Threads

Daniel Bosk[1]

Department of Information and Communication Systems (ICS),
Mid Sweden University, Sundsvall.

thread.tex 231 2018-02-05 08:55:44Z jimahl

Mittuniversitetet
MID SWEDEN UNIVERSITY

Overview

1. Threads
   - What is a thread?
   - Benefits
   - Issues
2. Implementation of threads
   - Types of threads
   - The different models
3. Signals
   - What are signals?
4. Playing with Threads in Python
   - With race-condition
   - Without race-condition

Literature

This lecture covers threads. It gives an overview of Chapter 4 “Multithreaded Programming” in [SGG13a]

Overview

What is a thread?

- The process is what the operating system consider the smallest entity of execution.
- This contains the program code, the variables (data), etc.
- A process traditionally has only one *thread* of execution. Such a process is called a heavy-weight process.
- Now, however, we'll extend this with more threads of execution.

What is a thread?

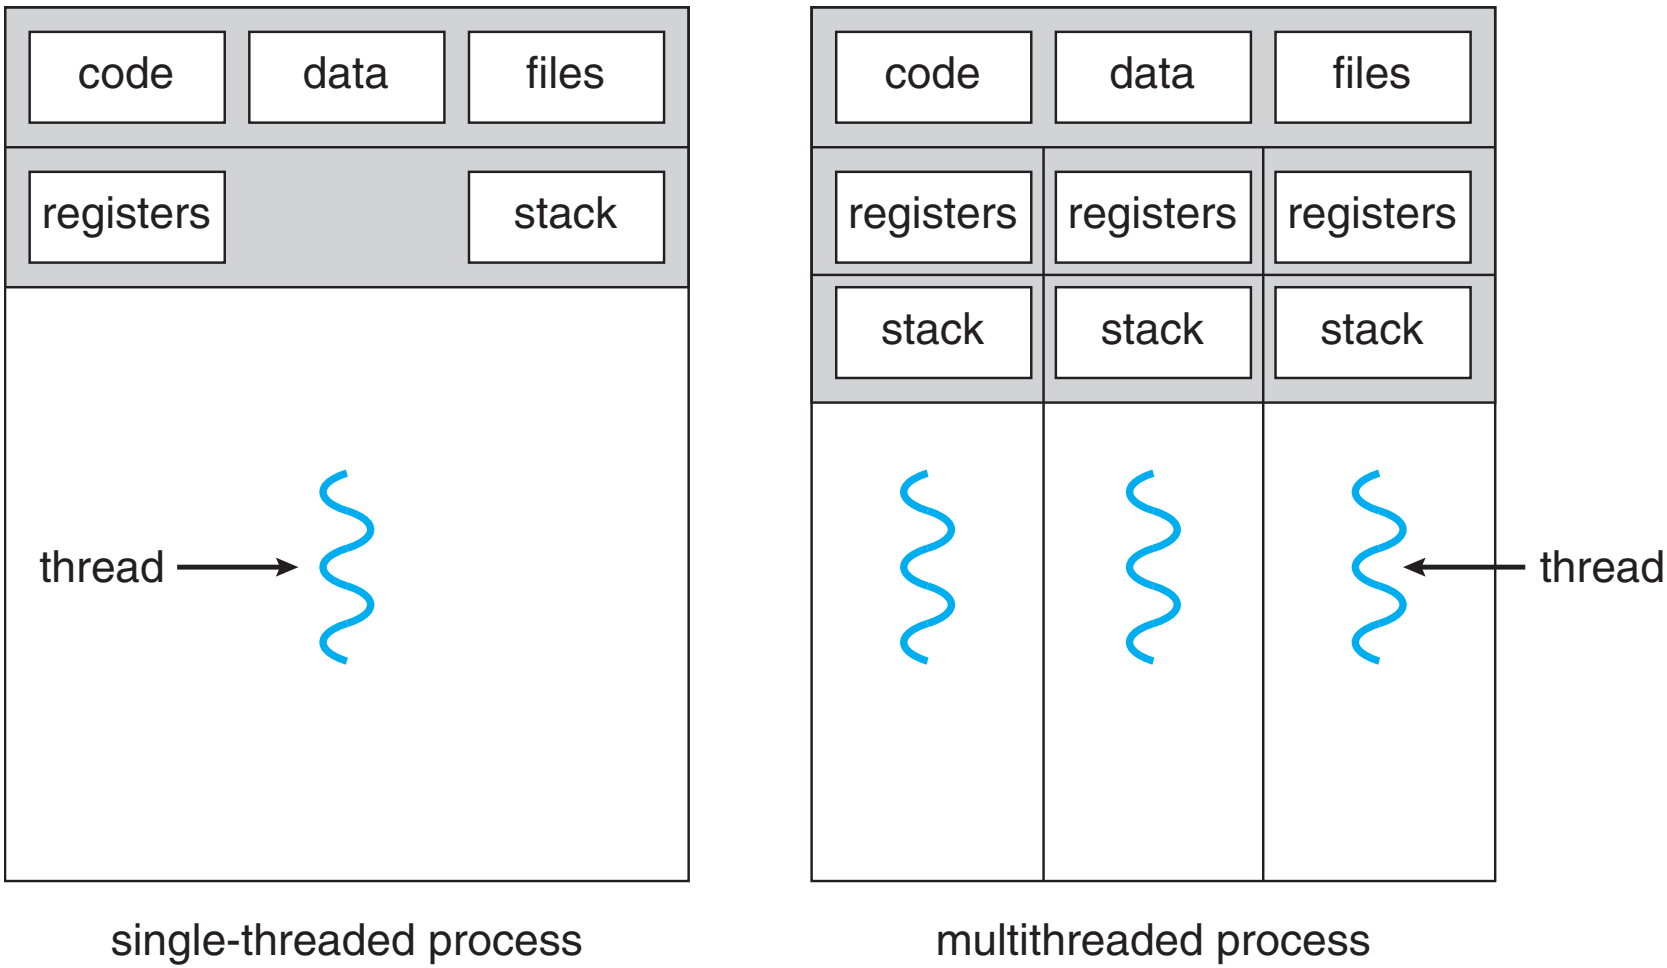

Figure: Single-threaded vs. multithreaded process. Image: [SGG13b].

What is a thread?

- Multiple threads of execution means we can do many things simultaneously in a process.
- Have to be careful though, you never know who is changing something in the process at a given time.

Benefits

- Responsiveness
- Resource sharing
- Economy (context switching, process creation)
- Scalability

Benefits
Scalability

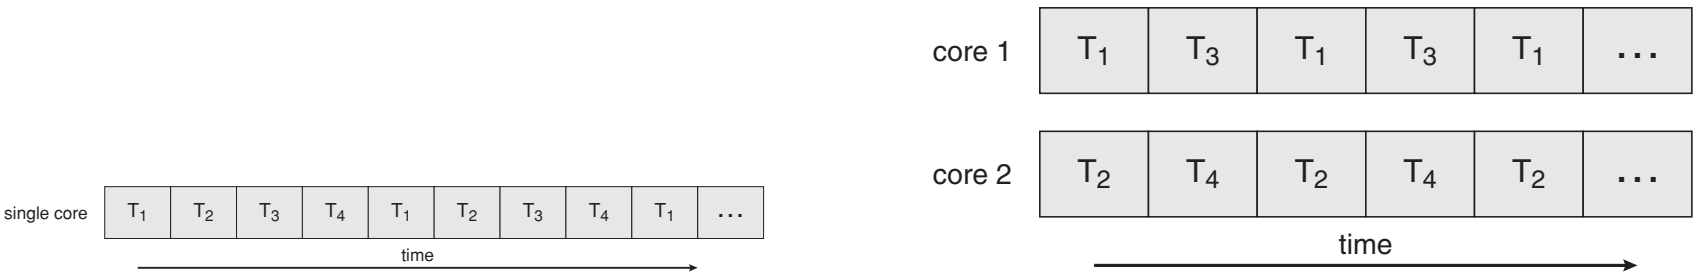

Figure: A single-core and multi-core execution of four threads. Image: [SGG13b].

Issues

- What happens with fork(2) and exec(2)? The calling thread only, or all threads?
- exec(2) usually replaces all threads.
- fork(2) can be different; in some cases only the calling thread is reasonable to fork, in other cases all of them.
- Programming becomes more complex, introducing many potential bugs.

Overview

Types of threads

User thread These execute in user-mode, and are invisible to the kernel.

Kernel thread These are part of the kernel. They do not necessarily execute all code in kernel-mode though.

Types of threads

- User threads are mapped to kernel threads.
- This mapping can be done in several ways:
  - One-to-one,
  - Many-to-one,
  - Many-to-many.

The different models

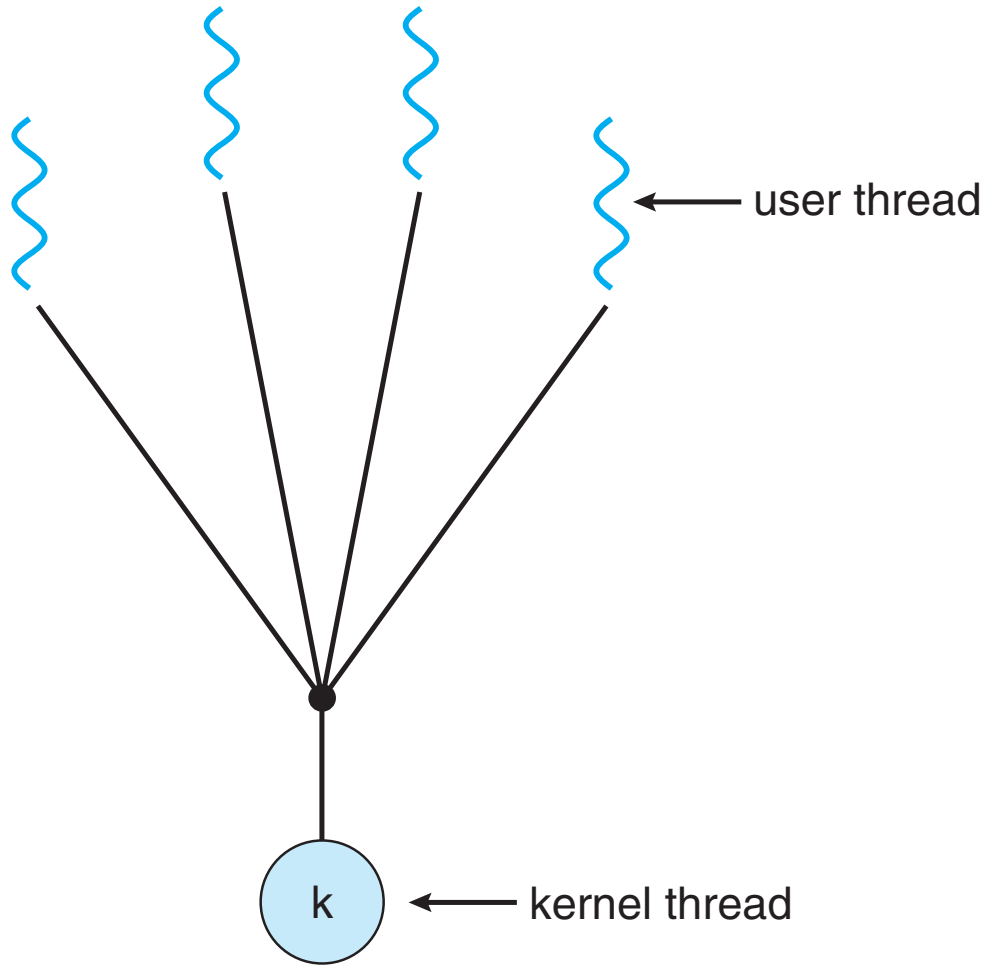

Figure: The many-to-one model. Image: [SGG13b].

The different models

- Can achieve responsiveness, but is not scalable.
- Bad in that if one thread blocks, e.g. through a system call, the whole process blocks, and thus all threads block.
- Can be combined with the one-to-one model though, then threads can be partitioned.

The different models

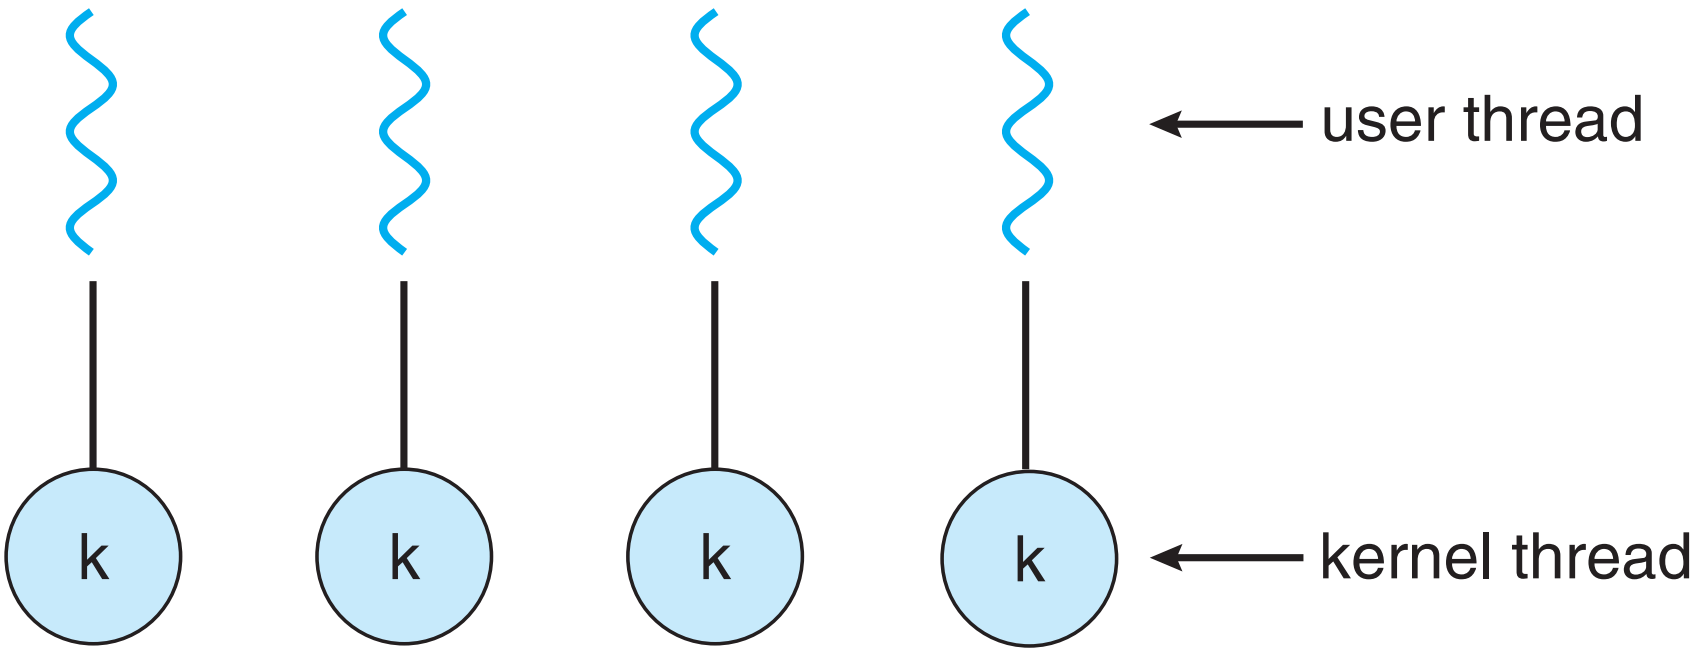

Figure: The one-to-one model. Image: [SGG13b].

The different models

- Good in that it's the OS handling all scheduling.
- Bad in that it requires as many kernel-threads as user-threads.

The different models

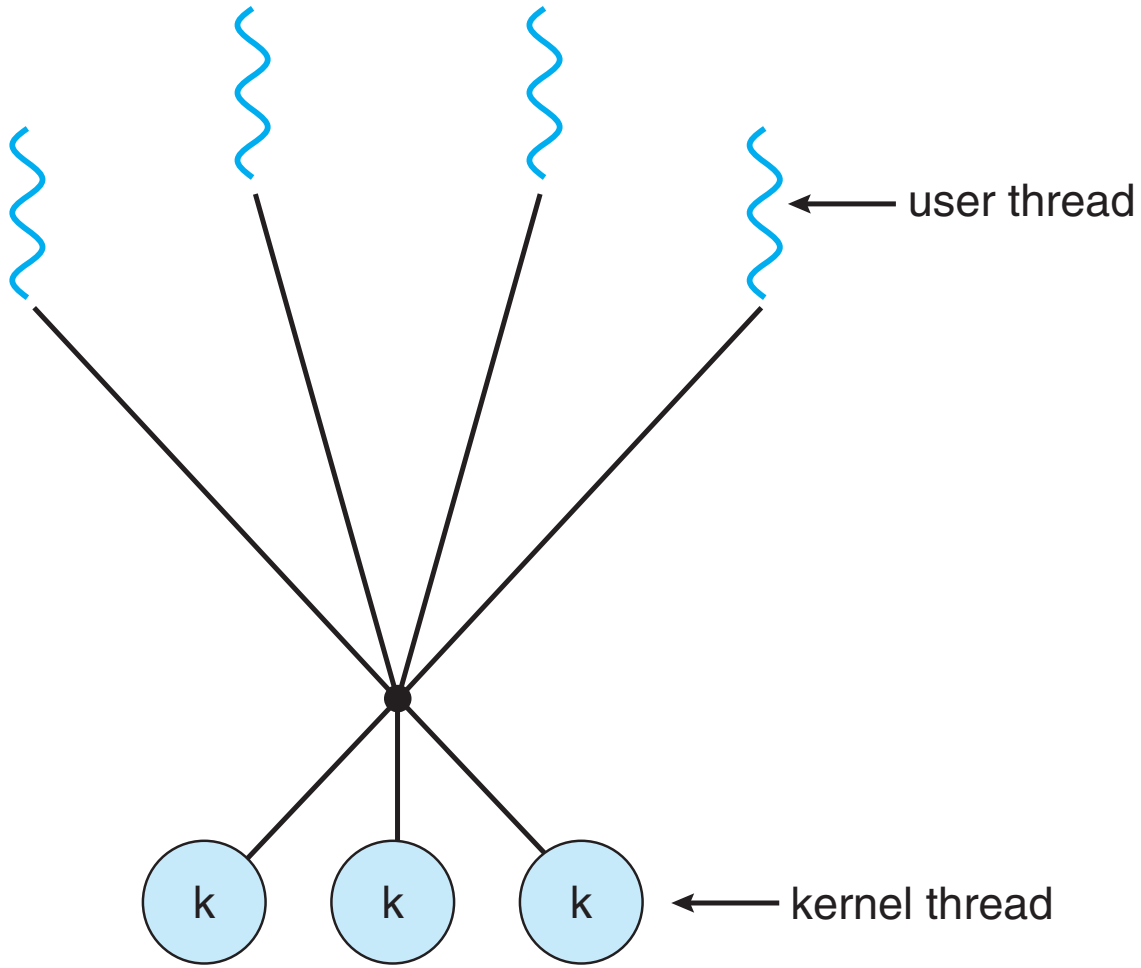

Figure: The many-to-many model. Image: [SGG13b].

The different models

- The many-to-many model requires a light-weight process (LWP).
- To the process this appears as a processor on which it can schedule its threads.
- I.e. the thread library does the scheduling of its threads, as is the case in many-to-one also.
- The kernel schedules all threads in the one-to-one model.

The different models

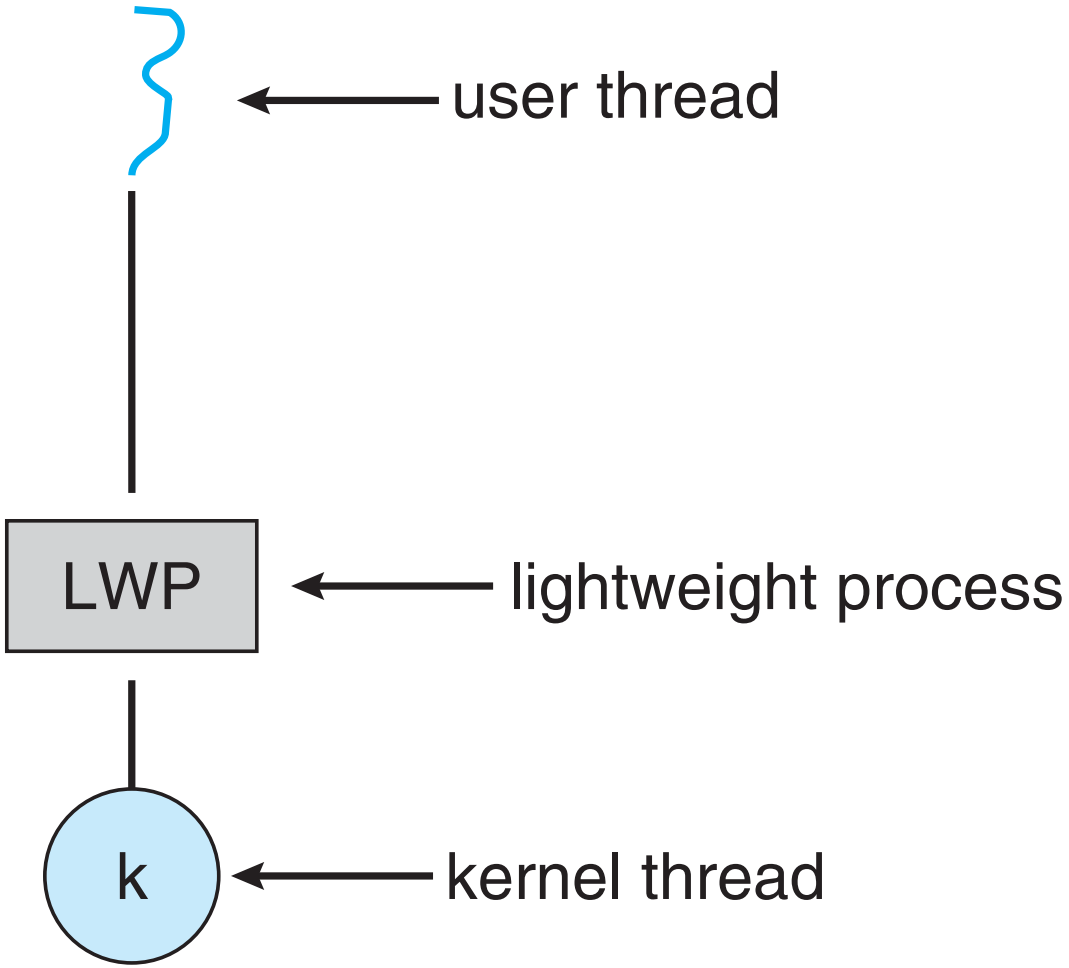

Figure: Structure including an LWP. Image: [SGG13b].

The different models

- The many-to-many model has the advantage that it has none of the problems of the other models.
- It can have an arbitrarily large amount of threads.
- If one thread blocks, i.e. one LWP is blocked, just schedule another LWP.

The different models

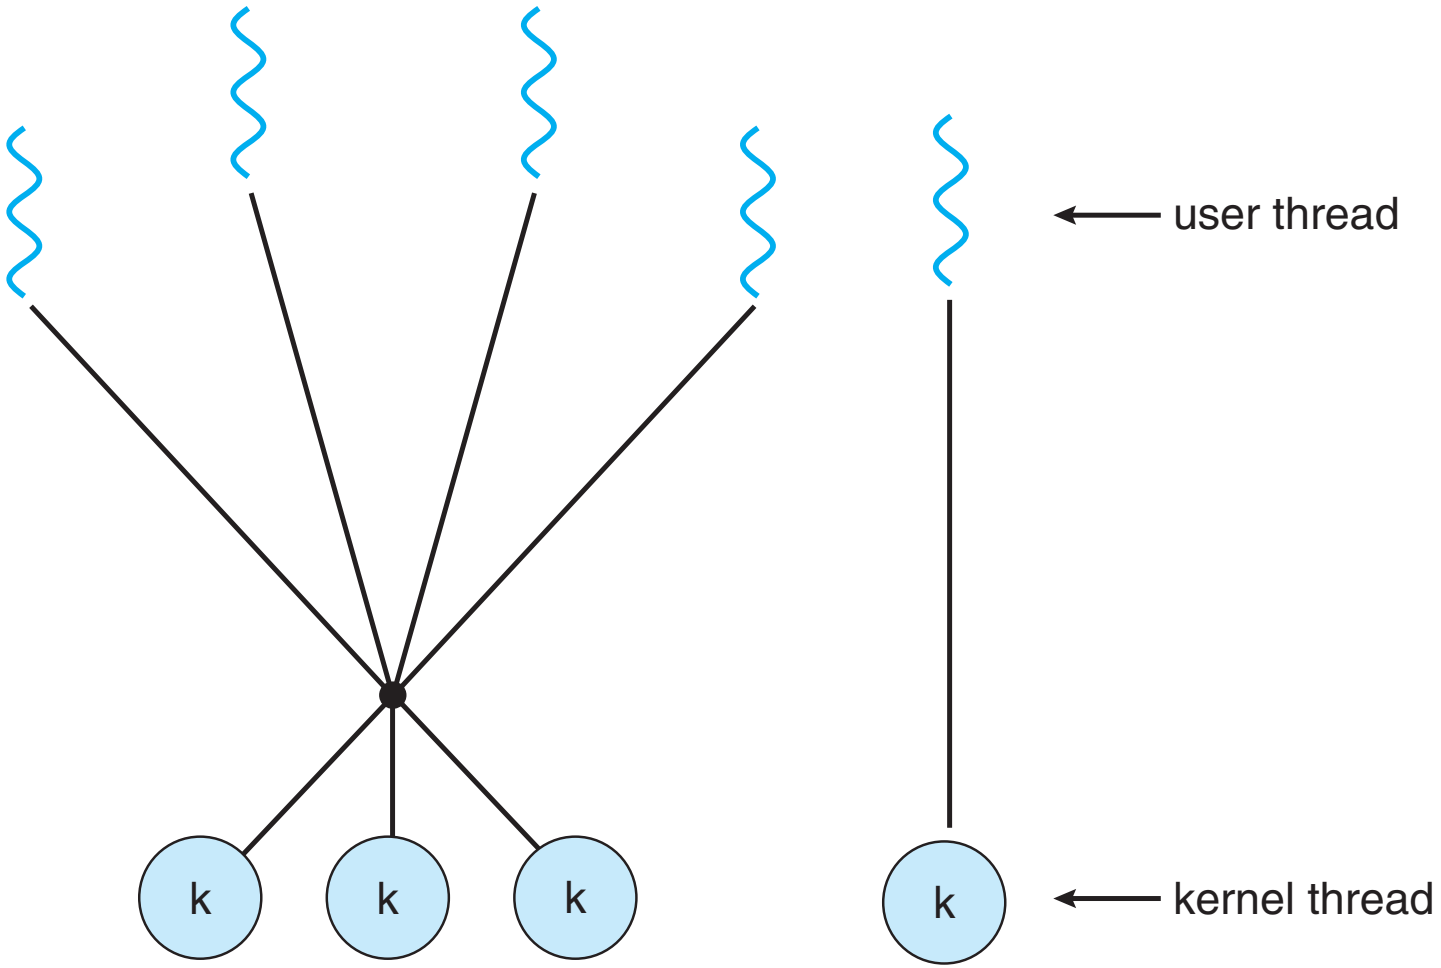

Figure: The two-level model. Image: [SGG13b].

Overview

What are signals?

- A signal is like an interrupt for a process.
- The OS sends the process a signal.
- The process must stop its current work to handle the signal, then it may return to previous work.
- Examples include “Division by zero”, “Illegal memory access” and hitting Ctrl+C keyboard combination.

What are signals?

1. Process executes division by zero.
2. CPU generates interrupt, calls OS interrupt handler.
3. OS interrupt handler notes the process currently using the CPU.
4. OS generates as signal to the process.
5. Process's signal handler executes.

What are signals?

- With a multithreaded process, which thread should handle a signal from the OS?
- This can usually be specified by the programmer.

Overview

With race-condition I

```
#!/usr/bin/env python3

# Author:   Daniel Bosk <daniel.bosk@miun.se>
# Date:     15 May 2012

import sys, threading, time

# function to be run in each separate thread
def test(thread_id, delay):
  for i in range(2):
    for j in range(5):
      # critical section
      # print values of i and j
      print( str(thread_id) + ": i=" + str(i) +
          ", j=" + str(j) )
      # remainder section
      # sleep for delay seconds
```

With race-condition II

```
17          time.sleep( float(delay) )
18
19  def main():
20      # default to using two threads.  this can be
            overridden by passing a number
21      # as argument from the command line.
22      # usage: race.py <nthreads>
23      nthreads = 2
24      if len(sys.argv) > 1:
25          nthreads = int( sys.argv[1] )
26
27      threads = []
28      for n in range( nthreads ):
29          # create a thread which runs the
                test-function above, documentation:
30          #
                http://docs.python.org/library/threading.html?
```

With race-condition III

```
31      t = threading.Thread( target=test,
32          args=("thread"+str(n), 1+float(n)/10) )
33      t.start()
34      threads.append( t )
35
36    print( "waiting␣..." )
37    # wait for all threads to finish
38    for t in threads:
39      t.join()
40    print( "done" )
41
42  if __name__ == "__main__":
43    main()
```

Without race-condition I

```
 1 #!/usr/bin/env python3
 2
 3 # Author:  Daniel Bosk <daniel.bosk@miun.se>
 4 # Date:    15 May 2012
 5
 6 import sys, threading, time
 7
 8 # function to be run in each separate thread
 9 def test(lock, thread_id, delay):
10   for i in range(2):
11     for j in range(5):
12       # entry section, wait for lock
13       lock.acquire()
14       # critical section
15       # print values of i and j
16       print( str(thread_id) + ": i=" + str(i) +
             ", j=" + str(j) )
```

Without race-condition II

```
17          # exit section, release lock
18          lock.release()
19          # remainder section
20          # sleep for delay seconds
21          time.sleep( float(delay) )
22
23  def main():
24      # prepare a lock for stdout, to synchronise
            output
25      stdout = threading.Lock()
26
27      # default to using two threads.  this can be
            overridden by passing a number
28      # as argument from the command line.
29      # usage: ./norace.py <nthreads>
30      nthreads = 2
31      if len(sys.argv) > 1:
```

Without race-condition III

```
32      nthreads = int( sys.argv[1] )
33
34    threads = []
35    for n in range( nthreads ):
36      # create a thread which runs the
            test-function above, documentation:
37      #
            http://docs.python.org/library/threading.html?
38      t = threading.Thread( target=test,
39          args=(stdout, "thread"+str(n),
                1+float(n)/10) )
40      t.start()
41      threads.append( t )
42
43    print( "waiting␣..." )
44    # wait for all threads to finish
45    for t in threads:
```

Without race-condition IV

```
46      t.join()
47    print( "done" )
48
49 if __name__ == "__main__":
50    main()
```

Referenser I

Abraham Silberschatz, Peter Baer Galvin, and Greg Gagne. *Operating System Concepts*. 9th ed. International Student Version. Hoboken, N.J.: John Wiley & Sons Inc, 2013.

Abraham Silberschatz, Peter Baer Galvin, and Greg Gagne. *Operating System Concepts*. 9th ed. Hoboken, N.J.: John Wiley & Sons Inc, 2013.