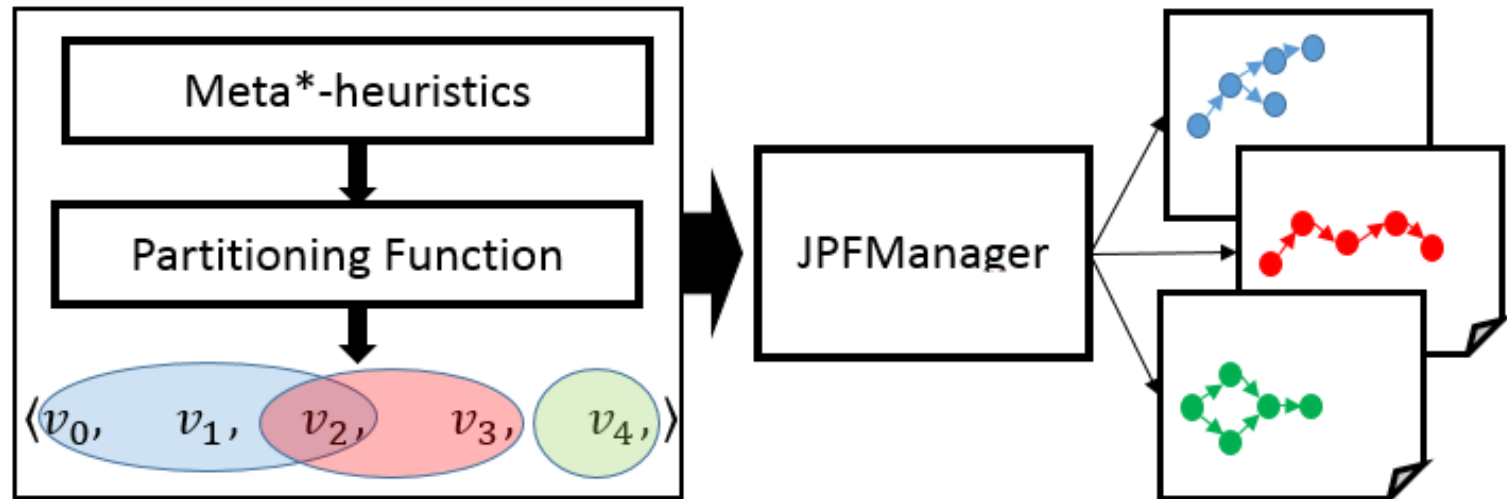


# Dynamic Metadata

- The annotations are calls to a metadata library.
- The annotation method itself does nothing, but calls to it are intercepted by a special listener.
- The listener decorates JPF states with annotations which can then be leveraged by a heuristic search.




# Multiobjective Search




# Experimental Results

# Benchmark Generation

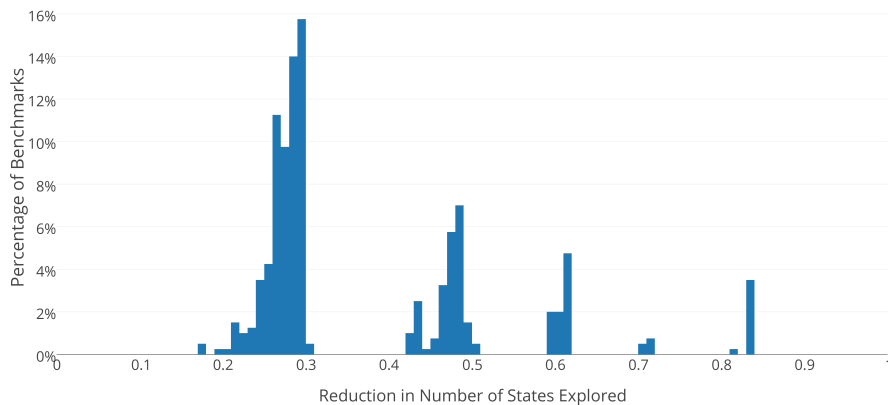


```
1 public void node5(){
2   int choice = Verify.getInt(0,1);
3   if(choice == 0) node7();
4   if(choice == 1) node8();
5 }
6 public void node7(){ read_loc_1(); node6(); }
7 public void node8(){ write_loc_1(); }
8 public int read_loc_1(){
9   synchronized(lock1){ return memloc1; }
10 }
11 public void write_loc_1(){
12   synchronized(lock1){ memloc1 = 1; }
13 }
```

# Branching Factor Tests




Relative Improvement of Ariadne over BFS



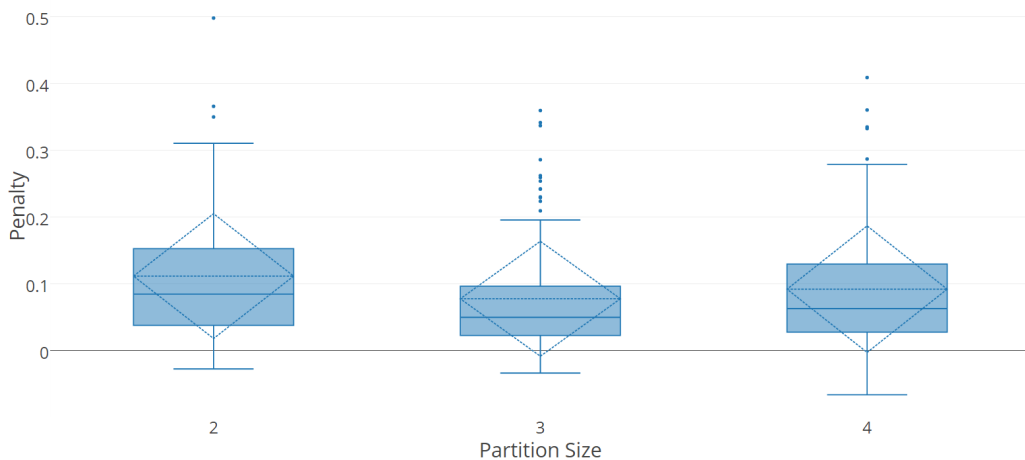
**Median Reduction vs. BFS:**

$$\frac{1}{e}$$

# Partitioning Tests



Average Performance Penalty Relative to Best Single Heuristic



## Performance Penalty

2-partitions: **11%**

3-partitions: **7%**

4-partitions: **9%**

# Elevator Benchmark Tests

Heuristic	States Explored, Path Length	Time (sec), Memory Consumed (gb)
BFS	427613 (44)	135 s (12.39 gb)
Interleaving	974806 (101)	350 s (31.77 gb)
Ariadne (single race, best)	<b>105120</b> (44)	<b>42 s (5.90 gb)</b>
Ariadne (single race, avg)	419242 (44)	136 s (7.97 gb)
Ariadne (partitioned, best)	<b>280818</b> (45)	<b>113.278 s (7.29 gb)</b>
Ariadne (partitioned, avg)	439473 (44)	138 s (7.90 gb)

4x fewer states, 263 gb required

1.52x fewer states, 79 gb required

Partitioned, multiobjective search allows for flexible parallelization.

# Ariadne: Conclusion

- Even given the imperfect nature of the static analysis (e.g. high false positive rate), we are able to find real bugs faster by integrating it into the search process.
- Dynamic metadata is a cost-effective vehicle for communicating information to the model checker.
- Multiobjective search allows us to efficiently allocate resources at scale with the volume of evidence we receive from the prior analysis.



Questions?