
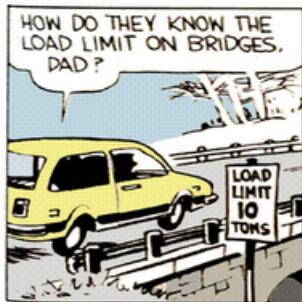


Security Engineering (9)

Email: christian.urban at kcl.ac.uk

Office: S1.27 (1st floor Strand Building)

Slides: KEATS (also homework is there)



Old-Fashioned Eng. vs. CS



bridges:

engineers can “look” at a bridge and have a pretty good intuition about whether it will hold up or not (redundancy; predictive theory)



code:

programmers have very little intuition about their code; often it is too expensive to have redundancy; not “continuous”

Trusting Computing Base

When considering whether a system meets some security objectives, it is important to consider which parts of that system are trusted in order to meet that objective (TCB).

Trusting Computing Base

When considering whether a system meets some security objectives, it is important to consider which parts of that system are trusted in order to meet that objective (TCB).

The smaller the TCB, the less effort it takes to get some confidence that it is trustworthy, by doing a code review or by performing some (penetration) testing.

CPU, compiler, libraries, OS, NP \neq P, random number generator, ...

Dijkstra on Testing

“Program testing can be a very effective way to show the presence of bugs, but it is hopelessly inadequate for showing their absence.”

unfortunately attackers exploit bugs (Satan's computer vs Murphy's)

Proving Programs to be Correct

Theorem: There are infinitely many prime numbers.

Proof ...

similarly

Theorem: The program is doing what it is supposed to be doing.


Long, long proof ...

This can be a gigantic proof. The only hope is to have help from the computer. 'Program' is here to be understood to be quite general (protocols, OS, ...).

Mars Pathfinder Mission 1997

- despite NASA's famous testing procedures, the lander crashed frequently on Mars
- a scheduling algorithm was not used in the OS

low priority



high priority



low priority



high priority



low priority




Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



locked a resource

low priority




Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



locked a resource

low priority




Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



low priority

locked a resource




Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



locked a resource

low priority



Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



medium pr.



locked a resource

low priority



Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



medium pr.



low priority

locked a resource



o

time



Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



medium pr.



low priority

locked a resource



o

time



Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority



medium pr.



locked a resource

low priority



Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority




medium pr.



locked a resource

low priority




Scheduling: You want to avoid that a high priority process is starved indefinitely.

high priority

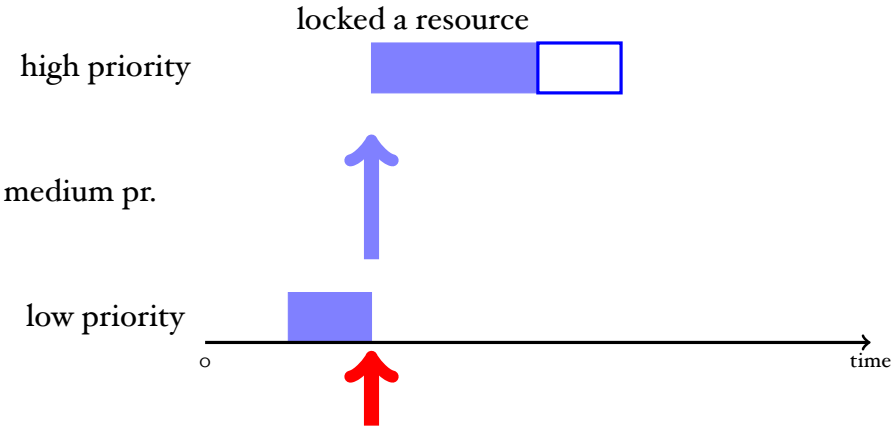


medium pr.

low priority



Scheduling: You want to avoid that a high priority process is starved indefinitely.



Scheduling: You want to avoid that a high priority process is starved indefinitely.

locked a resource


high priority



medium pr.



low priority




Scheduling: You want to avoid that a high priority process is starved indefinitely.

Priority Inheritance Scheduling


- Let a low priority process L temporarily inherit the high priority of H until L leaves the critical section unlocking the resource.
- Once the resource is unlocked L returns to its original priority level.

low priority



high priority


low priority



high priority




low priority




high priority



low priority




high priority




low priority



high priority




low priority



high priority




low priority



high priority




low priority



high priority



low priority




high priority




medium pr.



low priority





Scheduling: You want to avoid that a high priority process is staved indefinitely.


Priority Inheritance Scheduling

- Let a low priority process L temporarily inherit the high priority of H until L leaves the critical section unlocking the resource.
- Once the resource is unlocked L returns to its original priority level. **BOGUS**


Priority Inheritance Scheduling

- Let a low priority process L temporarily inherit the high priority of H until L leaves the critical section unlocking the resource.
- Once the resource is unlocked L returns to its original priority level. **BOGUS**
- ... L needs to switch to the highest **remaining** priority of the threads that it blocks.


this error is already known since around 1999




- by Rajkumar, 1991
- *“it resumes the priority it had at the point of entry into the critical section”*



- by Jane Liu, 2000
- *“The job J_1 executes at its inherited priority until it releases R ; at that time, the priority of J_1 returns to its priority at the time when it acquires the resource R .”*
- gives pseudo code and totally bogus data structures
- interesting part *“left as an exercise”*



- by Laplante and Ovaska, 2011 (\$113.76)
- “*when [the task] exits the critical section that caused the block, it reverts to the priority it had when it entered that section*”



- by Silberschatz, Galvin, and Gagne, 2013 (OS-bible)
- *“Upon releasing the lock, the [low-priority] thread will revert to its original priority.”*

Priority Scheduling

- a scheduling algorithm that is widely used in real-time operating systems
- has been “proved” correct by hand in a paper in 1983
- but this algorithm turned out to be incorrect, despite its “proof”

Priority Scheduling

- a scheduling algorithm that is widely used in real-time operating systems
- has been “proved” correct by hand in a paper in 1983
- but this algorithm turned out to be incorrect, despite its “proof”
- we corrected the algorithm and then **really** proved that it is correct
- we implemented this algorithm in a small OS called PINTOS (used for teaching at Stanford)
- our implementation was much more efficient than their reference implementation

Design of AC-Policies

Imagine you set up an access policy (root, lpd, users, staff, etc)

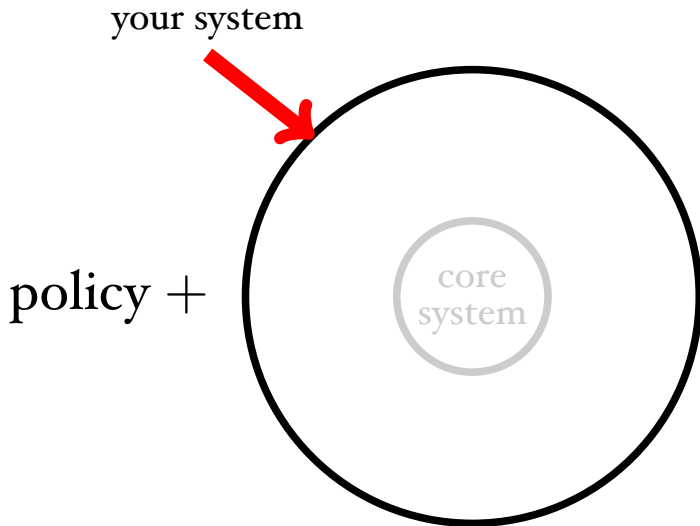
Design of AC-Policies

Imagine you set up an access policy (root, lpd, users, staff, etc)


“what you specify is what you get but not necessarily what you want...”

main work by Chunhan Wu (a PhD-student in Nanjing)


Testing Policies




Testing Policies




Testing Policies




Testing Policies




Testing Policies



Testing Policies



Testing Policies



Fuzzy Testing C-Compilers

- tested GCC, LLVM and others by randomly generating C-programs
- found more than 300 bugs in GCC and also many in LLVM (some of them highest-level critical)
- about CompCert:

“The striking thing about our CompCert results is that the middle-end bugs we found in all other compilers are absent. As of early 2011, the under-development version of CompCert is the only compiler we have tested for which Csmith cannot find wrong-code errors. This is not for lack of trying: we have devoted about six CPU-years to the task.”

Big Proofs in CS (2)

- in 2010, verification of a micro-kernel operating system (approximately 8700 loc)
 - used in helicopters and mobile phones
 - 200k loc of proof
 - 25 - 30 person years
 - found 160 bugs in the C code (144 by the proof)

“Real-world operating-system kernel with an end-to-end proof of implementation correctness and security enforcement”

Big Proofs in CS (2)

- in 2010, verification of a micro-kernel operating system (approximately 8700 loc)
 - used in helicopters and mobile phones
 - 200k loc of proof
 - 25 - 30 person years
 - found 160 bugs in the C code (144 by the proof)

“Real-world operating-system kernel with an end-to-end proof of implementation correctness and security enforcement”

unhackable kernel (New Scientists, September 2015)

Big Proofs in CS (3)

- verified TLS implementation
- verified compilers (CompCert, CakeML)
- verified distributed systems (Verdi)
- verified Oses or OS components
(seL4, CertiKOS, Ironclad Apps, ...)
- verified cryptography

How Did This Happen?

Lots of ways!

- better programming languages
 - basic safety guarantees built in
 - powerful mechanisms for abstraction and modularity
- better software development methodology
- stable platforms and frameworks
- better use of specifications

If you want to build software that works or is secure, it is helpful to know what you mean by “works” and by “is secure”!

Goal

Remember the Bridges example?

- Can we look at our programs and somehow ensure they are secure/bug free/correct?

Goal

Remember the Bridges example?

- Can we look at our programs and somehow ensure they are secure/bug free/correct?
- Very hard: Anything interesting about programs is equivalent to halting problem, which is undecidable.

Goal

Remember the Bridges example?

- Can we look at our programs and somehow ensure they are secure/bug free/correct?
- Very hard: Anything interesting about programs is equivalent to halting problem, which is undecidable.
- **Solution:** We avoid this “minor” obstacle by being as close as possible of deciding the halting problem, without actually deciding the halting problem. \Rightarrow yes, no, don't know (static analysis)