

SECURE EMBEDDED SYSTEMS NEED MICROKERNELS

Gernot Heiser

Embedded, Real-Time and Operating Systems Program
National ICT Australia

February 2004



Australian Government

Department of Communications,
Information Technology and the Arts

Australian Research Council

NICTA Members



Department of State and
Regional Development



NICTA Partners

EMBEDDED SYSTEM ARE UBIQUITOUS



EMBEDDED SYSTEM ARE UBIQUITOUS



But are they *Secure*?

SECURITY CHALLENGES



- Growing functionality
- Wireless connectivity
- Downloaded contents (entertainment)

SECURITY CHALLENGES



- Growing functionality
 - increasing software complexity
 - increased number of faults
 - increased likelihood of security faults
- Wireless connectivity
- Downloaded contents (entertainment)

SECURITY CHALLENGES



- Growing functionality
 - increasing software complexity
 - increased number of faults
 - increased likelihood of security faults
- Wireless connectivity
 - subject to attacks from outside (crackers)
- Downloaded contents (entertainment)

SECURITY CHALLENGES



- Growing functionality
 - increasing software complexity
 - increased number of faults
 - increased likelihood of security faults
- Wireless connectivity
 - subject to attacks from outside (crackers)
- Downloaded contents (entertainment)
 - subject to attacks from inside (viruses, worms)

SECURITY CHALLENGES



- Growing functionality
 - increasing software complexity
 - increased number of faults
 - increased likelihood of security faults
- Wireless connectivity
 - subject to attacks from outside (crackers)
- Downloaded contents (entertainment)
 - subject to attacks from inside (viruses, worms)
- Increasing dependence on embedded systems
 - increased exposure to embedded-systems security weaknesses

PRESENT APPROACHES



- Real-time executives

PRESENT APPROACHES



- Real-time executives
 - ★ suitable for systems with very limited functionality
 - ★ no internal protection

PRESENT APPROACHES



- Real-time executives
 - ★ suitable for systems with very limited functionality
 - ★ no internal protection
 - every small bug/failure is fatal
 - no defence against viruses, limited defence against crackers

PRESENT APPROACHES



- Real-time executives
 - ★ suitable for systems with very limited functionality
 - ★ no internal protection
 - every small bug/failure is fatal
 - no defence against viruses, limited defence against crackers
- Linux, Windows Embedded, ...

PRESENT APPROACHES



- Real-time executives
 - ★ suitable for systems with very limited functionality
 - ★ no internal protection
 - every small bug/failure is fatal
 - no defence against viruses, limited defence against crackers
- Linux, Windows Embedded, ...
 - ★ dubious or non-existent real-time capabilities
 - unsuitable for hard real-time systems

PRESENT APPROACHES



- Real-time executives
 - ★ suitable for systems with very limited functionality
 - ★ no internal protection
 - every small bug/failure is fatal
 - no defence against viruses, limited defence against crackers
- Linux, Windows Embedded, ...
 - ★ dubious or non-existent real-time capabilities
 - unsuitable for hard real-time systems
 - ★ huge code base (millions of lines of code)
 - excessive for small embedded system
 - too much code on which security of system is dependent

REQUIRES MINIMAL TRUSTED COMPUTING BASE (TCB)

TCB: The part of system that must be relied on for the correct operation of the system

REQUIRES MINIMAL TRUSTED COMPUTING BASE (TCB)

TCB: The part of system that must be relied on for the correct operation of the system

... otherwise the security issue is intractable!

REQUIRES MINIMAL TRUSTED COMPUTING BASE (TCB)

TCB: The part of system that must be relied on for the correct operation of the system

... otherwise the security issue is intractable!

MINIMAL TCB MEANS:

1. Split software into trusted and untrusted part

REQUIRES MINIMAL TRUSTED COMPUTING BASE (TCB)

TCB: The part of system that must be relied on for the correct operation of the system

... otherwise the security issue is intractable!

MINIMAL TCB MEANS:

1. Split software into trusted and untrusted part
2. Minimise the trusted part

REQUIRES MINIMAL TRUSTED COMPUTING BASE (TCB)


TCB: The part of system that must be relied on for the correct operation of the system

... otherwise the security issue is intractable!

MINIMAL TCB MEANS:



1. Split software into trusted and untrusted part
2. Minimise the trusted part
3. Make the trusted part *trustworthy*

MINIMISING THE TCB



System: traditional
embedded




MINIMISING THE TCB



System: traditional
embedded

Linux/
Windows

MINIMISING THE TCB




System: traditional
embedded


Linux/
Windows

Microkernel-
based


MINIMISING THE TCB



System: traditional
embedded




Linux/
Windows




Microkernel-
based


MINIMISING THE TCB



System: traditional
embedded
TCB: **all** code






Linux/
Windows



Microkernel-
based

MINIMISING THE TCB






System: traditional
embedded
all code

Linux/
Windows

Microkernel-
based

MINIMISING THE TCB




System: traditional
embedded
all code


Linux/
Windows

Microkernel-
based


MINIMISING THE TCB



System: traditional
embedded
TCB: **all** code






Linux/
Windows



Microkernel-
based

MINIMISING THE TCB




System: traditional
embedded
TCB: **all** code


Linux/
Windows
100,000's loc

Microkernel-
based


MINIMISING THE TCB



System: traditional
embedded
TCB: **all** code






System: Linux/
Windows
TCB: 100,000's loc



System: Microkernel-
based

MINIMISING THE TCB






System: traditional
embedded
TCB: **all** code

Linux/
Windows
100,000's loc

Microkernel-
based

MINIMISING THE TCB






System: traditional
embedded
TCB: **all** code

Linux/
Windows
100,000's loc

Microkernel-
based

MINIMISING THE TCB




System: traditional
embedded
TCB: **all** code

Linux/
Windows
100,000's loc

Microkernel-
based
10,000's loc

MINIMISING THE TCB



System: traditional
embedded
TCB: **all** code

Linux/
Windows
100,000's loc

Microkernel-
based
10,000's loc

Small kernel \Rightarrow small TCB \Rightarrow more trustworthy TCB!

MAKING THE TCB TRUSTWORTHY



TRADITIONAL APPROACHES

ADVANCED APPROACHES

MAKING THE TCB TRUSTWORTHY



TRADITIONAL APPROACHES

- Testing
- Code inspection

ADVANCED APPROACHES

MAKING THE TCB TRUSTWORTHY



TRADITIONAL APPROACHES

- Testing
- Code inspection
- ✓ Smaller TCB \Rightarrow fewer remaining bugs
 - ✗ ... but no *assurance* of absence of bugs!

ADVANCED APPROACHES

MAKING THE TCB TRUSTWORTHY



TRADITIONAL APPROACHES

- Testing
- Code inspection
- ✓ Smaller TCB \Rightarrow fewer remaining bugs
 - ✗ ... but no *assurance* of absence of bugs!

ADVANCED APPROACHES

- *Prove* the correctness of the TCB
 - ✓ assurance of absence of bugs

MAKING THE TCB TRUSTWORTHY



TRADITIONAL APPROACHES

- Testing
- Code inspection
- ✓ Smaller TCB \Rightarrow fewer remaining bugs
 - ✗ ... but no *assurance* of absence of bugs!

ADVANCED APPROACHES

- *Prove* the correctness of the TCB
 - ✓ assurance of absence of bugs
 - ✗ unfeasible for 100,000s loc
 - ✓ feasible for very small microkernel

MICROKERNEL VERIFICATION



MICROKERNEL VERIFICATION



PROOF OF CORRECTNESS OF MICROKERNEL

PROOF OF CORRECTNESS OF MICROKERNEL

1. Formal specification of microkernel API

→ prerequisite for all formal work on kernel

PROOF OF CORRECTNESS OF MICROKERNEL

1. Formal specification of microkernel API

→ prerequisite for all formal work on kernel

2. Prove security (isolation) properties of API

→ confinement, data flow, interference, covert storage channels

PROOF OF CORRECTNESS OF MICROKERNEL

1. Formal specification of microkernel API

→ prerequisite for all formal work on kernel

2. Prove security (isolation) properties of API

→ confinement, data flow, interference, covert storage channels

3. Prove correctness of microkernel implementation

→ *refine* API spec into compilable source code

PROOF OF CORRECTNESS OF MICROKERNEL

1. Formal specification of microkernel API

→ prerequisite for all formal work on kernel

2. Prove security (isolation) properties of API

→ confinement, data flow, interference, covert storage channels

3. Prove correctness of microkernel implementation

→ *refine* API spec into compilable source code

COMPLETE TEMPORAL MODEL OF MICROKERNEL

PROOF OF CORRECTNESS OF MICROKERNEL

1. Formal specification of microkernel API

→ prerequisite for all formal work on kernel

2. Prove security (isolation) properties of API

→ confinement, data flow, interference, covert storage channels


3. Prove correctness of microkernel implementation

→ *refine* API spec into compilable source code

COMPLETE TEMPORAL MODEL OF MICROKERNEL


- exhaustive measurement of latencies of kernel operations
 - prerequisite for real-time analysis of whole system

A SAMPLE SYSTEM




- Sensitive part of system has small TCB

A SAMPLE SYSTEM



- Sensitive part of system has small TCB
- Standard API supported by de-privileged Linux server
 - full binary compatibility with native Linux


A SAMPLE SYSTEM




- Sensitive part of system has small TCB
- Standard API supported by de-privileged Linux server
 - full binary compatibility with native Linux
- Compromised legacy system cannot interfere with trusted part

- Microkernel operational


- 10,000 lines of code
- highly efficient



- Microkernel operational
 - 10,000 lines of code
 - highly efficient
- Basic infrastructure exists
 - core services
 - Linux server, drivers



- Microkernel operational
 - 10,000 lines of code
 - highly efficient
- Basic infrastructure exists
 - core services
 - Linux server, drivers
 - in commercial deployment



- Microkernel operational
 - 10,000 lines of code
 - highly efficient
- Basic infrastructure exists
 - core services
 - Linux server, drivers
 - in commercial deployment
- Verification in progress
 - successful pilot project for verification of API “slice”
 - full-scale correctness proof project commencing
 - temporal analysis commencing

