

Unpredictable Behavior of Concurrent Unit Tests

- Thread switching is non-deterministic and machine-specific \rightarrow Unit tests may pass on one run, fail on the next
- **Success of unit tests does not indicate correct concurrent behavior**
- **Project Goal:** Develop an open-source framework for reliable unit testing of **concurrent Java applications**

Schedule-Based Unit Testing

- Execution using a fixed schedule \rightarrow Deterministic and machine-independent execution
- Execution using all different schedules → Success of unit tests implies correct concurrent behavior
- Unit testing not performance critical \rightarrow Overhead more acceptable in unit testing than in acceptance testing

Test-Driven Development

Unit Tests

- Written before the application code
- Must pass before code can be committed to the repository \rightarrow Fewer bugs enter the repository
- Prevent old bugs from reappearing \rightarrow Before a bug is fixed, a unit test exhibiting that bug is added
- Catch bugs early in development
- Current tools: JUnit, Ant

Our Experience with DrJava

- Unit tests used extensively \rightarrow 30% of code lines are unit tests \rightarrow 60% of code paths are tested
- **Production quality software**
- **Constantly changing development** team that includes sophomores
- Unit tests proved effective with single thread of control
- Major portion of DrJava is concurrent \rightarrow Existing tools not effective

Unit Testing for Concurrent Programs



	CK
Atomic	C
Block Block	(

Generation Technique 2: Exhaustive search \rightarrow Enumeration of all arrangements of atomic blocks covers all possible program behaviors \rightarrow Fewer arrangements of atomic blocks

than of instructions

Schedule 1	Atomic Block		Atomic Block	Atomic Block		
Schedule 2	Atomic Block	ļ	Atomic Block	Ato	mic Block	
: Schedule n	Atomic Block		Atomic Block		Atomic Block	
Eiguro, Enumorating arrangements of atomic blocks						

Figure: Enumerating arrangements of atomic blocks

- Modified class files will run on any Java VM \rightarrow Portable and open-source
- Bytecode can be dynamically compiled by an embedded JIT compiler \rightarrow Faster than interpretation

Variable X	Locks held	Variable Y	Locks held
Access 1	A C	Access 1	A C
Access 2	C	Access 2	C
Access 3	BC	Access 3	B
Intersection	С	Intersection	

X is guarded, intersection is non-empty

Y is not properly guarded, intersection is empty

Figure: Cumulative intersection of locks held

Testing Distributed Programs

- Extend framework to distributed Java programs using RMI (Remote Method Invocation)
- Data cannot be shared except using RMI operations \rightarrow Treat RMI operations as additional synchronization points
- Challenge: Make all machines observe the dependencies in the replayed schedule

Corky Cartwright, Mathias Ricken {cork | mgricken }@rice.edu



A Framework for Concurrent Unit Testing

- Open-source extension of JUnit \rightarrow Portable and accessible
- Schedule-based execution \rightarrow Deterministic, reliable results
- **Detection of unguarded access** \rightarrow Valuable debugging information
- Validation of schedule suites \rightarrow Unit testing remains automated
- Partial execution of test suite \rightarrow Faster non-concurrent testing
- Testing of distributed programs \rightarrow Desirable for testing of DrJava