Mutation Operators for Concurrent Java (J2SE 5.0)*[†]

Technical Report 2006-520

Jeremy S. Bradbury, James R. Cordy, Juergen Dingel

School of Computing, Queen's University
Kingston, Ontario, Canada
{*bradbury, cordy, dingel*}*@cs.queensu.ca*

November 2006

Abstract

The current version of Java (J2SE 5.0) provides a high level of support for concurreny in comparison to previous versions. For example, programmers using J2SE 5.0 can now achieve synchronization between concurrent threads using explicit locks, semaphores, barriers, latches, or exchangers. Furthermore, built-in concurrent data structures such as hash maps and queues, built-in thread pools, and atomic variables are all at the programmer's disposal.

We are interested in using mutation analysis to evaluate, compare and improve quality assurance techniques for concurrent Java programs. Furthermore, we believe that the current set of method mutation operators and class operators proposed in the literature are insufficient to evaluate concurrent Java source code because the majority of operators do not directly mutate the portions of code responsible for synchronization. In this paper we will provide an overview of concurrency constructs in J2SE 5.0 and a new set of concurrent mutation operators. We will justify the operators by categorizing them with an existing bug pattern taxonomy for concurrency. Most of the bug patterns in the taxonomy have been used to classify real bugs in a benchmark of concurrent Java applications.

1 Introduction

As a result of advances in hardware technology (e.g. multi-core processors) a number of practioners and researchers have advocated the need for concurrent software development [SL05]. Unfortunately, developing correct concurrent code is much more difficult than developing correct sequential code. The difficulty in programming concurrently is due to the many different, possibly unexpected, executions of the program. Reasoning about all possible interleavings in a program and ensuring that interleavings do not contain bugs is non-trivial. Edward A. Lee discussed concurrency bugs in a recent paper [Lee06]:

> *"I conjecture that most multithreaded-general purpose applications are so full of concurrency bugs that - as multicore architectures become commonplace - these bugs will begin to show up as system failures."*

*This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

[†] This technical report is an extended version of a paper with the same title in the *Proc. of the* 2^{nd} *Workshop on Mutation Analysis (Mutation 2006)*

The presence of bugs in concurrent code can have serious consequences including deadlock, starvation, livelock, dormancy, and incoincidence (calls occurring at the wrong time) [LSW05].

We are interested in using mutation to evaluate, compare, and improve quality assurance techniques for concurrent Java. The use of mutation with Java has been proposed in previous work – for instance the MuJava tool [MOK05]. MuJava includes two general types of mutation operators for Java: method level operators [KO91, MOK05] and class level operators [MKO02]. The method level operators include modifications to statements (e.g., statement deletion) and modifications to operands and operators in expressions (e.g., arithmetic operator insertion). The class level operators are related to inheritance (e.g., super keyword deletion), polymorphism (e.g., cast type change), and Java-specifc features. In general, the method and class level mutation operators do not directly mutate the synchronization portions of the source code in Java (J2SE 5.0) that handle concurrency. Furthermore, we conjecture that additional operators are needed in order to provide a more comprehensive set of operators that can truly reflect the types of bugs that often occur in concurrent programs. In this paper we present a set of concurrent operators for Java (J2SE 5.0). We believe our new set of concurrency mutation operators used in conjunction with existing method and class level operators provide a more comprehensive set of mutation metrics for the comparison and improvement of quality assurance testing and analysis for concurrency.

In the next section (Section 2) we will provide an overview of the support for concurrency in Java (J2SE 5.0). In Section 3 we provide an overview of real concurrency bug patterns which we will use to classify our concurrency mutation operators and demonstrate that the set of operators is both comprehensive and representative of real bugs. The set of mutation operators for concurrency and the bug pattern classification are presented in Section 5. Finally in Section 6 we provide our conclusions and an overview of our future work on using our new mutation operators.

2 Java Concurrency

Threads. Java concurrency is built around the notion of multi-threaded programs. The Java documentation defines a thread as "...a thread of execution in a program."[2] A typical thread is created and then started using the start() method and will be terminated once it has finished running. While a thread is alive it can often alternate between being runnable and not runnable. A number of methods exist that can affect the status of a thread:

- sleep(): will cause the current thread to become not runnable for a certain amount of time.
- yield(): will cause the current thread that is running to pause (temporarily).
- join(): will cause the caller thread to wait for a target thread to terminate.
- wait(): will cause the caller thread to wait until a condition is satisfied. Another thread notifies the caller that a condition is satisfied using the notify() or notifyAll() method.

Synchronization. Prior to J2SE 5.0, Java provided support for concurrency primarily through the use of the synchronized keyword. Java supports both synchronization methods and synchronization blocks. Additionally, synchronization blocks can be used in combination with implicit monitor locks.

Other Concurrency Mechanisms. In J2SE 5.0, additional mechanisms to support concurrency were added as part of `java.util.concurrent`[1]:

[2] `java.lang.Thread` documentation

[1] definitions of mechanisms and methods from the `java.util.concurrent` and the `java.util.concurrent.locks` documentation

- Explicit Lock (with Condition): Provides the same semantics as the implicit monitor locks but provides additional functionality such as timeouts during lock acquisition.
  - lock(), lockInterruptibly(), tryLock(): lock acquisition methods.
  - unlock(): lock release method.
  - await(), awaitNanos(), awaitUniterruptibly(), awaitUntil(): will cause a thread to wait (similar to wait() method).
  - signal(), signalAll(): will awaken waiting threads (similar to notify() and notifyAll() methods).
- Semaphore: Maintains a set of permits that restrict the number of threads accessing a resource. A Semaphore with one permit acts the same as a Lock.
  - acquire(), acquireUninterruptibly(), tryAcquire(): permit acquisition methods, some of which block until a permit is available.
  - release(): permit release method that will send a permit back to the semaphore.
- Latch: Allows threads from a set to wait until other threads complete a set of operations.
  - await(): will cause current thread to wait until latch has finished counting down or until the thread is interrupted.
  - countDown(): will decrement the latch count.
- Barrier: A point at which threads from a set wait until all other threads reach the point.
  - await(): used by a set of threads to wait until all other threads in the set have invoked the await() method.
- Exchanger: Allows for the exchange of objects between two threads at a given synchronization point.

Built-in Concurrent Data Structures. To reduce the overhead of developing concurrent data structures, J2SE 5.0 provides a number of collection types including ConcurrentHashMap and five different BlockingQueues.

Built-in Thread Pools. J2SE 5.0 provides a built-in FixedThreadPool and an unbounded CachedThreadPool.

Atomic Variables. The `java.util.concurrent.atomic` package includes a number of atomic variables that can be used in place of synchronization: AtomicInteger, AtomicIntegerArray, AtomicLong, AtomicLongArray, AtomicBoolean, AtomicReference and AtomicReferenceArray. Each atomic variable type contains new methods to support concurrency. For example, AtomicInteger contains methods such as addAndGet(), getAndSet() and others.

3 Bug Patterns for Java Concurrency

Farchi, Nir, and Ur have developed a bug pattern taxonomy for Java concurrency [FNU03]. The bug patterns are based on common mistakes programmers make when developing concurrent code in practice. Furthermore, the taxonomy has been expanded and used to classify bugs in an existing public domain concurrency benchmark maintained by IBM Research [EU04]. The benchmark contains 40 programs ranging in size from 57 to 17000 loc. Programs in the benchmark are from a variety of sources including student created programs, tool developer programs, open source programs, and a commercial product. In our attempt to develop a comprehensive set of concurrency mutation operators we will later classify our operators with respect to the bug patterns taxonomy. Since this bug pattern taxonomy was developed prior to J2SE 5.0 we have had to add some additional patterns that occur in concurrency constructs not available at the time the taxonomy was proposed. We distinguish between the original bug patterns(*), the added bug patterns also used in the benchmark classification(**) and new patterns that we are including ($^+$):

- **Nonatomic operations assumed to be atomic bug pattern.*** *"...an operation that "looks" like one operation in one programmer model (e.g., the source code level of the programming language). but actually consists of several unprotected operations at the lower abstraction levels" [FNU03].* In this paper we also include nonatomic floating point operations** in this pattern.

- **Two-state access bug pattern.*** *"Sometimes a sequence of operations needs to be protected but the programmer wrongly assumes that separately protecting each operation is enough" [FNU03].*

- **Wrong lock or no lock bug pattern.*** *"A code segment is protected by a lock but other threads do not obtain the same lock instance when executing. Either these other threads do not obtain a lock at all or they obtain some lock other than the one used by the code segment" [FNU03].*

- **Double-checked lock bug pattern.*** *"When an object is initialized, the thread local copy of the objects field is initialized but not all object fields are necessarily written to the heap. This might cause the object to be partially initialized while its reference is not null" [FNU03].*

- **The sleep() bug pattern.*** *"The programmer assumes that a child thread should be faster than the parent thread in order that its results be available to the parent thread when it decides to advance. Therefore, the programmer sometimes adds an 'appropriate' sleep() to the parent thread. However, the parent thread may still be quicker in some environment. The correct solution would be for the parent thread to use the join() method to explicitly wait for the child thread" [FNU03].*

- **Losing a notify bug pattern.*** *"If a notify() is executed before its corresponding wait(), the notify() has no effect and is "lost" ... the programmer implicitly assumes that the wait() operation will occur before any of the corresponding notify() operations" [FNU03].*

- **Other missing or nonexistent signals.**$^+$ This pattern generalizes the losing a notify bug pattern to all other signals. For example, at a barrier the await() method has to be called by a set number of threads before the program can proceed. If an await() from one thread never occurs then all of threads at the barrier may be stuck waiting.

- **Notify instead of notify all bug pattern.**** If a notify() is executed instead of notifyAll() then threads with some of its corresponding wait() calls will not be notified [LDG$^+$04].

- **A "blocking" critical section bug pattern.*** *"A thread is assumed to eventually return control but it never does" [FNU03].*

- **The orphaned thread bug pattern.*** *"If the master thread terminates abnormally, the remaining threads may continue to run, awaiting more input to the queue and causing the system to hang" [FNU03].*

- **The interference bug pattern.**** A pattern in which *"...two or more concurrent threads access a shared variable and when at least one access is a write, and the threads use no explicit mechanism to prevent the access from being simultaneous."* [LSW05]. The interference bug pattern can also be generalized from classic data race interference to include high level data races** which deal *"...with accesses to sets of fields which are related and should be accessed atomically" [AHB03].*

- **The deadlock (deadly embrace) bug pattern.**** *"...a situation where two or more processes are unable to proceed because each is waiting for one of the others to do something in a deadlock cycle ... For example, this occurs when a thread holds a lock that another thread desires and vice-versa"* [LSW05].

- **Starvation bug pattern.**$^+$ This bug occurs when their is a failure to *"...allocate CPU time to a thread. This may be due to scheduling policies..."* [Lea00]. For example, an unfair lock acquisition scheme might cause a thread never to be scheduled.

- **Resource exhaustion bug pattern.**$^+$ *"A group of threads together hold all of a finite number of resources. One of them needs additional resources but no other thread gives one up"* [Lea00].

- **Incorrect count initialization bug pattern.**$^+$ This pattern occurs when there is an incorrect initialization in a barrier for the number of parties that must be waiting for the barrier to trip, or an incorrect initialization of the number of threads required to complete some action in a latch, or an incorrect initialization of the number of permits in a semaphore.

4 Related Work: Existing Mutation Operators for Java

Currently, there are two main groups of operators for mutating Java source code: method and class mutation operators. As previously stated we believe the existing operators are complementary to our concurrency mutation operators.

4.1 Method Mutation Operators

Method level operators [KO91, MOK05] have been used in previous mutation tools for other programming languages besides Java (e.g., the Mothra tool set for mutating Fortran programs **??**). These operators are applied to statements, operands and operators (see Table 0(a)). Operators applied to statements perform actions such as modification, replacement, and deletion. Operators applied to operands primarily are replacements. Operators applied to operators include insertion, deletion, and replacement. The sufficient set of method level mutation operators in Table 0(a) have been implemented in the muJava tool [OMK04, MOK05]. The sufficient set of operators was defined with respect to an empirical study of method operators with a set of Fortran programs [OLR$^+$96]. Further research is also empirically studying this problem [NA06].

4.2 Class Mutation Operators

A set of class level operators [MKO02] were developed and implemented in the muJava tool [OMK04, MOK05] are used for the creation of mutants that results from object oriented bugs. The class level operators are related to inheritance, polymorphism, and Java-specific object oriented features (see Table 0(b)). Operators related to inheritance include access modifier changes, overriding method changes, and changes in reference to parent classes. Operators related to polymorphism include changes to class type, type casting, and overloaded methods. Operators related to Java-specific features including inserting and deleting the this and static keywords.

5 Concurrent Mutation Operators

We propose five categories of mutation operators for concurrent Java: modify parameters of concurrent methods, modify the occurrence of concurrency method calls (removing, replacing and exchanging), modify keywords (addition and removal), switch concurrent objects, and modify critical regions (shift, expand,

(a) Method mutation operators [KO91, OLR^{+}96]

	Operator	Description
Method Level Mutation Operators (sufficient set)	AOR	*Arithmetic Operator Replacement*
	AOI	*Arithmetic Operator Insertion*
	AOD	*Arithmetic Operator Deletion*
	ROR	*Relational Operator Replacement*
	COR	*Conditional Operator Replacement*
	COI	*Conditional Operator Insertion*
	COD	*Conditional Operator Deletion*
	SOR	*Shift Operator Replacement*
	LOR	*Logical Operator Replacement*
	LOI	*Logical Operator Insertion*
	LOD	*Logical Operator Deletion*
	ASR	*Assignment Operator Replacement*
Method Level Mutation Operators (not part of sufficient set)	AAR	*Array Reference for Array Reference Replacement*
	ABS	*Absolute Value Insertion*
	ACR	*Array Reference for Constant Replacement*
	ASR	*Array Reference for Scalar Variable Replacement*
	CAR	*Constant for Array Reference Replacement*
	CNR	*Comparable Array Name Replacement*
	CRP	*Constant Replacement*
	DER	*Do Statement End Replacement*
	DSA	*Data Statement Alterations*
	GLR	*Goto Label Replacement*
	RSR	*Return Statement Replacement*
	SAN	*Statement Analysis*
	SAR	*Scalar Variable for Array Reference Replacement*
	SCR	*Scalar for Constant Replacement*
	SDL	*Statement Deletion*
	SRC	*Source Constant Replacement*
	SVR	*Scalar Variable Replacement*

(b) Class mutation operators used in muJava [MKO02]

	Operator	Description
Class Mutation Operators (Inheritance)	AMC	*Access Modifier Change*
	IHD	*Hiding Variable Deletion*
	IHI	*Hiding Variable Insertion*
	IOD	*Overridding method deletion*
	IOP	*Overridding method calling position change*
	IOR	*Overridding method rename*
	ISI	**`super`** *keyword insertion*
	ISD	**`super`** *keyword deletion*
	IPC	*Explicit call to a parent's constructor deletion*
Class Mutation Operators (Polymorphism)	PNC	*new method call with child class type*
	PMD	*Member variable declaration with parent class type*
	PPD	*Parameter variable declaration with child class type*
	PCI	*Type cast operator insertion*
	PCC	*Cast type change*
	PCD	*Type cast operator deletion*
	PRV	*Reference assignment with other comparable variable*
	OMR	*Overloading method contents replace*
	OMD	*Overloading method deletion*
	OAC	*Arguments of overloading method call change*
Class Mutation Operator (Java-Specific Features)	JTI	`this` *keyword insertion*
	JTD	`this` *keyword deletion*
	JSI	`static` *modifier insertion*
	JSD	`static` *modifier deletion*
	JID	*Member variable intialization deletion*
	JDC	*Java-supported default constructor creation*
	EOA	*Reference assignment and content assignment replacement*
	EOC	*Reference comparison and content comparison replacement*
	EAM	*Accessor method change*
	EMM	*Modifier method change*

Table 1: Existing mutation operators

shrink and split). The relationship between these general operator categories and the concurrency mechanisms provided in J2SE 5.0 is presented in Table 2 – which demonstrates that the operators provide coverage over the J2SE 5.0 concurrency mechanisms.

A complete list of the operators we will be presenting in this section is provided in Table 3. The mutant operators are designed specifically to represent mistakes that programmers may make when implementing concurrency. Therefore, many of the operators are specific only to concurrency methods, objects and keywords. We have tried to use context and knowledge about Java concurrency to make the operators as specific as possible in order to make concurrency mutation analysis more feasible by reducing the total number of mutants produced.

Readers familiar with method and class level mutation operators will notice that some of our mutation operators are special cases of existing mutation operators while others are new operators that have not been

Java (J2SE 5.0) Concurrency Mutation Operator Categories	*Threads*	*Synchronization methods*	*Synchronization statements*	*Synchronization with implicit monitor locks*	*Explicit locks*	*Semaphores*	*Barriers*	*Latches*	*Exchangers*	*Built-in concurrent data structures (e.g. queues)*	*Built-in thread pools*	*Atomic variables (e.g. LongInteger)*
Modify Parameters of Concurrent Methods	✓	–	✓	✓	✓	✓	✓	✓	–	–	–	–
Modify the Occurrence of Concurrency Method Calls	✓	–	–	–	✓	✓	✓	✓	–	–	–	✓
Modify Keyword	–	✓	✓	✓	✓	–	–	–	–	–	–	–
Switch Concurrent Objects	–	–	–	–	✓	✓	✓	✓	✓	✓	✓	–
Modify Concurrent Region	–	✓	✓	✓	✓	✓	–	–	–	–	–	–

Table 2: The relationship between new mutation operators for concurrency and the concurrency features provided by J2SE 5.0

Operator Category	Concurrency Mutation Operators for Java (J2SE 5.0)
Modify Parameters of Concurrent Methods	M*X*T – Modify Method-X Time *(wait(), sleep(), join(), and await() method calls)*
	MSP - Modify Synchronized Block Parameter
	ESP - Exchange Synchronized Block Parameters
	MSF - Modify Semaphore Fairness
	M*X*C - Modify Permit Count in Semaphore and Modify Thread Count in Latches and Barriers
	MBR - Modify Barrier Runnable Parameter
Modify the Occurrence of Concurrency Method Calls	RT*X*C – Remove Thread Method-X Call *(wait(), join(), sleep(), yield(), notify(), notifyAll() Methods)*
	RC*X*C – Remove Concurrency Mechanism Method-X Call *(methods in Locks, Semaphores, Latches, Barriers, etc.)*
	RNA - Replace NotifyAll() with Notify()
	RJS - Replace Join() with Sleep()
	ELPA - Exchange Lock/Permit Acquisition
	EAN - Exchange Atomic Call with Non-Atomic
Modify Keyword	ASTK – Add Static Keyword to Method
	RSTK – Remove Static Keyword from Method
	ASK - Add Synchronized Keyword to Method
	RSK - Remove Synchronized Keyword from Method
	RSB - Remove Synchronized Block
	RVK - Remove Volatile Keyword
	RFU - Remove Finally Around Unlock
Switch Concurrent Objects	R*X*O - Replace One Concurrency Mechanism-X with Another *(Locks, Semaphores, etc.)*
	EELO - Exchange Explicit Lock Objects
Modify Critical Region	SHCR - Shift Critical Region
	SKCR - Shrink Critical Region
	EXCR – Expand Critical Region
	SPCR - Split Critical Region

Table 3: Concurrency mutation operators for Java

previously proposed. Other related work from the concurrency bug detection community includes a set of 18 hand-created concurrency mutants [LDG$^+$04] for a previous version of Java that did not contain many of the concurrency mechanisms available in J2SE 5.0. We have compared our comprehensive set of operators with this work and found that our operators in combination with the method and class level operators subsume the manual mutants used in the previous work. The idea of using mutation for concurrency was also suggested by Ghosh who proposed two mutation operators (RSYNCHM and RSYNCHB) for removing the synchronized keyword from methods and removing synchronized blocks [Gho02]. The operators proposed by Ghosh are equivalent to the Remove Synchronized Keyword from Method (RSK) and Remove Synchronized Block (RSB) operators presented in this paper.

5.1 Modify parameters of concurrent method

These operators involve modifying the parameters of methods for thread and concurrency classes. Some of the method level mutation operators that modify operands are similar to the operators proposed here.

5.1.1 MXT - Modify Method-X Timeout

The MXT operator can be applied to the wait(), sleep(), and join() method calls (introduced in Section 2) that include an optional timeout parameter. For example, in Java a call to wait() with the optional timeout parameter will cause a thread to no longer be runnable until a condition is satisfied or a timeout has occurred. The MXT replaces the timeout parameter, t, of the wait() method by some appropriately chosen fraction or multiple of t (e.g., $t/2$ and $t * 2$). We could replace the timeout parameter by a variable of an equivalent type however since we know that the parameter represents a time value it is just as meaningful to mutate the method to both increase and decrease the time by a factor of 2.

Original Code:

```
long time = 10000;
try{
    wait(time);
} catch ...
```

MXT Mutant for wait():

```
long time = 10000;
try{
    wait(time*2);  //or replace with time/2
} catch ...
```

The MXT operator with the wait() method is most likely to result in an interference bug or a data race. The MXT operator with the sleep() and join() methods is most likely to result in the sleep() bug pattern. For example, in a situation where a sleep() or join() is used by a caller thread to wait for another thread, reducing the time may cause the caller thread to not wait long enough for the other thread to complete.

The MXT operator can also be applied to the optional timeout parameter in await() method calls. Both barriers and latches have an await() method. In barriers the await() method is used to cause a thread to wait until all threads have reached the barrier. In latches the await() method is used by threads to wait until the latch has finished counting down, that is until all operations in a set are complete. For example:

Original Code:

```
CountDownLatch latch1
   =new CountDownLatch(1);
...
long time = 50;
latch1.await(time,TimeUnit.MILLISECONDS);
...
```

MXT Mutant for await():

```
CountDownLatch latch1
   =new CountDownLatch(1);
...
long time = 50;
latch1.await(time/2,TimeUnit.MILLISECONDS);
//or replace time with time*2
...
```

The MXT operator when applied to an await() method call will most likely result in an interference bug.

5.1.2 MSP - Modify Synchronized Block Parameter

Common parameters for a synchronized block include the this keyword, indicating that synchronization occurs with respect to the instance object of the class, and implicit monitor objects. If the keyword this or an object is used as a parameter for a synchronized block we can replace the parameter by another object or the keyword this. For example:

Original Code:

```
private Object lock1 = new Object();
private Object lock2 = new Object();
....
public void methodA(){
    synchronized(lock1){ ... }
}
...
```

MSP Mutant:

```
private Object lock1 = new Object();
private Object lock2 = new Object();
...
public void methodA(){
    synchronized(lock2){ ... }
}
...
```

Another MSP Mutant:

```
private Object lock1 = new Object();
private Object lock2 = new Object();
...
public void methodA(){
    synchronized(this){ ... }
}
...
```

The MSP operator will result in the wrong lock bug pattern.

5.1.3 ESP - Exchange Synchronized Block Parameters

If a critical region is guarded by multiple synchronized blocks with implicit monitor locks the ESP operator exchanges two adjacent lock objects. For example:

Original Code:

```
private Object lock1 = new Object();
private Object lock2 = new Object();
....
public void methodA(){
    synchronized(lock1){
        synchronized(lock2){ ... }
    }
}
...
public void methodB(){
    synchronized(lock1){
        synchronized(lock2){ ... }
    }
}
...
```

ESP Mutant:

```
private Object lock1 = new Object();
private Object lock2 = new Object();
....
public void methodA(){
    //switched lock1 and lock2 in methodA
    synchronized(lock2){
        synchronized(lock1){ ... }
    }
}
...
public void methodB(){
    synchronized(lock1){
        synchronized(lock2){ ... }
    }
}
...
```

The ESP mutation operator can result in a wrong lock bug because exchanging two adjacent locks will cause the locks to be acquired at incorrect times for incorrect critical regions. The ESP operator can also cause a classic deadlock (via deadly embrace) bug to occur as is the case in the above example.

5.1.4 MSF - Modify Semaphore Fairness

Recall in Section 2 that a semaphore maintains a set of permits for accessing a resource. In the constructor of a Semaphore there is an optional parameter for a boolean fairness setting. When the fairness setting is not used the default fairness value is false which allows for unfair permit acquisition. If the fairness parameter is a constant then the MSF operator is a special case of the Constant Replacement (CRP) method level operator and replaces a true value with false and a false value with true. In the case that a boolean variable is used as a parameter we simply negate it.

A potential consequence of expecting a semaphore to be fair when in fact it is not is that there is a potential for starvation because no guarantees about permit acquisition ordering can be given. In fact, when a semaphore is unfair any thread that invokes the Semaphore's acquire() method to obtain a permit may receive one prior to an already waiting thread - this is known as barging[3].

Original Code:

```
int permits = 10;
private final Semaphore sem
    = new Semaphore (permits, true);
...
```

MSF Mutant:

```
int permits = 10;
private final Semaphore sem
    = new Semaphore (permits, false);
...
```

5.1.5 M*X*C - Modify Concurrency Mechanism-X Count

The M*X*C operator is applied to parameters in three of Java's concurrency mechanisms: Semaphores, Latches, and Barriers. A latch allows a set of threads to countdown a set of operations and a barrier allows a set of threads to wait at a point until a number of threads reach that point. The count being modified in Semaphores is the set of permits, and in Latches and Barriers it is the number of threads. We will next provide an example of the M*X*C operator for Semaphores, Latches, and Barriers.

The constructor of the Semaphore class has a parameter that refers to the maximum number of available permits that are used to limit the number of the threads accessing the shared resource. Access is acquired using the acquire() method and released using the release() method. Both the acquire() and release() method calls have optional count parameters referring to the number of permits being acquired or released. The M*X*C operator modifies the number of permits, p, in calls to these methods by decrementing (p--) and incrementing (p++) it by 1. For example:

Original Code:

```
int permits = 10;
private final Semaphore sem
    = new Semaphore (permits,true);
...
```

M*X*C Mutant for a Semaphore:

```
int permits = 10;
private final Semaphore sem
    = new Semaphore (permits--,true);
...
```

A potential bug that can occur from modifying permit counts in Semaphores. In the above example if the total number of permits had been one then decrementing the number of permits by 1 would have lead to a situation where no permits were ever available. Another bug could occur if we increased the number of permits acquired by the acquire() method but did not increase the count in the release() method which could

[3] `java.util.concurrent` documentation

eventually exhaust the resources. In this case we could end up with a blocking critical section bug once all of the permits were held but not released.

Similar to the Semaphore constructors permit count, The constructor of the concurrent latch class CountDownLatch has a thread count parameter that can also be incremented and decremented. For example:

Original Code:

```
int i = 10;
CountDownLatch latch1
    = new CountDownLatch(i);
...
```

M*X*C Mutant for a Latch:

```
int i = 10;
CountDownLatch latch1
    = new CountDownLatch(i--);
...
```

The M*X*C operator can also increment and decrement the thread count parameter in the constructor of the concurrent barrier class CyclicBarrier. For example:

Original Code:

```
int i=10;
CyclicBarrier barrier1
    = new CyclicBarrier(i,
    new Runnable(){
        public void run(){
        }
    });
...
```

M*X*C Mutant for a Barrier:

```
int i=10;
CyclicBarrier barrier1
    = new CyclicBarrier(i++,
    new Runnable(){
        public void run(){
        }
    });
...
```

A potential bug that can occur from modifying the number of threads in Latches and Barriers is resource exhaustion.

5.1.6 MBR - Modify Barrier Runnable Parameter

The CyclicBarrier constructor has a parameter that is an optional runnable thread that can happen after all the threads complete and reach the barrier. The MBR operator modifies the runnable thread parameter by removing it if it is present. This is a special case of the method level mutation operator, statement deletion (SDL). For example:

Original Code:

```
int i=10;
CyclicBarrier barrier1
    = new CyclicBarrier(i,
    new Runnable(){
        public void run(){
        }
    });
...
```

MBR Mutant:

```
int i=10;
CyclicBarrier barrier1
    = new CyclicBarrier(i);
    //runnable thread parameter removed‘
...
```

An example of a bug caused by the MBR operator is missed or nonexistent signals if some signal calls were present in the runnable thread.

5.2 Modify the occurrence of concurrency method calls: remove, replace, and exchange

This class of operators is primarily interested in modifying calls to thread methods and methods of concurrency mechanism classes. Examples of modifications include removal of a method call and replacement or exchange of a method call with a different but similar method call. The operators that remove method calls are special cases of the method level operator: Statement Deletion (SDL).

5.2.1 RT*X*C - Remove Thread Method-X Call

The RT*X*C operator removes calls to the following methods: wait(), join(),sleep(), yield(), notify(), and notifyAll(). Removing the wait() method can cause potential interference, removing the join() and sleep() methods can cause the sleep() bug pattern, and removing the notify() and notifyAll() method calls is an example of losing a notify bug. We will now provide an example of the RT*X*C operator used to remove a wait() method call.

Original Code:

```
try{
    wait();
} catch ...
```

RT*X*C Mutant for wait():

```
try{
    //removed wait();
} catch ...
```

5.2.2 RC*X*C - Remove Concurrency Mechanism Method-X Call

The RC*X*C operator can be applied to the following concurrency mechanisms: Locks (lock(), unlock()), Condition (signal(), signalAll()), Semaphore (acquire(), release()), Latch(countDown(), and ExecutorService(e.g., submit()).

Let us consider a specific application of the RC*X*C operator in a ReentrantLock or a ReentrantReadWriteLock with a call to the unlock() method. The RC*X*C operator removes this call thus the lock is not released causing an example of a blocking critical section bug. For example:

Original Code:

```
private Lock lock1 = new ReentrantLock();
...
lock1.lock();
try{
   ...
} finally{
   lock1.unlock();
}
...
```

RC*X*C Mutant for a Lock:

```
private Lock lock1 = new ReentrantLock();
...
lock1.lock();
try{
   ...
} finally{
 //removed lock1.unlock();
}
...
```

The RC*X*C operator can also be used to remove calls to the acquire() and release() methods for a Semaphore. On the one hand, if an acquire() call is removed interference may occur. On the other hand, if a release() call is removed a blocking critical section bug might be the result.

Original Code:

```
int permits = 10;
private final Semaphore sem = new Semaphore (permits, true);
...
sem.acquire();
...
sem.release();
...
```

RC*X*C Mutant for a Semaphore:

```
int permits = 10;
private final Semaphore sem
    = new Semaphore (permits, true);
...
//removed sem.acquire();
...
sem.release();
...
```

Another RC*X*C Mutant for a Semaphore:

```
int permits = 10;
private final Semaphore sem
    = new Semaphore (permits, true);
...
sem.acquire();
...
//removed sem.release();
...
```

Due to the similar nature of applying the RC*X*C operator for other concurrency mechanisms we will not provide any additional examples.

5.2.3 RNA - Replace NotifyAll() with Notfiy()

The RNA operator replaces a notifyAll() with a notify() and is an example of the notify instead of notify all bug pattern.

Original Code:

```
...notifyAll();...
```

RNA Mutant:

```
...notify();...
```

5.2.4 RJS - Replace Join() with Sleep()

The RJS operator replaces a join() with a sleep() and is an example of the sleep() bug pattern.

Original Code:

```
...join();...
```

RJS Mutant:

```
...sleep(10000);...
```

5.2.5 ELPA - Exchange Lock/Permit Acquistion

In a Semaphore the acquire(), acquireUninterruptibly() and tryAcquire() methods can be used to obtain one or more permits to access a shared resource. The ELPA operator exchanges one method for another which can lead to potential timing changes as well as starvation. For example, an acquire() method will try and obtain one or more permits and will block and wait until the permit or permits become available. If the thread that invoked the acquire() method is interrupted it will no longer continue to block and wait. If the acquire() method invocation is changed to acquireUninterruptibly() it will behave exactly the same except it can no longer be interupted. Thus in situations where the semaphore is unfair or if for other reasons the number of requested permits never becomes available the thread that invoked the acquireUninterruptibly() will stay dormant and wait. If an acquire() method invocation is changed to a tryAcquire() then a permit will be acquired if one is available otherwise the thread will not block and wait. tryAcquire() will acquire a permit or permits unfairly even if the fairness setting is set to fair. Use of tryAcquire() may cause starvation for threads waiting for permits.

Original Code:

```
int permits = 10;
private final Semaphore sem = new Semaphore (permits, true);
...
sem.acquire();
...
```

ELPA Mutant:

```
int permits = 10;
private final Semaphore sem = new Semaphore (permits, true);
...
sem.acquireUninterruptibly();
...
```

Another ELPA Mutant:

```
int permits = 10;
private final Semaphore sem = new Semaphore (permits, true);
...
sem.tryAcquire();
...
```

The ELPA operator can also be applied to the lock(), lockInterruptibly(), tryLock() method calls with Locks.

5.2.6 SAN - Switch Atomic Call with Non-Atomic

A call to the getAndSet() method in an atomic variable class is replaced by a call to the get() method and a call to the set() method. The effect of this replacement is that the combined get and set commands are no longer atomic. For example:

Original Code:

```
AtomicInteger int1 = 15;
...
int oldVal = int1.getandSet(40);
...
```

SAN Mutant:

```
AtomicInteger int1 = 15;
...
int oldVal = int1.get();
int1.set(40);
...
```

5.3 Modify keywords: add and remove

We consider what happens when we add and remove keywords such as static, synchronized,volatile, and finally.

5.3.1 ASTK - Add Static Keyword to Method

The static keyword used for a synchronized method indicates that the method is synchronized using the class object not the instance object. The ASTK operator adds static to non-static synchronized methods and causes synchronization to occur on the class object instead of the instance object. The ASTK operator is an example of the wrong lock bug pattern.

Original Code:

```
public synchronized void aMethod(){
    ... }
```

ASTK Mutant:

```
public static synchronized void aMethod(){
    ... }
```

5.3.2 RSTK - Remove Static Keyword from Method

The RSTK operator removes static from static synchronized methods and causes synchronization to occur on the instance object instead of the class object. Similar to the ASTK operator, the RSTK operator is an examples of the wrong lock bug pattern.

Original Code:

```
public static synchronized void bMethod(){
      ... }
```

RSTK Mutant:

```
public synchronized void bMethod(){
    ... }
```

5.3.3 ASK - Add Synchronized Keyword to Method

The synchronized keyword is added to a non-synchronized method in a class that has synchronized methods or statements. The ASK operator has the potential to cause a deadlock, for example, if a critical region already exists inside the method.

Original Code:

```
public void aMethod(){
    ... }
```

ASK Mutant:

```
public synchronized void aMethod(){
    ... }
```

5.3.4 RSK - Remove Synchronized Keyword from Method

The synchronized keyword is important in defining concurrent methods and the omission of this keyword is a plausible mistake that a programmer might make when writing concurrent source code. The RSK operator removes the synchronized keyword from a synchronized method and causes a potential no lock bug. For example:

Original Code:

```
public synchronized void aMethod(){ ... }
```

RSK Mutant:

```
public void aMethod(){ ... }
```

5.3.5 RSB - Remove Synchronized Block

Similar to the RSK operator, the RSB operator removes the synchronized keyword from around a statement block which can cause a no lock bug. For example:

Original Code:

```
    synchronized(this){
     <statement_c1>
    }
```

RSB Mutant:

```
    //synchronized(this) is removed
    <statement_c1>
     ...
```

5.3.6 RVK - Remove Volatile Keyword

The volatile keyword is used with a shared variable and prevents operations on the variable from being reordered in memory with other operations. In the below example we remove the volatile keyword from a shared long variable. If a long variable, which is 64-bit, is not declared volatile then reads and writes will be treated as two 32-bit operations instead of one operation. Therefore, the RVK operator can cause a situation where a nonatomic operation is assumed to be atomic. For example:

Original Code:

```
volatile long x;
```

RVK Mutant:

```
long x;
```

5.3.7 RFU - Remove Finally Around Unlock

The finally keyword is important in releasing explicit locks. In the below example, finally ensures that the unlock() method call will occur after a try block regardless of whether or not an exception is thrown. If finally is removed the unlock() will not occur in the presence of an exception and cause a blocking critical section bug.

Original Code:

```
private Lock lock1 = new ReentrantLock();
...
lock1.lock();
try{
    ...
} finally{
   lock1.unlock();
}
...
```

RFU Mutant:

```
private Lock lock1 = new ReentrantLock();
...
lock1.lock();
try{
    ...
}
lock1.unlock();
...
```

5.4 Switch concurrent objects

When multiple instances of the same concurrent class type exist we can replace one concurrent object with the other.

5.4.1 R*X*O - Replace One Concurrency Mechanism-X with Another

When two instances of the same concurrency mechanism exist we replace a call to one with a call to the other. For example, consider the replacement of Lock method calls:

Original Code:

```
private Lock lock1 = new ReentrantLock();
private Lock lock2 = new ReentrantLock();
...
lock1.lock();
...
```

R*X*O Mutant for Locks:

```
private Lock lock1 = new ReentrantLock();
private Lock lock2 = new ReentrantLock();
...
//should be call to lock1.lock()
lock2.lock();
...
```

We can also apply the R*X*O operator when 2 or more objects exist of type Semaphore, CountDownLatch, CyclicBarrier, Exchanger, and more. For example consider the application of the R*X*O operator with two

Semaphores and two Barriers:

Original Code:

```
private final Semaphore sem1
   = new Semaphore(100, true);
private final Semaphore sem2
   = new Semaphore(50, true);
...
sem1.acquire();
...
```

RXO Mutant for Semaphores:

```
private final Semaphore sem1
   = new Semaphore(100, true);
private final Semaphore sem2
   = new Semaphore(50, true);
...
//should be call to sem1.acquire()
sem2.acquire();
...
```

Original Code:

```
final CyclicBarrier bar1
   = new CyclicBarrier(20,
                new Runnable() {...});
final CyclicBarrier bar2
   = new CyclicBarrier(20,
                new Runnable() {...});
...
bar1.await();
...
```

RXO Mutant for Barriers:

```
final CyclicBarrier bar1
    = new CyclicBarrier(20,
                new Runnable() {...});
final CyclicBarrier bar2
    = new CyclicBarrier(20,
                new Runnable() {...});
...
//should be call to bar1.await()
bar2.await();
...
```

5.4.2 EELO - Exchange Explicit Lock Object

We have already seen the exchanging of two implicit lock objects in a synchronized block and the potential for deadlock (Section 5.1.3). The EELO operator is identical only it exchanges two explicit lock object instances:

Original Code:

```
private Lock lock1 = new ReentrantLock();
private Lock lock2 = new ReentrantLock();
...
lock1.lock();
...
lock2.lock();
...
finally{
   lock2.unlock();
}
...
finally{
   lock1.unlock();
}
...
```

EELO Mutant:

```
private Lock lock1 = new ReentrantLock();
private Lock lock2 = new ReentrantLock();
...
lock2.lock();
...
lock1.lock();
...
finally{
   lock2.unlock();
}
...
finally{
   lock1.unlock();
}
...
```

5.5 Modify critical region : shift, expand, shrink and split

The modify critical region operators cause the modification of the critical region by moving statements both inside and outside the region and by dividing the region into multiple regions.

5.5.1 SHCR - Shift Critical Region

Shifting a critical region up or down can potentially cause interference bugs by no longer synchronizing access to a shared variable. An example of shifting a synchronized block up is provided below. The SHCR operator can also be applied to shift up or down critical regions using other concurrency mechanisms.

Original Code:

```
<statement n1>
<statement n2>
synchronized (this){
    //critical region
    <statement c1>
    <statement c2>
}
<statement n3>
<statement n4>
...
```

SHCR Mutant:

```
<statement n1>
<statement n2>
//critical region
<statement c1>
synchronized (this){
    <statement c2>
    <statement n3>
}
<statement n4>
...
```

The SHCR operator can also be used to shift the critical region created by an explicit lock. For example:

Original Code:

```
private Lock lock1 = new ReentrantLock();
...
public void m1 (){
<statement n1>
<statement n2>
lock1.lock();
try{
    //critical region
    <statement c1>
    <statement c2>
} finally{
    lock1.unlock();
}
<statement n3>
...
```

SHCR Mutant:

```
private Lock lock1 = new ReentrantLock();
...
public void m1 (){
<statement n1>
lock1.lock();
try{
    <statement n2>
    //critical region
    <statement c1>
} finally{
    lock1.unlock();
}
<statement c2>
<statement n3>
...
```

5.5.2 EXCR - Expand Critical Region

Expanding a critical region to include statements above and below the statements required to be in the critical region can cause performance issues by unnecessarily reducing the degree of concurrency. For example:

Original Code:

```
<statement n1>
<statement n2>
synchronized (this){
    //critical region
    <statement c1>
    <statement c2>
}
<statement n3>
<statement n4>
...
```

EXCR Mutant:

```
<statement n1>
synchronized (this){
    <statement n2>
    //critical region
    <statement c1>
    <statement c2>
    <statement n3>
}
<statement n4>
...
```

The EXCR operator can also cause correctness issues and consequences such as deadlock when an expanded critical region overlaps with or subsumes another critical region.

5.5.3 SKCR - Shrink Critical Region

Shrinking a critical region will have similar consequences (interference) to shifting a region since both the SHCR and SKCR operators move statements that require synchronization outside the critical section. Below we provide an example of the SKCR operator using a Lock.

Original Code:

```
private Lock lock1 = new ReentrantLock();
...
public void m1 (){
<statement n1>
lock1.lock();
try{
    //critical region
    <statement c1>
    <statement c2>
    <statement c3>
} finally{
    lock1.unlock();
}
<statement n2>
...
```

SKCR Mutant:

```
private Lock lock1 = new ReentrantLock();
...
public void m1 (){
<statement n1>
//critical region
<statement c1>
lock1.lock();
try{
    <statement c2>
} finally{
    lock1.unlock();
}
<statement c3>
<statement n2>
...
```

5.5.4 SPCR - Split Critical Region

Unlike the SHCR or SKCR operators, splitting a critical region into two regions will not cause statements to move outside of the critical region. However, the consequences of splitting a critical region into 2 regions is potentially just as serious since a split may cause a set of statements that were meant to be atomic to be nonatomic. For example, in between the two split critical regions another thread might be able to acquire the lock for the region and modify the value of shared variables before the second half of the old critical region is executed.

Original Code:

```
<statement n1>
synchronized (this){
    //critical region
    <statement c1>
    <statement c2>
}
<statement n2>
...
```

SPCR Mutant:

```
<statement n1>
synchronized (this){
    //critical region
    <statement c1>
}
synchronized (this){
    <statement c2>
}
<statement n2>
...
```

5.6 Summary

In the above subsections we have provided an overview of concurrency mutation operators for Java (J2SE 5.0). In our discussion of each operator we have briefly mentioned the bug pattern that relates to that operator. Table 4 provides a summary of this relationship and shows that the operators we propose are examples of real bug patterns. Overall almost all of the bug patterns are covered by the operators demonstrating that the proposed concurrency operators are not only representative but provide good coverage. The bug patterns that do not have mutation operators are typically more specific complex patterns and the development of general operators related to these patterns is not feasible.

Concurrency Bug Pattern	Mutation Operators
Nonatomic operations assumed to be atomic bug pattern	RVK, EAN
Two-stage access bug pattern	SPCR
Wrong lock or no lock bug pattern	MSP, ESP, EELO, SHCR, SKCR, EXCR, RSB, RSK, ASTK, RSTK, RCXC, RXO
Double-checked locking bug pattern	–
The sleep() bug pattern	MXT, RJS, RTXC
Losing a notify bug pattern	RTXC, RCXC
Notify instead of notify all bug pattern	RNA
Other missing or nonexistent signals bug pattern	MXC, MBR, RCXC
A “blocking” critical section bug pattern	RFU, RCXC
The orphaned thread bug pattern	–
The interference bug pattern	MXT, RTXC, RCXC
The deadlock (deadly embrace) bug pattern	ESP, EXCR, EELO, RXO, ASK
Starvation bug pattern	MSF, ELPA
Resource exhaustion bug pattern	MXC
Incorrect count initialization bug pattern	MXC

Table 4: Concurrency bug patterns vs. concurrency mutation operators

6 Conclusion

We have presented a set of concurrency mutation operators to be used as a metric in the comparison of different test suites and testing strategies for concurrent Java as well as different quality assurance tools for concurrency. Although we are primarily interested in concurrent mutation operators as comparative metrics we believe that these operators can also serve a role similar to method and class level mutation operators as both comparative metrics and coverage criteria. Our new concurrency operators should be viewed as a complement not a replacement for the existing operators used in tools like muJava. For example, using the concurrency operators can cause direct concurrency bugs while using the method and class level operators can cause indirect concurrency bugs.

We believe that our concurrency operators are comprehensive and representative of real bugs. We have justified the operators by comparing them to a set of bug patterns that have been used to identify real bugs in concurrent Java programs. Additionally, our classification of concurrency operators shows that the operators are well distributed across the majority of bug patterns.

Currently we are implementing our concurrency mutation operators in a source transformation language TXL [CDMS02]. Upon completion of our implementation we plan to validate the operators with our mutation analysis framework ExMAn [BCD06]. We are interested in using our concurrency operators with the programs in the IBM concurrency benchmark to compare concurrency testing and model checking.

References

[AHB03] C. Artho, K. Havelund, and A. Biere. High-level data races. In *The First International Workshop on Verification and Validation of Enterprise Information Systems (VVEIS'03)*, Apr. 2003.

[BCD06] Jeremy S. Bradbury, James R. Cordy, and Juergen Dingel. ExMAn: A generic and customizable framework for experimental mutation analysis. In *Proc. of the 2nd Workshop on Mutation Analysis (Mutation 2006)*, Nov. 2006.

[CDMS02] James R. Cordy, Thomas R. Dean, Andrew J. Malton, and Kevin A. Schneider. Source transformation in software engineering using the TXL transformation system. *J. of Information and Software Technology*, 44(13):827–837, 2002.

[EU04] Yaniv Eytani and Shmuel Ur. Compiling a benchmark of documented multi-threaded bugs. In *Workshop on Parallel and Distributed Testing and Debugging (PADTAD 2004), proceedings of IPDPS*, page 266a. IEEE Computer Society, 2004.

[FNU03] Eitan Farchi, Yarden Nir, and Shmuel Ur. Concurrent bug patterns and how to test them. In *IPDPS '03: Proceedings of the 17th International Symposium on Parallel and Distributed Processing*, page 286.2, Washington, DC, USA, 2003. IEEE Computer Society.

[Gho02] S. Ghosh. Towards measurement of testability of concurrent object-oriented programs using fault insertion: a preliminary investigation. In *Proc. of the* 2^{nd} *IEEE International Workshop on Source Code Analysis and Manipulation (SCAM 2002)*, pages 17–25, 2002.

[KO91] K. N. King and A. Jefferson Offutt. A Fortran language system for mutation-based software testing. *Softw. Pract. Exper.*, 21(7):685–718, 1991.

[LDG$^+$04] Brad Long, Roger Duke, Doug Goldson, Paul A. Strooper, and Luke Wildman. Mutation-based exploration of a method for verifying concurrent Java components. In *Proc. of Workshop on Parallel and Distributed Systems: Testing and Debugging (PADTAD 2004)*, Apr. 2004.

[Lea00] Doug Lea. *Concurrent Programming in Java*$_{TM}$ *Second Edition*. Addison Wesley, 2000.

[Lee06] E.A. Lee. The problem with threads. *Computer*, 39(5):33– 42, May 2006.

[LSW05] Brad Long, Paul Strooper, and Luke Wildman. A method for verifying concurrent java components based on an analysis of concurrency failures. In *Concurrency Computat.: Pract. Exper.*, volume (submitted), pages 1–7, 2005.

[MKO02] Yu-Seung Ma, Yong-Rae Kwon, and Jeff Offutt. Inter-class mutation operators for Java. In *Proc. of the* 13^{th} *International Symposium on Software Reliability Engineering*, pages 352–363. IEEE Computer Society Press, Nov. 2002.

[MOK05] Yu-Seung Ma, Jeff Offutt, and Yong-Rae Kwon. MuJava : An automated class mutation system. *Journal of Software Testing, Verification and Reliability*, 15(2):97–133, Jun. 2005.

[NA06] Akbar S. Namin and James Andrews. Finding sufficient mutation operators via variable reduction. In *Proc. of* 2^{nd} *Workshop on Mutation Analysis*, Nov. 2006.

[OLR$^+$96] A. Jefferson Offutt, Ammei Lee, Gregg Rothermel, Roland H. Untch, and Christian Zapf. An experimental determination of sufficient mutant operators. *ACM Trans. Softw. Eng. Methodol.*, 5(2):99–118, 1996.

[OMK04] Jeff Offutt, Yu-Seung Ma, and Yong-Rae Kwon. An experimental mutation system for Java. In *Proc. of the Workshop on Empirical Research in Software Testing (WERST'2004)*. SIGSOFT Software Engineering Notes, 29(5):1–4, ACM Press, 2004.

[SL05] Herb Sutter and James Larus. Software and the concurrency revolution. *Queue*, 3(7):54–62, 2005.