

Verification of System Calls in PikeOS

Christoph Baumann, Holger Blasum, Thorsten Bormer

Saarland Univ., SYSGO AG, Univ. of Koblenz-Landau

19.05.2009

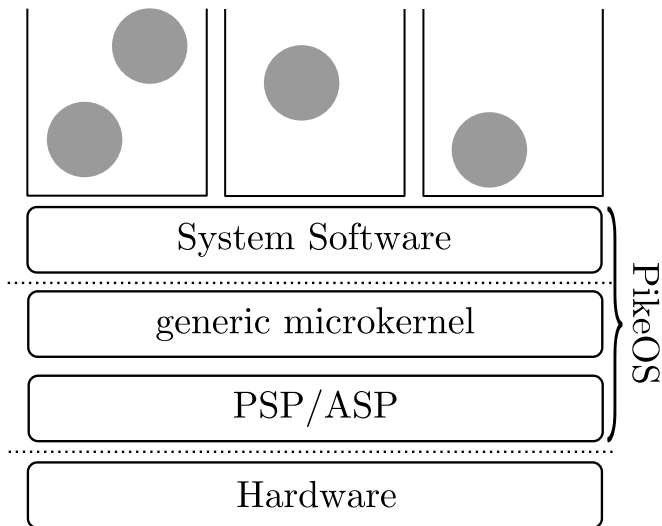
Our Goal

The goal of our subproject is to ...

- verify a “real-world” microkernel (PikeOS)
- using the VCC toolchain
- ~→ case study for large scale verification, VCC

- microkernel for use in safety and security-critical systems developed by SYSGO AG for multiple architectures
- verification target: PowerPC architecture, the snapshot under analysis is of compact size: 10% assembly language, rest in C

PikeOS – System Architecture



Verification of the PikeOS kernel

Verification Progress

- Done: helper functions only visible inside the kernel, sequential setting
- Goal: specify externally visible behavior of the kernel, concurrently executing

First target: system calls

Tasks

- abstract model of kernel state; prove refinement relation
- model and verify (inline) assembly instructions in VCC
- specify and verify sequential execution of system call
- adapt verification to concurrent setting

Verification of the PikeOS kernel

Prove that VCC methodology fits to concurrency model in PikeOS

- prove that scheduling operation is not visible to current thread (separation properties)
- prove functional correctness of the scheduler
- adapt specification to VCC technicalities

Verification of Hardware-related Layers

- model of PPC hardware as VCC ghost structure
 - introduce one VCC spec function for each assembly instruction
 - replace assembly instructions by C spec. functions
- ⇒ (inline) assembly can be verified as usual with VCC

Requirement Specification: p4_fast_set_prio

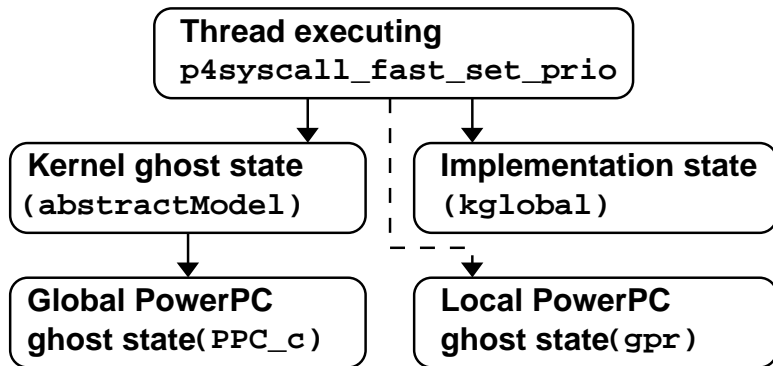
From the kernel reference manual:

“This function sets the current thread’s priority to newprio. Invalid or too high priorities are limited to the caller’s task MCP. Upon success, a call to this function returns the current thread’s priority before setting it to newprio.”

Implementation: p4_fast_set_prio(-helper)

```
1 P4_prio_t p4_runner_changeprio
2 (P4k_thrinfo_t *proc, P4_prio_t newprio)
3 {
4     P4_prio_t oldprio; P4_cpureg_t oldstat;
5
6     oldstat = p4arch_disable_int();
7     oldprio = proc->userprio;
8     proc->userprio = newprio;
9     proc->schedprio = newprio;
10    kglobal.kinfo->currprio = newprio;
11    p4arch_restore_int(oldstat);
12
13    return oldprio;
14 }
```

PikeOS Entities in our Verification Setup



Abstract Kernel Model

```
1 spec( struct absModel_str {
2     bool interruptsEnabled;
3     invariant(interruptsEnabled ==
4         (PPC_c.msr.fld.EE == 1))
5
6     struct P4k_thrinfo_t *currentThread;
7     invariant(currentThread != NULL)
8
9     invariant(keeps(currentThread, &PPC_c))
10 } abstractModel; )
```



Specification: p4_runner_changeprio

```
1 P4_prio_t p4_runner_changeprio
2 (P4k_thrinfo_t *proc, P4_prio_t newprio)
3   requires(proc ==
4             abstractModel.currentThread)
5   ensures(proc->schedprio == newprio && ...)
6   returns(old(proc->userprio))
7
8   maintains(wrapped(...))
9   writes(...)
10 {
11   ...
12 }
```

Results – Sequential Setting

- abstract model of PikeOS
- proof of refinement relation between abstract model and concrete state
- proof of sequential behavior of first system calls in terms of abstract model

Concurrent Setting



Consequences for the Specification of setprio

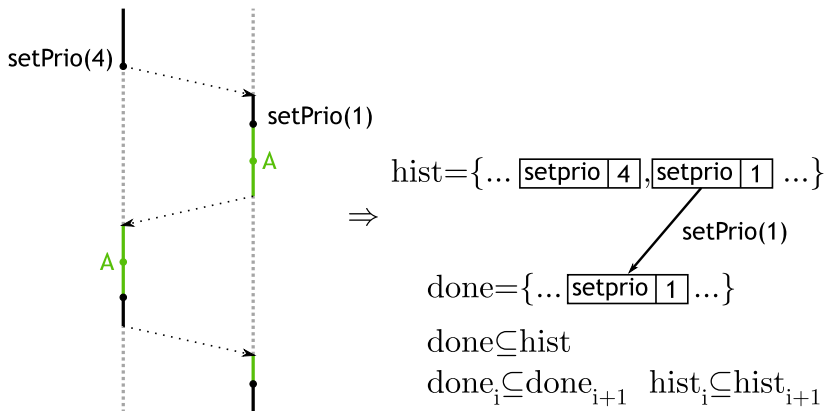
In the concrete implementation, other threads may interfere after atomic block

- one of the threads wins, but we don't know which
- only a rather weak invariant can be shown without further information

⇒ introduce history for system calls

Concurrent Specification of setprio

Idea: each invocation of a system call is recorded in a history



Definition of thread data structure

```
1 typedef struct update {
2     int id;
3     int value;
4 } update, *pUpdate;
5
6 typedef struct thread {
7     volatile int prio;
8
9     spec(volatile ptrset done;);
10    spec(volatile ptrset hist;);
11 } thread, *pThread;
```

Invariants of thread

```
1 //hist contains only updates
2 invariant(forall(obj_t o; set_in(o, hist)
3 ==> is(o, update) && set_in(o, owns(this))))
4
5 //D is a subset of H
6 invariant(set_subset(done, hist))
7
8 //H only increases
9 invariant(set_subset(old(hist), hist))
10
11 //D only increases
12 invariant(set_subset(old(done), done))
```

Invariants of thread

We execute each update from hist, but only once:

```
1 invariant(unchanged(prio) ||  
2   exists(update *u; set_in((obj_t) u, done)  
3       && !set_in((obj_t) u, old(done))  
4       && prio == u->value))
```

Implementation of setprio for VCC

```
1 void setPrio(pThread t, int v)
2 {
3     atomic(...) {
4         t->prio = v;
5     }
6
7     atomic(...) {}
8 }
```

Concurrent Specification of setprio

```
1 void setPrio(pThread t, int v
2     spec(update *up))
3     maintains(up->value == v)
4
5     requires(set_in((obj_t) up, t->hist)
6         && !set_in((obj_t) up, t->done))
7
8     ensures(exists(update *u;
9         set_in((obj_t) u, t->done)
10            && !set_in((obj_t) u, old(t->done))
11            && t->prio == u->value))
12
13     ensures(set_in((obj_t) up, t->done))
```

Conclusion

Results

- verification of concurrent system call with histories
- analysis of PikeOS concurrency model w.r.t VCC model

Further Work

- solve remaining technicalities
- extend verification to other system calls
- prove concurrency model