



Data-Race Free Java



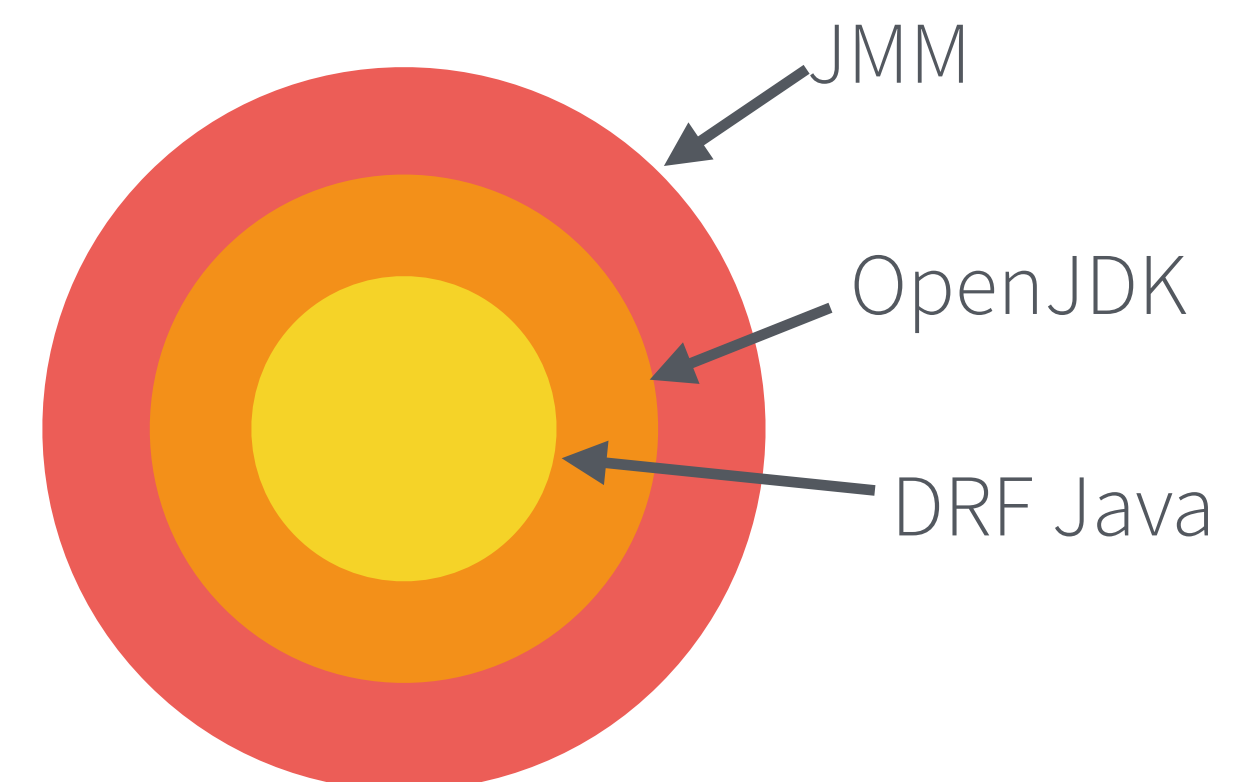
Tobias Wrigstad, PLISS 2026

An Introduction to Data-Race Free Java

- What are we doing: **reinterpreting the JMM** (— changing the way Java programs execute!)
 1. We want to make it **simpler** to write a **correct** concurrent program
 2. We want to **data-races** to throw **exceptions**
 3. We want to **remove behaviour** not fully defined by the JMM
- Our goal is to be backwards-compatible

Correctly synchronised programs should **behave exactly the same** as they do in Java today

All **incorrectly** synchronised programs should exhibit behaviour that is **permitted** by the **current JMM**



The core idea: enforced thread isolation + communication

- When a thread starts, it gets its own **private snapshot** of the current system

There is **no shared state**

- Synchronisation actions in the JMM are reinterpreted as **communication** with other threads

Taking a lock means **updating the private snapshot** with changes **published by other threads**

Releasing a lock means **publishing private changes** so that others **can update** their snapshots with these change

- Thread **isolation** gives us structural **freedom from read-write races**

Communication can detect **write-write races**


Write-write races are detected at **well-defined points in time**: synchronisation actions

Leakage

Very simple example:

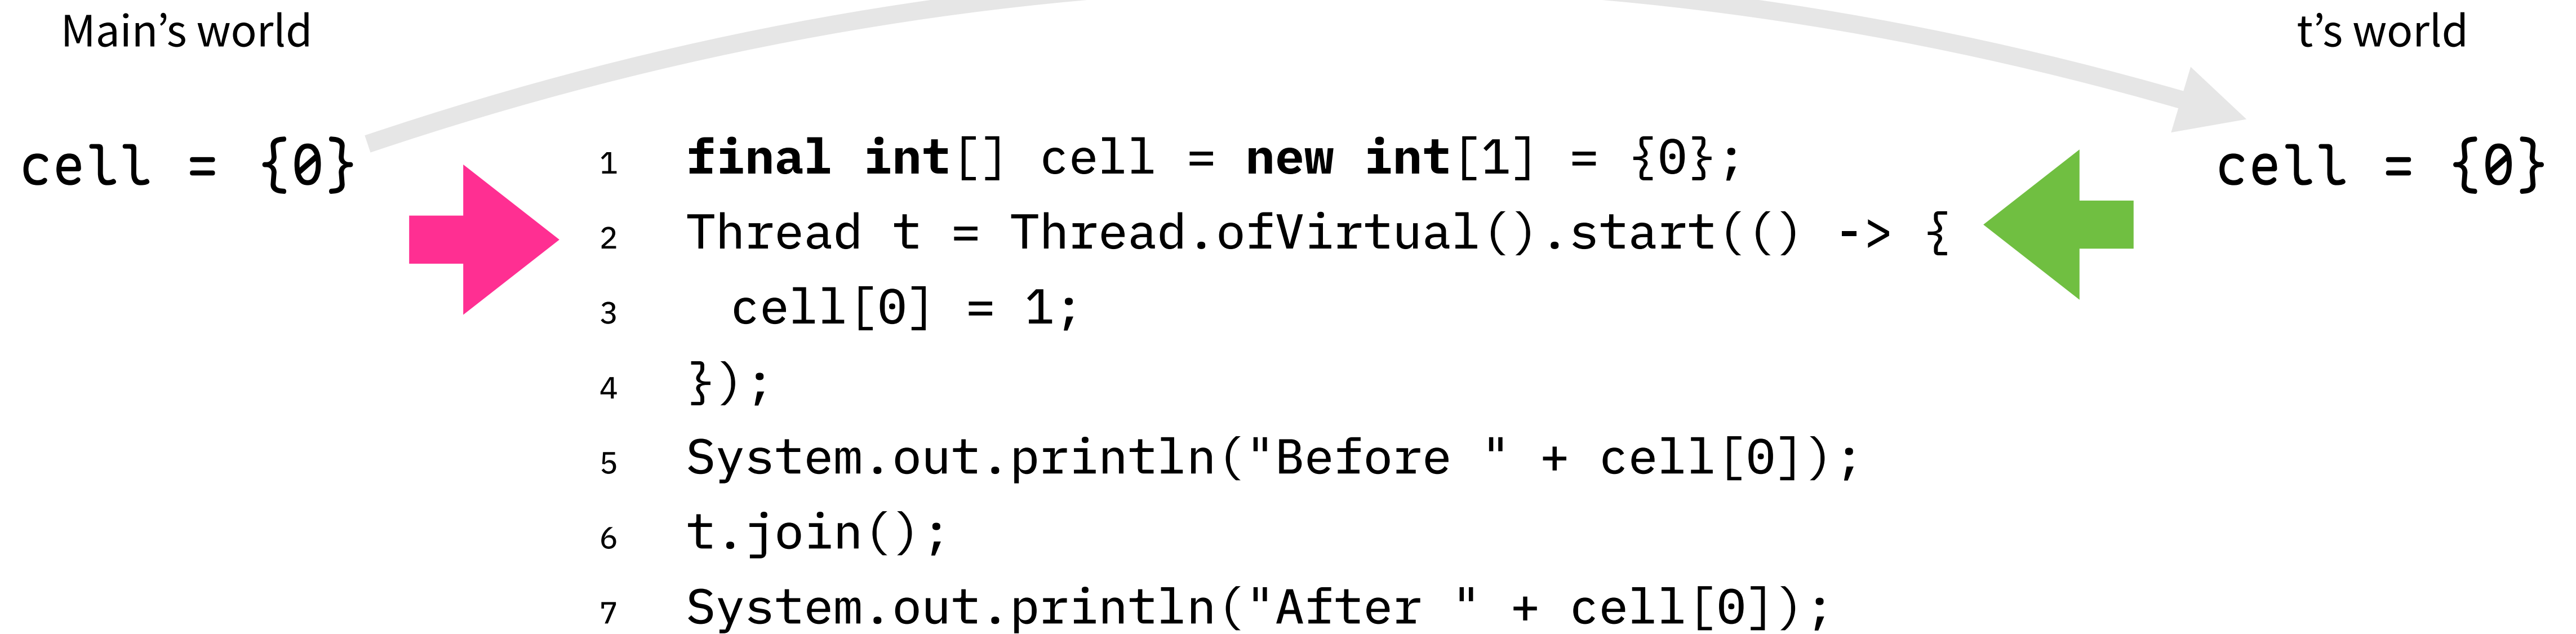
Main's world

cell = {0}



```
1 final int[] cell = new int[1] = {0};  
2 Thread t = Thread.ofVirtual().start(() -> {  
3     cell[0] = 1;  
4 });  
5 System.out.println("Before " + cell[0]);  
6 t.join();  
7 System.out.println("After " + cell[0]);
```


Very simple example:




Very simple example:

Main's world

cell = {0}



```
1 final int[] cell = new int[1] = {0};
2 Thread t = Thread.ofVirtual().start(() -> {
3     cell[0] = 1;
4 });
5 System.out.println("Before " + cell[0]);
6 t.join();
7 System.out.println("After " + cell[0]);
```




t's world

cell = {1}


Very simple example:

Main's world

cell = {0}



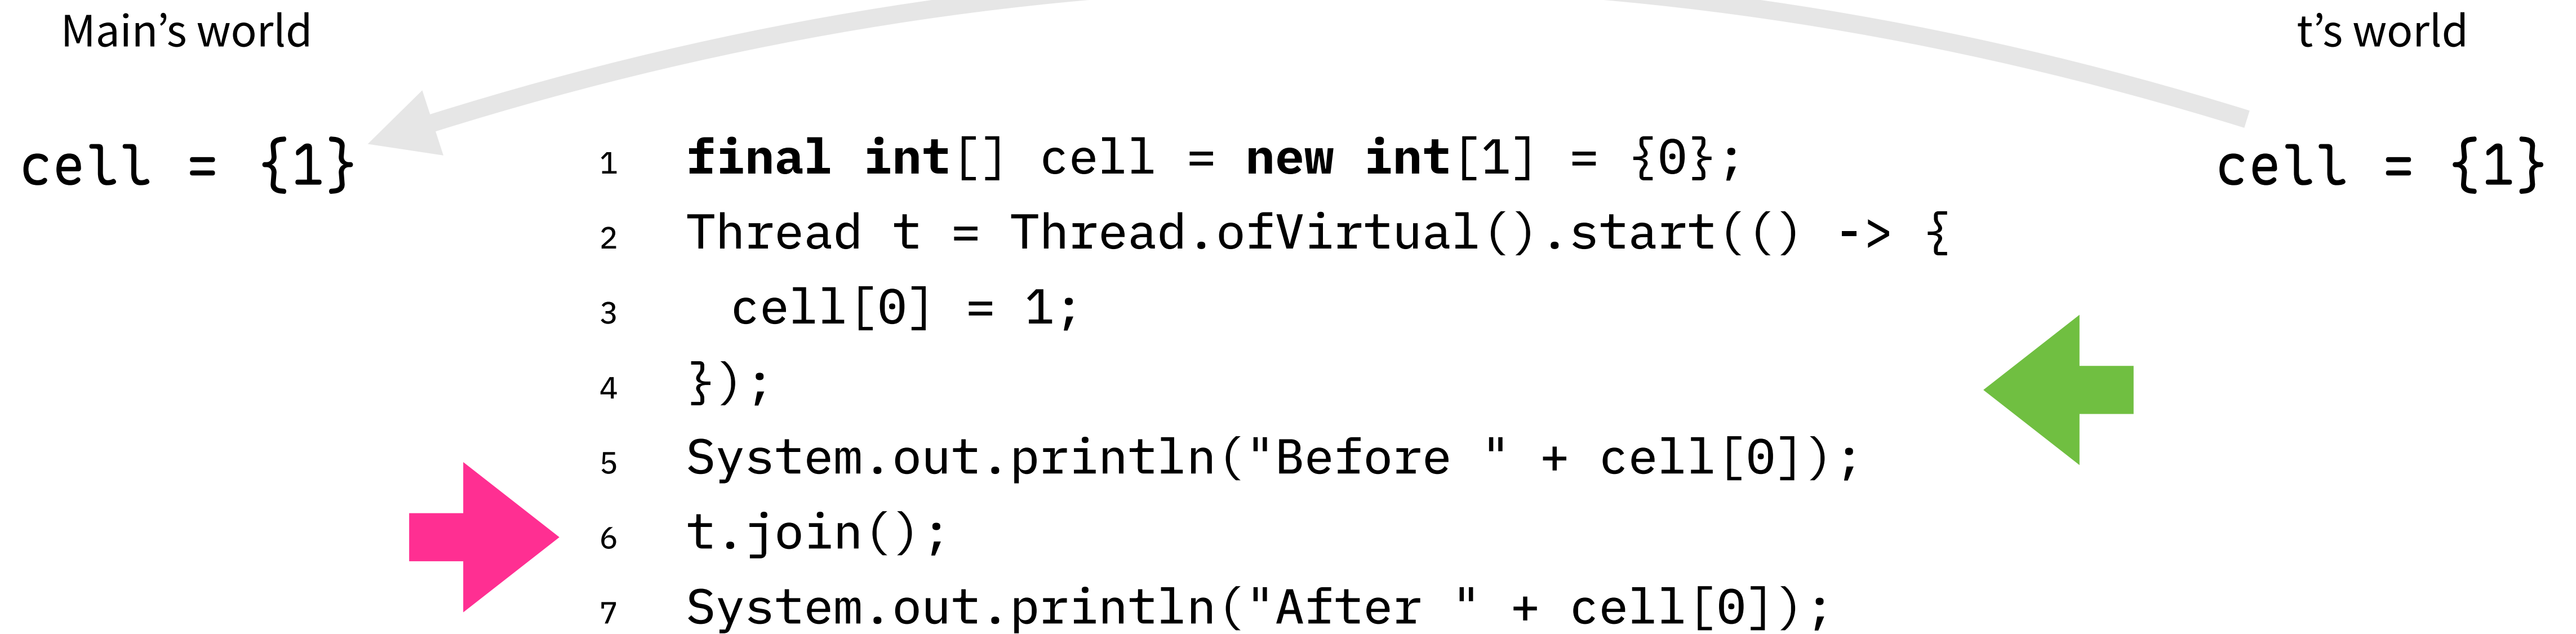
```
1 final int[] cell = new int[1] = {0};
2 Thread t = Thread.ofVirtual().start(() -> {
3     cell[0] = 1;
4 });
5 System.out.println("Before " + cell[0]);
6 t.join();
7 System.out.println("After " + cell[0]);
```



t's world

cell = {1}

Very simple example:




Very simple example:

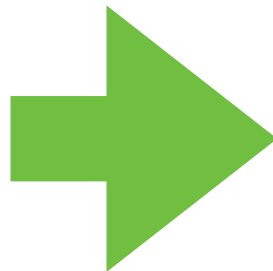
Main's world

```
cell = {1}
```

```
1  final int[] cell = new int[1] = {0};  
2  Thread t = Thread.ofVirtual().start(() -> {  
3    cell[0] = 1;  
4  });  
5  System.out.println("Before " + cell[0]);  
6  t.join();  
7  System.out.println("After " + cell[0]);
```



Simpler to write a correct concurrent program



```

1  final class Logger {
2    private static Logger instance;
3    private String path;
4    private Logger(String path) {
5      this.path = path;
6    }
7
8    public static Logger getLogger() {
9      if (instance == null) {
10     synchronized (Logger.class) {
11       if (instance == null) {
12         instance = new Logger("f.log");
13       }
14     }
15     return instance;
16   }

```

Original →

```

1  instance =
2    new Logger("f.log");

```


After inlining →

```

1  var tmp = new Logger();
2  tmp.path = "f.log";
3  instance = tmp;

```

After optimisation →



```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

```

1  final class Logger {
2    private static Logger instance;
3    private String path;
4    private Logger(String path) {
5      this.path = path;
6    }
7
8    public static Logger getLogger() {
9      if (instance == null) {
10     synchronized (Logger.class) {
11       if (instance == null) {
12         instance = new Logger("f.log");
13       }
14     }
15     return instance;
16   }

```

Original →

```

1  instance =
2    new Logger("f.log");

```

After inlining →

```

1  var tmp = new Logger();
2  tmp.path = "f.log";
3  instance = tmp;


```

After optimisation →

```

1  instance = new Logger();
2  instance.path = "f.log";

```



Fixing the program

```

1  final class Logger {
2    private static Logger instance; volatile
3    private String path; final
4    private Logger(String path) {
5      this.path = path;
6    }
7
8    public static Logger getLogger() {
9      if (instance == null) {
10     synchronized (Logger.class) {
11       if (instance == null) {
12         instance = new Logger("f.log");
13       }
14     }
15     return instance;
16   }

```

Original →

```

1  instance =
2    new Logger("f.log");

```

After inlining →

```

1  var tmp = new Logger();
2  tmp.path = "f.log";
3  instance = tmp;

```

After optimisation →


```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

`instance = null`



```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15         }
16     }

```

`instance = null`

After optimisation →


```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

`instance = null`



```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15         }
16     }

```

`instance = null`

After optimisation →


```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = null



```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15             return instance;
16         }
17     }
18 }

```

instance = null

After optimisation →

```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = null




```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15             return instance;
16         }
17     }
18 }

```

instance = null



After optimisation →



```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = 0x...




```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15         }
16     }

```

instance = null



After optimisation →



```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = 0x...




```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15             return instance;
16         }
17     }
18 }

```

instance = null



After optimisation →



```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = 0x...




```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15             return instance;
16         }
17     }
18 }

```

instance = null



After optimisation →



```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = 0x...




```


1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15             return instance;
16         }
17     }
18 }

```

instance = null



After optimisation →



```

1  instance = new Logger();
2  instance.path = "f.log";

```


Simpler to write a correct concurrent program

instance = 0x...


```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15         }
16     }

```



instance = null



After optimisation →

```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program



instance = 0x...

```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15         }
16     }

```

instance = null



After optimisation →

```

1  instance = new Logger();
2  instance.path = "f.log";

```

Simpler to write a correct concurrent program

instance = 0x...



```

1  final class Logger {
2      private static Logger instance;
3      private String path;
4      private Logger(String path) {
5          this.path = path;
6      }
7
8      public static Logger getLogger() {
9          if (instance == null) {
10             synchronized (Logger.class) {
11                 if (instance == null) {
12                     instance = new Logger("f.log");
13                 }
14             }
15         }
16     }

```



instance = 0x...



After optimisation →

```

1  instance = new Logger();
2  instance.path = "f.log";

```

What DRF Java Is and Is Not

DRF Java is not "just" dynamic data-race detection

Structurally guaranteed to be free from read-write races

Several programs that are considered incorrectly synchronised under the JMM are correctly synchronised in DRF Java

– As exemplified by the singleton example (double-checked locking without volatiles)

DRF Java is not software transactional memory

No imposed nesting

We can do e.g. hand-over-hand locking

There is no rollback

...

Semantic Overview

We are exploring two different models in parallel that differ in how threads communicate changes with each other

Linear model

All synchronisation interacts with a "global store"

Faster convergence — aka fails faster

More leakage

Branching model

Locks, volatile fields, and threads have their own stores

Permits more divergence

Less leakage


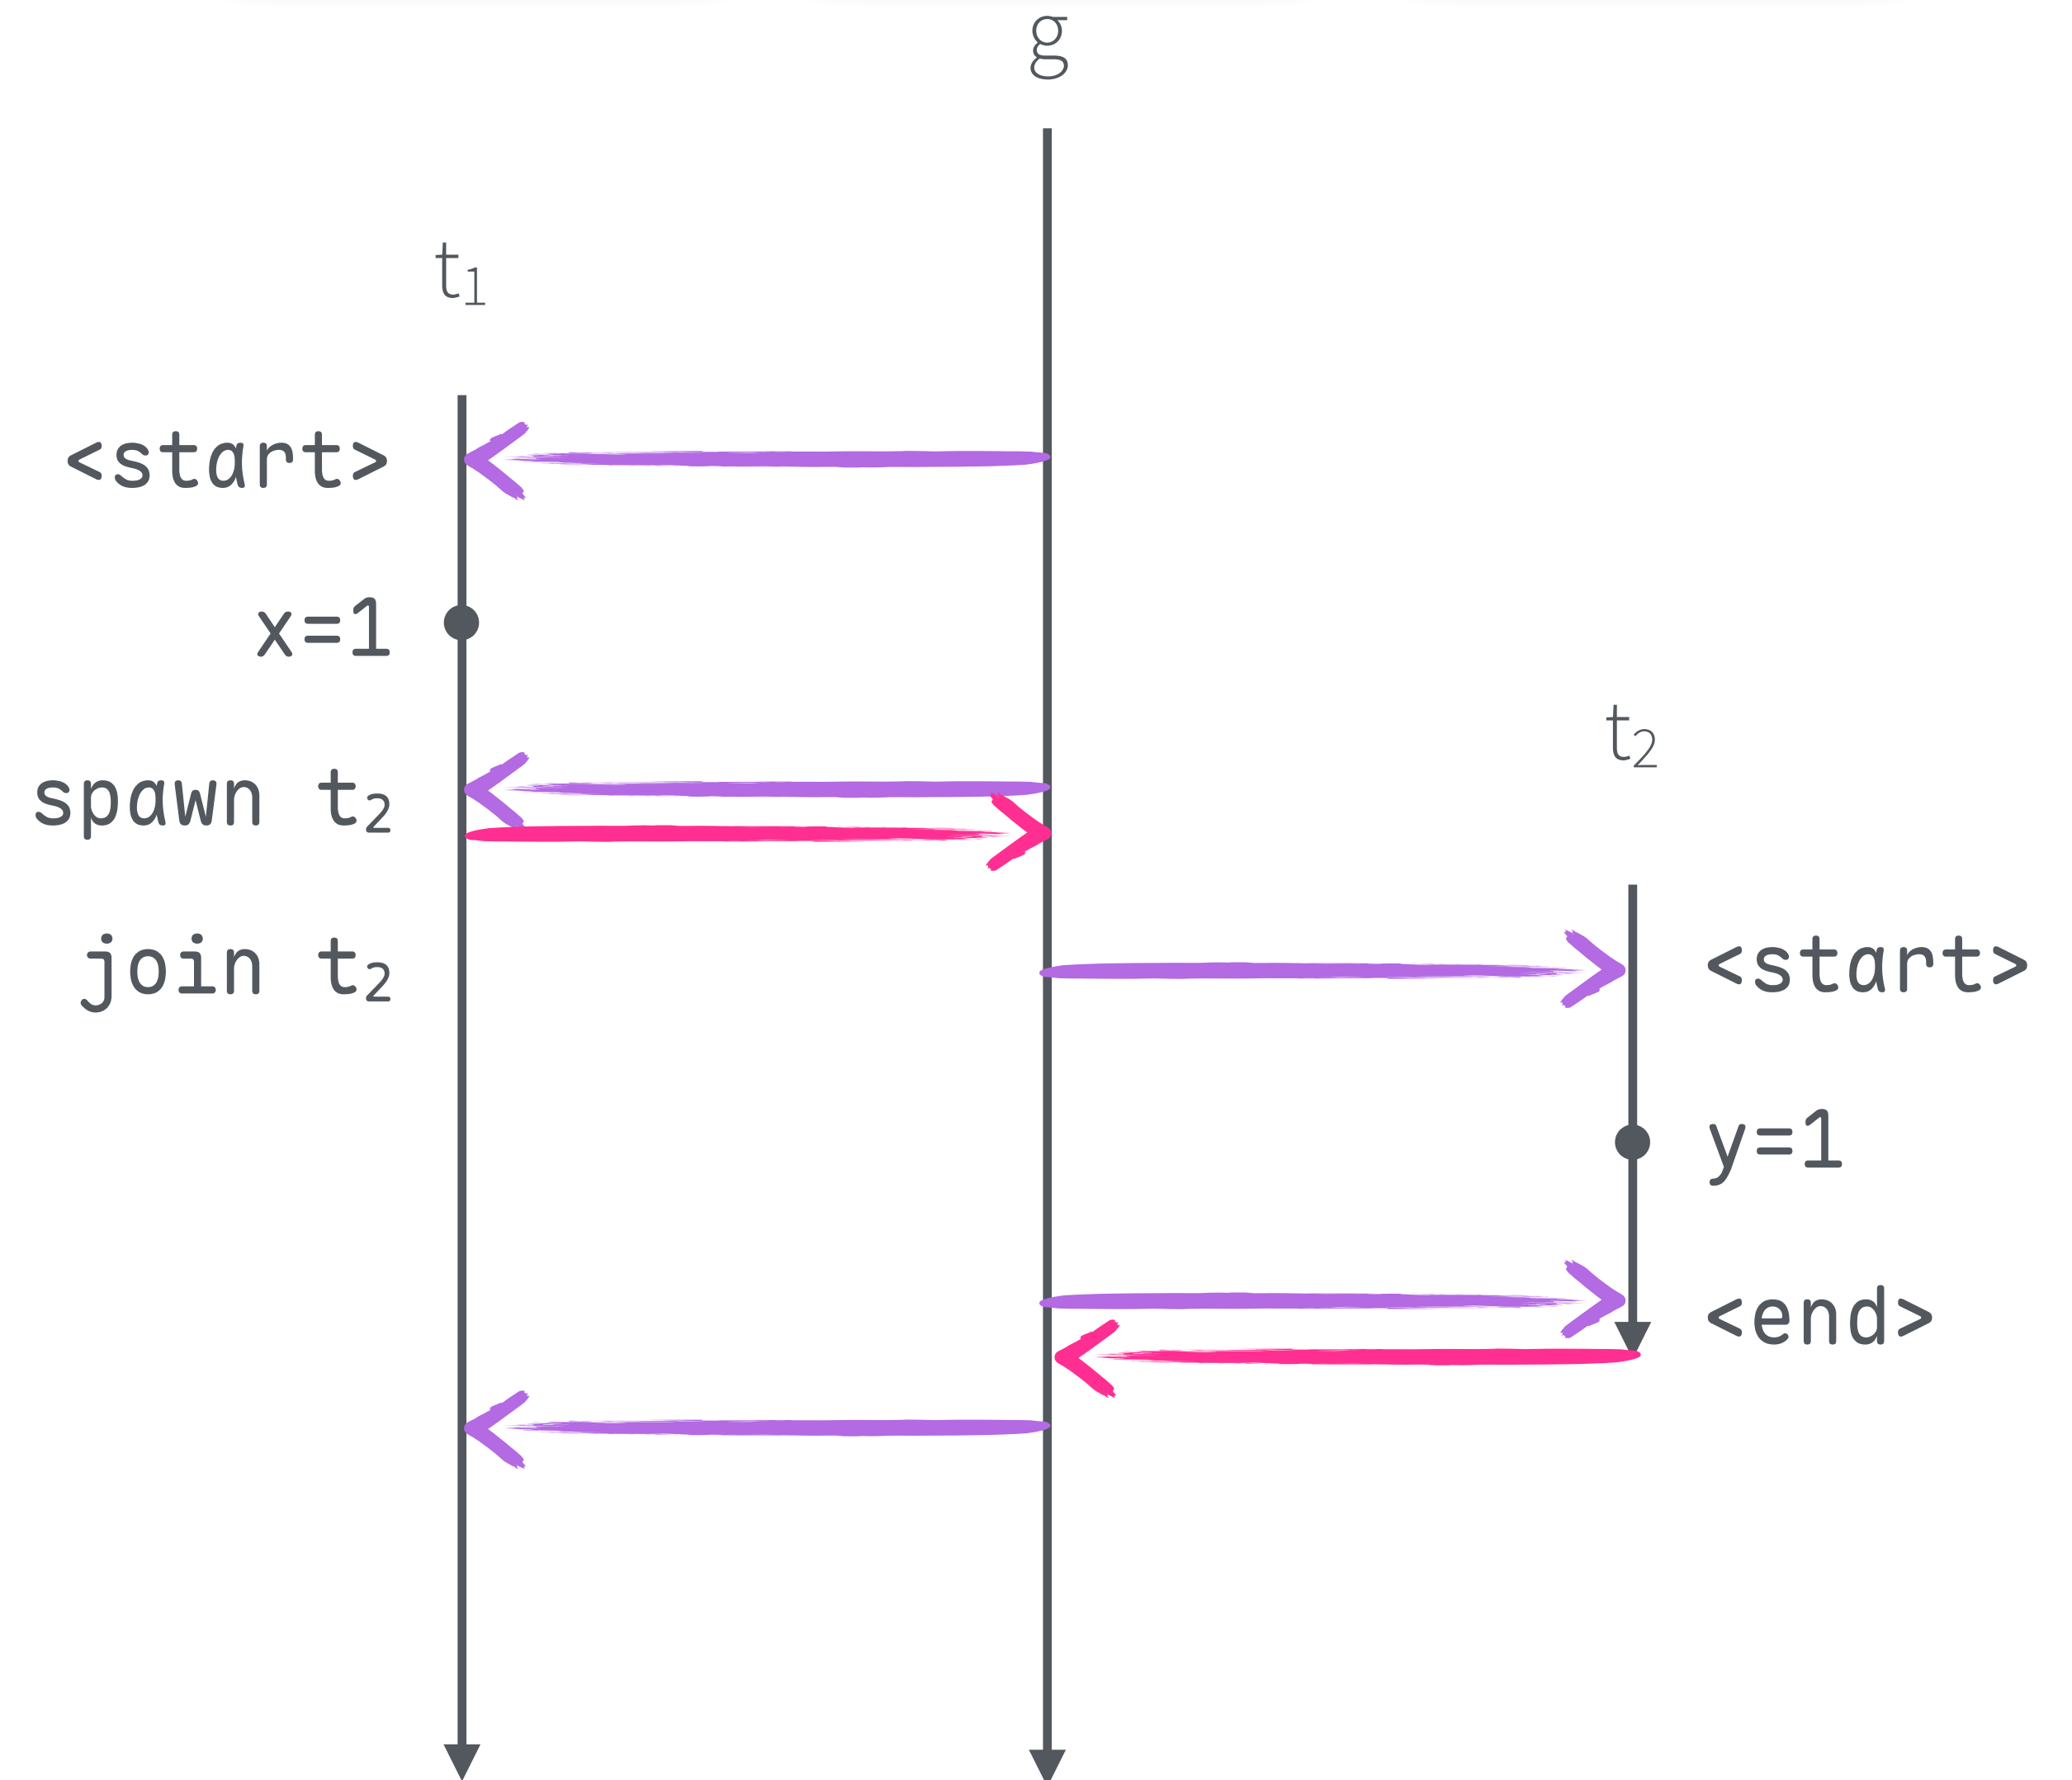
(Ignoring thread interrupt for now)

Sync. Action		Re-interpretation
Object	Action	
thread	spawn t' join t'	
lock	lock l unlock l	
volatile	read v write v	

Sync. Action	Re-interpretation
start	
end	

Diagrammatic Notation Primer



t_1	t_2
<pre><start> x = 1 spawn t2 join t2</pre>	<pre><start> y = 1 <end></pre>



Detecting Conflicts

t_1	t_2
<start> $x = 1$ <end>	<start> $x = 1$ <end>

Counter
3

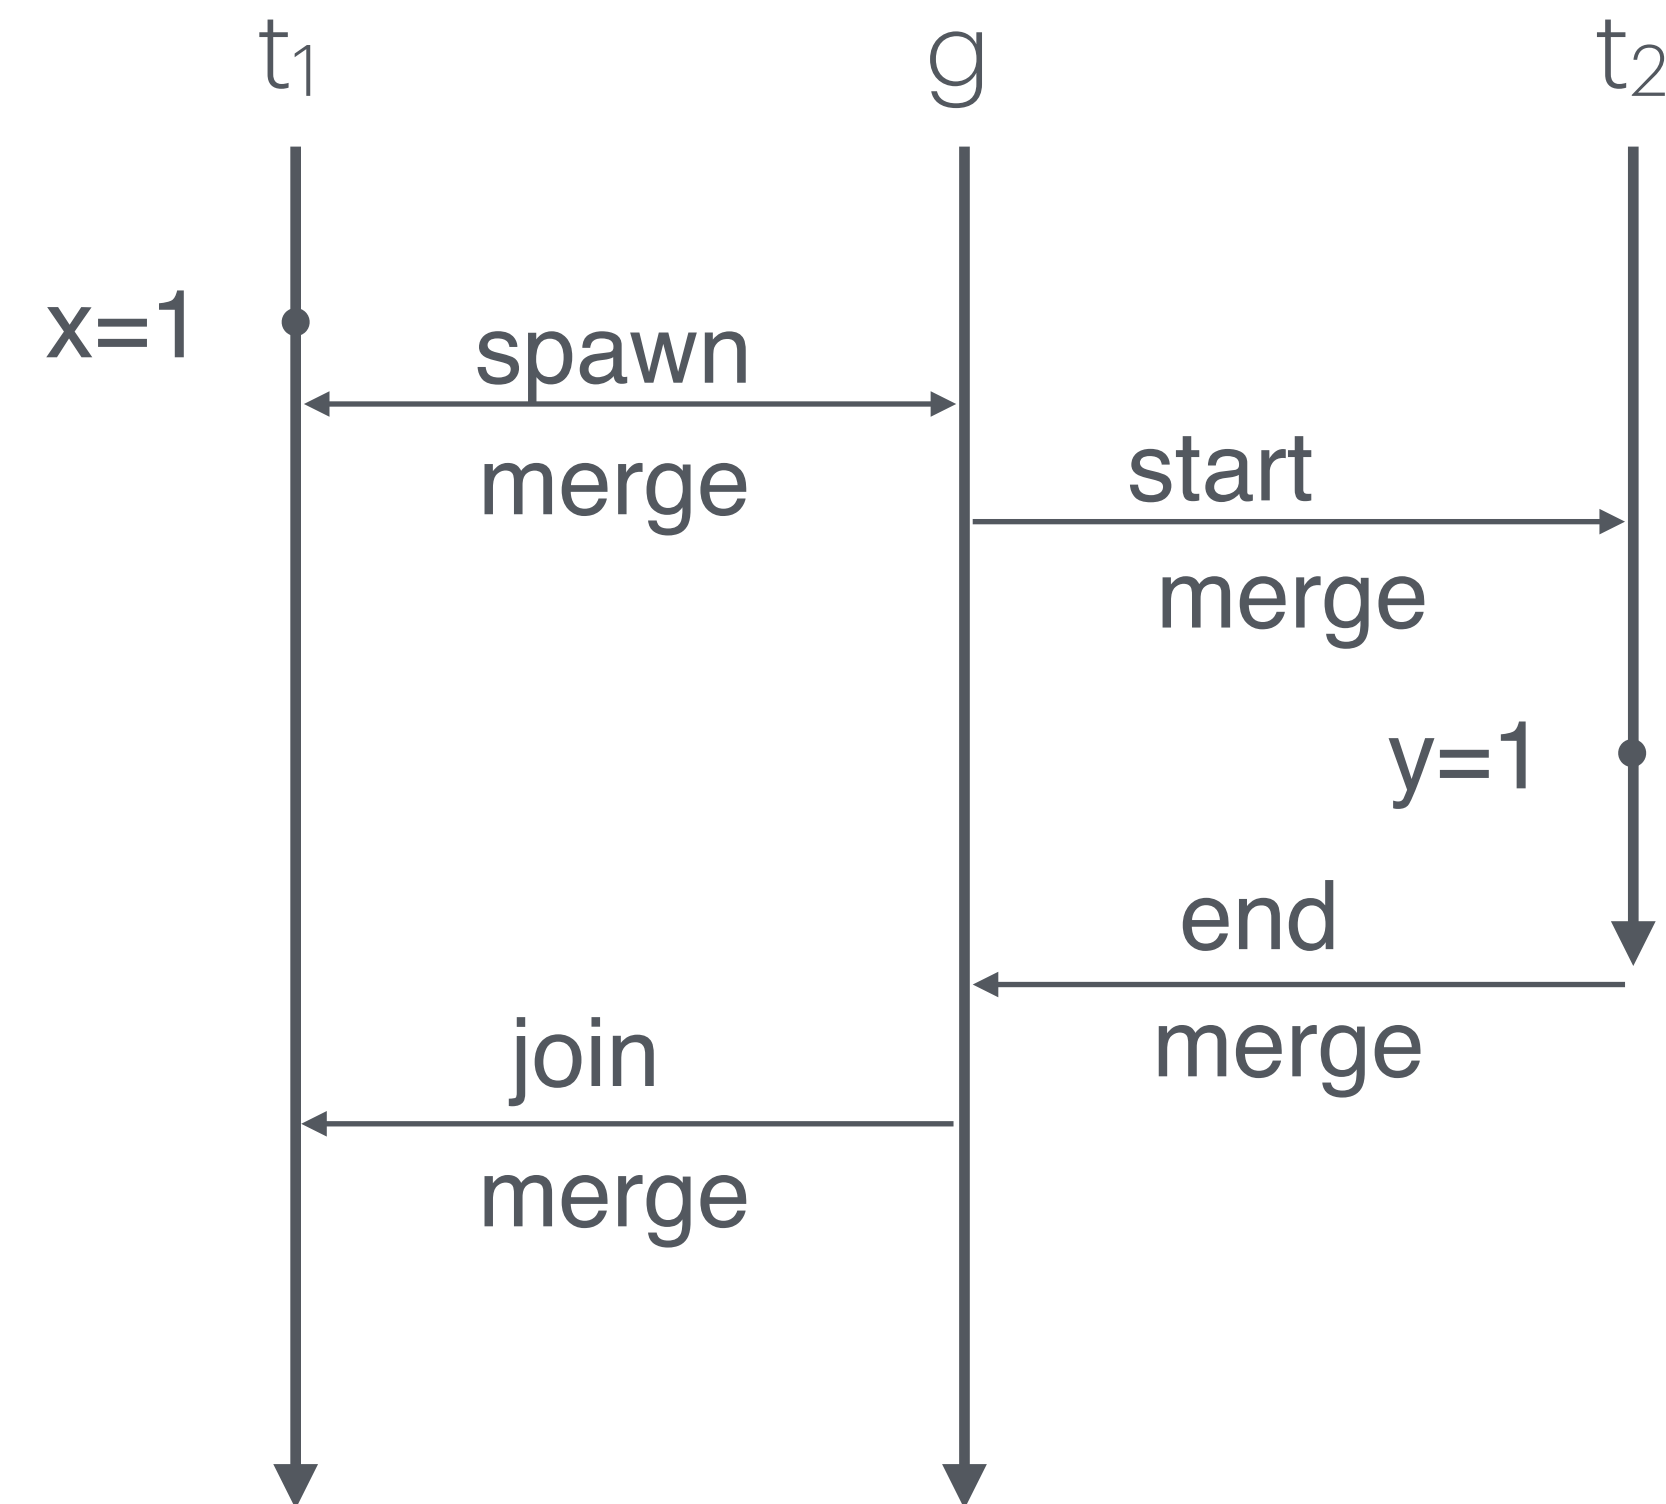


X 2 \neq 1 **X**

Example 1

Linear — correctly synchronised, changes different parts of the state

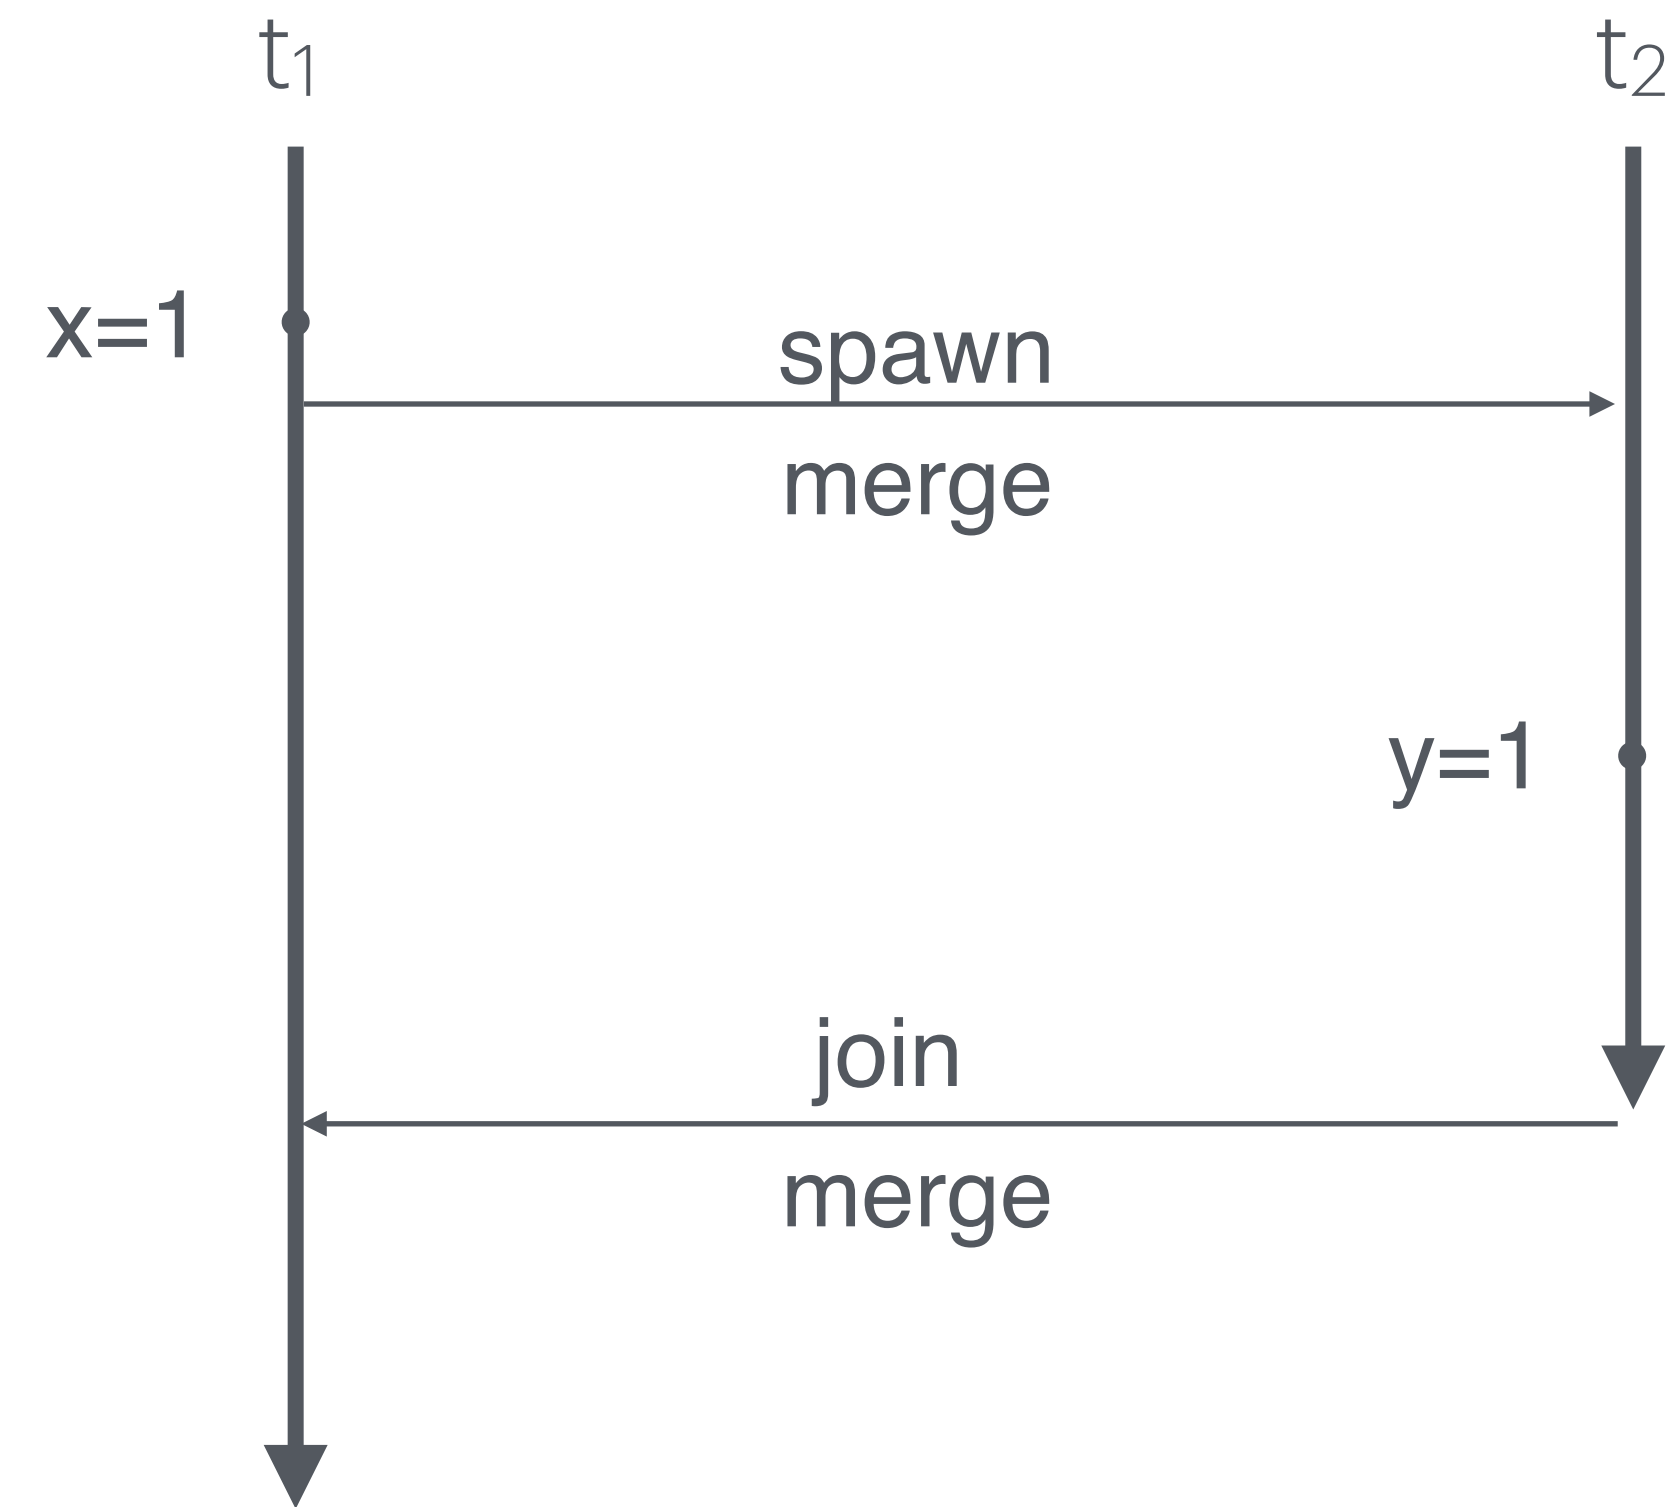
t_1	t_2
$x = 1$ spawn t_2 join t_2	$y = 1$



Example 1

Branching — correctly synchronised, changes different parts of the state

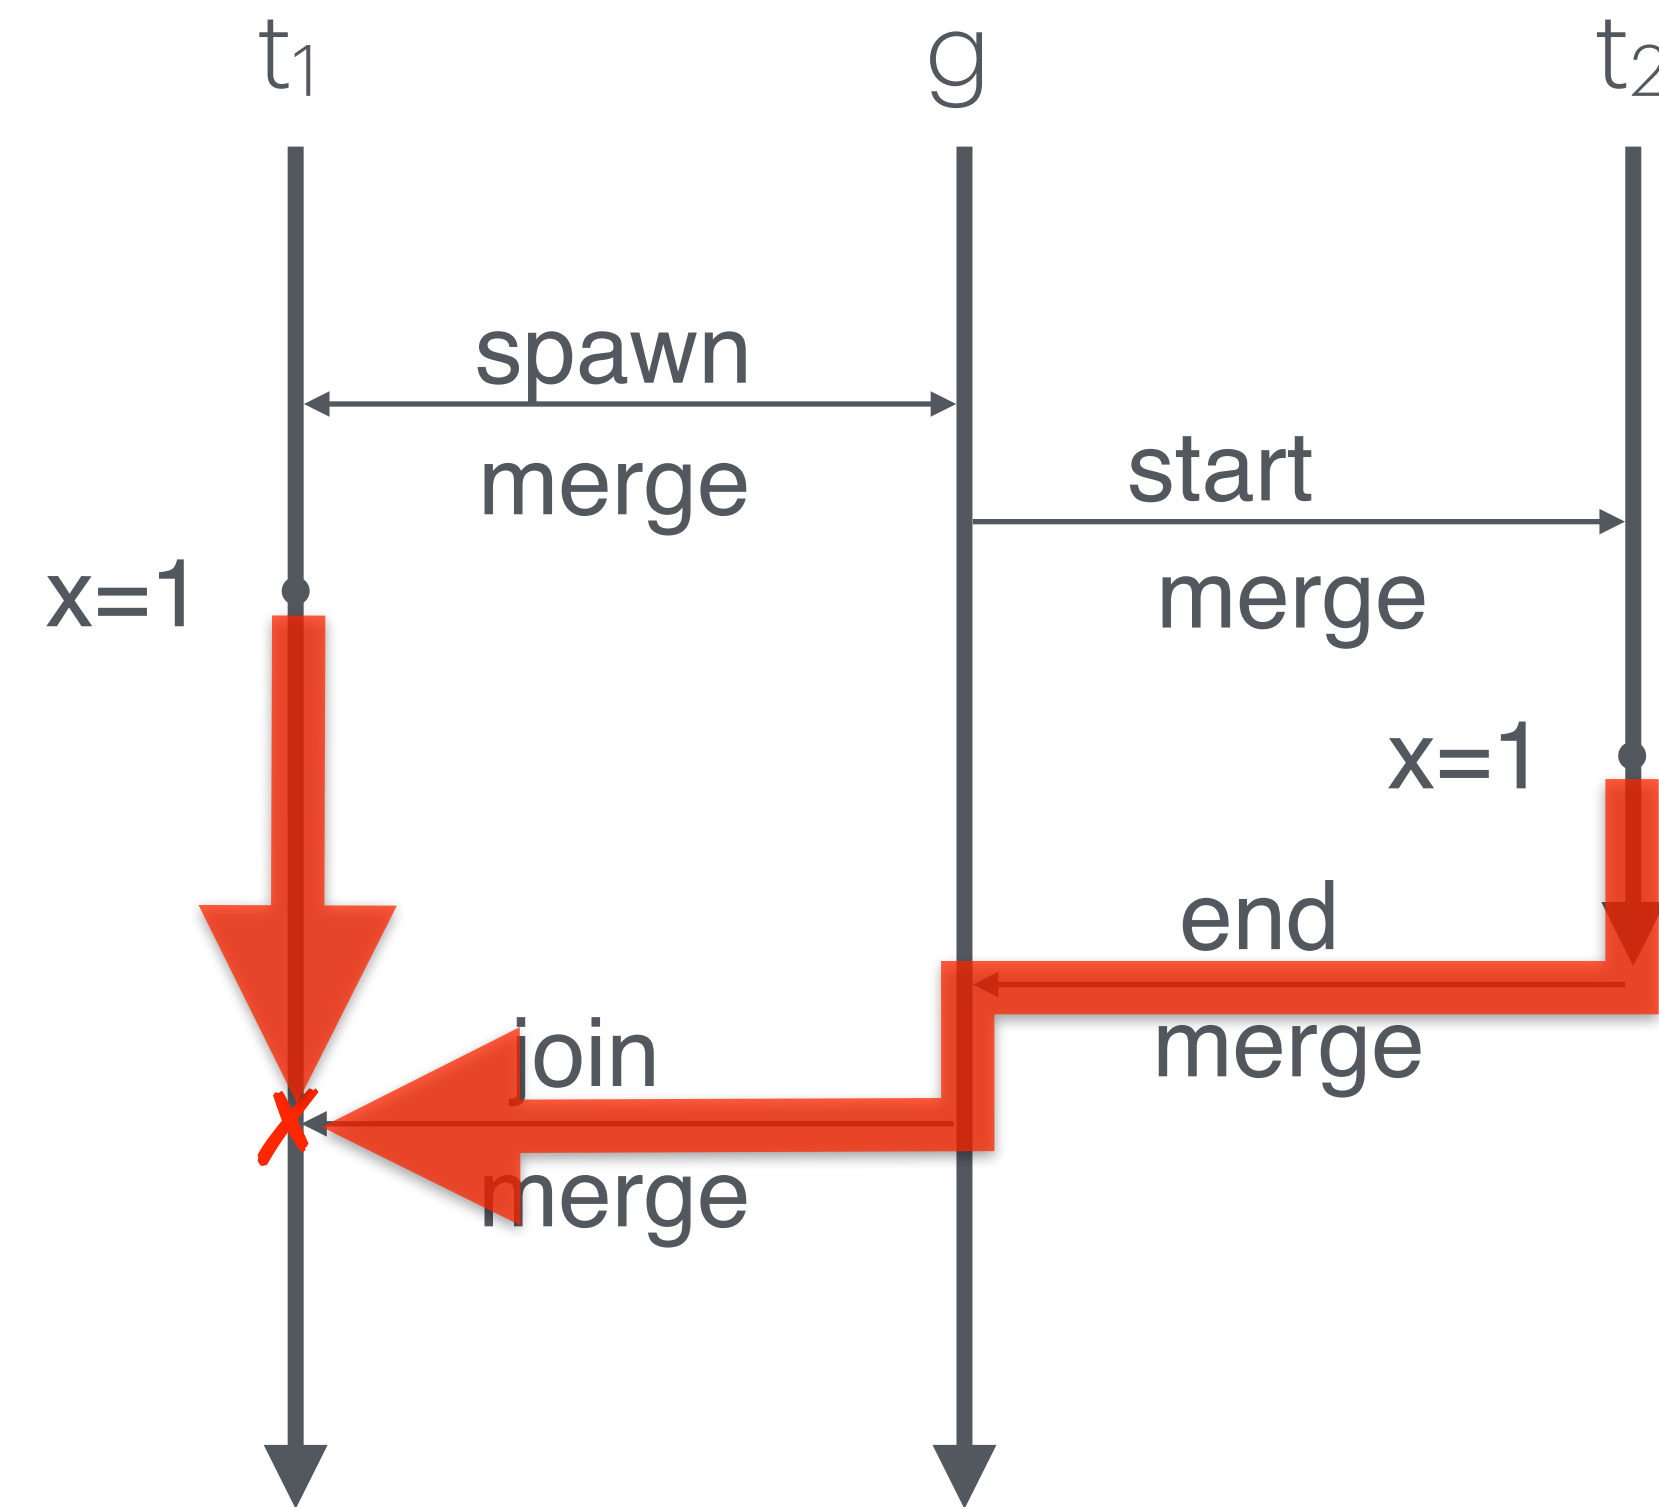
t_1	t_2
$x = 1$ spawn t_2 join t_2	$y = 1$



Example 2

Linear — **incorrectly** synchronised, write-write race on x

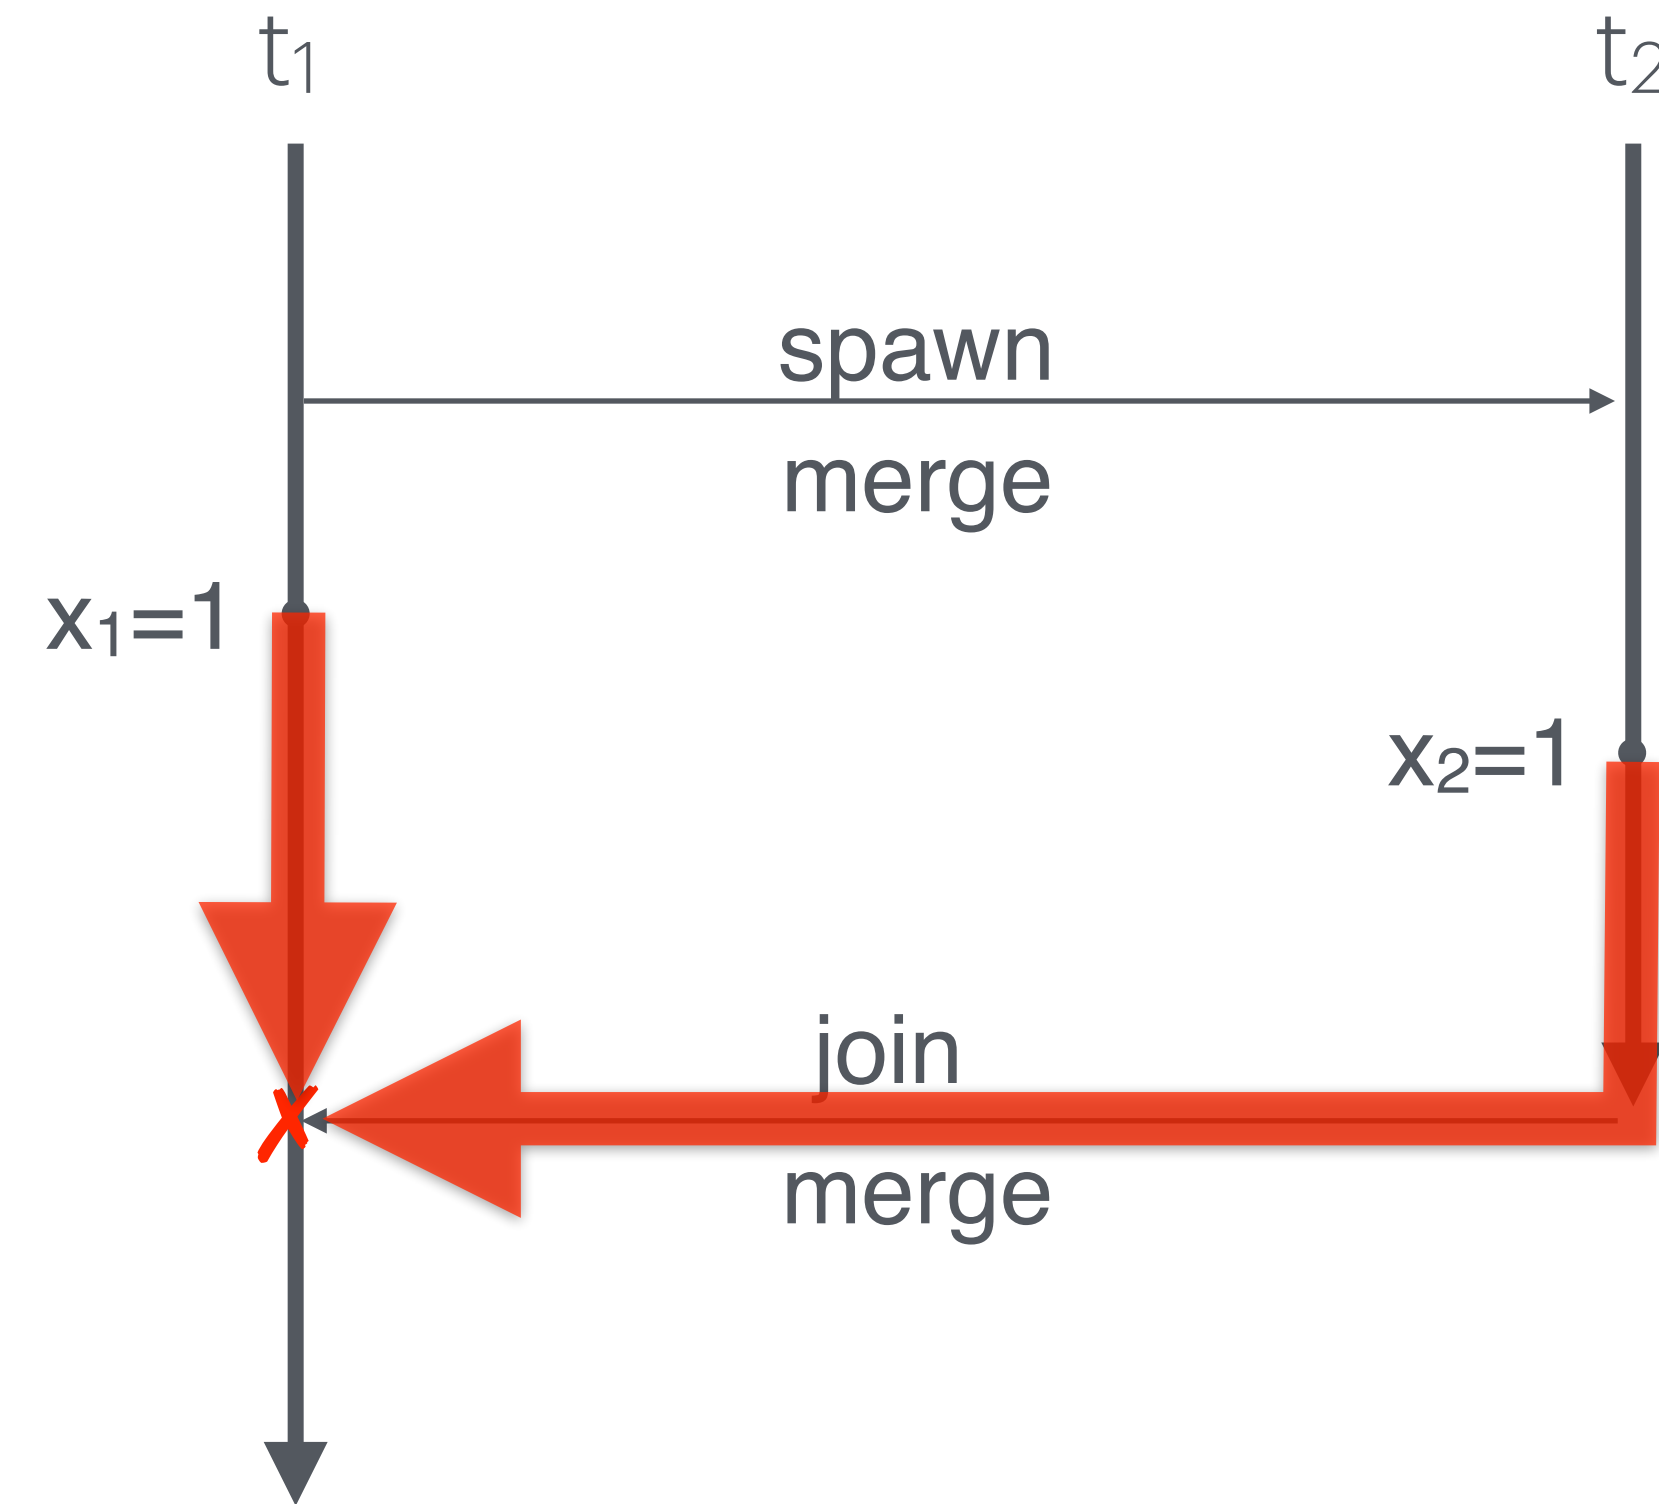
t ₁	t ₂
spawn t ₂ x = 1 join t ₂	x = 1



Example 2

Branching — **incorrectly** synchronised, write-write race on x


t_1	t_2
spawn t_2 $x = 1$ join t_2	$x = 1$



Example 3

Linear — correctly synchronised, lock l orders writes to x

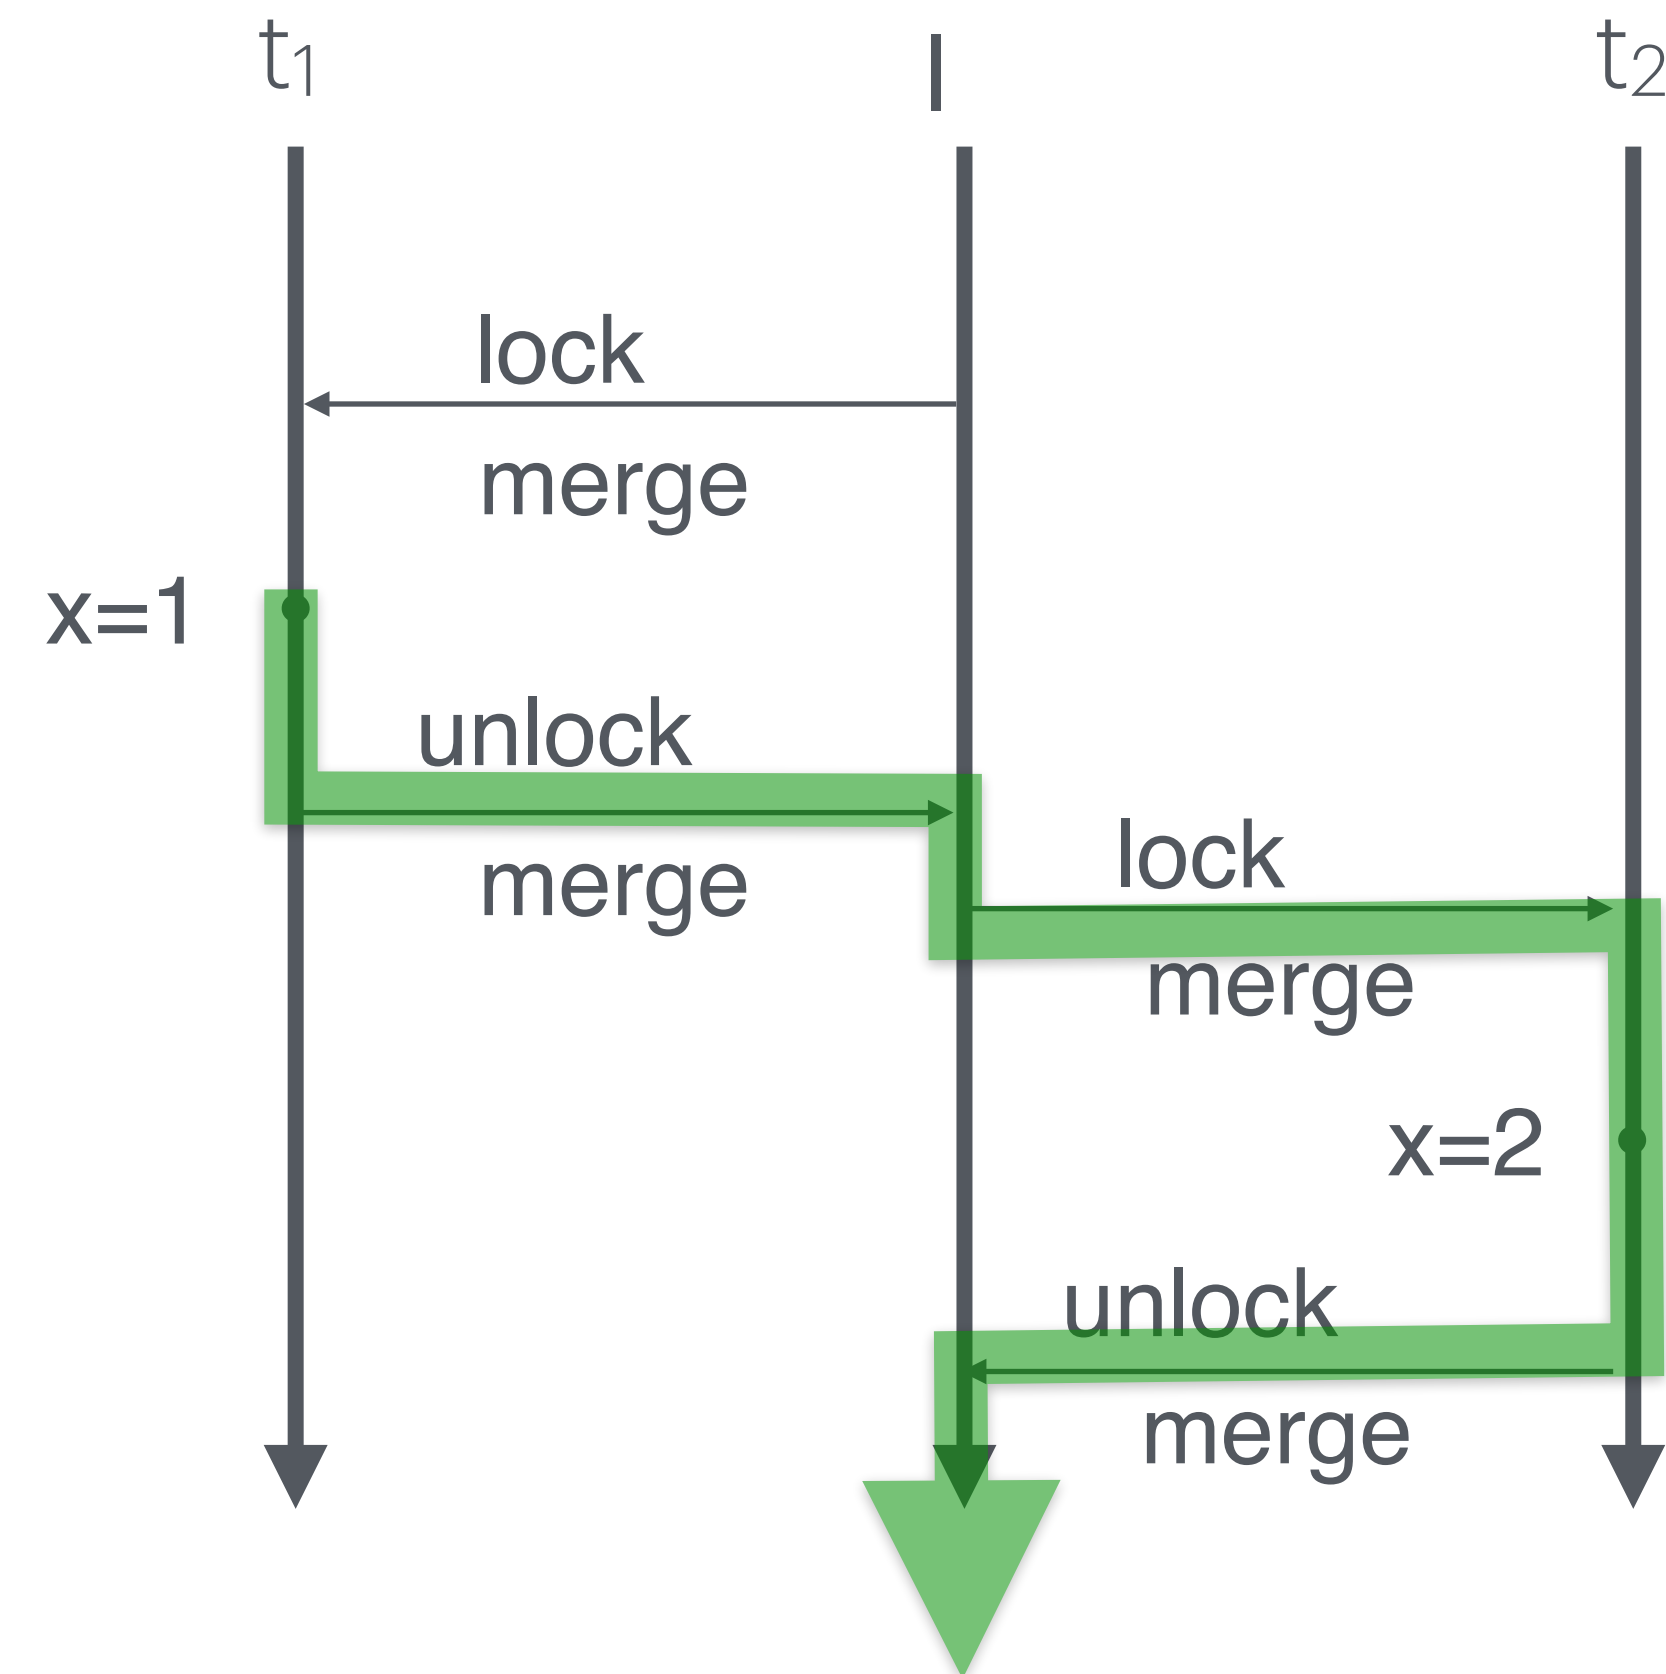
t ₁	t ₂
lock l x = 1 unlock l	lock l x = 2 unlock l



Example 3

Branching — correctly synchronised, lock l orders writes to x


t ₁	t ₂
lock l x = 1 unlock l	lock l x = 2 unlock l



Example 4

Linear — **incorrectly** synchronised, write-write race on x under **different** locks ("lucky" with timing)

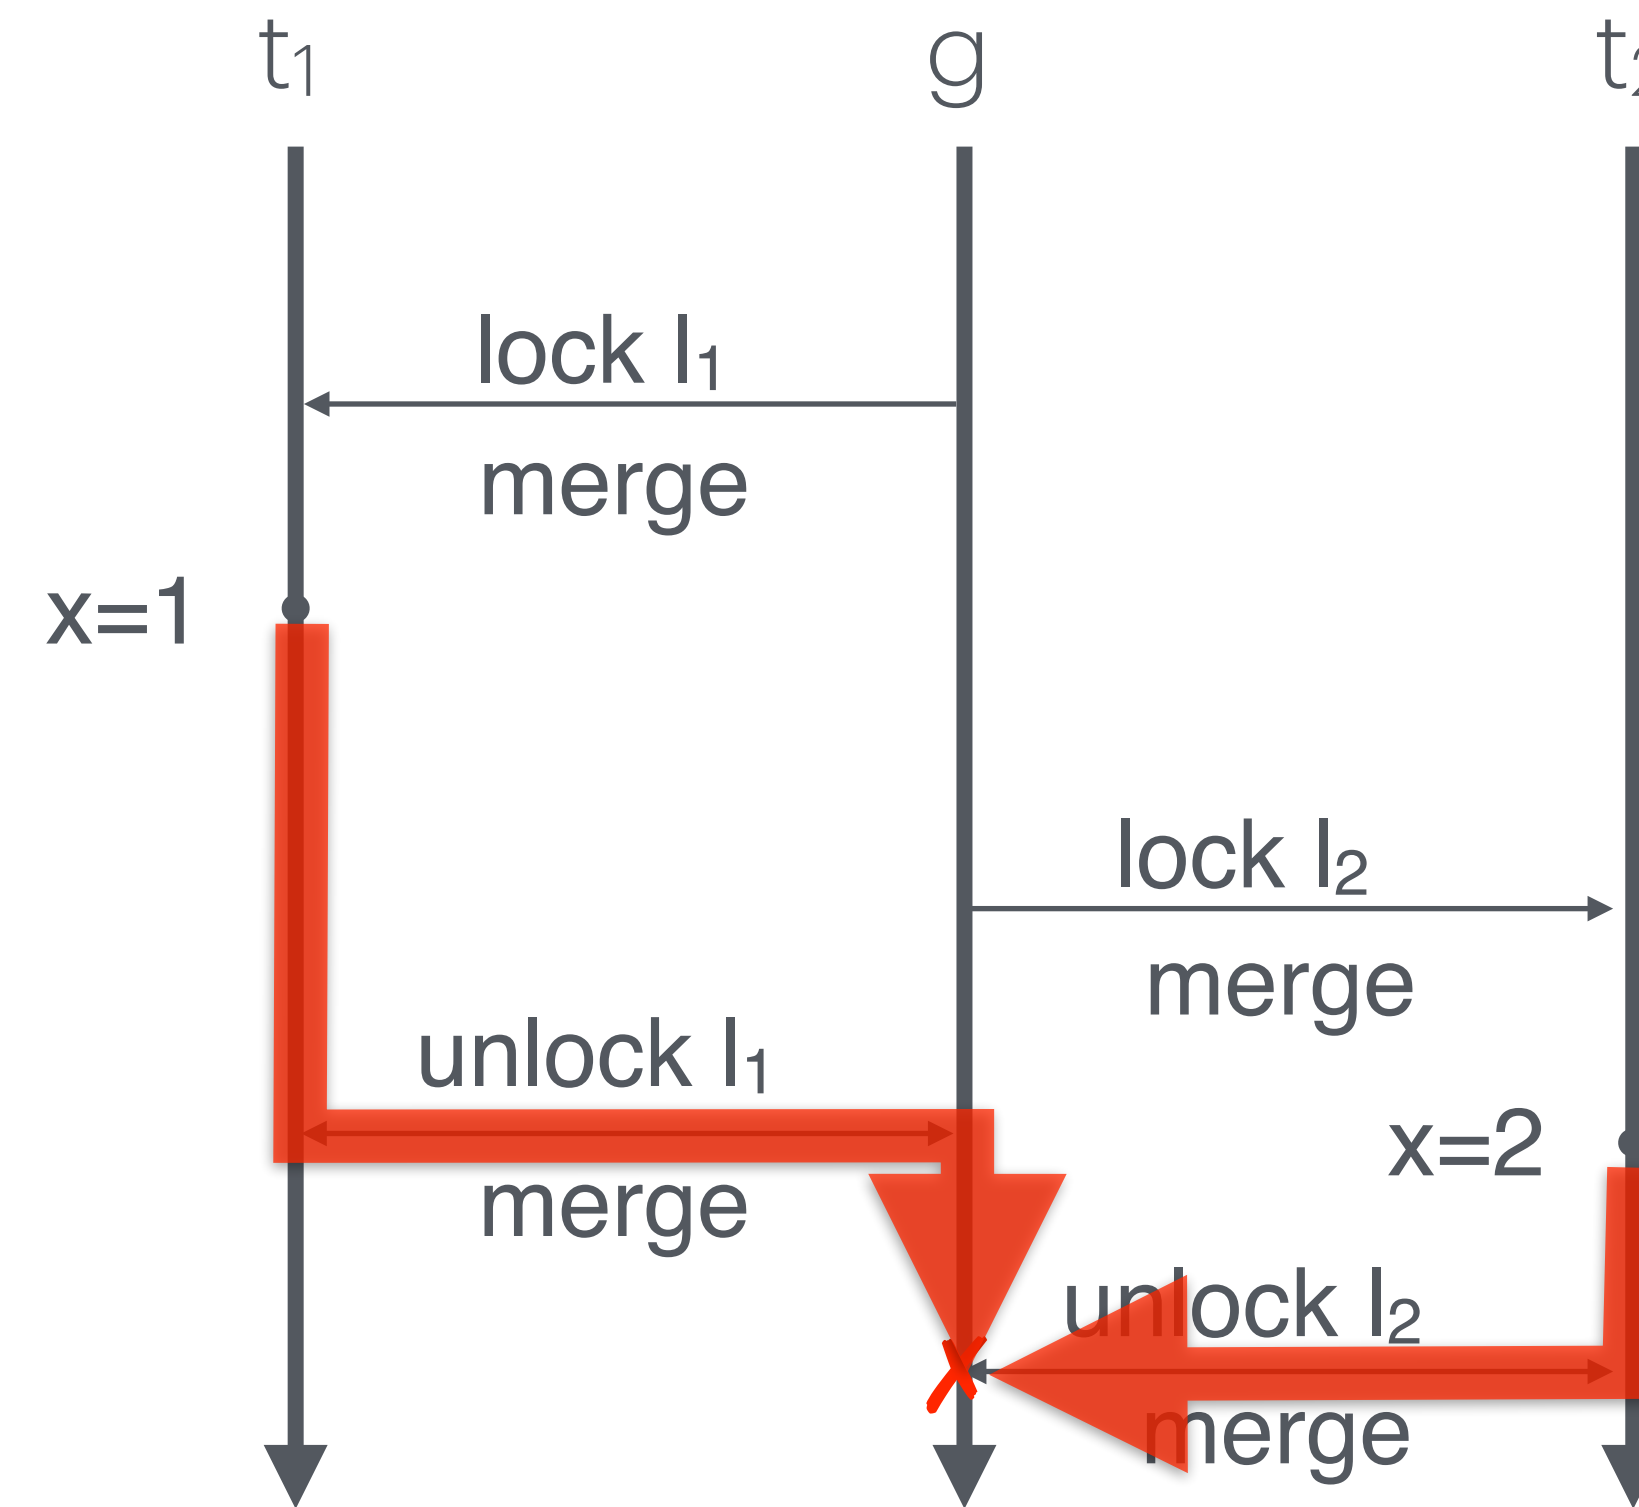
t_1	t_2
lock l_1 $x = 1$ unlock l_1	lock l_2 $x = 2$ unlock l_2



Example 4

Linear — **incorrectly** synchronised, write-write race on x under **different** locks (not "lucky")


t ₁	t ₂
lock l ₁ x = 1 unlock l ₁	lock l ₂ x = 2 unlock l ₂



Example 4

Branching — **incorrectly** synchronised, write-write race on x under **different** locks (thread isolated)


t ₁	t ₂
lock l ₁ x = 1 unlock l ₁	lock l ₂ x = 2 unlock l ₂



Example 5

Branching — **incorrectly** synchronised, write-write race on x under **different** locks (thread join)

t ₁	t ₂
lock l ₁ x = 1 unlock l ₁ join t ₂	lock l ₂ x = 2 unlock l ₂



Some Open Design Choices

- Example 4 shows that two threads that never synchronise can hold diverging world views
- Example 5 showed that two such threads joining (or synchronising) will eventually raise a conflict
- In Java, there is a synthetic join of all threads when the VM tears down

Should we raised a conflict exception here or not?

- Our approach: a "ConcurrencyManager"

Gives programmer some control over what should happen on a conflict

Not intended for fine-grained conflict resolution

The Leakage Problem

(Initially $x = 0$)

t_1	t_2
lock l_1 $x = 1$ unlock l_1	lock l_2 print(x) unlock l_2

(Initially $x = 0, y = 0$)

t_1	t_2
lock l_1 $x = 1$ lock l_2 $y = 2$ unlock l_2	lock l_2 print(x) unlock l_2

Status

Completed

Informal model


Diagrams for communication

Toy tool for writing litmus test programs that simulate execution

Pen & Paper + Mechanised formalism of our models

On-going

Implementation in the OpenJDK 26 interpreter



$$\begin{array}{c}
 \frac{T(t) = (\mathcal{N}, e, \emptyset, c)}{\langle T; H; L \rangle \longrightarrow \langle T[t \mapsto_{ts} \mathcal{R}]; H; L \rangle} \text{START} \qquad \frac{T(t) = (\mathcal{R}, r, \delta, c)}{\langle T; H; L \rangle \xrightarrow{\text{commit}_t} \langle T'; H'; L \rangle} \text{END} \\
 \\
 \frac{T(t) = (\mathcal{R}, e, \delta, c) \quad e \xrightarrow{\text{read } i} K \quad v = \text{read}(H, \delta, c, i)}{\langle T; H; L \rangle \longrightarrow \langle T[t \mapsto_e K(v)]; H; L \rangle} \text{READ} \qquad \frac{T(t) = (\mathcal{R}, e, \delta, c) \quad e \xrightarrow{\text{write } i \ v} K \quad (H', \delta') = \text{write}(H, \delta, i, v) \quad T' = T[t \mapsto (\mathcal{R}, K(\epsilon), \delta', c)]}{\langle T; H; L \rangle \longrightarrow \langle T'; H'; L \rangle} \text{WRITE} \\
 \\
 \frac{T(t) = (\mathcal{R}, e, \delta, c) \quad e \xrightarrow{\text{spawn } e_s} K \quad T' = T[t \mapsto_e K(t')][t' \mapsto (\mathcal{N}, e_s, \emptyset, 0)] \quad \langle T'; H; L \rangle \xrightarrow{\text{push}_{t,t'}} \langle T''; H'; L \rangle \quad t' \notin \text{dom}(T)}{\langle T; H; L \rangle \longrightarrow \langle T''; H'; L \rangle} \text{SPAWN} \\
 \\
 \frac{T(t) = (\mathcal{R}, e, \delta, c) \quad e \xrightarrow{\text{join } t'} K \quad T(t') = (\mathcal{T}, r, _ , _) \quad \langle T; H; L \rangle \xrightarrow{\text{pull}_{t,t'}} \langle T'; H'; L \rangle}{\langle T; H; L \rangle \longrightarrow \langle T'[t \mapsto_e K(r)]; H'; L \rangle} \text{JOIN} \\
 \\
 \frac{T(t) = (\mathcal{R}, e, \delta, c) \quad e \xrightarrow{\text{lock } \ell} K \quad L(\ell).ls = \text{free} \quad \langle T; H; L \rangle \xrightarrow{\text{pull}_{t,\ell}} \langle T'; H'; L \rangle}{\langle T; H; L \rangle \longrightarrow \langle T'[t \mapsto_e K(\epsilon)]; H'; L[\ell \mapsto_{ls} \text{held}(t)] \rangle} \text{LOCK} \\
 \\
 \frac{T(t) = (\mathcal{R}, e, \delta, c) \quad e \xrightarrow{\text{unlock } \ell} K \quad L(\ell).ls = \text{held}(t) \quad \langle T; H; L \rangle \xrightarrow{\text{push}_{t,\ell}} \langle T'; H'; L' \rangle}{\langle T; H; L \rangle \longrightarrow \langle T'[t \mapsto_e K(\epsilon)]; H'; L'[\ell \mapsto_{ls} \text{free}] \rangle} \text{UNLOCK}
 \end{array}$$

(b) Concurrency Semantics.

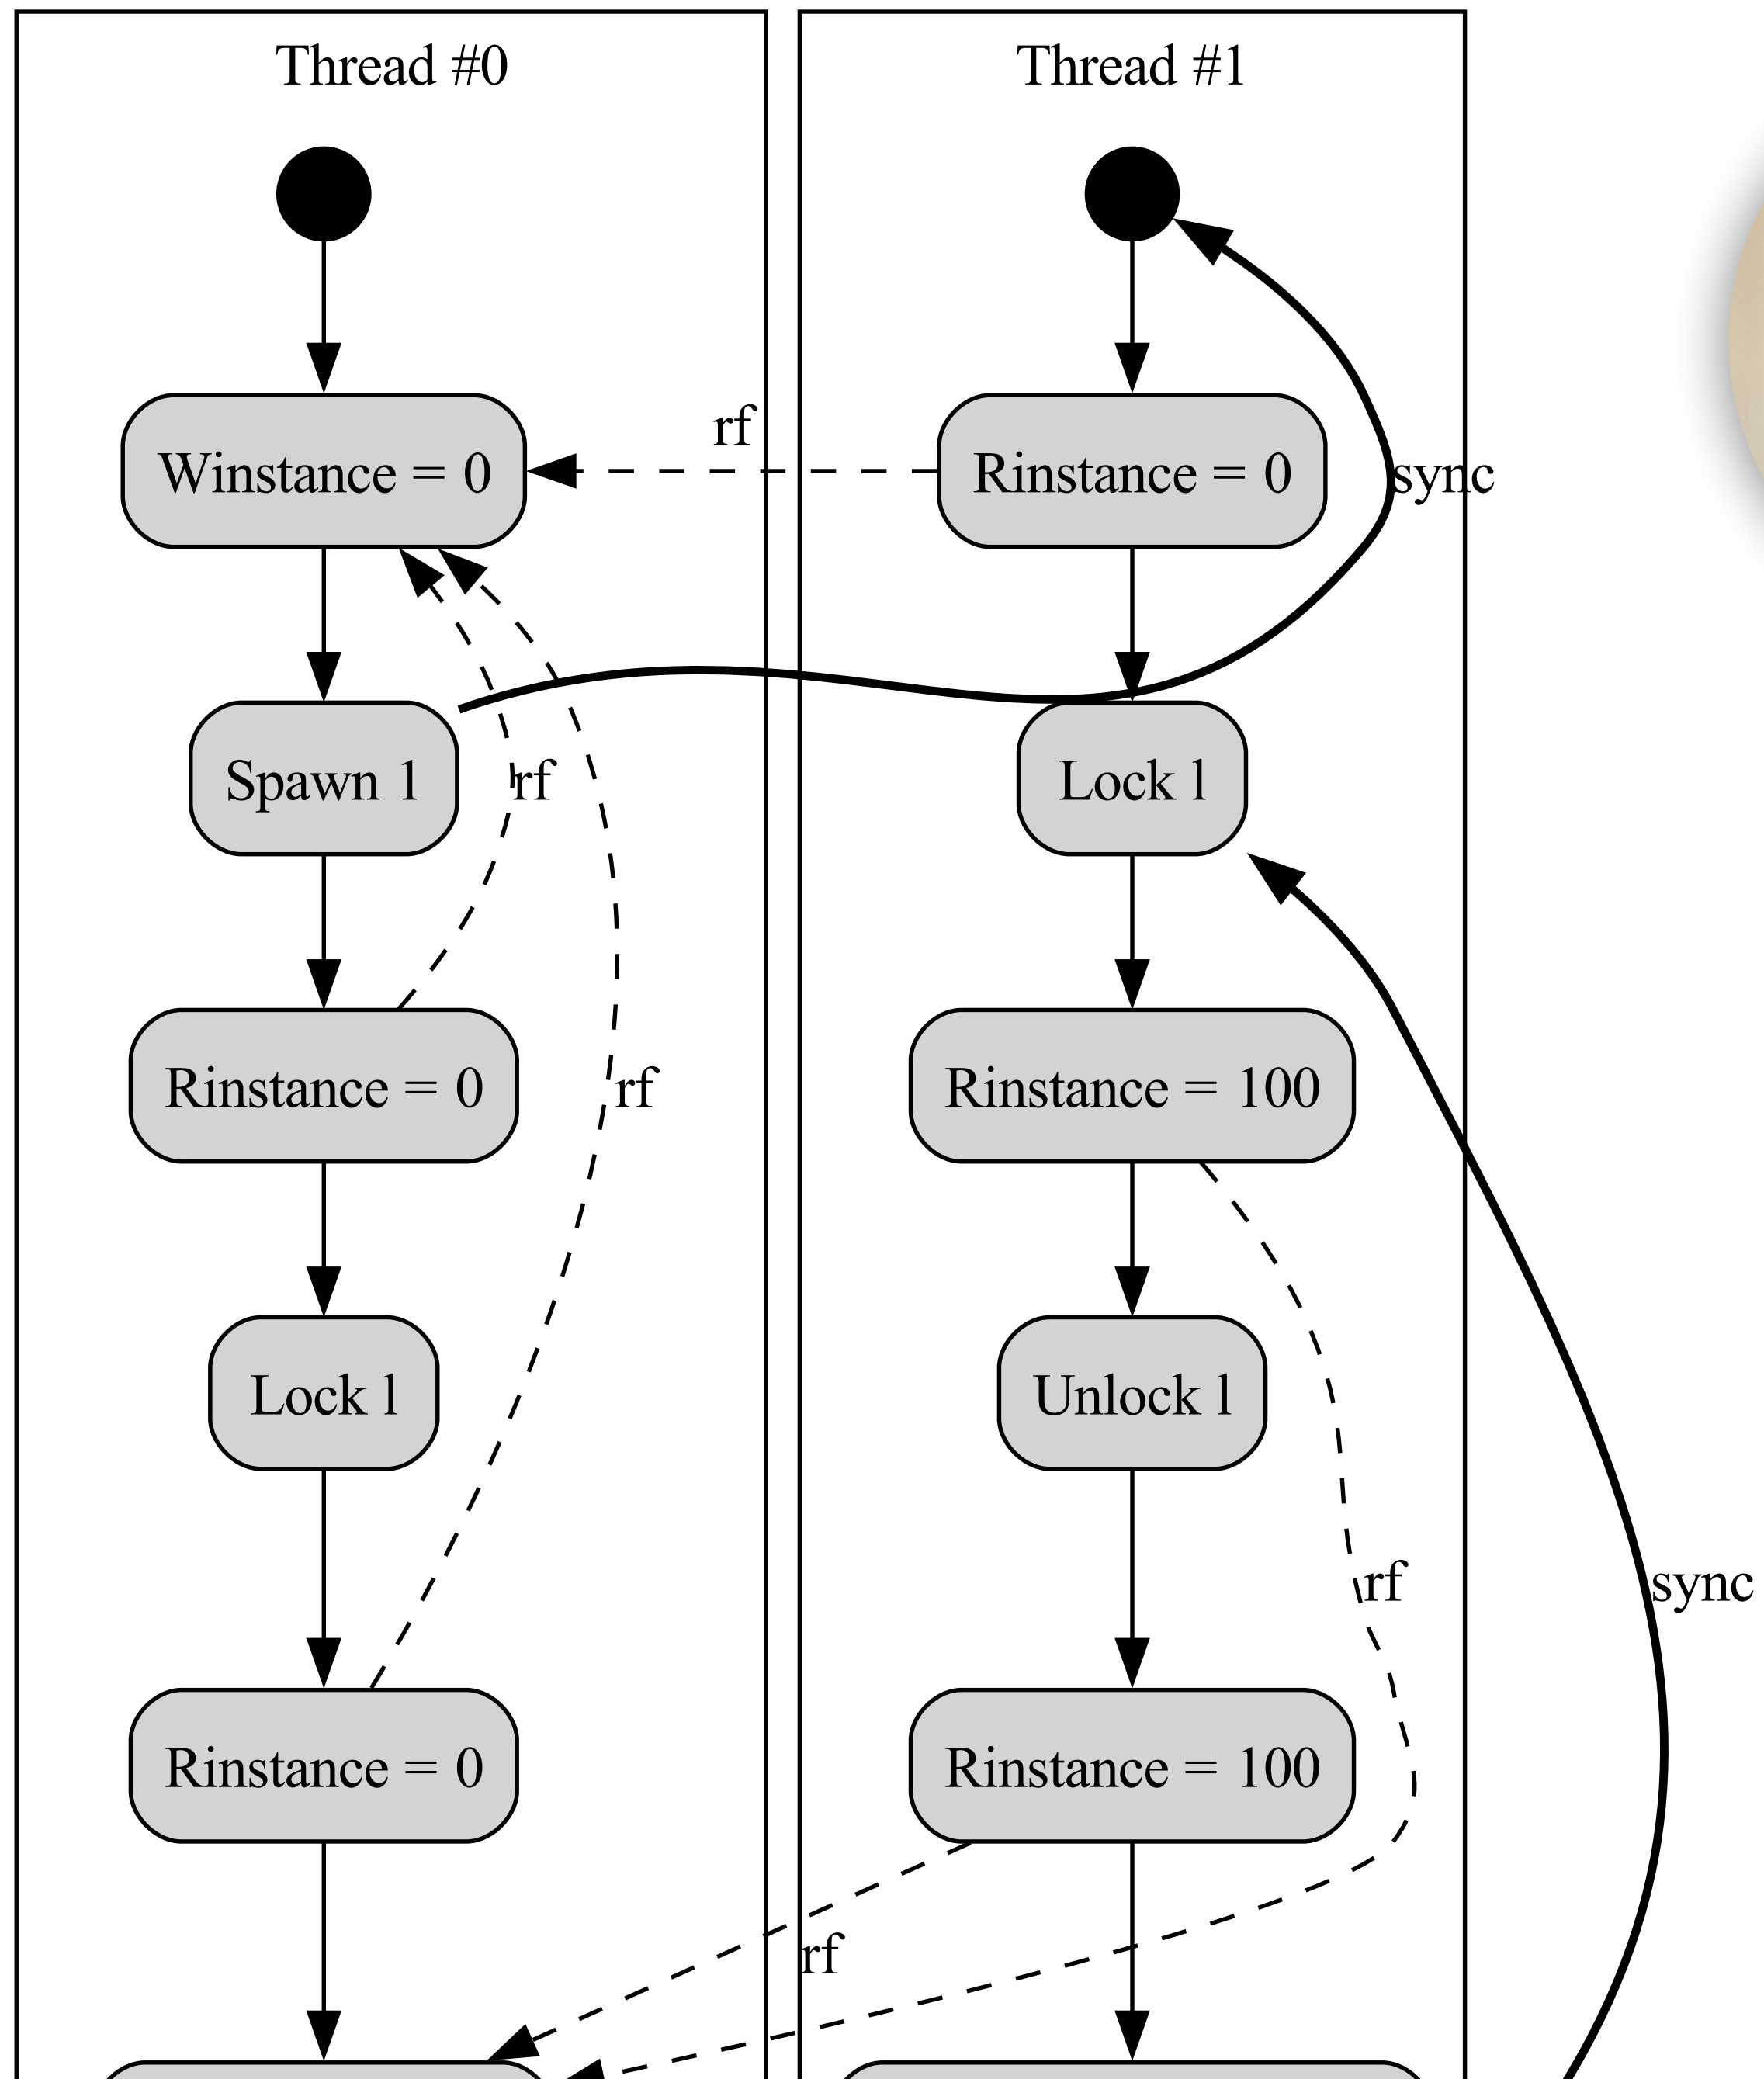
```
// Construct the singleton pattern:
// ...
```

```
instance = 0;
```

```
$t1 = spawn {
  if (instance == 0) {
    lock l;
    if (instance == 0) {
      instance = 100;
    }
    unlock l;
  }
  assert(instance == 100);
};
```

```
if (instance == 0) {
  lock l;
  if (instance == 0) {
    instance = 100;
  }
  unlock l;
}
assert(instance == 100);
```

```
join $t1;
assert(instance == 100);
```



Can this be sufficiently performant?

- We do not know!
- Next step: instrument OpenJDK to measure various events to give an idea of the overhead costs
- This also enables some optimisations that can be used to off-set the costs

141

Fixing the program

```

1 final class Logger {
2   private static Logger instance; volatile
3   private String path; final
4   private Logger(String path) {
5     this.path = path;
6   }
7
8   public static Logger getLogger() {
9     if (instance == null) {
10      synchronized (Logger.class) {
11        if (instance == null) {
12          instance = new Logger("f.log");
13        }
14      }
15    }
16    return instance;
17  }
18 }

```

Original →

```

1 instance =
2   new Logger("f.log");

```

After inlining →

```

1 var tmp = new Logger();
2 tmp.path = "f.log";
3 instance = tmp;

```

After reordering

```

1 instance = new Logger();
2 instance.path = "f.log";

```

