



Threads vs. Events


TAME – Event Style Programming with Threads

TAME

- expressive abstractions for event-based programming
- implemented via source-source translation
- avoids stack ripping
- type safety and composability via templates

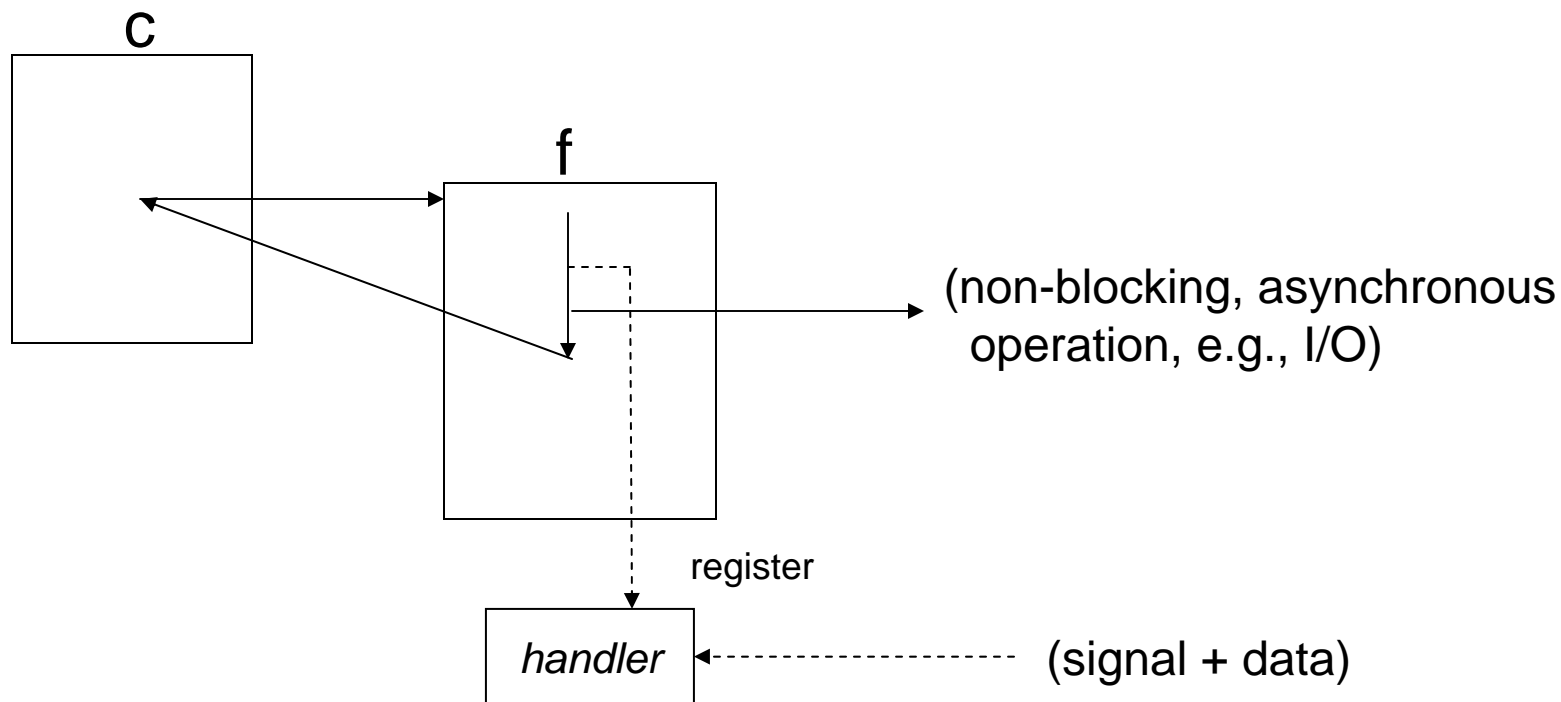
M. Krohn, E. Kohler, M.F. Kaashoek, “Events Can Make Sense,”
USENIX Annual Technical Conference, 2007, pp. 87-100.

A typical thread programming problem



Problem: the thread becomes blocked in the called routine (f) and the caller (c) is unable to continue even if it logically is able to do so.

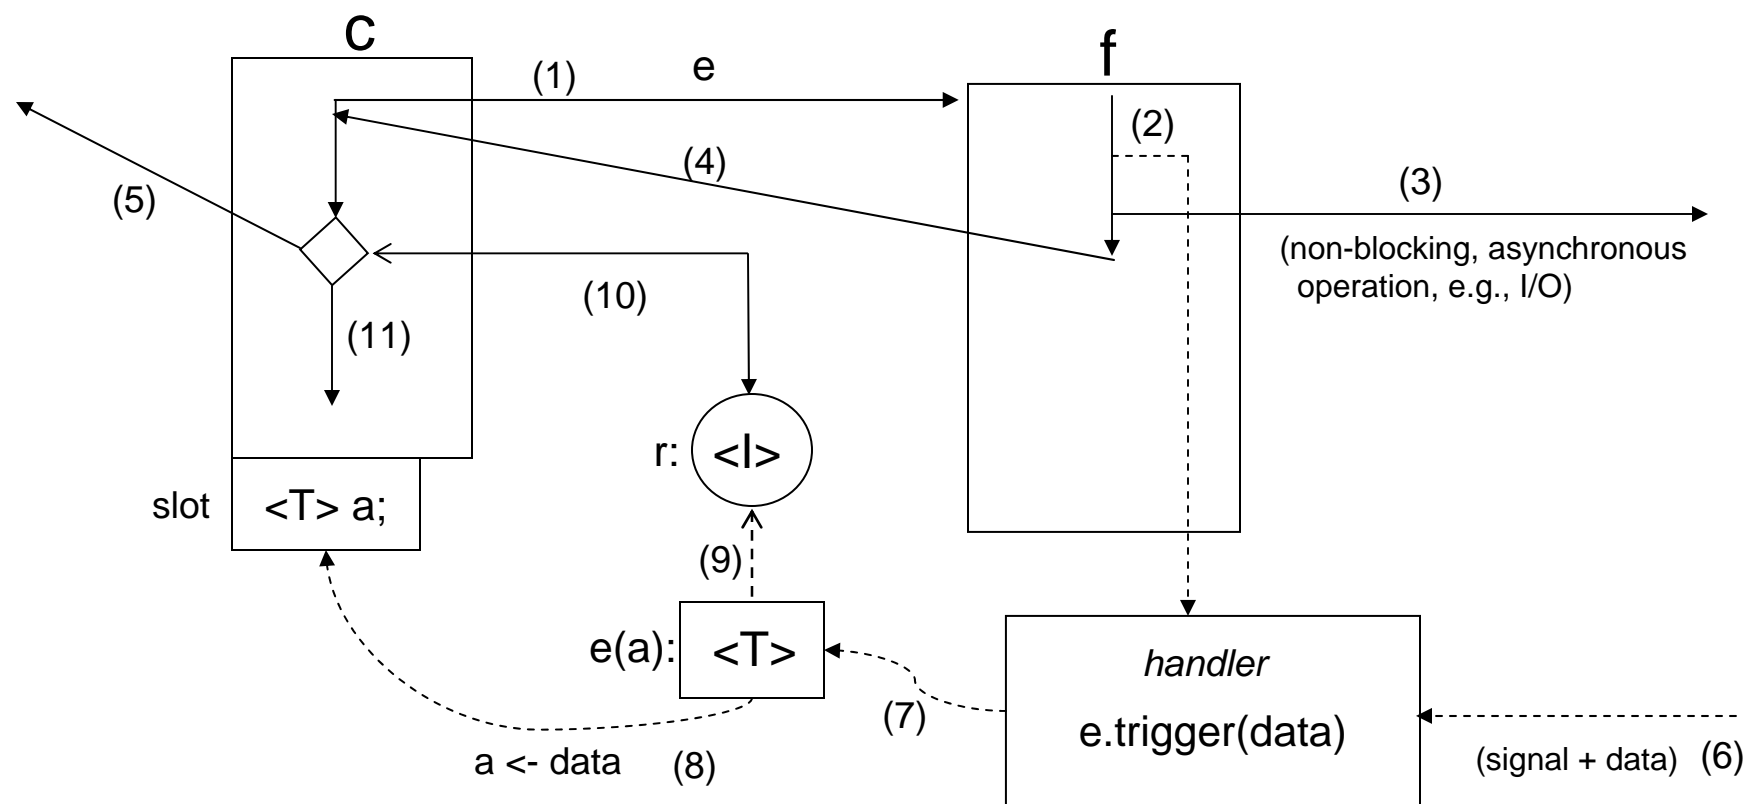
A partial solution



Issues

- Synchronization: how does the caller know when the signal has occurred without busy-waiting?
- Data: how does the caller know what data resulted from the operation?

A "Tame" solution



Tame Primitives

Classes	Keywords & Language Extensions	Functions & Methods
<p><code>event<></code></p> <ul style="list-style-type: none"> • A basic event. <p><code>event<T></code></p> <ul style="list-style-type: none"> • An event with a single <i>trigger value</i> of type <i>T</i>. This value is set when the event occurs; an example might be a character read from a file descriptor. Events may also have multiple trigger values of types $T_1 \dots T_n$. <p><code>rendezvous<I></code></p> <ul style="list-style-type: none"> • Represents a set of outstanding events with event IDs of type <i>I</i>. Callers name a rendezvous when they block, and unblock on the triggering of any associated event. 	<p><code>twait(r[i]);</code></p> <ul style="list-style-type: none"> • A wait point. Block on explicit rendezvous <i>r</i>, and optionally set the event ID <i>i</i> when control resumes. <p><code>tamed</code></p> <ul style="list-style-type: none"> • A return type for functions that use <code>twait</code>. <p><code>tvars { ... }</code></p> <ul style="list-style-type: none"> • Marks safe local variables. <p><code>twait { statements; }</code></p> <ul style="list-style-type: none"> • Wait point syntactic sugar: block on an implicit rendezvous until all events created in <i>statements</i> have triggered. 	<p><code>mkevent(r,i,s);</code></p> <ul style="list-style-type: none"> • Allocate a new event with event ID <i>i</i>. When triggered, it will awake rendezvous <i>r</i> and store trigger value in slot <i>s</i>. <p><code>mkevent(s);</code></p> <ul style="list-style-type: none"> • Allocate a new event for an implicit <code>twait{}</code> rendezvous. When triggered, store trigger value in slot <i>s</i>. <p><code>e.trigger(v);</code></p> <ul style="list-style-type: none"> • Trigger event <i>e</i>, with trigger value <i>v</i>. <p><code>timer(to,e); wait_on_fd(fd,rw,e);</code></p> <ul style="list-style-type: none"> • Primitive event interface for timeouts and file descriptor events, respectively.

Figure 2: Tame primitives for event programming in C++.

An example

```

1 void multidns(dnsname name[], ipaddr a[], int n) {
2     for (int i = 0; i < n; i++)
3         a[i] = gethostbyname(name[i]);
4 }

```

```

1 tamed multidns_tame(dnsname name[], ipaddr a[],
2                     int n, event<> done) {
3     tvars { int i; }
4     for (i = 0; i < n; i++)
5         twait { gethost_ev(name[i], mkevent(a[i])); }
6     done.trigger();
7 }

```

```
tamed gethost_ev(dsname name, event<ipaddr> e);
```

Variations on control flow

```
1  tamed multidns_par(dnsname name[], ipaddr a[],
                      int n, event<> done) {
2      twait {
3          for (int i = 0; i < sz; i++)
4              gethost_ev(name[i], mkevent(a[i]));
5      }
6      done.trigger();
7  }
```

parallel control
flow

```
1  tamed multidns_win(dnsname name[], ipaddr a[],
                      int n, event<> done) {
2      tvars { int sent(0), recv(0); rendezvous<> r; }
3      while (recv < n)
4          if (sent < n && sent - recv < WINDOWSIZE) {
5              gethost_ev(name[sent], mkevent(r,a[sent]));
6              sent++;
7          } else {
8              twait(r);
9              recv++;
10         }
11     done.trigger();
12 }
```

window/pipeline
control flow

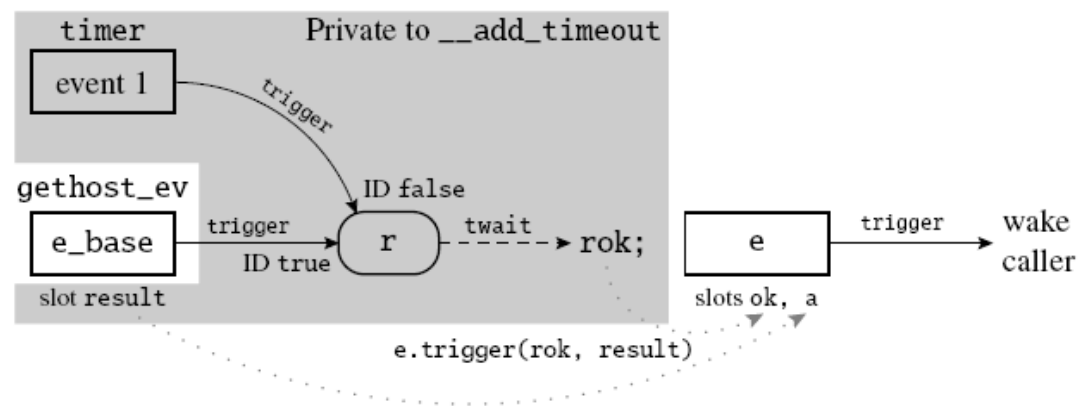
Event IDs & Composability

```

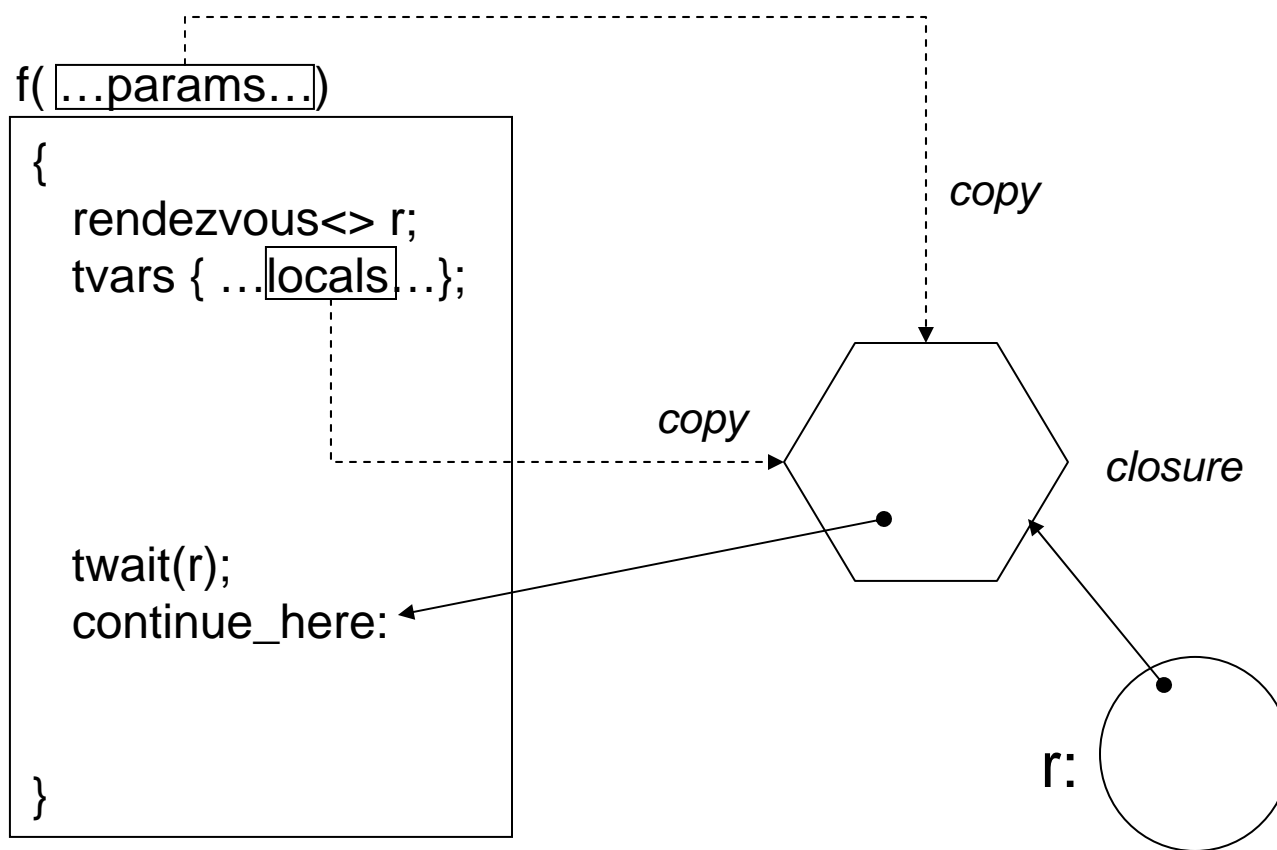
1  template <typename T> tamed
   __add_timeout(event<T> &e_base, event<bool, T> e) {
2      tvars { rendezvous<bool> r; T result; bool rok; }
3      timer(TIMEOUT, mkevent(r, false));
4      e_base = mkevent(r, true, result);
5      twait(r, rok);
6      e.trigger(rok, result);
7      r.cancel();
8  }

9  template <typename T> event<T> add_timeout(event<bool, T> e) {
10     event<T> e_base;
11     __add_timeout(e_base, e);
12     return e_base;
13 }

```



Closures



Smart pointers and reference counting insure correct deallocation of events, rendezvous, and closures.

Performance (relative to Capriccio)

	Capriccio	Tame
Throughput (connections/sec)	28,318	28,457
Number of threads	350	1
Physical memory (kB)	6,560	2,156
Virtual memory (kB)	49,517	10,740

Figure 7: Measurements of Knot at maximum throughput. Throughput is averaged over the whole one-minute run. Memory readings are taken after the warm-up period, as reported by ps.