Abstract

How to implement reliable system software with transactions. Theory and practice.
Error handling and concurrency are usually the most complicated to implement and test.
In this presentation we'll examine the I/O code of two example programs implemented in C. We'll first look at basic problems and afterwards how transactions can help to solve these.
On the practical side, we'll talk about the software *picotm*, a system-level transaction manager for POSIX systems.
Reimplementing the example programs on top of picotm will make them thread-safe and less error prone.
Picotm can handle arbitrary resources. In the presentation's final part, we'll look at the functionality that is currently provided, such as transactional memory, C string and memory functions, memory allocation, file-descriptor I/O, and others.

System-Level Transactions with *picotm*

Thomas Zimmermann

June 21, 2017

Handling Errors

Transitioning through consistent states; doing I/O in between.

```
1   int fd0, fd1; /* file descriptors */
2
3   char ibuf[100]; /* input buffer */
4   char obuf[100]; /* output buffer */
5
6   while (true) {
7
8       wait_for_input();
9
10      fill_input_buffer(ibuf);
11
12      compute_output_buffer(ibuf, obuf);
13
14      write(fd0, obuf, sizeof(obuf));
15      write(fd1, obuf, sizeof(obuf));
16  }
```

Handling Concurrency

Two threads writing concurrently to the same file.

```
1   int fd; /* file descriptor */
2
3   void thread_1_func() {
4       char obuf[100]; /* output buffer */
5
6       compute_output_buffer_t1(obuf); /* obuf = "42" */
7
8       pwrite(fd, obuf, sizeof(obuf), 256);
9   }
10
11  void thread_2_func() {
12      char obuf[100]; /* output buffer */
13
14      compute_output_buffer_t2(obuf); /* obuf = "23" */
15
16      pwrite(fd, obuf, sizeof(obuf), 256);
17  }
```

What We Actually Wanted

- To fix our examples, we have to ensure a number of constraints.

What We Actually Wanted

- To fix our examples, we have to ensure a number of constraints.

Atomic Our second example shall not write partial strings.

What We Actually Wanted

- To fix our examples, we have to ensure a number of constraints.

Atomic Our second example shall not write partial strings.

Consistent Our first example shall output consistent data to both files.

What We Actually Wanted

- To fix our examples, we have to ensure a number of constraints.

Atomic	Our second example shall not write partial strings.
Consistent	Our first example shall output consistent data to both files.
Isolated	Threads in our second example shall not interfere.

What We Actually Wanted

- To fix our examples, we have to ensure a number of constraints.

Atomic	Our second example shall not write partial strings.
Consistent	Our first example shall output consistent data to both files.
Isolated	Threads in our second example shall not interfere.
Durable	Bonus point: Once we made a write, it should not disappear.

What We Actually Wanted

- To fix our examples, we have to ensure a number of constraints.

Atomic	Our second example shall not write partial strings.
Consistent	Our first example shall output consistent data to both files.
Isolated	Threads in our second example shall not interfere.
Durable	Bonus point: Once we made a write, it should not disappear.

- So what we actually wanted are transactions!

Introducing Transactions

- Transactional semantics mandate

Atomicity	Do everything, or nothing.
Consistency	Work on consistant data.
Isolation	Isolate transactions from each other.
Durability	Effects of committed transactions do not disappear.

Introducing Transactions

- Transactional semantics mandate

Atomicity	Do everything, or nothing.
Consistency	Work on consistant data.
Isolation	Isolate transactions from each other.
Durability	Effects of committed transactions do not disappear.

- Many databases around, but hardly anything for arbitrary software.

Introducing Transactions

- Transactional semantics mandate

Atomicity	Do everything, or nothing.
Consistency	Work on consistant data.
Isolation	Isolate transactions from each other.
Durability	Effects of committed transactions do not disappear.

- Many databases around, but hardly anything for arbitrary software.
- Enter *picotm* ***drum rolls***

Put into Practice with *picotm*

- *picotm* is a transaction manager for C applications and firmware

Put into Practice with *picotm*

- *picotm* is a transaction manager for C applications and firmware

Basic C interface of *picotm*

```
1  picotm_begin
2      /* execution phase; put your code here */
3
4  picotm_commit  /* commit phase; provided by picotm */
5
6      /* recovery phase; put your error handling here */
7  picotm_end
```

Handling Errors, transactionally

Transitioning through consistent states; doing transactional I/O in between.

```
1   int fd0, fd1; /* file descriptors */
2
3   char ibuf[100]; /* input buffer */
4   char obuf[100]; /* output buffer */
5
6   while (true) {
7
8       wait_for_input();
9
10      picotm_begin
11
12          fill_input_buffer_tx(ibuf);
13
14          compute_output_buffer_tx(ibuf, obuf);
15
16          write_tx(fd0, obuf, sizeof(obuf));
17          write_tx(fd1, obuf, sizeof(obuf));
18
19      picotm_commit
20
21          if (picotm_error_is_non_recoverable()) {
22              notice_admin_and_abort();
23          } else {
24              handle_error_and_retry();
25              picotm_restart();
26          }
27
28      picotm_end
29  }
```

Transaction Log

- *picotm* keeps a log of all operations that are
  - delayed until commit time, or
  - reverted during a rollback.

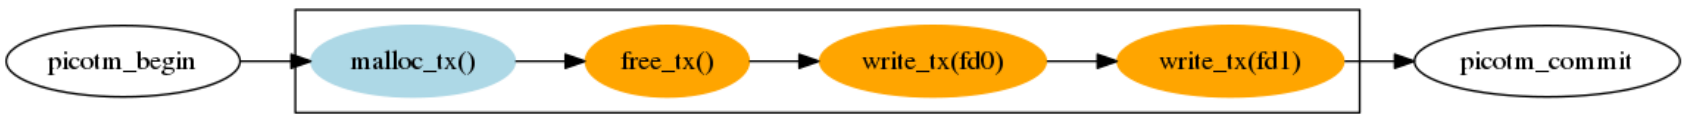

Figure: The complete transaction log for example 1. Delayed operations are displayed in Orange, revertable operations are in Light Blue.

Handling Concurrency, transactionally

Two threads writing transactionally to the same file.

```
1  int fd; /* file descriptor */
2
3  void thread_1_func()
4  {
5      picotm_begin
6          char obuf[100]; /* output buffer */
7          compute_output_buffer_t1_tx(obuf); /* obuf = "42" */
8          pwrite_tx(fd, obuf, sizeof(obuf), 256);
9      picotm_commit
10         if (picotm_error_is_non_recoverable()) {
11             notice_admin_and_abort();
12         } else {
13             handle_error_and_retry();
14             picotm_restart();
15         }
16     picotm_end
17 }
18
19 void thread_2_func()
20 {
21     picotm_begin
22         char obuf[100]; /* output buffer */
23         compute_output_buffer_t2_tx(obuf); /* obuf = "23" */
24         pwrite_tx(fd, obuf, sizeof(obuf), 256);
25     picotm_commit
26         [...]
27     picotm_end
28 }
```

Modules of *picotm*

- All application functionality is provided by modules
- Modules can be combined as needed
- New modules can be added

Module interface for interacting with *picotm*.

```
1  /* Register a module */
2  unsigned long
3  picotm_register_module(picotm_module_lock_function lock,
4                         picotm_module_unlock_function unlock,
5                         picotm_module_is_valid_function is_valid,
6                         picotm_module_apply_function apply,
7                         picotm_module_undo_function undo,
8                         picotm_module_apply_events_function apply_events,
9                         picotm_module_undo_events_function undo_events,
10                        picotm_module_update_cc_function update_cc,
11                        picotm_module_clear_cc_function clear_cc,
12                        picotm_module_finish_function finish,
13                        picotm_module_uninit_function uninit);
14
15 /* Append event to transaction log */
16 void
17 picotm_append_event(unsigned long module, unsigned long op, uintptr_t cookie);
18
19 /* Inform picotm about an error */
20 void
21 picotm_recover_from_error(const struct picotm_error* error);
```

Transactional Memory

- load_tx()
- store_tx()
- privatize_tx()

String and Memory helpers

- memcpy_tx(), memcmp_tx(), etc.
- strcpy_tx(), strcmp_tx(), etc.

Memory Allocation

- malloc_tx()
- free_tx()

File I/O

- open_tx(), close_tx()
- read_tx(), write_tx()
- pread_tx(), pwrite_tx()

Others

- errno
- C Standard Math Library
  - Math functions
  - Floating-Point environment
  - Floating-Point exceptions
- Some VFS support

Summary

- Transactional code is safer and less error prone than traditional one.
- Implement error handling and concurency control *exactly* once.
- *picotm* is available as Open Source at
  - picotm.org
- More information, tutorials, background on my blog at
  - transactionblog.org
  - twitter.com/transactionblog
- Or reach out to me via
  - tdz@users.sourceforge.net

picotm.org