

Isolating Failure-Inducing Thread Schedules

by Jong-Deok Choi - IBM T.J. Watson Research Center
and Andreas Zeller - Saarland University

Marlene Hochrieser
FMV - Institute for Formal Models and Verification

Introduction to the problem

Consider a multi threaded application is running on the same input several times and fails occasionally.

- How do I reproduce a failure?

Thread switches are non-deterministic.

- How do I isolate the error?

Thread schedules may be composed of thousands of thread switches

- How do I find failing runs?

Is it possible to test different thread switch sequences with the same input?

Introduction to the solution

- Record a test run
- Deterministic replay
- Generate test cases
- Isolate failure causes by using delta debugging
- Relate causes to errors

Contents

- Code Example
- *Jalapeño - DejaVu*
- Capturing and Playbacking
- Isolating failure with Delta Debugging
- Generating altered schedules
- Case study - Ray tracer
- Related work
- Conclusion and future work
- Literature

Code example

```
1 class IntQueue {
2     // The queue holds integers in the range
3     // of [1..numberOfElements - 1]
4     static final int numberOfElements = 100;
5
6     // link[N] is N's successor in the queue
7     int link[] = new int[numberOfElements];
8
9     int head; // First element of queue
10    int tail; // Last element of queue
11
12    // Constructor
13    IntQueue() {
14        head = 0;
15        tail = 0;
16        for (int i = 0; i < numberOfElements;
17             i++) {
18            link[i] = 0;
19        }
20    }
21
22    // Enqueue ELEM.
23    public void enqueue(int elem) {
24        link[elem] = 0;
25
26        if (head == 0)
27            head = elem;
28        else {
29            synchronized (this) {
30                link[tail] = elem;
31            }
32        }
33        tail = elem;
34    }
35
36    // Return first element of queue.
37    // No error checking.
38    public int dequeue() {
39        int elem = head;
40        if (elem == tail)
41            tail = 0;
42
43        synchronized (this) {
44            head = link[head];
45        }
46
47        return elem;
48    }
49
50    // Print elements of queue
51    public void print() {
52        for (int e = head; e != 0; e = link[e])
53            System.out.print(e + " ");
54        System.out.println();
55    }
56 }
57
58 }
```

Three test runs with same input - Passing / Failing

<p>Clock</p> <p>1</p> <p>2</p> <p>3</p> <p>4</p> <p>5</p> <p>6</p> <p>7</p> <p>8</p> <p>9</p> <p>10</p> <p>11</p> <p>12</p> <p>13</p> <p>14</p> <p>15</p> <p>16</p>	<p>Thread A</p> <pre> enqueue(11) 26 link[elem] = 0; // link[11] = 0 28 if (head == 0) // 0 == 0 29 head = elem; // head = 11 36 tail = elem; // tail = 11 </pre>	<p>Thread B</p> <p>1 →</p> <pre> dequeue() 42 elem = head // elem = 11 43 if (elem == tail) // 11 == 11 44 tail = 0; // 47 head = link[head]; // head = 0 50 return elem; // return 11 </pre>	<p>Thread C</p> <p>2 →</p> <pre> enqueue(95) 26 link[elem] = 0; // link[95] = 0 28 if (head == 0) // 0 == 0 29 head = elem; // head = 95 36 tail = elem; // tail = 95 </pre>
<p>Clock</p> <p>1</p> <p>2</p> <p>3</p> <p>4</p> <p>5</p> <p>6</p> <p>7</p> <p>8</p> <p>9</p> <p>10</p> <p>11</p> <p>12</p> <p>13</p> <p>14</p> <p>15</p> <p>16</p>	<p>Thread A</p> <pre> enqueue(11) 26 link[elem] = 0; // link[11] = 0 28 if (head == 0) // 0 == 0 29 head = elem; // head = 11 </pre> <p>36 tail = elem; // tail = 11</p>	<p>Thread B</p> <p>1 →</p> <pre> dequeue() 42 elem = head // elem = 11 </pre> <p>2 ←</p> <p>3 →</p> <pre> 43 if (elem == tail) // 11 == 11 44 tail = 0; // </pre>	<p>Thread C</p> <p>4 →</p> <pre> enqueue(95) 26 link[elem] = 0; // link[95] = 0 28 if (head == 0) // 11 == 0 32 link[tail] = elem; // link[0] = 95 36 tail = elem; // tail = 95 </pre> <p>5 ←</p> <pre> 47 head = link[head]; // head = 0 50 return elem; // return 11 </pre>

Three test runs with same input - Passing

Clock	Thread A	Thread B	Thread C
1	enqueue(11)		
2	26 link[elem] = 0; // link[11] = 0		
3	28 if (head == 0) // 0 == 0		
4	29 head = elem; // head = 11		
5		$\xrightarrow{1}$ dequeue()	
6		42 elem = head // elem = 11	
7	36 tail = elem; // tail = 11	$\xleftarrow{2}$	
8		$\xrightarrow{3}$	
9		43 if (elem == tail) // 11 == 11	
10		44 tail = 0; //	
11		47 head = link[head]; // head = 0	
12		50 return elem; // return 11	
13		...	
14			$\xrightarrow{4}$ enqueue(95)
15			26 link[elem] = 0; // link[95] = 0
16			28 if (head == 0) // 11 == 0
17			29 head = elem; // head = 95
			36 tail = elem; // tail = 95

Jalapeño - DeJaVu

Jalapeño - Research VM for Java servers, developed at the IBM T.J. Watson Research Center.

- Designed for scalability (SMP), high performance, sophisticated thread support, availability, rapid response,..
- Papers are available at <http://jikesrvm.org/> (redirected from IBM's research page)
- Overview of *Jalapeño*: see [1]

DeJaVu - **D**eterministic **J**ava **U**tility is a tool to *capture, alter, and replay* Thread schedules.

- Running as part of *Jalapeño* VM
- Making failures reproducible by capturing a non-deterministic run and replaying it deterministically.

Jalapeño and Scheduling

Using implementation of the garbage collector to do scheduling.

Quasi-preemptive scheduling

Safe point - *is a program location where the compiler that created the method body is able to describe where all the live references exit.*

Yield point - *is a safe point located at a method prologue (such as function invocation or at a loop back-edge).*

Thread switching takes place only when a running thread has reached a *yield point*.

Capturing and playbacking

- *Dejavu* is using the *yield points* as global clock values
- During recording *Dejavu* stores global clock values of thread switches
- During playback a thread switch is forced at every *yield point*
- Different sequences of global clock values can be generated
- *Dejavu* will use them to force a thread switch
- Can be used for test generation


Capturing thread schedule - another approach [2]

```

class Test {
    static public volatile int f = 0;
        // shared variable
    static public volatile int g = 20;
        // shared variable
    static public void main(String argv[]) {
        int j; // local variable
        MyThread t1 = new MyThread();
        t1.start();
        j = 20;
        System.out.println("f = " + f
            + " j = " + j);
    }
}

class MyThread extends Thread {
    public void run() {
        int k; // local variable
        k = 5;
        Test.f = Test.f + k;
        Test.g = Test.g - k;
    }
}

```



- Thread schedule is a sequence of time intervals (time slices), containing thread schedule information
- Logical thread schedules: Generate equivalence classes for different order of access to shared variable
- Use logical thread schedule information with physical thread schedule

Isolating failure with Delta Debugging

A thread schedule is defined as a list of n clock times: t_1, \dots, t_n

Padded form of example results:

$\langle 6, 12, 17, 17, 17 \rangle$ - passing thread schedule in the *padded* form

$\langle 5, 7, 8, 10, 15 \rangle$ - failure-inducing thread schedule

$\langle 5, 7, 8, 13, 17 \rangle$ - passing thread schedule

stest is defined to return:

✓ if the queue holds the value 95

✗ if the queue is empty

? otherwise

Difference decomposition

- Difference of two schedules T_{ν} and $T_{\mathbf{x}}$ are defined as $\delta : \tau \rightarrow \tau$ with $\delta(T_{\nu}) = (T_{\mathbf{x}})$

- Set of all differences:

$$C = \tau^{\tau}$$

- Decompose δ in number of thread switch changes δ_i

- A schedule difference δ between two schedules

$T_{\nu} = \langle t_{\nu_1}, \dots, t_{\nu_n} \rangle$ and $T_{\mathbf{x}} = \langle t_{\mathbf{x}_1}, \dots, t_{\mathbf{x}_n} \rangle$ is defined as $\delta = \delta_1 \circ \delta_2 \circ \dots \circ \delta_n$ where each

$\delta_i : \tau \rightarrow \tau$ maps t_{ν_i} to $t_{\mathbf{x}_i}$; that is $\delta_i(T_{\nu}) = \langle t_{1\nu}, \dots, t_{i-1\nu}, t_{\mathbf{x}_i}, t_{\nu_{i+1}}, \dots, t_{\nu_{i+n}} \rangle$

$\circ : C \times C \rightarrow C$ is defined as $(\delta \circ \delta_j)(T) = \delta_i(\delta_j(T))$

Atomic decomposition

$\delta_i = \delta_{i,1} \circ \delta_{i,2} \circ \dots \circ \delta_{i,|t_{v_i}-t_{x_i}|}$ where each $\delta_{i,j}$ is defined as $\delta_{i,j}(T_{\checkmark}) = \delta_{i,j}(\langle t_{v_1}, t_{v_2}, \dots, t_{v_n} \rangle) = \langle t_{v_1}, t_{v_2}, \dots, t_{v_{i-1}}, t'_{v_i}, t_{v_{i+1}}, \dots, t_{v_n} \rangle$

where t'_{v_i} is the value altered by $\delta_{i,j}$; that is

→ $t_{v_i} + 1$ if $t_{v_i} < t_{x_i}$

→ $t_{v_i} - 1$ if $t_{v_i} > t_{x_i}$

Tests	$\delta_{1,1}$	$\delta_{2,1}\delta_{2,2}\delta_{2,3}\delta_{2,4}\delta_{2,5}$	$\delta_{3,1}\delta_{3,2}\delta_{3,3}\delta_{3,4}\delta_{3,5}\delta_{3,6}\delta_{3,7}\delta_{3,8}\delta_{3,9}$	$\delta_{4,1}\delta_{4,2}\delta_{4,3}\delta_{4,4}\delta_{4,5}\delta_{4,6}\delta_{4,7}$	$\delta_{5,1}\delta_{5,2}$	Schedule	Outcome
T_{\checkmark}	$\langle 6, 12, 17, 17, 17 \rangle$	✓
T_{\times}	□	□ □ □ □ □	□ □ □ □ □ □ □ □ □	□ □ □ □ □ □ □ □	□ □	$\langle 5, 7, 8, 10, 15 \rangle$	✗
(1)	□	□ □ □ □ □	□ □ □ □ □ □ □ □ □	$\langle 5, 7, 8, 17, 17 \rangle$	✓
(2)	□	□ □ □ □ □	□ □ □ □ □ □ □ □ □	□ □ □ □ □ □ □ □	. .	$\langle 5, 7, 8, 10, 17 \rangle$	✗
(3)	□	□ □ □ □ □	□ □ □ □ □ □ □ □ □	□ □ □ □	$\langle 5, 7, 8, 13, 17 \rangle$	✓
(4)	□	□ □ □ □ □	□ □ □ □ □ □ □ □ □	□ □ □ □ □ □ .	. .	$\langle 5, 7, 8, 11, 17 \rangle$	✓
Result				□			

Generating altered schedules with fuzzy approach


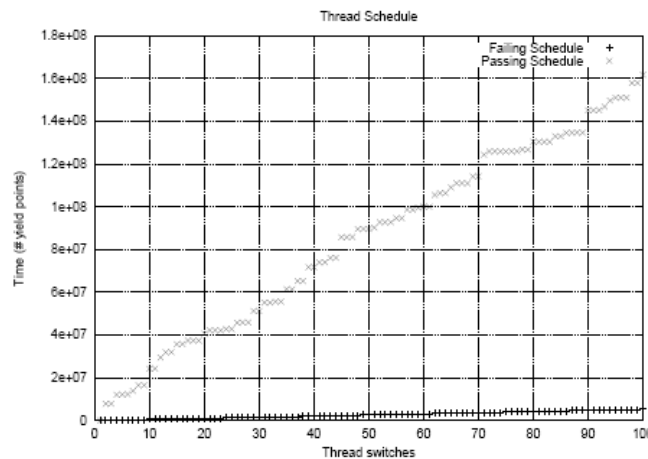
- Failing test run → passing test run
- Passing test run → failing test run

- Start from existing schedule
- Generate similar test cases to optimize determining differences
- Use simple Gaussian distribution centered around t :

From given schedule $T = \langle t_1, t_2, \dots, t_n \rangle$ generate $T = \langle f(t_1), f(t_2), \dots, f(t_n) \rangle$, where $f(t)$ is a perturbation function. - Widen distribution until an altered schedule is found

Case study - Ray tracer

```
25 public class Scene { ...
44     private static int ScenesLoaded = 0;
45     (more methods...)
81     private
82     int LoadScene(String filename) {
83         int OldScenesLoaded = ScenesLoaded;
84         (more initializations...)
91         infile = new DataInputStream(...);
92         (more code...)
130        ScenesLoaded = OldScenesLoaded + 1;
131        System.out.println(" " +
132            ScenesLoaded + " scenes loaded.");
132    }
133    ...
134 }
135 ...
733 }
```



- Each ray-tracing thread calls *LoadScene* → race condition → no update for *ScenesLoaded*
 - Record of failing schedule T ✘ contained 3770 thread switches
 - Using fuzzy approach it took 66 test to generate a passing schedule
 - More than a million *yield points* → 3,842,577,240 atomic deltas have been applied
 - Outcome: Amount of time was larger for T ✓
-
- Error isolated at *yield point* 59,772,127 → back tracing for a given set of *yield points* has been implemented

Related work

Article about *Debugging concurrent processes*, presented by [5]:

- Alter thread schedules manually
- Main idea of the paper *Isolating Failure-Inducing Thread Schedules*, but solution for automated testing.

Testing alternate schedules by [4]:

- Idea is to manipulate in order to get different schedules
- Using sleep or priorities
- Focus on coverage rather than on isolating failure causes

Conclusion and future work

- Method that automatically isolates the failure-inducing difference(s) between a passing and a failing schedule
- Purely experimental / analysis of the program in question is not required
- Use capturing, replaying, and isolating thread schedule as integrated part of testing and debugging of concurrent applications
- Use delta debugging
- Focus on cause-effect chains and other circumstances
- Integration of static analysis, dynamic analysis, and automated experiments
- More case studies

Literature

- [1] B. Alpern, C. R. Attanasio, J. J. Barton, M. G. Burke, P. Cheng, J.-D. Choi, A. Cocchi, S. J. Fink, D. Grove, M. Hind, S. F. Hummel, D. Lieber, V. Litvinov, M. F. Mergen, T. Ngo, J. R. Russell, V. Sarkar, M. J. Serrano, J. C. Shepherd, S. E. Smith, V. C. Sreedhar, H. Srinivasan, , and J. Whaley. The jalapeno virtual machine. IBM Systems Journal, Vol 39, No 1, 2000.
- [2] Jong-Deok Choi and Harini Srinivasan. Deterministic replay of java multithreaded applications. In SPDT 98: Proceedings of the SIGMETRICS symposium on Parallel and distributed tools, pages 4859, New York, NY, USA, 1998. ACM Press.
- [3] Jong-Deok Choi and Andreas Zeller. Isolating failure-inducing thread schedules. In ISSTA 02: Proceedings of the 2002 ACM SIGSOFT international symposium on Software testing and analysis, pages 210220, New York, NY, USA, 2002. ACM Press.
- [4] Y. Nir G.Ratsaby O. Edelstein, E. Farchi and S. Ur. Multithreaded java program test generation. IBM Systems Journal, 41(1):111-125, Februar 2002.
- [5] J. M. Stone. Debugging concurrent processes: A case study., June 1988.