

LOOM: Bypassing Races in Live Applications with Execution Filters

Jingyue Wu, Heming Cui, Junfeng Yang
Columbia University

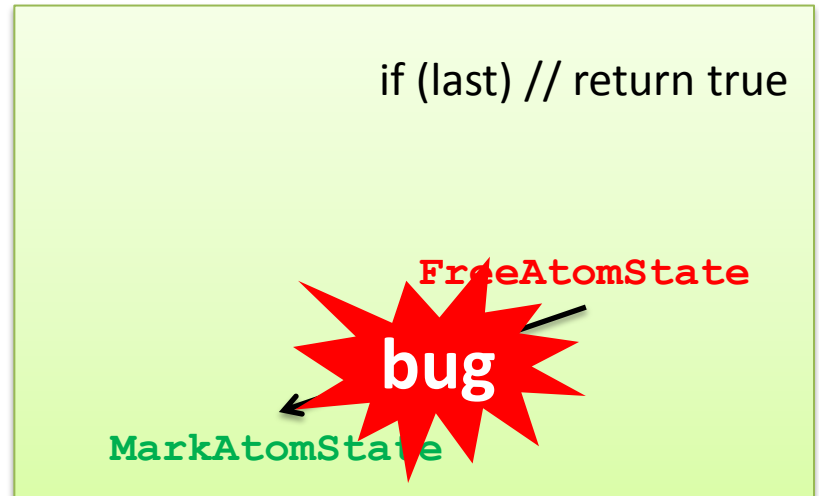
Mozilla Bug #133773

```
void js_DestroyContext(  
    JSContext *cx) {  
    JS_LOCK_GC(cx->runtime);  
    MarkAtomState(cx);  
    if (last) { // last thread?  
        ...  
        FreeAtomState(cx);  
        ...  
    }  
    JS_UNLOCK_GC(cx->runtime);  
}
```

A buggy interleaving

Non-last Thread

Last Thread



Complex Fix

```
void js_DestroyContext() { void js_ForceGC(bool last) }
  if (last) { { gcLevel = 1;
                gcLock.release();
                restart:
                MarkAtomState();
                gcLock.acquire();
                if (gcLevel > 1) {
                  gcLevel = 1;
                  gcLock.release();
                  goto restart;
                }
                gcLevel = 0;
                gcPoke = false;
                gcLock.release();
            }
    state = LANDING; gcPoke = true;
    if (requestDepth == 0) js_GC(last);
    js_BeginRequest(); }
    while (gcLevel > 0) void js_GC(bool last) {
        JS_AWAIT_GC_DONE(); if (state == LANDING &&
        js_ForceGC(true); !last)
        while (gcPoke) return;
        js_GC(true); gcLock.acquire();
        FreeAtomState(); if (!gcPoke) {
        } else { gcLock.release();
                gcLevel = 0;
                gcPoke = true; return;
                gcLock.release();
                js_GC(false); }
        if (gcLevel > 0) { }
    }
    gcLevel++;
    while (gcLevel > 0)
        JS_AWAIT_GC_DONE();
        gcLock.release();
        return;
}
void js_BeginRequest() {
    while (gcLevel > 0)
        JS_AWAIT_GC_DONE();
}
```

- 4 functions; 3 integer flags
- Nearly a month
- Not the only example

LOOM: Live-workaround Races

- Execution filters: temporarily filter out buggy thread interleavings

```
void js_DestroyContext(JSContext *cx) {  
    MarkAtomState(cx) ;  
    if (last thread) {  
        ...  
        FreeAtomState(cx) ;  
        ...  
    }  
}
```

**A mutual-exclusion
execution filter to bypass
the race on the left**

```
js_DestroyContext <> self
```

- Declarative, easy to write

LOOM: Live-workaround Races

- Execution filters: temporarily filter out buggy thread interleavings
- Installs execution filters to live applications
 - Improve server availability
 - STUMP [PLDI '09], Ginseng [PLDI '06], KSplice [EUROSYS '09]
- Installs execution filters safely
 - Avoid introducing errors
- Incurs little overhead during normal execution

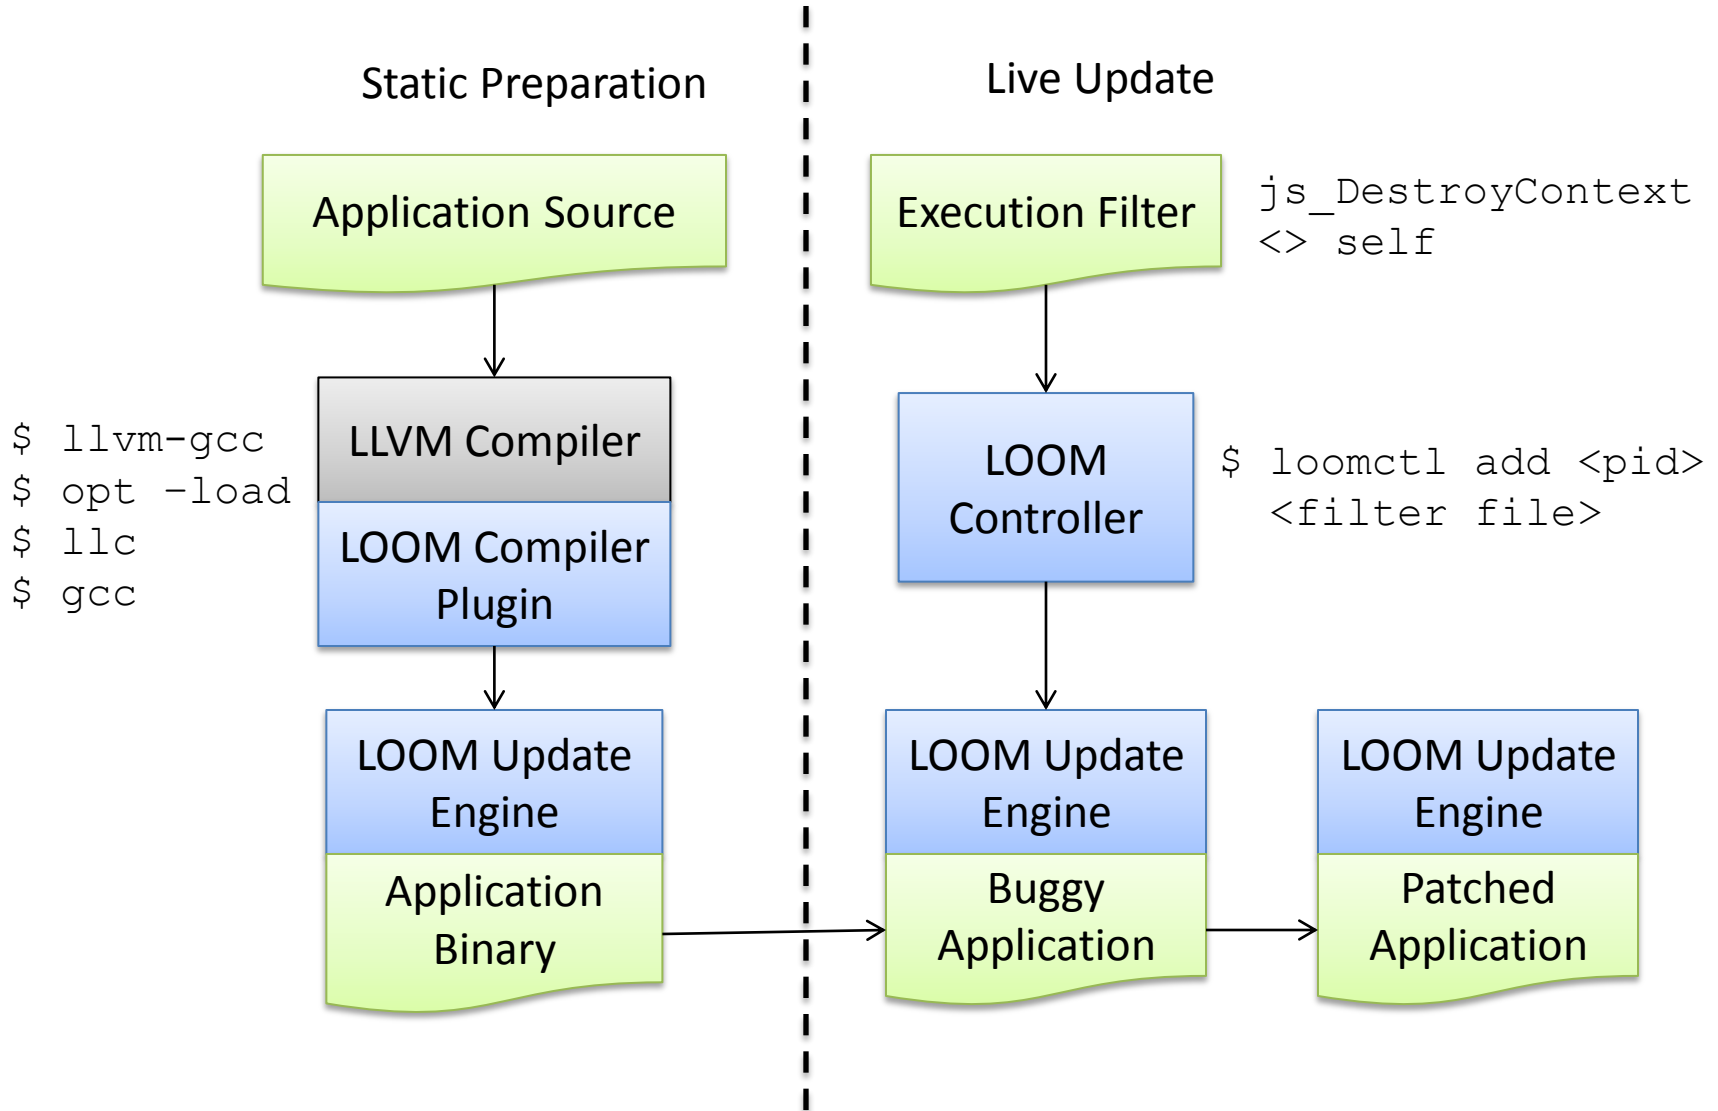
Summary of Results

- We evaluated LOOM on nine real races.
 - Bypasses all the evaluated races safely
 - Applies execution filters immediately
 - Little performance overhead (< 5%)
 - Scales well with the number of application threads (< 10% with 32 threads)
 - Easy to use (< 5 lines)

Outline

- Architecture
 - Combines static preparation and live update
- Safely updating live applications
- Reducing performance overhead
- Evaluation
- Conclusion

Architecture




Outline


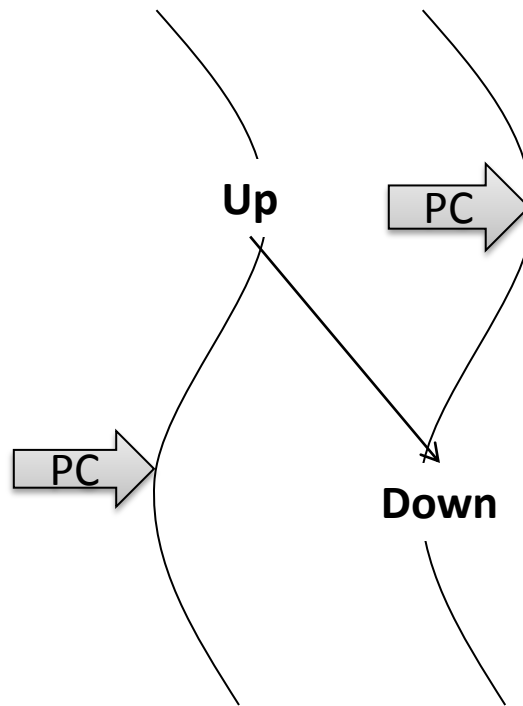
- Architecture
 - Combines static preparation and live update
- **Safely updating live applications**
- Reducing performance overhead
- Evaluation
- Conclusion

Safety: Not Introducing New Errors

Mutual Exclusion




Order Constraints




Evacuation Algorithm

1. Identify the dangerous region using static analysis
2. Evacuate threads that are in the dangerous region
3. Install the execution filter


Unsafe to update




“Evacuate”




Safe to update



Install Filter

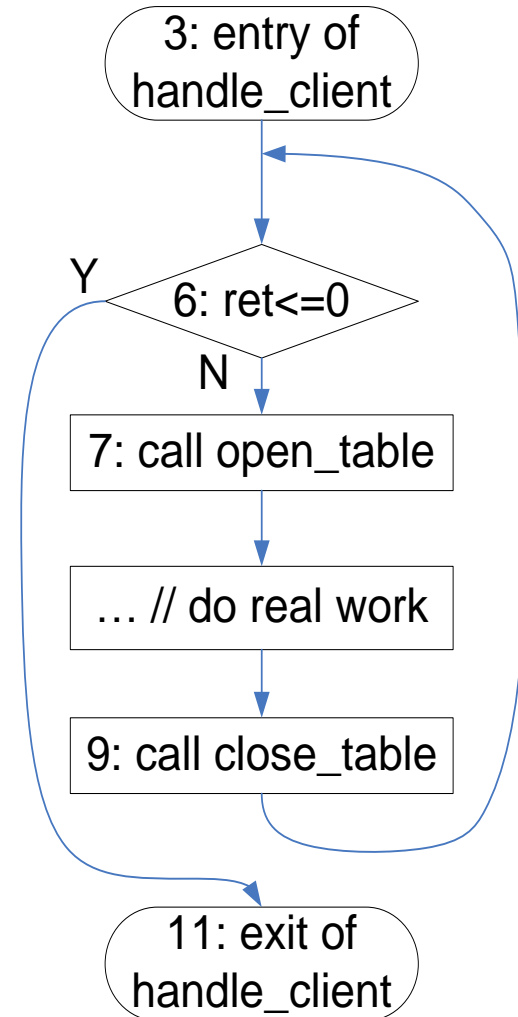


Updated

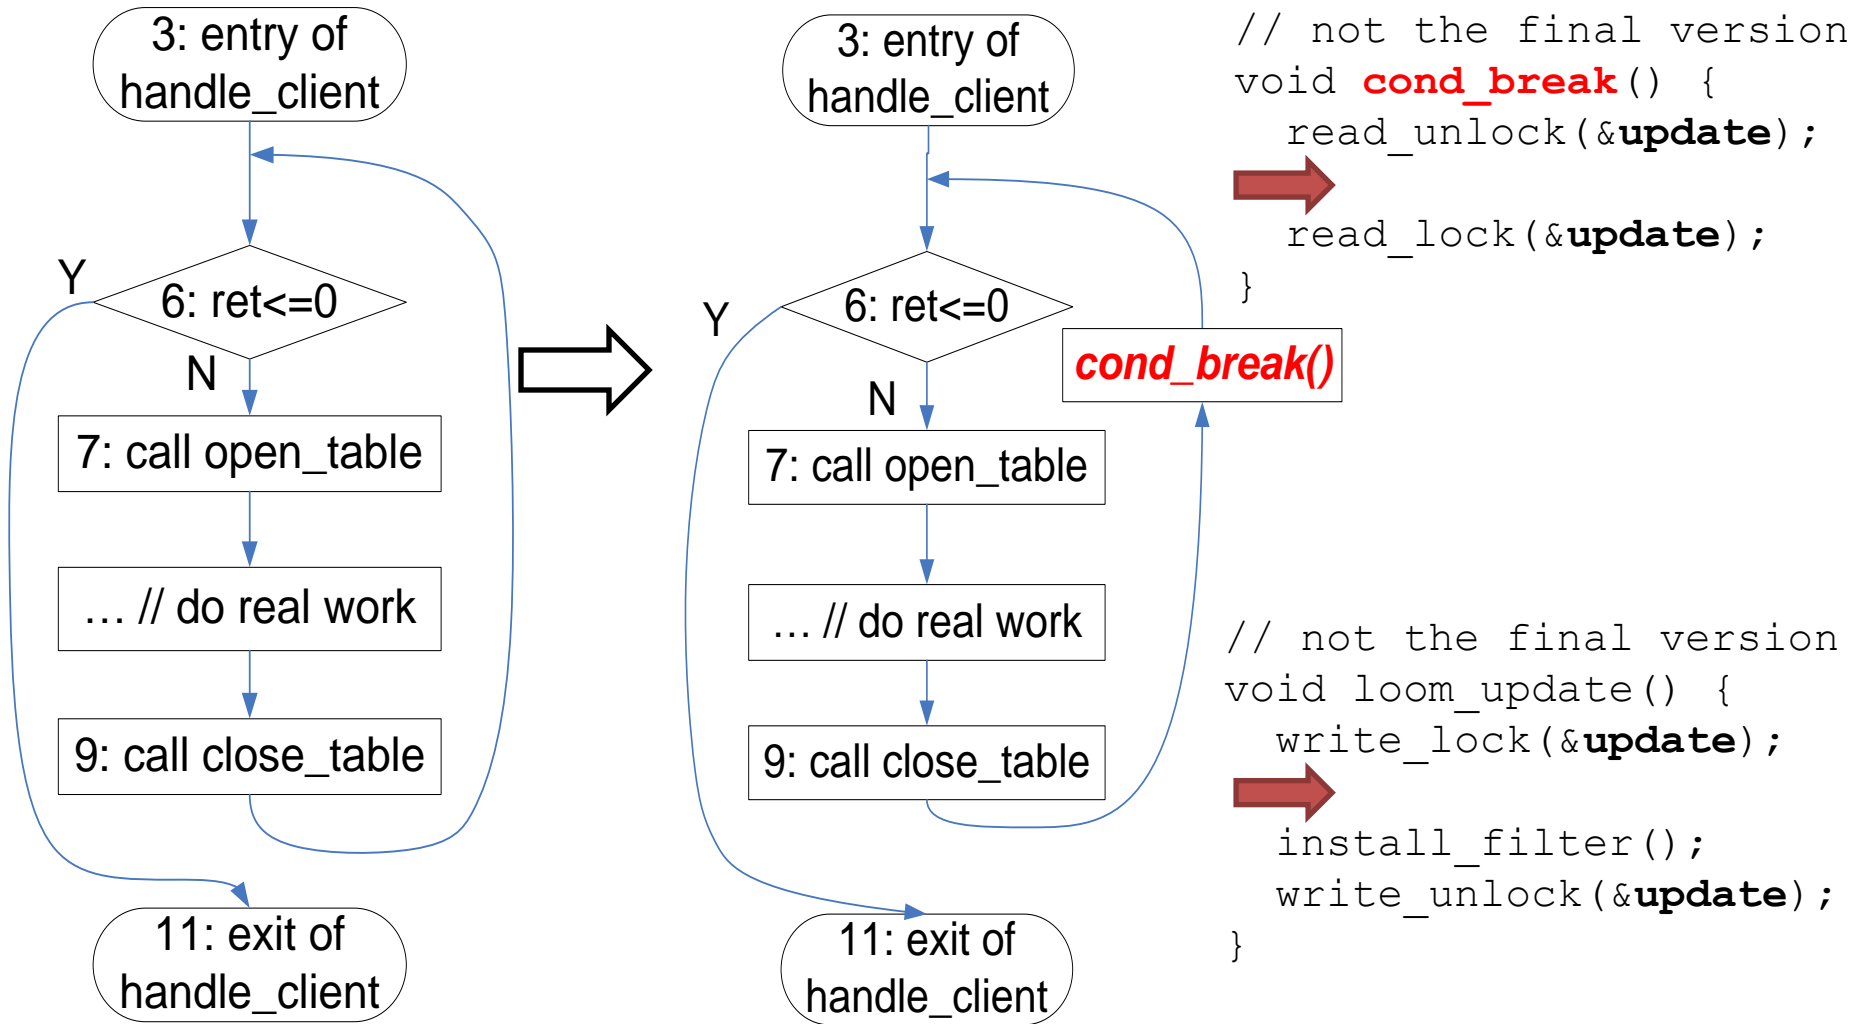


Control Application Threads

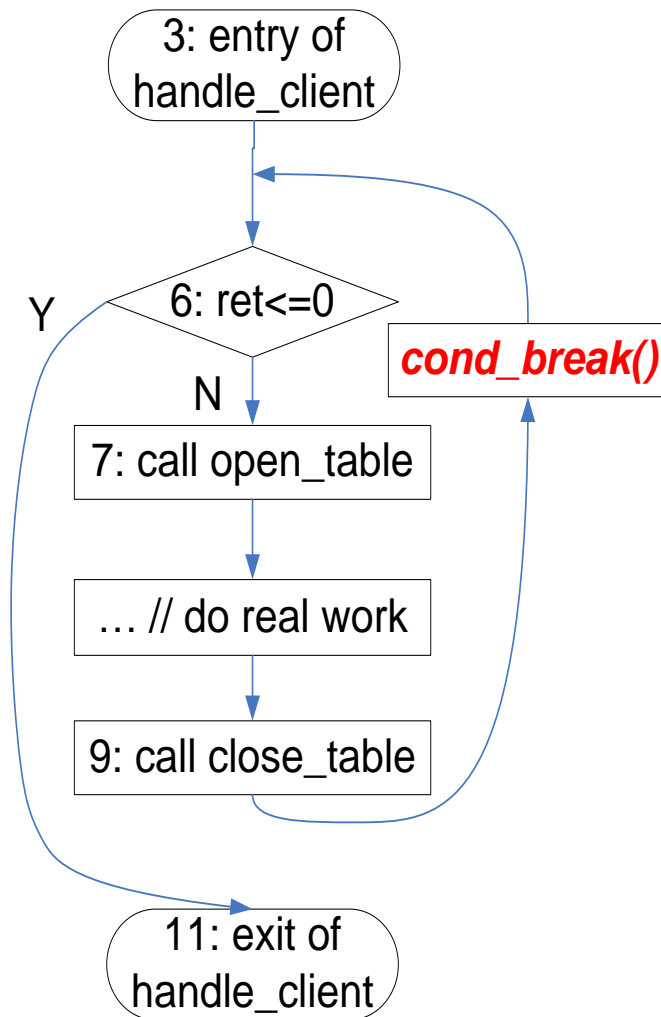
```
1 : // database worker thread
2 : void handle_client(int fd) {
3 :     for(;;) {
4 :         struct client_req req;
5 :         int ret = recv(fd, &req, ...);
6 :         if(ret <= 0) break;
7 :         open_table(req.table_id);
8 :         ... // do real work
9 :         close_table(req.table_id);
10:    }
11: }
```



Control Application Threads (cont'd)



Pausing Threads at Safe Locations



```
    cmpl 0x0, 0x845208c
    je 0x804b56d
```

```
void cond_break() {
    if (wait[backedge_id]) {
        read_unlock(&update);
        while (wait[backedge_id]);
        read_lock(&update);
    }
}
```

Annotations: A white arrow points to the `wait` variable in the `if` statement. A red box highlights the `wait` variable in the `while` loop. A red arrow points to the `while` loop.

```
void loom_update() {
    identify_safe_locations();
    for each safe backedge E
        wait[E] = true;
    write_lock(&update);
    install_filter();
    for each safe backedge E
        wait[E] = false;
    write_unlock(&update);
}
```

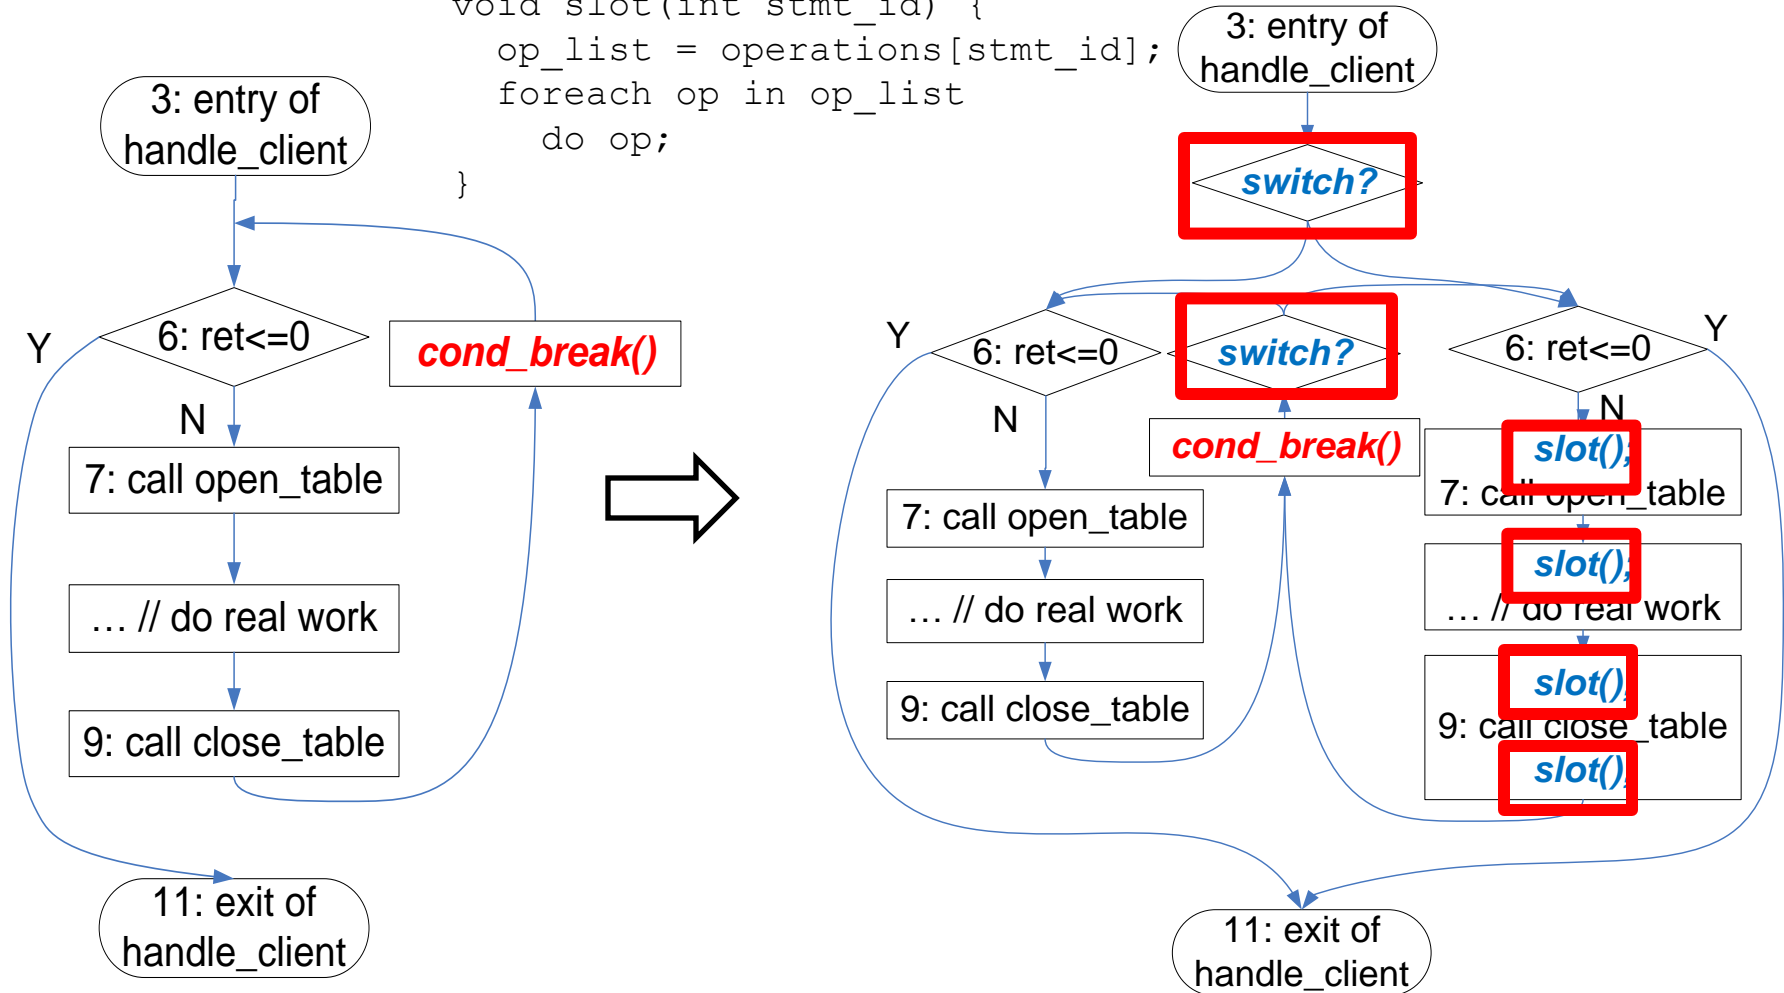
Annotations: Red arrows point to the `identify_safe_locations()` call, the `wait[E] = true;` line, the `wait[E] = false;` line, and the `write_unlock(&update);` call.

Outline


- Architecture
 - Combines static preparation and live update
- Safely updating live applications
- **Reducing performance overhead**
- Evaluation
- Conclusion

Hybrid Instrumentation

```
void slot(int stmt_id) {  
    op_list = operations[stmt_id];  
    foreach op in op_list  
        do op;  
}
```

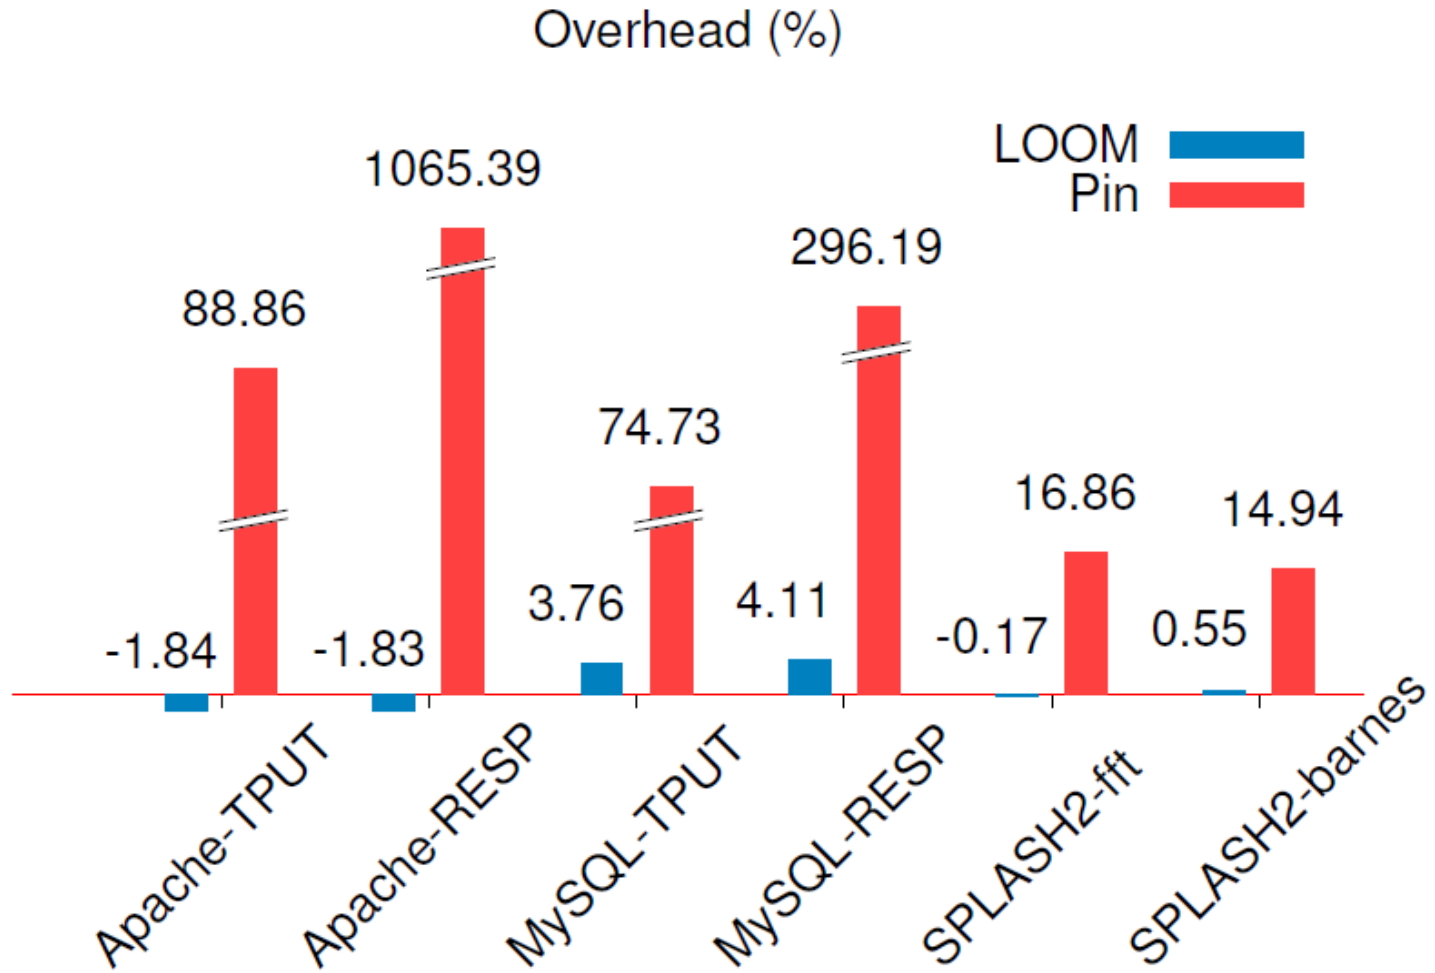


Bare Instrumentation Overhead



Performance overhead < 5%

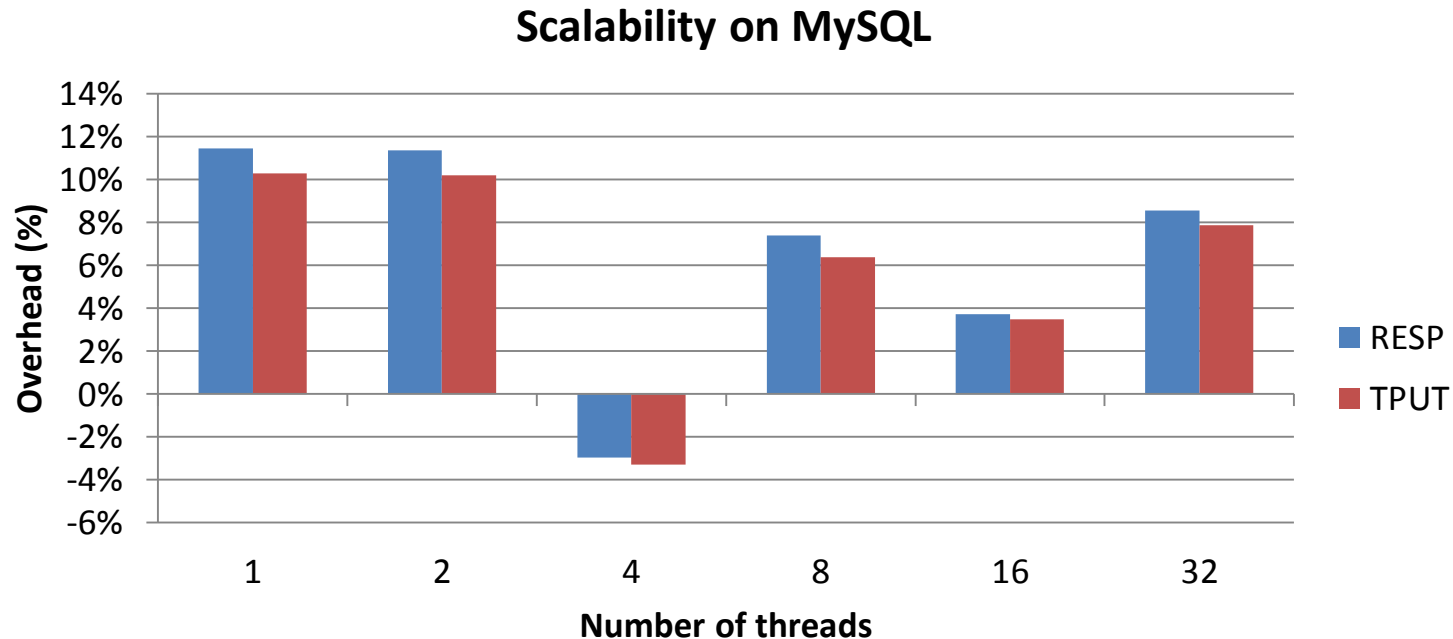
Bare Instrumentation Overhead



Performance overhead < 5%

Scalability

- 48-core machine with 4 CPUs; Each CPU has 12 cores.
- Pin the server to CPU 0, 1, 2, and the client to CPU 3.



Performance overhead does not increase

Conclusion

- LOOM: A live-workaround system designed to quickly and safely bypass races
 - **Execution filters:** easy to use and flexible (< 5 lines)
 - **Evacuation algorithm:** safe
 - **Hybrid instrumentation:** fast (overhead < 5%) and scalable (overhead < 10% with 32 threads)
- Future work
 - Generic hybrid instrumentation framework
 - Extend the idea to other classes of errors

Questions?