


# A Study of Dynamic Software Update Quiescence for Multithreaded Programs

---

Christopher M. Hayden, Karla Saur,  
Michael Hicks, Jeffrey S. Foster



# Update Timing


- Well-defined *update points* make it easier to reason about update correctness

```
1 void *thread_entry(void *arg) {  
2     /* thread init code */  
3     while (1) {  
4         dsu_update(); /* update point */  
5         /* loop body: typically handles a single program event */  
6     }  
7 }
```


- Good candidates are *quiescent points* in loops which have little in-flight state

# DSU and Threading

- *Timeliness* in multithreaded updates:
  - Full quiescence – all threads hit update point



- Concern - Updating at only specific points has the risk of delaying an update for too long, even *indefinitely*



# Goals & Approach:

- Questions:
  - Quick full quiescence in multithreaded programs?
  - What blocking calls impede quick quiescence?
- Created library: QBench
  - Interrupt blocking to facilitate quiescence with minimal delay
  - *Measures time* from update request to full quiescence
  - Idioms we develop in QBench we can roll into DSU systems

# Update at Quiescent Points

- Update point 'qbench\_update':
  - No update requested: call is a no-op
  - Update requested: calling thread blocks

```
1 void *thread_entry(void *arg) {  
2     /* thread init code */  
3     while (1) {  
4         qbench_update(); /* update point */  
5         /* loop body: typically handles a single program event */  
6     }  
7 }
```


- Request an update by sending a SIGUSR2 signal
  - QBench installs a signal handler indicating update requested.

# Threats to Quiescence

- Blocking calls in our experiments:

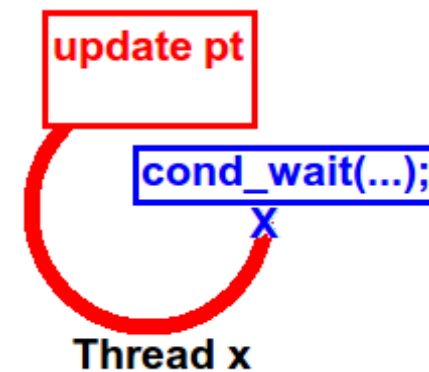
I/O:

Socket blocking on data



Condition Variables:

Threads sharing a mutex



# Blocking on I/O

Under normal circumstances an accept call will block until a connection is accepted.

```
1 void *thread_entry(void *arg) {  
2     /* thread init code */  
3     while (1) {  
4  
5         res = accept(sockfd, addr, addrlen);  
6  
7  
8         /* ... handle connection */  
9     }  
10 }
```

# Blocking on I/O

Under normal circumstances an accept call will block until a connection is accepted.

```
1 void *thread_entry(void *arg) {
2     /* thread init code */
3     while (1) {
4         qbench_update();
5         res = accept(sockfd, addr, addrlen);
6
7
8         /* ... handle connection */
9     }
10 }
```



# Blocking on I/O

Under normal circumstances an accept call will block until a connection is accepted.

```
1 void *thread_entry(void *arg) {
2     /* thread init code */
3     while (1) {
4         qbench_update();
5         res = accept(sockfd, addr, addrlen);
6         if (res == -1 && errno == EINTR)
7             continue;
8         /* ... handle connection */
9     }
10 }
```

A signal will interrupt *accept*, return -1, and set errno to EINTR.



# Blocking on I/O

Under normal circumstances an accept call will block until a connection is accepted.

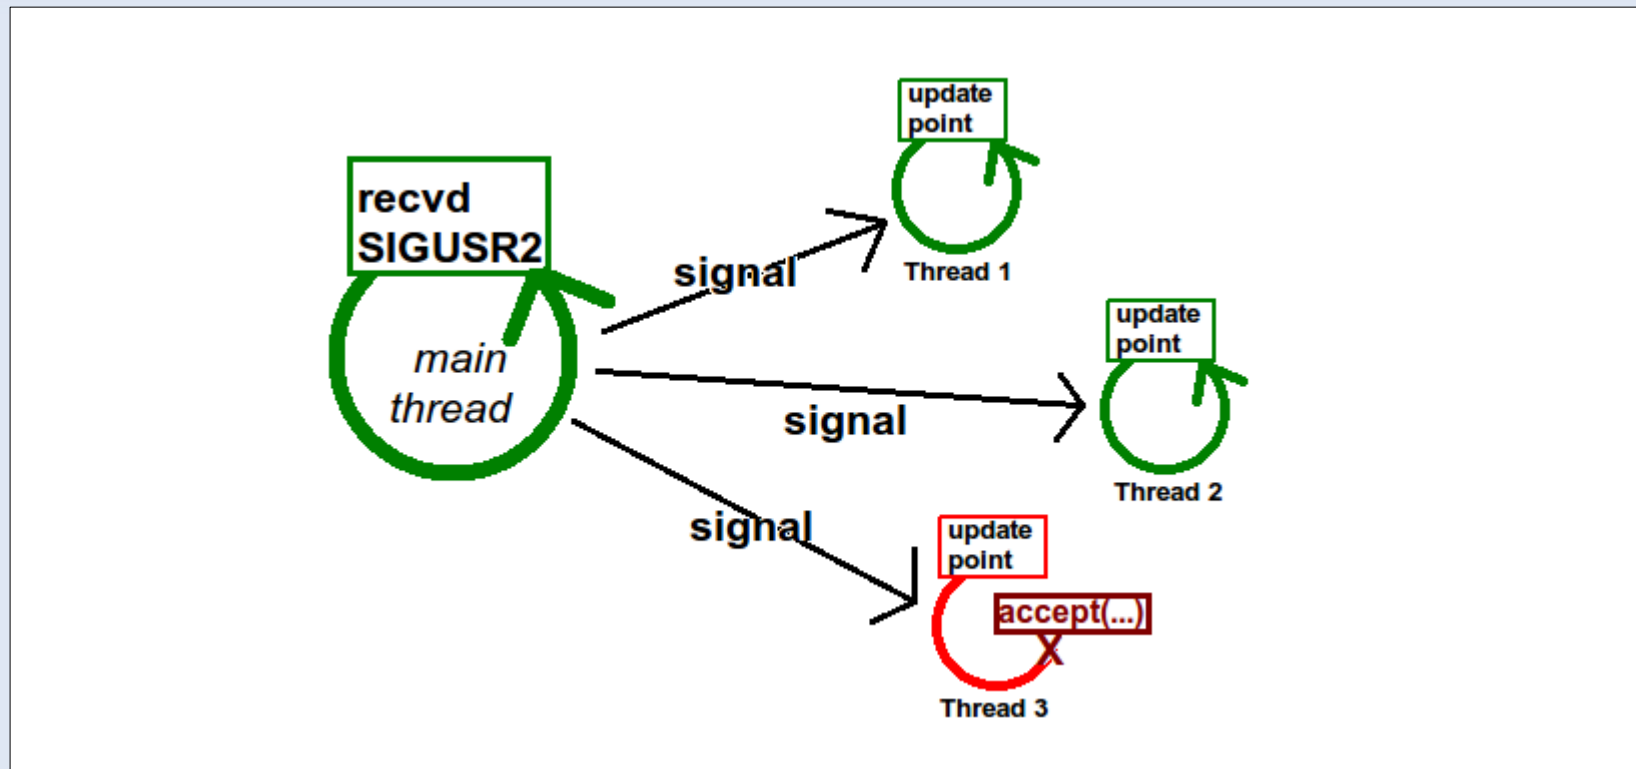
```
1 void *thread_entry(void *arg) {
2     /* thread init code */
3     while (1) {
4         qbench_update();
5         res = accept(sockfd, addr, addrlen);
6         if (res == -1 && errno == EINTR)
7             continue;
8         /* ... handle connection */
9     }
10 }
```

A signal will interrupt *accept*, return -1, and set *errno* to *EINTR*.

Returns to top of even loop to immediately hit update point

# UNIX Signals

- Signals are usually handled by main thread
  - Main thread signals all threads not blocked by condition variables



# Blocking on Condition Variables

Programmers guard against spurious wake-ups by placing `pthread_cond_wait` in a loop

```
1 void *thread_entry(void *arg) {
2     /* thread init code */
3     while (1) {
4         qbench_update();
5         pthread_mutex_lock(&mutex);
6         while (!input_is_ready()){
7             pthread_cond_wait(&cond,&mutex);
8         }
9         pthread_mutex_unlock(&mutex);
10
11
12     /* ... handle connection */
13 }
14 }
```

# Blocking on Condition Variables

Programmers guard against spurious wake-ups by placing `pthread_cond_wait` in a loop

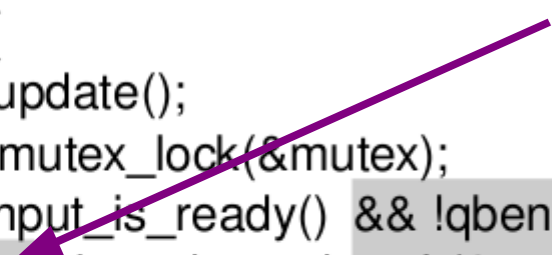
```
1 void *thread_entry(void *arg) {
2     /* thread init code */
3     while (1) {
4         qbench_update();
5         pthread_mutex_lock(&mutex);
6         while (!input_is_ready() && !qbench_update_requested()) {
7             qbench_pthread_cond_wait(&cond, &mutex);
8         }
9         pthread_mutex_unlock(&mutex);
10        if (qbench_update_requested())
11            continue; /* reaches qbench_update */
12        /* ... handle connection */
13    }
14 }
```

# Blocking on Condition Variables

Programmers guard against spurious wake-ups by placing `pthread_cond_wait` in a loop

```
1 void *thread_entry(void *arg) {
2     /* thread init code */
3     while (1) {
4         qbench_update();
5         pthread_mutex_lock(&mutex);
6         while (!input_is_ready() && !qbench_update_requested()) {
7             qbench_thread_cond_wait(&cond, &mutex);
8         }
9         pthread_mutex_unlock(&mutex);
10        if (qbench_update_requested())
11            continue; /* reaches qbench_update */
12        /* ... handle connection */
13    }
14 }
```

Allows thread to be signaled for update even when waiting on a condition variable



# Blocking on Condition Variables

Programmers guard against spurious wake-ups by placing `pthread_cond_wait` in a loop

```
1 void *thread_entry(void *arg) {
2   /* thread init code */
3   while (1) {
4     qbench_update();
5     pthread_mutex_lock(&mutex);
6     while (!input_is_ready() && !qbench_update_requested()) {
7       qbench_thread_cond_wait(&cond, &mutex);
8     }
9     pthread_mutex_unlock(&mutex);
10    if (qbench_update_requested())
11      continue; /* reaches qbench_update */
12    /* ... handle connection */
13  }
14 }
```

Allows thread to be signaled for update even when waiting on a condition variable

Reports true if an update request is signaled

# Blocking on Condition Variables

Programmers guard against spurious wake-ups by placing `pthread_cond_wait` in a loop

```
1 void *thread_entry(void *arg) {
2   /* thread init code */
3   while (1) {
4     qbench_update();
5     pthread_mutex_lock(&mutex);
6     while (!input_is_ready() && !qbench_update_requested()) {
7       qbench_thread_cond_wait(&cond, &mutex);
8     }
9     pthread_mutex_unlock(&mutex);
10    if (qbench_update_requested())
11      continue; /* reaches qbench_update */
12    /* ... handle connection */
13  }
14 }
```

Allows thread to be signaled for update even when waiting on a condition variable

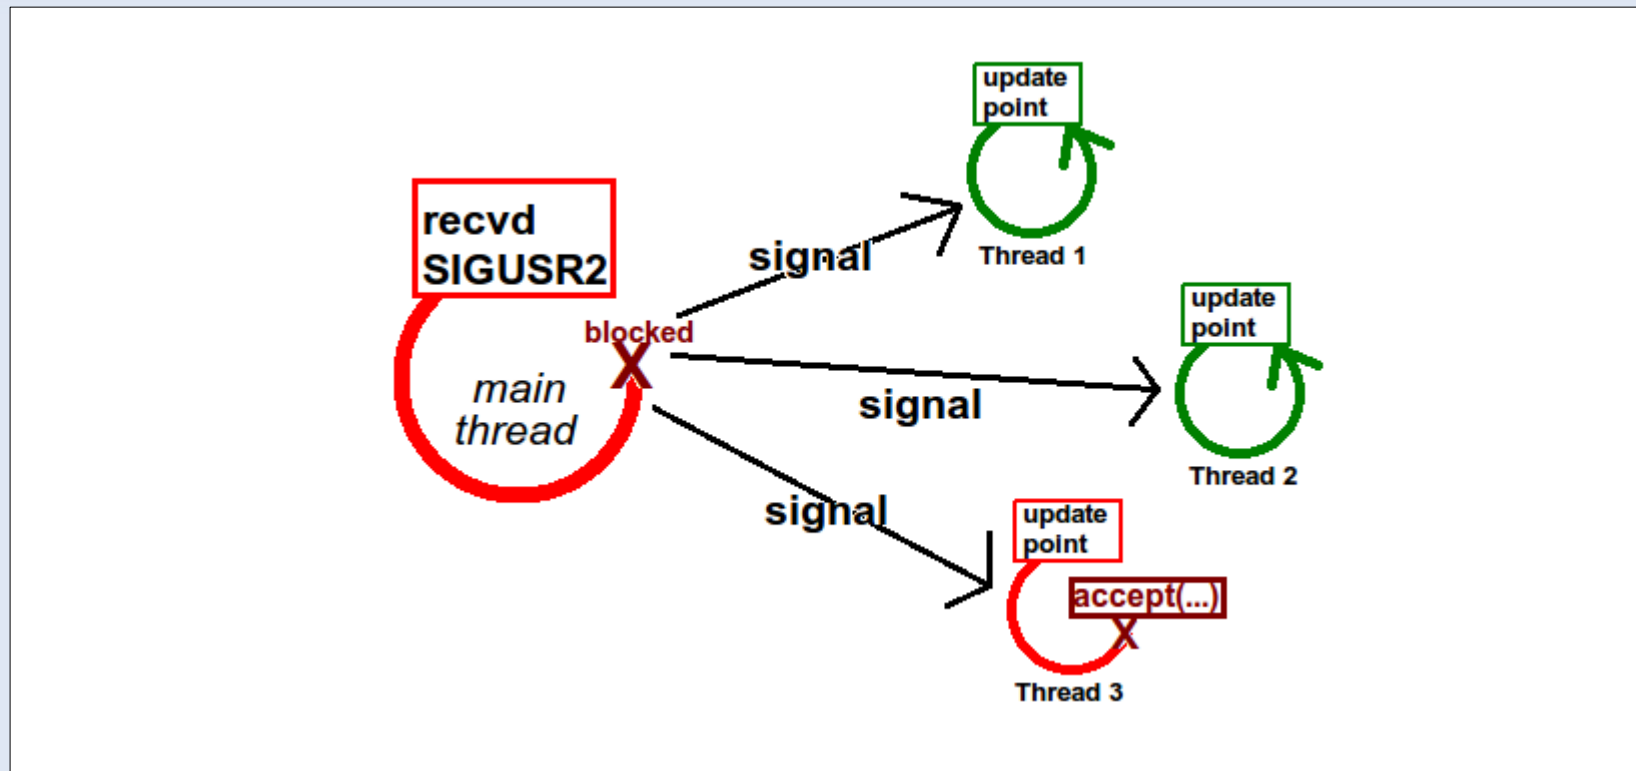
Reports true if an update request is signaled

Returns to top of even loop to immediately hit update point



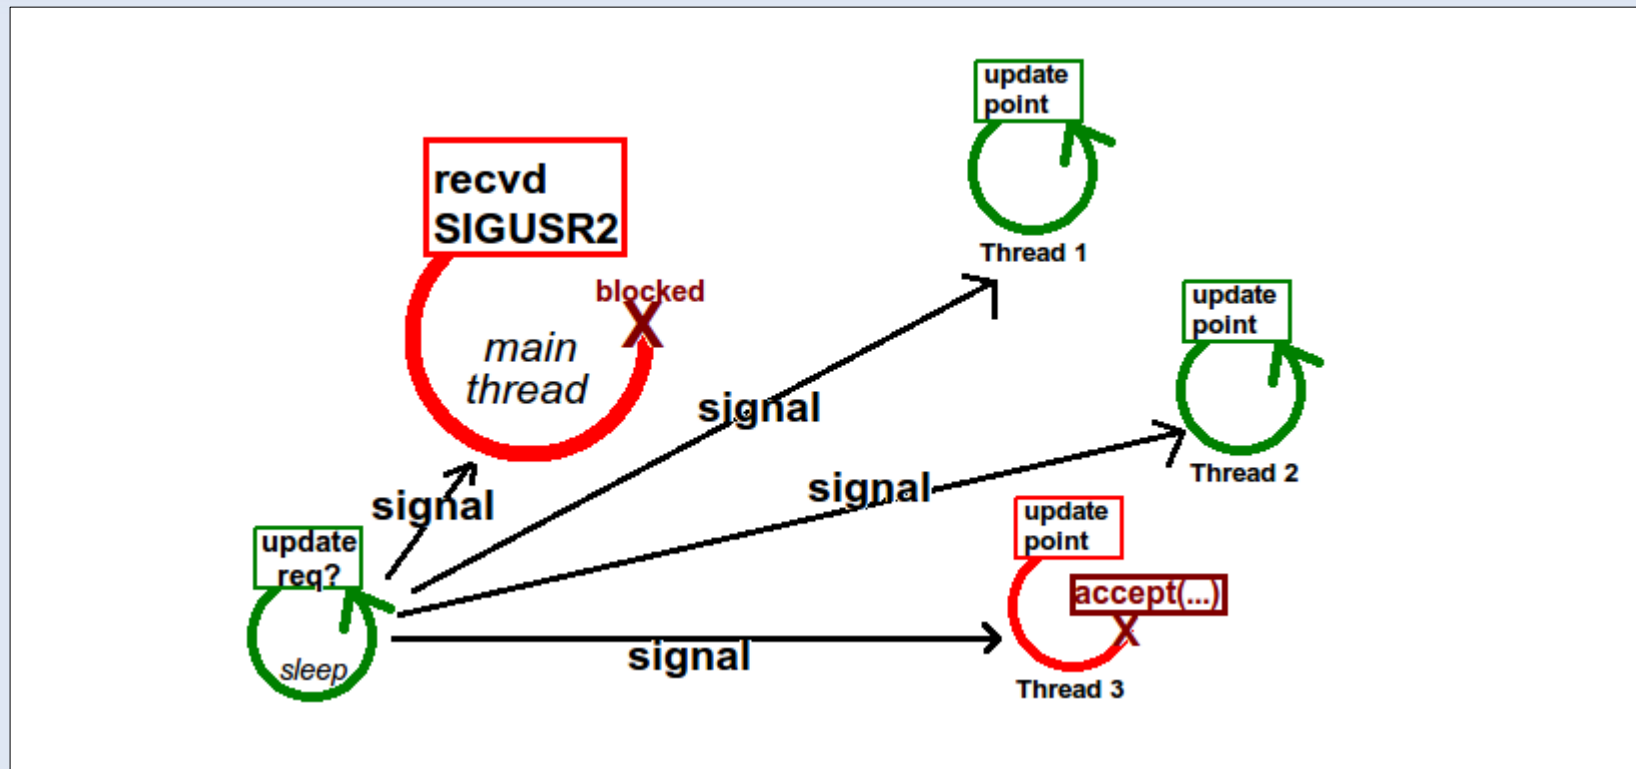
# Waking a Blocked Thread

- Condition Variables: Another thread must be available to signal



# Waking a Blocked Thread

- Condition Variables: Another thread must be available to signal



# Experiments

- We chose programs covering a wide range of domains.
- On average, 22 lines of code changed (including update points).
- Manual changes: changes beyond adding calls to QBench.

<b>Program</b>	<b>LoC Total</b>	<b># of Threads</b>	<b>Upd Points</b>	<b>Changed LoC (†)</b>	<b>Required Manual Chgs</b>
<i>httpd-2.2.22</i>	232651	$2 + c^*$ , $c = 3$	5	7 (5)	3 (Cond. Var. Loop)
<i>icecast-2.3.2</i>	17038	6	12	3 (3)	1 (Thread Sleeps)
<i>iperf-2.0.5</i>	3996	$3 + n^\circ$ , $n = 1$	5	8 (3)	1 (Cond. Var. Loop)
<i>memcached-1.4.13</i>	9404	$2 + c^*$ , $c = 4$	4	27 (4)	2 (libevent changes)
<i>space-tyrant-0.354</i>	8721	$3 + 2n^\circ$ , $n = 5$	6	8 (6)	1 (Thread Sleeps)
<i>suricata-1.2.1</i>	260344	$8 + c^*$ , $c = 3$	7	11 (6)	1 (libpcap break)

\*Configurable:  $c$  workers    $^\circ$ Varies by  $n$  connected clients    $^\dagger$ Calls to QBench excluding update

# Results

- Two Workloads:
  - Server idle (i.e., no connected clients)
  - Performing program-dependent work
- Nearly all programs quiesced in under 1ms
- Some would not quiesce without changes

Program	w/Load (ms)		w/o Load (ms)	
	All Chgs	UpdPt only	All Chgs	UpdPt only
<i>httpd-2.2.22</i>	0.185	0.230	0.123	0.150
<i>icecast-2.3.2</i>	105.152	954.32	107.558	986.265
<i>iperf-2.0.5</i>	0.193	DNQ	0.169	DNQ
<i>memcached-1.4.13</i>	0.166	DNQ	0.155	DNQ
<i>space-tyrant-0.354</i>	0.426	20.583	0.078	20.304
<i>suricata-1.2.1</i>	0.503	68.098	0.378	DNQ

DNQ = Does Not Quiesce

# Results

- Two Workloads:
  - Server idle (i.e., no connected clients)
  - Performing program-dependent work
- Nearly all programs quiesced in under 1ms
- Some would not quiesce without changes

Program	w/Load (ms)		w/o Load (ms)	
	All Chgs	UpdPt only	All Chgs	UpdPt only
<i>httpd-2.2.22</i>	0.185	0.230	0.123	0.150
<i>icecast-2.3.2</i>	105.152	954.32	107.558	986.265
<i>iperf-2.0.5</i>	0.193	DNQ	0.169	DNQ
<i>memcached-1.4.13</i>	0.166	DNQ	0.155	DNQ
<i>space-tyrant-0.354</i>	0.426	20.583	0.078	20.304
<i>suricata-1.2.1</i>	0.503	68.098	0.378	DNQ

DNQ = Does Not Quiesce

# Results

- Two Workloads:
  - Server idle (i.e., no connected clients)
  - Performing program-dependent work
- Nearly all programs quiesced in under 1ms
- Some would not quiesce without changes

Program	w/Load (ms)		w/o Load (ms)	
	All Chgs	UpdPt only	All Chgs	UpdPt only
<i>httpd-2.2.22</i>	0.185	0.230	0.123	0.150
<i>icecast-2.3.2</i>	105.152	954.32	107.558	986.265
<i>iperf-2.0.5</i>	0.193	DNQ	0.169	DNQ
<i>memcached-1.4.13</i>	0.166	DNQ	0.155	DNQ
<i>space-tyrant-0.354</i>	0.426	20.583	0.078	20.304
<i>suricata-1.2.1</i>	0.503	68.098	0.378	DNQ

DNQ = Does Not Quiesce

# Summary & Future Work

- Demonstrated multithreaded quiescence **quickly** and with **little implementation complexity** for many programs with fixed update points
- Time to quiescence ranged from 0.155 to 107.558 ms; **most were below 1 ms**
- We plan to integrate the multi-threaded quiescent functionality back into Kitsune