

Provable Security

Henrik Karlsson (henrik10@kth.se)

Supervisors: Mads Dam & Roberto Guanciale
KTH Royal Institute of Technology


October 27, 2022




- 1 Example Application & Proof-of-Concept
- 2 Simply Secure Separation Kernel (S3K)
- 3 S3K Process Scheduling
- 4 Multicore HoIBA & Kernel Verification
- 5 Research plan

Example Application & Proof-of-Concept


Application




Application



Application




- Monitor – Monitors system & handles commands
- Functions – Interact with environment and output data
- TCP/IP – Handle TCP/IP communication
- Crypto engine – Encrypts & Decrypts Packages



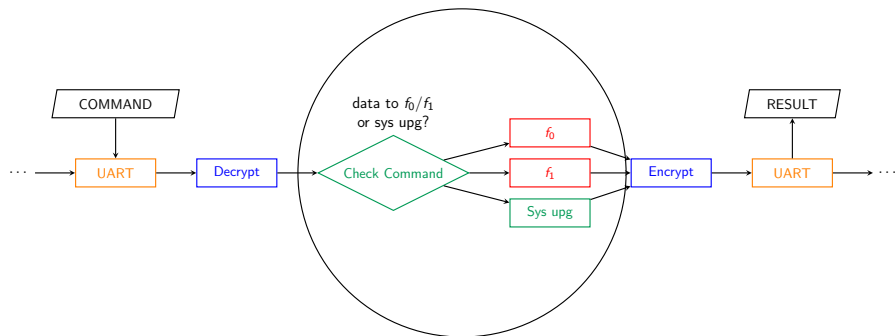
- Monitor – Manage apps & handles commands
- Dummy Functions f_0, f_1 – Process and output data
- UART – Handle UART communication
- Crypto engine – Encrypts & Decrypts Packages

Proof-of-Concept



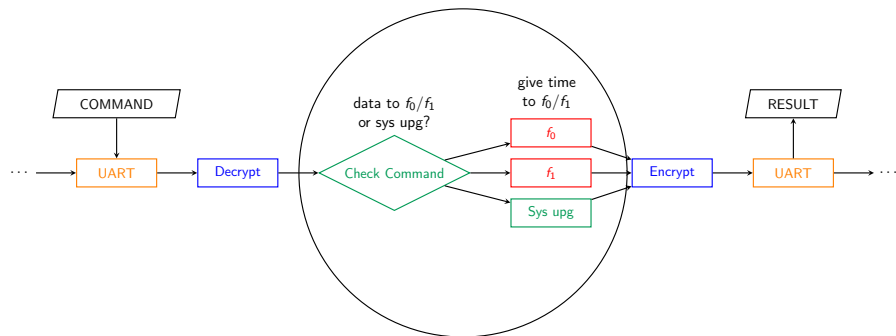
- Monitor – Manage apps & handles commands
- Dummy Functions f_0, f_1 – Process and output data
- UART – Handle UART communication
- Crypto engine – Encrypts & Decrypts Packages

Proof-of-Concept



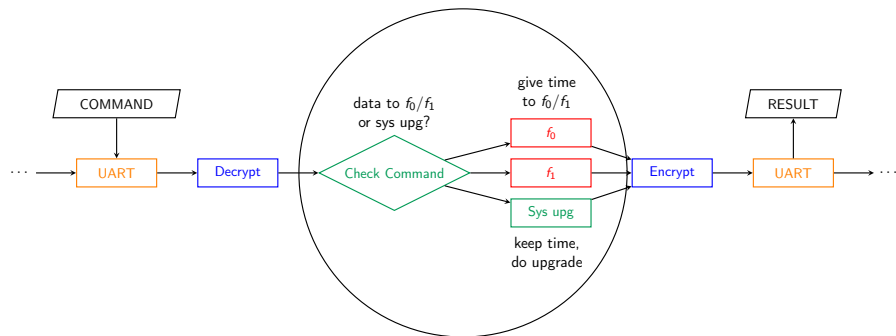
- Monitor – Manage apps & handles commands
- Dummy Functions f_0, f_1 – Process and output data
- UART – Handle UART communication
- Crypto engine – Encrypts & Decrypts Packages

Proof-of-Concept



- Monitor – Manage apps & handles commands
- Dummy Functions f_0, f_1 – Process and output data
- UART – Handle UART communication
- Crypto engine – Encrypts & Decrypts Packages

Proof-of-Concept



- Monitor – Manage apps & handles commands
- Dummy Functions f_0, f_1 – Process and output data
- UART – Handle UART communication
- Crypto engine – Encrypts & Decrypts Packages

Simply Secure Separation Kernel (S3K)

Simply Secure Separation Kernel (S3K)

- Designed and implemented a separation kernel with ...

¹MPU - Memory Protection Unit, protects physical memory.

Simply Secure Separation Kernel (S3K)

- Designed and implemented a separation kernel with ...
 - ▶ Memory Protection and Management

¹MPU - Memory Protection Unit, protects physical memory.

Simply Secure Separation Kernel (S3K)

- Designed and implemented a separation kernel with ...
 - ▶ Memory Protection and Management
 - ▶ Secure Time Management for Real-Time Systems

¹MPU - Memory Protection Unit, protects physical memory.

- Designed and implemented a separation kernel with ...
 - ▶ Memory Protection and Management
 - ▶ Secure Time Management for Real-Time Systems
 - ▶ Secure Inter-Process Communication (IPC)

¹MPU - Memory Protection Unit, protects physical memory.

- Designed and implemented a separation kernel with ...
 - ▶ Memory Protection and Management
 - ▶ Secure Time Management for Real-Time Systems
 - ▶ Secure Inter-Process Communication (IPC)
 - ▶ Process Monitoring

¹MPU - Memory Protection Unit, protects physical memory.

Simply Secure Separation Kernel (S3K)

- Designed and implemented a separation kernel with ...
 - ▶ Memory Protection and Management
 - ▶ Secure Time Management for Real-Time Systems
 - ▶ Secure Inter-Process Communication (IPC)
 - ▶ Process Monitoring
- Targeting standard RISC-V 64-bit (RV64IMA) with MPU¹

¹MPU - Memory Protection Unit, protects physical memory.

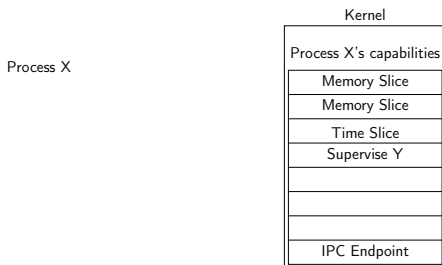
Simply Secure Separation Kernel (S3K)

- Designed and implemented a separation kernel with ...
 - ▶ Memory Protection and Management
 - ▶ Secure Time Management for Real-Time Systems
 - ▶ Secure Inter-Process Communication (IPC)
 - ▶ Process Monitoring
- Targeting standard RISC-V 64-bit (RV64IMA) with MPU¹
- ~ 2000 lines of C/Assembly

¹MPU - Memory Protection Unit, protects physical memory.

Previously mentioned features are implemented using capabilities


- Capability = object in kernel representing a resource



Capabilities

Previously mentioned features are implemented using capabilities


- Capability = object in kernel representing a resource
- Process owning a capability has access to corresponding resource



Capabilities

Previously mentioned features are implemented using capabilities

- Capability = object in kernel representing a resource
- Process owning a capability has access to corresponding resource
- Process can derive new capabilities from existing capabilities



- Memory Slice – Manage access to a physical memory region.
 - ▶ PMP – Configure RISC-V's MPU, grants memory access.


- Memory Slice – Manage access to a physical memory region.
 - ▶ PMP – Configure RISC-V's MPU, grants memory access.
- Time Slice – Manage and grant execution time on a core.

- Memory Slice – Manage access to a physical memory region.
 - ▶ PMP – Configure RISC-V's MPU, grants memory access.
- Time Slice – Manage and grant execution time on a core.
- Channels – Manage IPC channels and endpoints.
 - ▶ Receiver/Sender – Unidirectional IPC channel.
 - ▶ Server/Client – Bidirectional IPC channel.

- Memory Slice – Manage access to a physical memory region.
 - ▶ PMP – Configure RISC-V's MPU, grants memory access.
- Time Slice – Manage and grant execution time on a core.
- Channels – Manage IPC channels and endpoints.
 - ▶ Receiver/Sender – Unidirectional IPC channel.
 - ▶ Server/Client – Bidirectional IPC channel.
- Supervisor – Manage a set of processes.


S3K Process Scheduling

What is Process Scheduling?




- Cores can only run one process at the time.

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.
- Context switch \implies overhead.

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.
- Context switch \implies overhead.
- Core multiplexing technique = Process Scheduling.

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.
- Context switch \implies overhead.
- Core multiplexing technique = Process Scheduling.
- Best scheduler? Performance vs. Safety vs. Security.

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.
- Context switch \implies overhead.
- Core multiplexing technique = Process Scheduling.
- Best scheduler? Performance vs. Safety vs. Security.
 - ▶ Processor utilization? Supercomputer, Desktop, ...

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.
- Context switch \implies overhead.
- Core multiplexing technique = Process Scheduling.
- Best scheduler? Performance vs. Safety vs. Security.
 - ▶ Processor utilization? Supercomputer, Desktop, ...
 - ▶ Meet deadline? Airplane, Car, Routers, Industrial Machines, ...

What is Process Scheduling?




- Cores can only run one process at the time.
- Multiplexing the Core – main duty of the kernel.
- Context switch \implies overhead.
- Core multiplexing technique = Process Scheduling.
- Best scheduler? Performance vs. Safety vs. Security.
 - ▶ Processor utilization? Supercomputer, Desktop, ...
 - ▶ Meet deadline? Airplane, Car, Routers, Industrial Machines, ...
 - ▶ Prevent side-channels attacks? Secure servers, routers, ...

S3K Scheduling




- Modified RR – Minor frames defined by time slice capabilities.

S3K Scheduling



- Modified RR – Minor frames defined by time slice capabilities.
- Fair – Process with execution time gets execution time.


S3K Scheduling



- Modified RR – Minor frames defined by time slice capabilities.
- Fair – Process with execution time gets execution time.
- Predictable – Process knows time slice capabilities, thus their execution time.

²Flush cache, branch predictors, etc., support is hardware dependant.


S3K Scheduling



- Modified RR – Minor frames defined by time slice capabilities.
- Fair – Process with execution time gets execution time.
- Predictable – Process knows time slice capabilities, thus their execution time.
- Temporal Isolation – A process's execution time depends only on its capabilities.

²Flush cache, branch predictors, etc., support is hardware dependant.


S3K Scheduling



- Modified RR – Minor frames defined by time slice capabilities.
- Fair – Process with execution time gets execution time.
- Predictable – Process knows time slice capabilities, thus their execution time.
- Temporal Isolation – A process's execution time depends only on its capabilities.
- Low-overhead – Scheduling decided by a lookup table.

²Flush cache, branch predictors, etc., support is hardware dependant.


S3K Scheduling



- Modified RR – Minor frames defined by time slice capabilities.
- Fair – Process with execution time gets execution time.
- Predictable – Process knows time slice capabilities, thus their execution time.
- Temporal Isolation – A process's execution time depends only on its capabilities.
- Low-overhead – Scheduling decided by a lookup table.
- Dynamic – Process can alter their time slices.

²Flush cache, branch predictors, etc., support is hardware dependant.


Time Slice Capability



- hartid – ID of a hardware thread.³
- begin – start of a time slice
- free – start of minor frame
- end – end of a time slice and minor frame

³Hardware Thread – Logically separate processor.

Proof-of-Concept with Time Slices



Proof-of-Concept with Time Slice

Monitor has the initial time slice

(capabilities)



(scheduling)




Proof-of-Concept with Time Slice

Monitor derives capability B
(only create slices from *free* to *end*)

(capabilities)




(scheduling)



Proof-of-Concept with Time Slice

Monitor derives capability C

(capabilities)




(scheduling)



Proof-of-Concept with Time Slice

Monitor derives capability D

(capabilities)



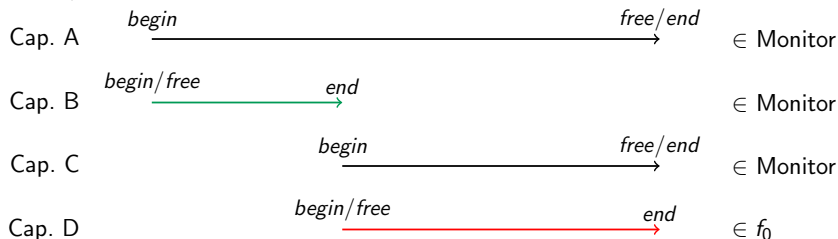
(scheduling)



Proof-of-Concept with Time Slice

Monitor sends capability D to F0
(using IPC or supervisor capability)

(capabilities)




(scheduling)



Proof-of-Concept with Time Slice

F0 deletes capability D
(Core idle from *free* to *end*)

(capabilities)



(scheduling)




Proof-of-Concept with Time Slice

Monitor call revoke on capability C

(deletes children & resets C)

(capabilities)




(scheduling)



Proof-of-Concept with Time Slice

Monitor derives capability D again


(capabilities)





(scheduling)




Multicore HoBA & Kernel Verification








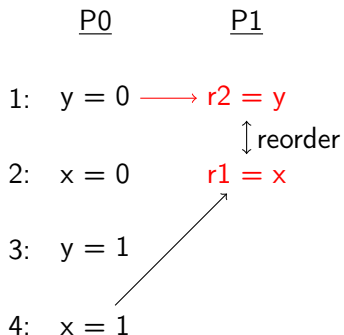
Memory operations may be reordered in RISC-V.



Writes may be reordered.




Reads may be reordered.




Verification of multicore RISC-V need all reorderings.
Multicore HoBA provides this!

Separation Kernel Example



Separation Kernel Example



Research Plan

- Dec. 2022 – Complete proof-of-concept
- Jan./Feb. 2023 – Evaluation and publication of kernel with proof-of-concept

⁴Worst-case execution time

- Dec. 2022 – Complete proof-of-concept
- Jan./Feb. 2023 – Evaluation and publication of kernel with proof-of-concept
- Spring 2023
 - ▶ Publication of multicore HoIBA
 - ▶ Measurements of WCET⁴ and jitter of non-preemptive kernel parts
 - ▶ Implement secure interrupts, optimized scheduler, and 32-bit kernel version

⁴Worst-case execution time

- Dec. 2022 – Complete proof-of-concept
- Jan./Feb. 2023 – Evaluation and publication of kernel with proof-of-concept
- Spring 2023
 - ▶ Publication of multicore HoIBA
 - ▶ Measurements of WCET⁴ and jitter of non-preemptive kernel parts
 - ▶ Implement secure interrupts, optimized scheduler, and 32-bit kernel version
- Summer 2023
 - ▶ Model and proofs on some concurrent kernel code using multicore HoIBA
 - ▶ Finish sequential high-level model of kernel

⁴Worst-case execution time


- Dec. 2022 – Complete proof-of-concept
- Jan./Feb. 2023 – Evaluation and publication of kernel with proof-of-concept
- Spring 2023
 - ▶ Publication of multicore HoIBA
 - ▶ Measurements of WCET⁴ and jitter of non-preemptive kernel parts
 - ▶ Implement secure interrupts, optimized scheduler, and 32-bit kernel version
- Summer 2023
 - ▶ Model and proofs on some concurrent kernel code using multicore HoIBA
 - ▶ Finish sequential high-level model of kernel
- Autumn 2023 – Finish concurrent high-level model of kernel

⁴Worst-case execution time

Questions?

Multicore Scheduling

Where does the blue process execute?




Process runs on one core, for determinism,⁵ we need priority rules.

- currently running core
- core with smallest ID.

⁵non-determinism may leak information

Multicore Scheduling

Where does the blue process execute?




Process runs on one core, for determinism,⁵ we need priority rules.

- currently running core
- core with smallest ID.

⁵non-determinism may leak information

Multicore Scheduling

Where does the blue process execute?



Process runs on one core, for determinism,⁵ we need priority rules.

- currently running core
- longest quantum (not implemented)
- core with smallest ID

⁵non-determinism may leak information