

Operating Systems Research at UNSW and NICTA

and
Opportunities for Students

COMP9242 2004/S2 Week 14

OVERVIEW

- What is NICTA?
- Embedded systems research agenda
- Present projects
- Opportunities

NATIONAL ICT AUSTRALIA (NICTA)

- National Centre of Excellence in Information and Communications Technology (ICT)
- Created in October 2002 by Australian Government

NATIONAL ICT AUSTRALIA (NICTA)

- National Centre of Excellence in Information and Communications Technology (ICT)
- Created in October 2002 by Australian Government
- Members:
 - ★ University of New South Wales (UNSW), Sydney
 - ★ Australian National University (ANU), Canberra
 - ★ NSW and ACT governments

NATIONAL ICT AUSTRALIA (NICTA)

- National Centre of Excellence in Information and Communications Technology (ICT)
- Created in October 2002 by Australian Government
- Members:
 - ★ University of New South Wales (UNSW), Sydney
 - ★ Australian National University (ANU), Canberra
 - ★ NSW and ACT governments
- Locations:
 - ★ Sydney (UNSW campus and ATP technology park)
 - ★ Canberra (ANU campus)
 - ★ Melbourne (University of Melbourne Campus)
 - ★ Brisbane (in progress)

NICTA'S FOUR PILLARS

- Research
- Education
- Linkages
- Commercialisation

NICTA'S FOUR PILLARS

- Research

- conduct world-class ICT research that makes an impact
- “use-inspired fundamental research”

- Education

- Linkages

- Commercialisation

NICTA'S FOUR PILLARS

- Research

- conduct world-class ICT research that makes an impact
- “use-inspired fundamental research”

- Education

- produce world-class PhD graduates
- students enrolled at one of the member/partner universities

- Linkages

- Commercialisation

NICTA'S FOUR PILLARS

- Research

- conduct world-class ICT research that makes an impact
- “use-inspired fundamental research”

- Education

- produce world-class PhD graduates
- students enrolled at one of the member/partner universities

- Linkages

- collaborate with top research institutions around the world

- Commercialisation

NICTA'S FOUR PILLARS

- Research

- conduct world-class ICT research that makes an impact
- “use-inspired fundamental research”

- Education

- produce world-class PhD graduates
- students enrolled at one of the member/partner universities

- Linkages

- collaborate with top research institutions around the world

- Commercialisation

- turn research into products
- focus on
 - local small and medium-sized enterprises (SMEs)
 - multinational corporations (MNCs)

NICTA STRUCTURE

- Presently \approx 100 researchers, 150 PhD students

NICTA STRUCTURE

- Presently \approx 100 researchers, 150 PhD students
- Researchers belong to *Research Programs*
 - aligned with discipline areas (\approx 5–10 researchers)
 - lifetime 5–10 years

NICTA STRUCTURE

- Presently \approx 100 researchers, 150 PhD students
- Researchers belong to *Research Programs*
 - aligned with discipline areas (\approx 5–10 researchers)
 - lifetime 5–10 years
- *Priority Challenges* drive research
 - PC1: trusted wireless networks
 - PC2: from data to knowledge

NICTA STRUCTURE

- Presently \approx 100 researchers, 150 PhD students
- Researchers belong to *Research Programs*
 - aligned with discipline areas (\approx 5–10 researchers)
 - lifetime 5–10 years
- *Priority Challenges* drive research
 - PC1: trusted wireless networks
 - PC2: from data to knowledge
- *Projects* focused on specific outcomes
 - collaborative or client-focused
 - lifetime 1–5 years

NICTA STRUCTURE

- Presently \approx 100 researchers, 150 PhD students
- Researchers belong to *Research Programs*
 - aligned with discipline areas (\approx 5–10 researchers)
 - lifetime 5–10 years
- *Priority Challenges* drive research
 - PC1: trusted wireless networks
 - PC2: from data to knowledge
- *Projects* focused on specific outcomes
 - collaborative or client-focused
 - lifetime 1–5 years
- *International Science Advisory Group*
 - Richard Newton (UCB), Shankar Sastry (UCB), Raj Reddy (CMU), Rodney Brooks (MIT), Jeff Ullman (Stanford), Gunnar Bjurel (SIKS), Gilles Kahn (INRIA), Ya-Qin. Zhang (Microsoft), ...

EMBEDDED, REAL-TIME AND OPERATING SYSTEMS (ERTOS) PROGRAM

- One of presently 14 Research Programs in NICTA
 - 7 FTE PhD-qualified researchers (10 total)
 - 5 FTE research engineers / research assistants
 - 20 PhD students (10 core OS topics)
- Competencies in
 - operating systems, microkernels
 - networking
 - real-time systems
 - reconfigurable computing
 - programming languages and compiler front-ends

EMBEDDED SYSTEMS IN AUSTRALIA

- Australian embedded systems industry landscape
 - ★ little industrial research
 - ★ innovation concentrated in SMEs
 - ★ little confidence in locally-developed technology
 - ★ operating in niche markets

EMBEDDED SYSTEMS IN AUSTRALIA

- Australian embedded systems industry landscape
 - ★ little industrial research
 - ★ innovation concentrated in SMEs
 - ★ little confidence in locally-developed technology
 - ★ operating in niche markets
- Implications:
 - ★ no scope for ASICs, use COTS hardware components
 - ★ main focus is on software

EMBEDDED SYSTEMS IN AUSTRALIA

- Australian embedded systems industry landscape
 - ★ little industrial research
 - ★ innovation concentrated in SMEs
 - ★ little confidence in locally-developed technology
 - ★ operating in niche markets
- Implications:
 - ★ no scope for ASICs, use COTS hardware components
 - ★ main focus is on software
 - ★ reconfigurable hardware will be increasingly important
 - prototyping
 - small series
 - flexibility

NICTA PRIORITY CHALLENGE: TRUSTED WIRELESS NETWORKS

To enable greater confidence, freedom, and capability through improved efficiency, reliability, and security of all wireless environments.

NICTA PRIORITY CHALLENGE: TRUSTED WIRELESS NETWORKS

To enable greater confidence, freedom, and capability through improved efficiency, reliability, and security of all wireless environments.

- Strongly based on embedded systems technology

NICTA PRIORITY CHALLENGE: TRUSTED WIRELESS NETWORKS

To enable greater confidence, freedom, and capability through improved efficiency, reliability, and security of all wireless environments.

- Strongly based on embedded systems technology
- Issues relevant to ERTOS:
 - ★ efficiency
 - ★ reliability
 - ★ security
 - ★ cost

ERTOS VISION

To develop methodologies, tools, components and systems that will deliver reliable, trustworthy and inexpensive embedded systems software

ERTOS VISION

To develop methodologies, tools, components and systems that will deliver reliable, trustworthy and inexpensive embedded systems software

ERTOS research is to be driven by applications

- to identify common challenges
- to provide generic systems software

EMBEDDED SYSTEM

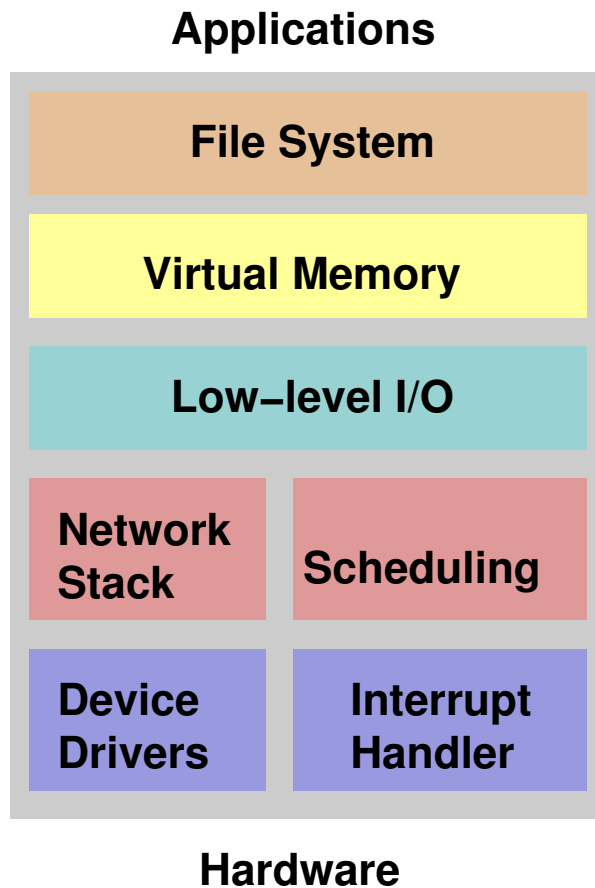


Computer system that is part of a larger system

GENERAL-PURPOSE VS. EMBEDDED SYSTEM

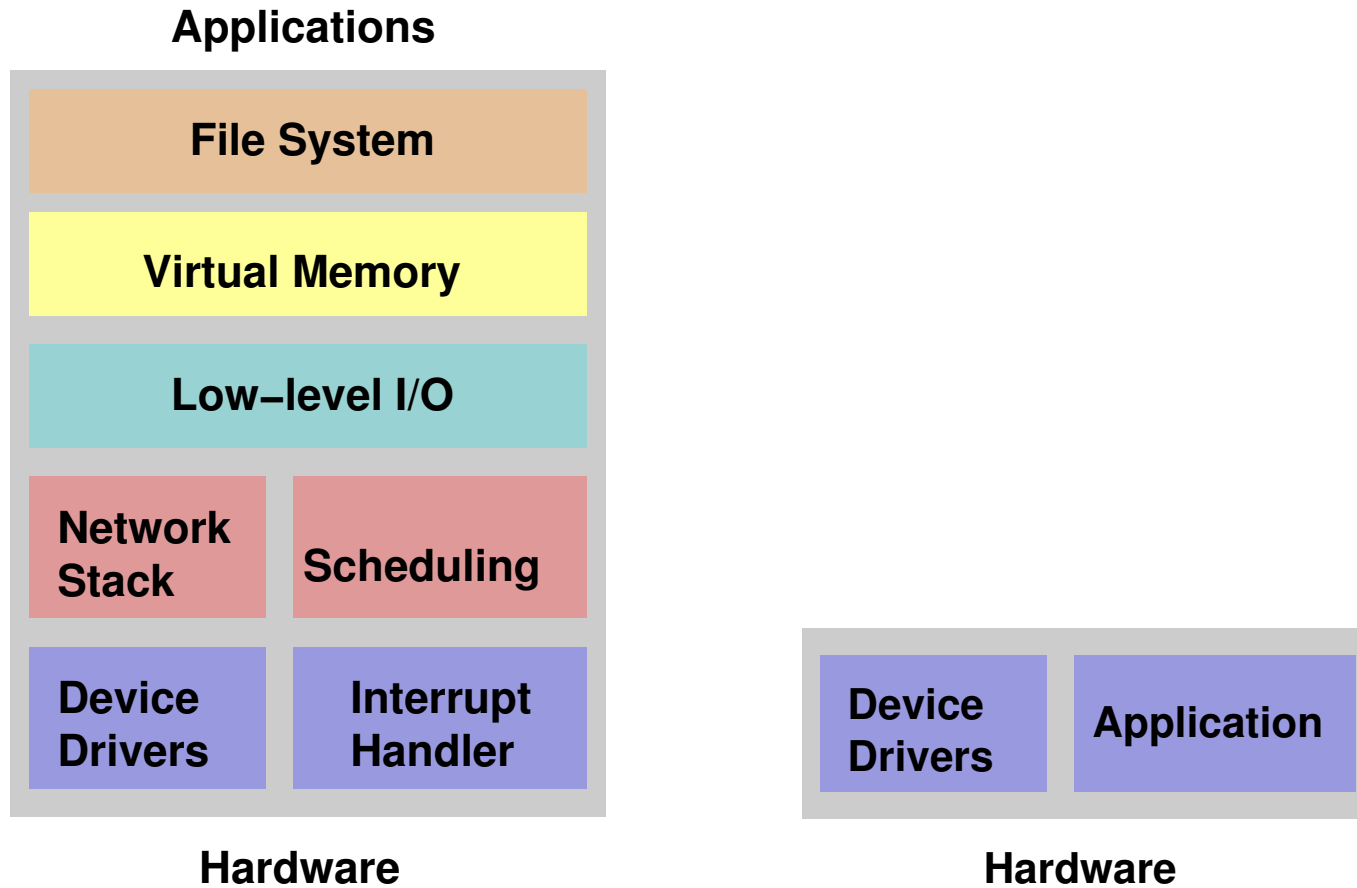
- Traditional model of general-purpose and embedded systems

GENERAL-PURPOSE VS. EMBEDDED SYSTEM



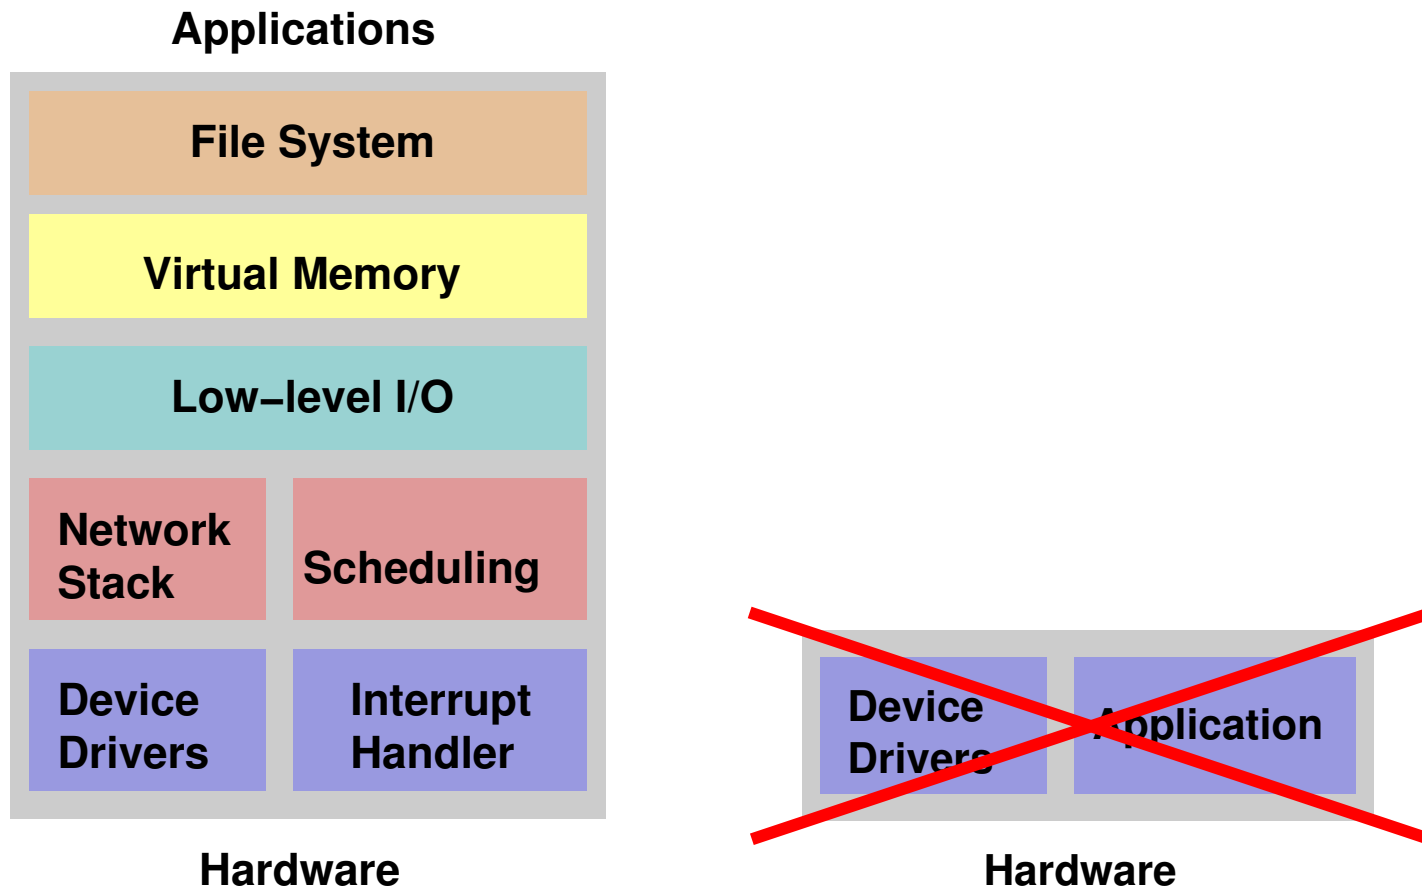
- Traditional model of general-purpose and embedded systems

GENERAL-PURPOSE VS. EMBEDDED SYSTEM



- Traditional model of general-purpose and embedded systems

GENERAL-PURPOSE VS. EMBEDDED SYSTEM



- Traditional model of general-purpose and embedded systems
 - No longer true for complex and networked embedded systems
 - Embedded systems increasingly face same challenges as desktops!

EMBEDDED SYSTEMS REQUIREMENTS

- Robustness, trustworthiness, security:

EMBEDDED SYSTEMS REQUIREMENTS

- Robustness, trustworthiness, security:

★ achieved by:

- exhaustive testing?
- systematic code inspection?
- formal methods?

EMBEDDED SYSTEMS REQUIREMENTS

- Robustness, trustworthiness, security:
 - ★ achieved by:
 - exhaustive testing?
 - systematic code inspection?
 - formal methods?
 - ★ requires minimal *trusted computing base* (TCB):
 - minimise exposure to bugs/faults
 - minimise exposure to attacks (internal and external)
 - support poorly-scaling verification methods

EMBEDDED SYSTEMS REQUIREMENTS

- Robustness, trustworthiness, security:
 - ★ achieved by:
 - exhaustive testing?
 - systematic code inspection?
 - formal methods?
 - ★ requires minimal *trusted computing base* (TCB):
 - minimise exposure to bugs/faults
 - minimise exposure to attacks (internal and external)
 - support poorly-scaling verification methods
- Low cost:

EMBEDDED SYSTEMS REQUIREMENTS

- Robustness, trustworthiness, security:
 - ★ achieved by:
 - exhaustive testing?
 - systematic code inspection?
 - formal methods?
 - ★ requires minimal *trusted computing base* (TCB):
 - minimise exposure to bugs/faults
 - minimise exposure to attacks (internal and external)
 - support poorly-scaling verification methods
- Low cost:
 - ★ achieved by good software-engineering support
 - modularisation, encapsulation
 - debugging support
 - standard tools

TRUSTED COMPUTING BASE

- What is it?

TCB is part of system that must be trusted to function correctly for correct operation of the system

TRUSTED COMPUTING BASE

- What is it?

TCB is part of system that must be trusted to function correctly for correct operation of the system

- What does it contain?

- ★ kernel (obviously)

- everything running in privileged mode can bypass security

TRUSTED COMPUTING BASE

- What is it?

TCB is part of system that must be trusted to function correctly for correct operation of the system

- What does it contain?

- ★ kernel (obviously)

- everything running in privileged mode can bypass security

- ★ device drivers

- DMA can bypass protection

- can mount DoS attacks

TRUSTED COMPUTING BASE

- What is it?

TCB is part of system that must be trusted to function correctly for correct operation of the system

- What does it contain?

- ★ kernel (obviously)

- everything running in privileged mode can bypass security

- ★ device drivers

- DMA can bypass protection

- can mount DoS attacks

- ★ services that control resources

TRUSTED COMPUTING BASE

- What is it?

TCB is part of system that must be trusted to function correctly for correct operation of the system

- What does it contain?

- ★ kernel (obviously)

- everything running in privileged mode can bypass security

- ★ device drivers

- DMA can bypass protection

- can mount DoS attacks

- ★ services that control resources

- ★ *everything* on MMU-less processors

TRUSTED COMPUTING BASE

- What is it?

TCB is part of system that must be trusted to function correctly for correct operation of the system

- What does it contain?

- ★ kernel (obviously)

- everything running in privileged mode can bypass security

- ★ device drivers

- DMA can bypass protection

- can mount DoS attacks

- ★ services that control resources

- ★ *everything* on MMU-less processors

- ★ compilers and tools...

MINIMISING THE SIZE OF THE TCB

- ... means first of all:
 - ★ use a *memory-management unit* (MMU)

MINIMISING THE SIZE OF THE TCB

- ... means first of all:
 - ★ use a *memory-management unit* (MMU)
 - ★ minimal size of the kernel!

MINIMISING THE SIZE OF THE TCB

- ... means first of all:
 - ★ use a *memory-management unit* (MMU)
 - ★ minimal size of the kernel!
- Kernel = part of system that executes in privileged mode
 - include in kernel only what is *required to build secure systems on top*

MINIMISING THE SIZE OF THE TCB

- ... means first of all:
 - ★ use a *memory-management unit* (MMU)
 - ★ minimal size of the kernel!
- Kernel = part of system that executes in privileged mode
 - include in kernel only what is *required to build secure systems on top*
 - if it does not *require* privileged mode, it *must not* be in kernel
 - use a *microkernel*

MONOLITHIC VS. MICROKERNEL

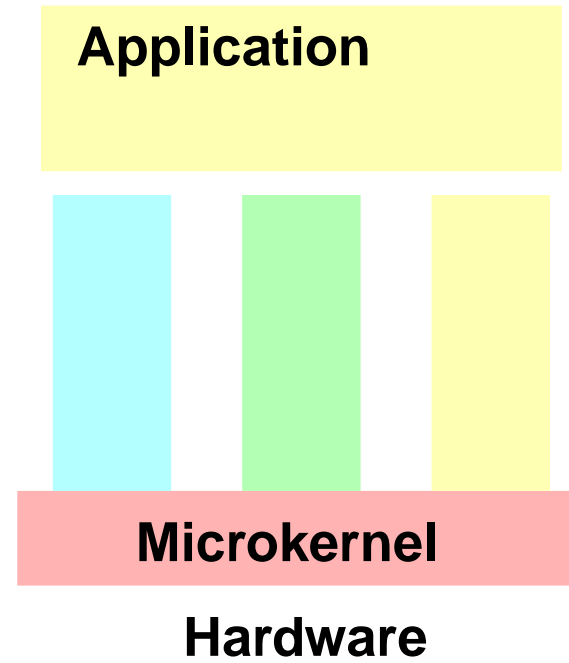
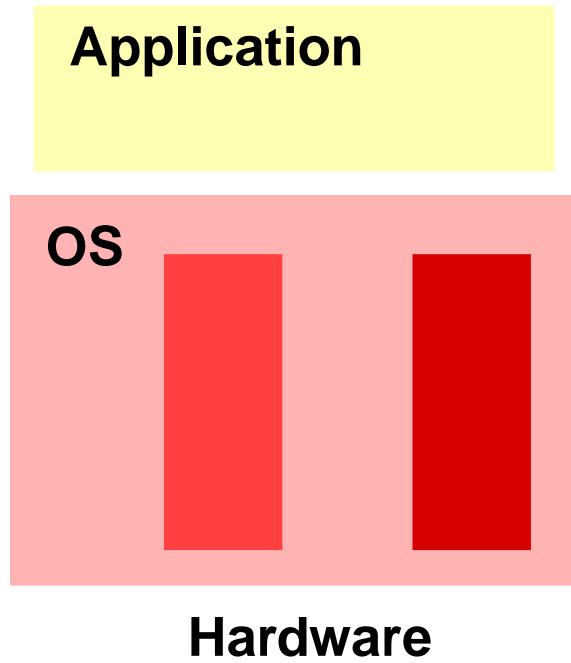
MONOLITHIC VS. MICROKERNEL

Application

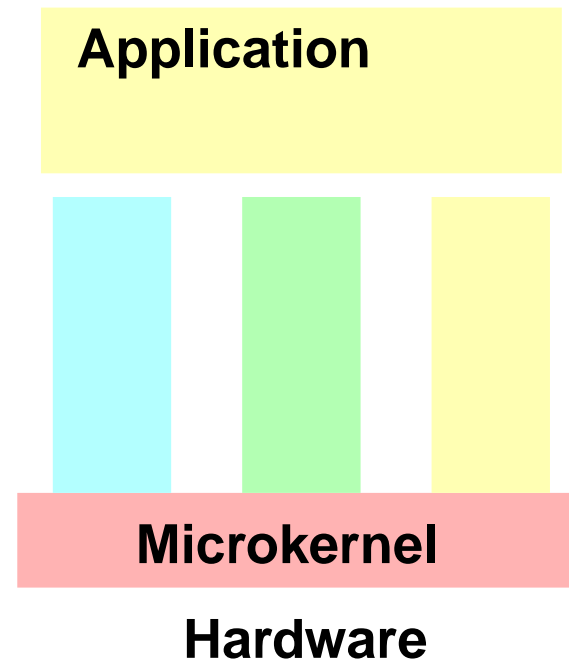
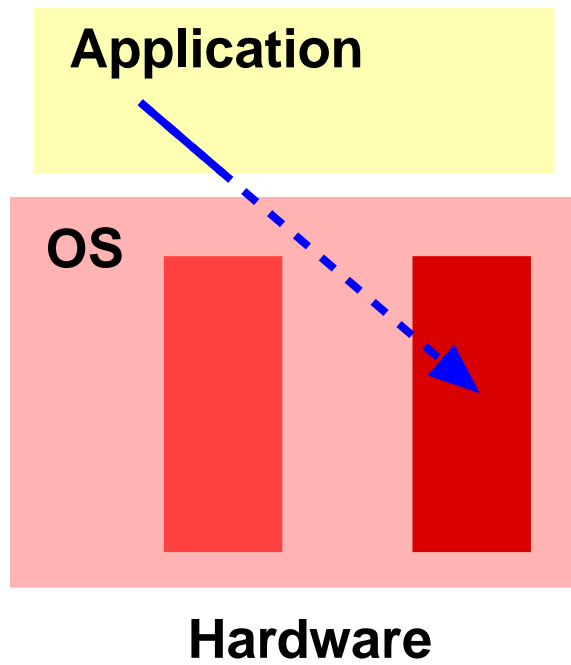


Hardware

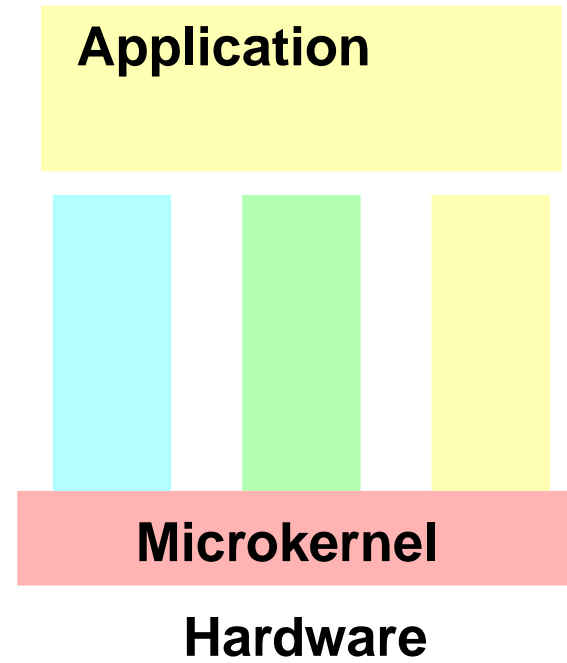
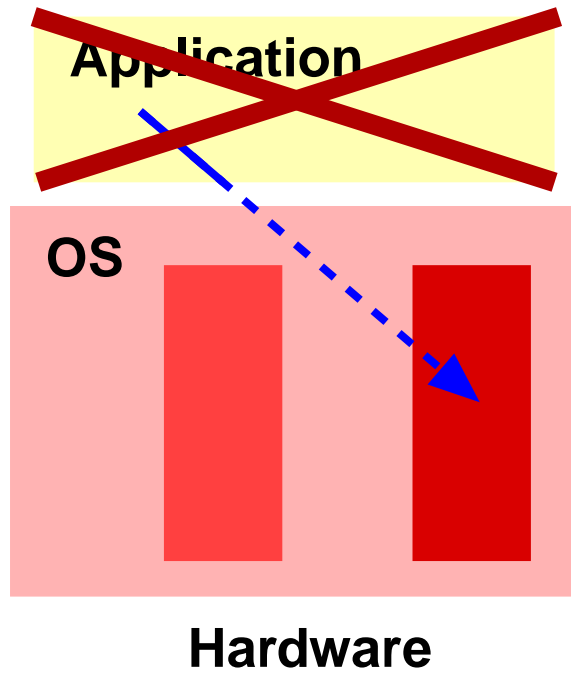
MONOLITHIC VS. MICROKERNEL



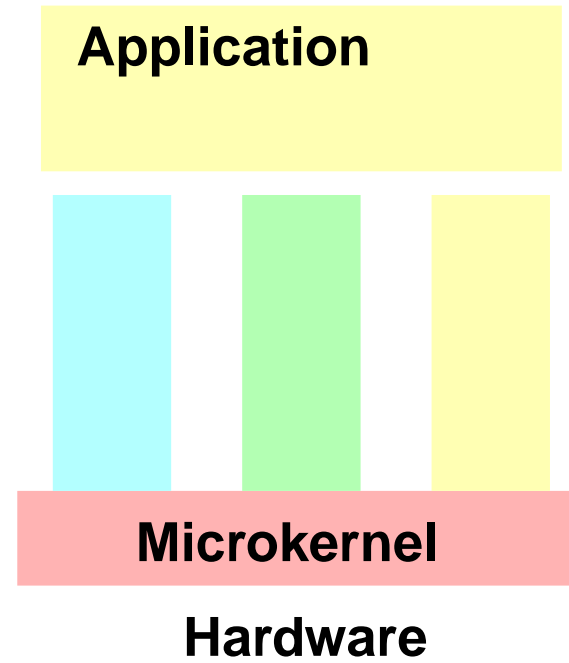
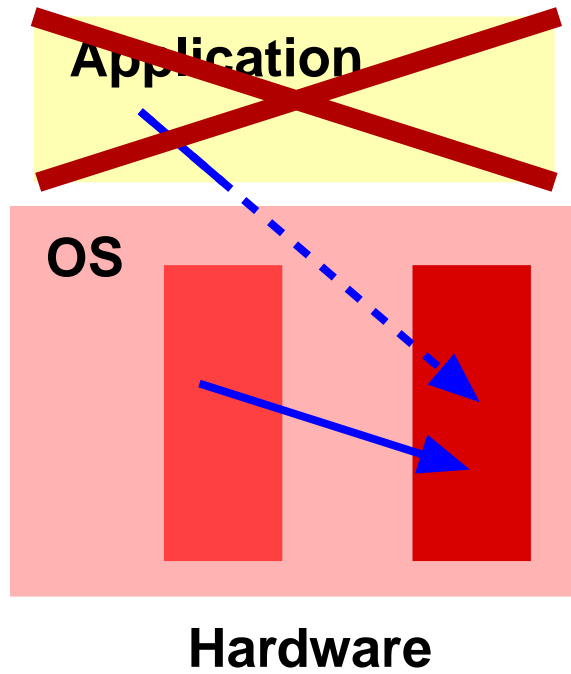
MONOLITHIC VS. MICROKERNEL



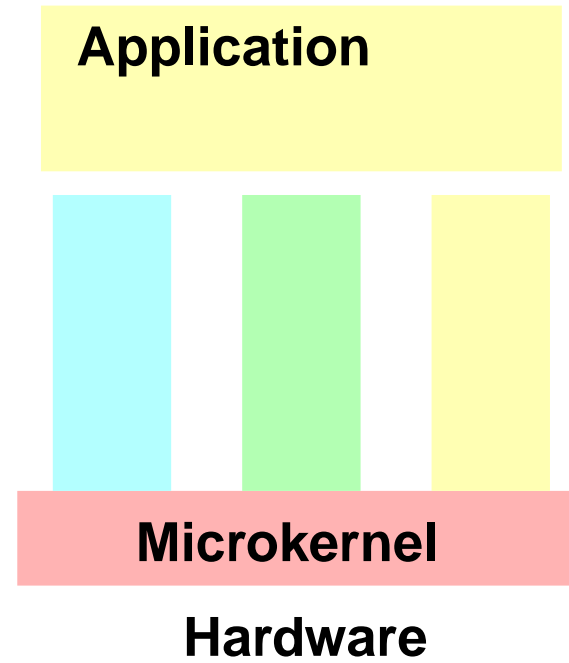
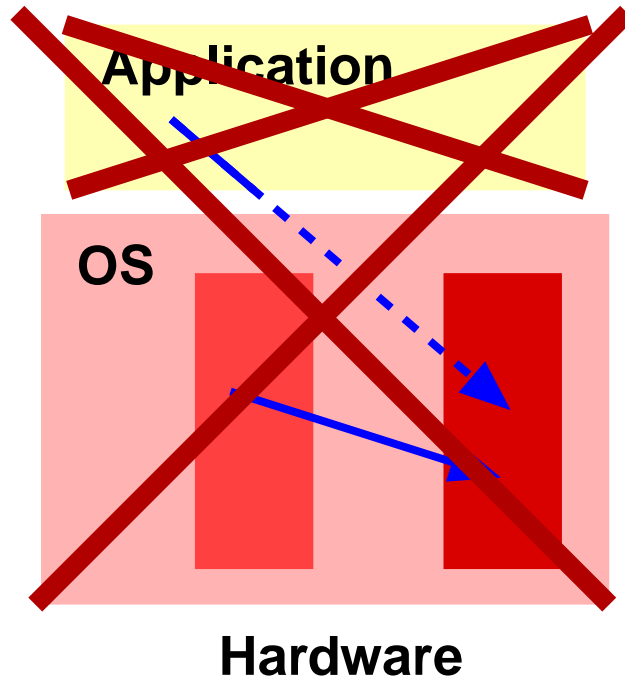
MONOLITHIC VS. MICROKERNEL



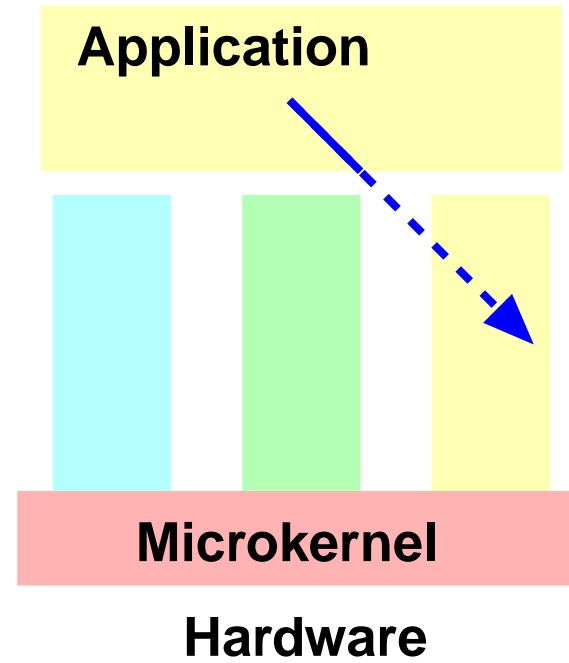
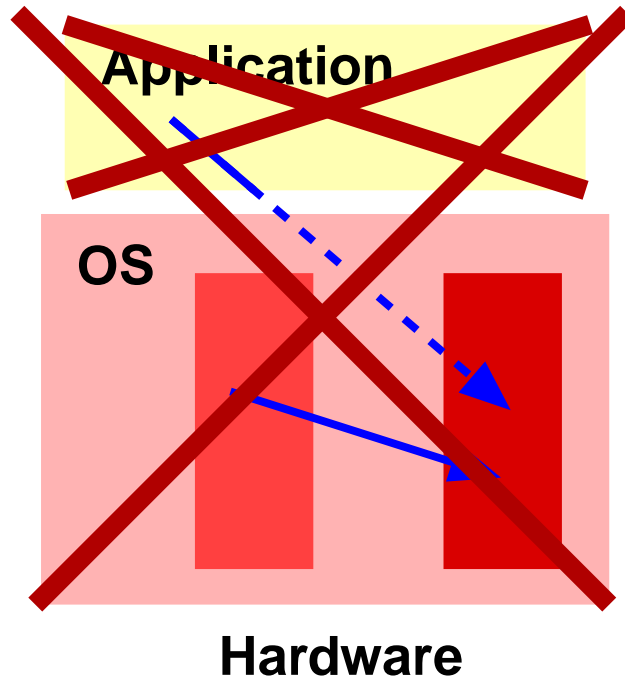
MONOLITHIC VS. MICROKERNEL



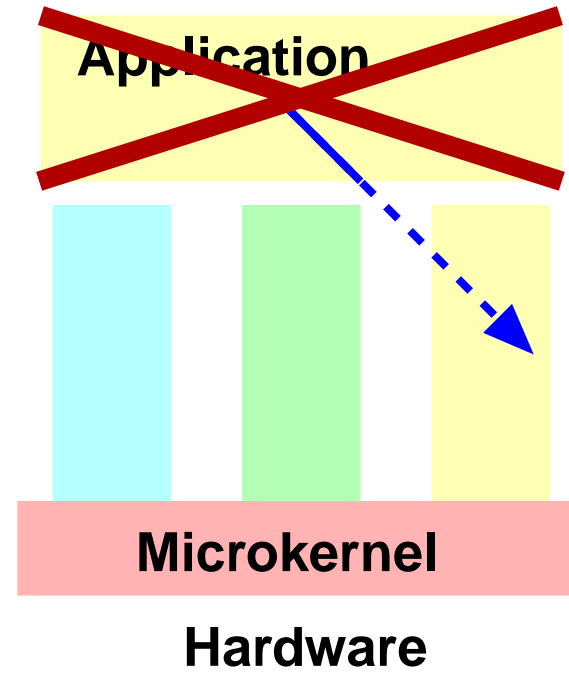
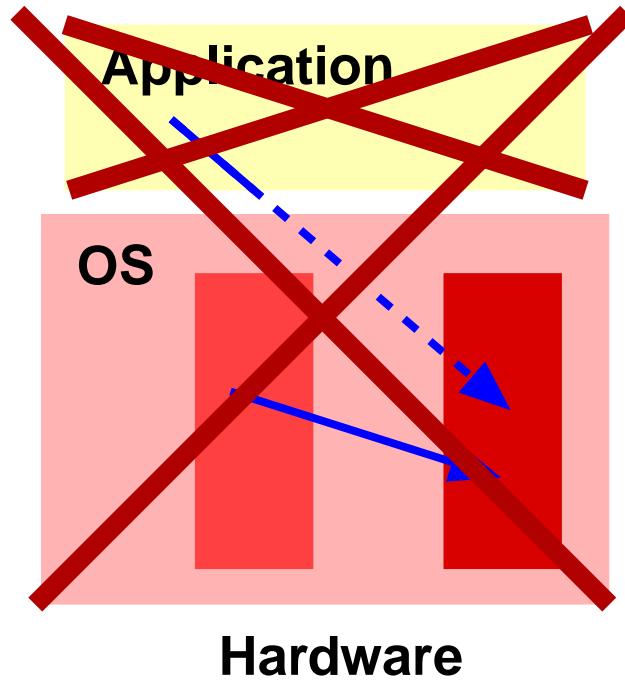
MONOLITHIC VS. MICROKERNEL



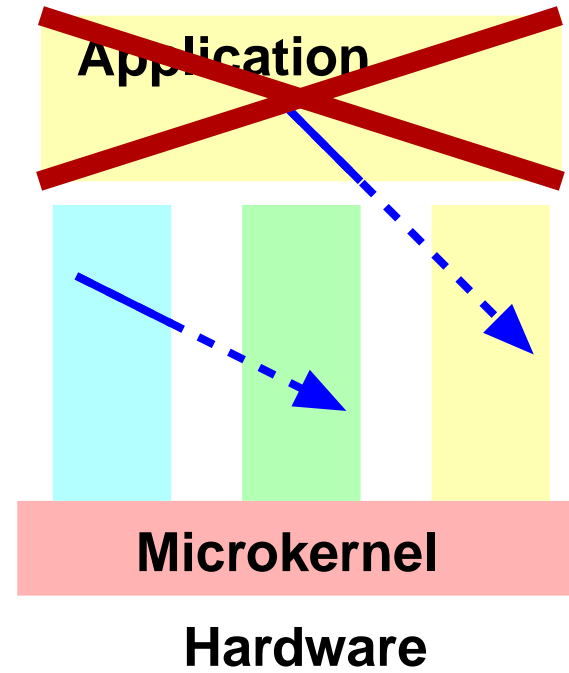
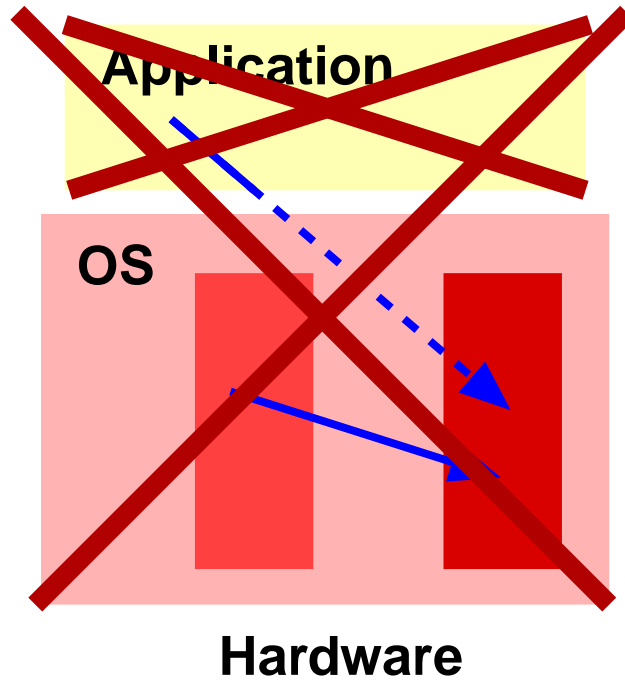
MONOLITHIC VS. MICROKERNEL



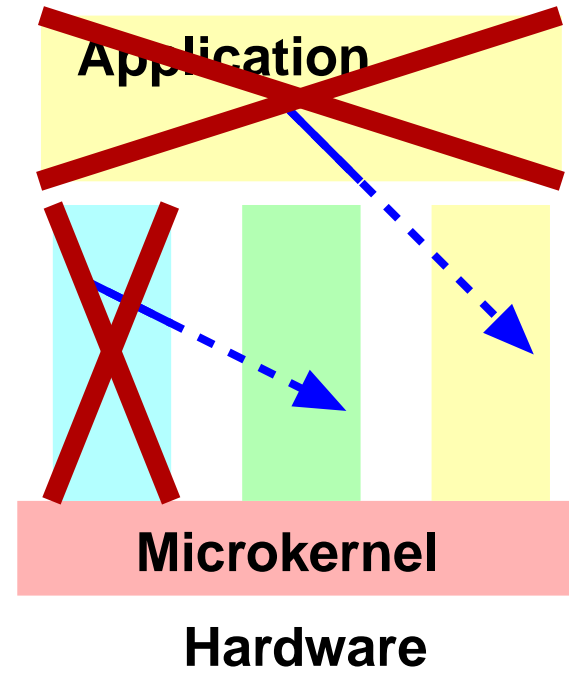
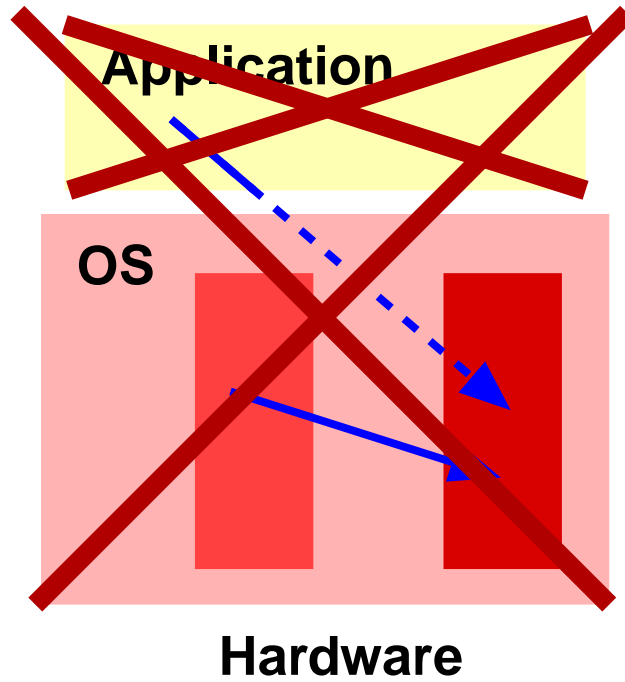
MONOLITHIC VS. MICROKERNEL



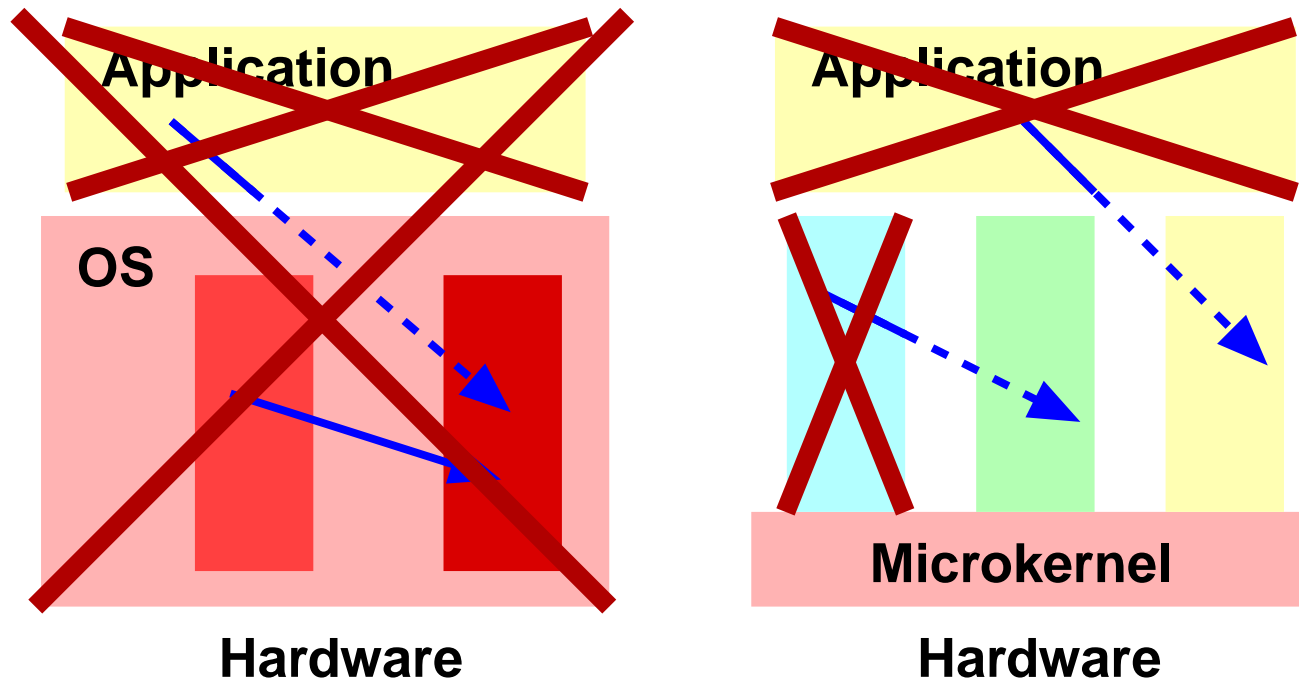
MONOLITHIC VS. MICROKERNEL



MONOLITHIC VS. MICROKERNEL

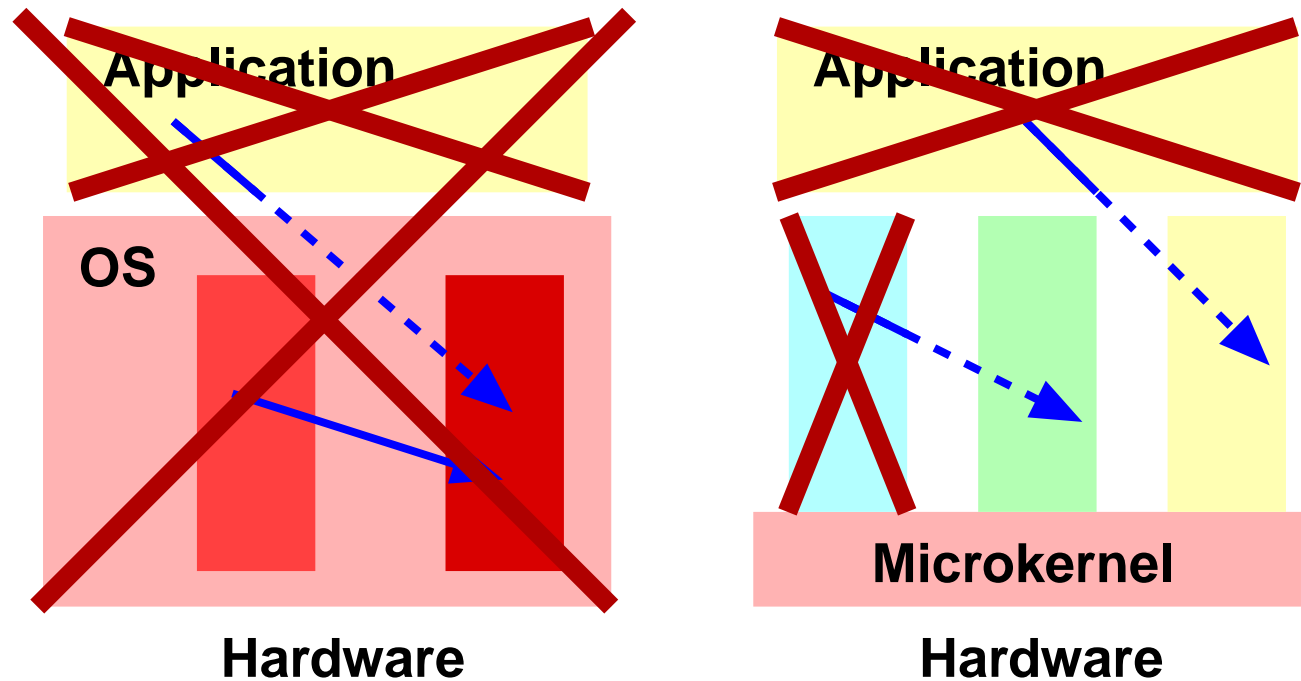


MONOLITHIC VS. MICROKERNEL



- Microkernel as small, simple, efficient lowest software layer
 - Provides address spaces for isolation/encapsulation/protection

MONOLITHIC VS. MICROKERNEL

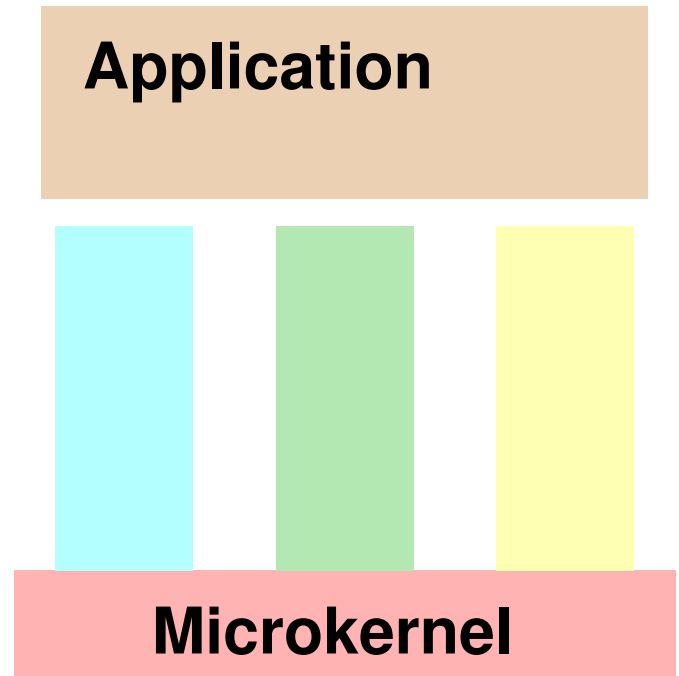


- Microkernel as small, simple, efficient lowest software layer
 - Provides address spaces for isolation/encapsulation/protection
 - Provides basis of modularity/componentisation and re-use
 - Provides basis for configurability/minimality
 - Supports standard tools for system development/debugging
 - Provides flexible resource management at application level

MICROKERNEL

SUPPORTS:

- fault isolation
- hot swapping / hot upgrade
- enforcement of security policies
- resource management for system services
- software engineering:
component architectures
- formal verification



PROJECT: L4 MICROKERNEL

- Very small and fast second-generation microkernel
 - IPC 20 times faster than Mach

PROJECT: L4 MICROKERNEL

- Very small and fast second-generation microkernel
 - IPC 20 times faster than Mach
- Originally developed by Jochen Liedtke (GMD, IBM) ≈95
 - ★ written in x86 assembler
 - ★ later implemented on MIPS (UNSW), Alpha (Dresden, UNSW)

PROJECT: L4 MICROKERNEL

- Very small and fast second-generation microkernel
 - IPC 20 times faster than Mach
- Originally developed by Jochen Liedtke (GMD, IBM) ≈95
 - ★ written in x86 assembler
 - ★ later implemented on MIPS (UNSW), Alpha (Dresden, UNSW)
- L4Ka::Pistachio implementation (Karlsruhe, UNSW) 2001–3
 - ★ written in C++ subset (with assembler fast path)
 - ★ portable:
 - x86, 32-bit PowerPC, Itanium (Karlsruhe)
 - ARM, MIPS, Alpha, 64-bit PowerPC (UNSW/NICTA)
 - UltraSPARC in progress (NICTA)

PROJECT: L4 MICROKERNEL

- Very small and fast second-generation microkernel
 - IPC 20 times faster than Mach
- Originally developed by Jochen Liedtke (GMD, IBM) ≈95
 - ★ written in x86 assembler
 - ★ later implemented on MIPS (UNSW), Alpha (Dresden, UNSW)
- L4Ka::Pistachio implementation (Karlsruhe, UNSW) 2001–3
 - ★ written in C++ subset (with assembler fast path)
 - ★ portable:
 - x86, 32-bit PowerPC, Itanium (Karlsruhe)
 - ARM, MIPS, Alpha, 64-bit PowerPC (UNSW/NICTA)
 - UltraSPARC in progress (NICTA)
 - ★ in industrial use (ARM)

PISTACHIO PERFORMANCE: IPC

Architecture	port/ optimisation	C++		optimised	
		intra AS	inter AS	intra AS	inter AS
Pentium-3	UKa	180	367	113	305
Small Spaces	UKa				213
Pentium-4	UKa	385	983	196	416
Itanium 2	UKa/NICTA	508	508	36	36
cross CPU	UKa	7419	7410	N/A	N/A
MIPS64	NICTA/UNSW	276	276	109	109
cross CPU	NICTA/UNSW	3238	3238	690	690
PowerPC-64	NICTA/UNSW	330	518	200‡	200‡
Alpha 21264	NICTA/UNSW	440	642	≈70†	≈70†
ARM/XScale	NICTA/UNSW	250	11,400	120–140‡	10,000‡
FASS	NICTA/UNSW	340	340	120–140‡	120–140‡
UltraSPARC	NICTA/UNSW			100‡	100‡

† “Version 2” assembler kernel

‡ Guestimate!

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
- Approach: remove non-essential stuff

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
- Approach: remove non-essential stuff
 - long IPC!

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
- Approach: remove non-essential stuff
 - long IPC!
 - (excessive) virtual registers

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
- Approach: remove non-essential stuff
 - long IPC!
 - (excessive) virtual registers
 - timeouts

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
- Approach: remove non-essential stuff
 - long IPC!
 - (excessive) virtual registers
 - timeouts
 - virtually-addressed TCBs

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
- Approach: remove non-essential stuff
 - long IPC!
 - (excessive) virtual registers
 - timeouts
 - virtually-addressed TCBs
 - per-thread kernel stacks

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
 - Approach: remove non-essential stuff
 - long IPC!
 - (excessive) virtual registers
 - timeouts
 - virtually-addressed TCBs
 - per-thread kernel stacks
- ... without reducing functionality or performance!

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues
 - ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
 - ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
 - ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance
 - Approach: remove non-essential stuff
 - long IPC!
 - (excessive) virtual registers
 - timeouts
 - virtually-addressed TCBs \leftarrow *honours thesis*
 - per-thread kernel stacks
- ... without reducing functionality or performance!

L4: ONGOING WORK

SIMPLER/SMALLER IMPLEMENTATION

- Issues

- ★ L4Ka::Pistachio optimised for supporting high-end systems (servers)
- ★ large memory footprint (presently $\approx 100\text{kB}$)
 - should fit on smartcard
- ★ complexity (in-kernel page faults, preemption, code bloat)
 - makes it hard to validate implementation (formally or informally)
 - makes it hard to establish hard worst-case latencies
 - leads to more bugs
 - harder to optimise for high performance

- Approach: remove non-essential stuff

- long IPC!
- (excessive) virtual registers
- timeouts
- virtually-addressed TCBs \Leftarrow *honours thesis*
- per-thread kernel stacks \Leftarrow *summer project*

... without reducing functionality or performance!

L4: ONGOING WORK

SECURE-SYSTEMS API

- Goal
 - ★ satisfy high security requirements
 - ★ support *certifiably secure systems*

L4: ONGOING WORK

SECURE-SYSTEMS API

- Goal
 - ★ satisfy high security requirements
 - ★ support *certifiably secure systems*
 - requires *proof of security properties*
 - present API is known not to support this

L4: ONGOING WORK

SECURE-SYSTEMS API

- Goal
 - ★ satisfy high security requirements
 - ★ support *certifiably secure systems*
 - requires *proof of security properties*
 - present API is known not to support this
- Issues:
 - ★ information flow control
 - V2: clans & chiefs: clumsy, inflexible, inefficient
 - V4: redirectors: still clumsy, often inefficient

L4: ONGOING WORK

SECURE-SYSTEMS API

- Goal
 - ★ satisfy high security requirements
 - ★ support *certifiably secure systems*
 - requires *proof of security properties*
 - present API is known not to support this
- Issues:
 - ★ information flow control
 - V2: clans & chiefs: clumsy, inflexible, inefficient
 - V4: redirectors: still clumsy, often inefficient
 - ★ IPC addressing via thread IDs reveals *global* names
 - violates *need to know*

L4: ONGOING WORK

SECURE-SYSTEMS API

- Goal
 - ★ satisfy high security requirements
 - ★ support *certifiably secure systems*
 - requires *proof of security properties*
 - present API is known not to support this
- Issues:
 - ★ information flow control
 - V2: clans & chiefs: clumsy, inflexible, inefficient
 - V4: redirectors: still clumsy, often inefficient
 - ★ IPC addressing via thread IDs reveals *global* names
 - violates *need to know*
 - ★ Primitive kernel memory management
 - allows denial-of-service attacks
 - reveals internal state
- Joint work with Dresden, Johns Hopkins (Jonathan Shapiro)
 - 2 local students (1 ME, 1 honours)

HIGH-PERFORMANCE USER-LEVEL DRIVERS

- L4 device drivers are *always* outside the kernel (at user level)
 - ★ Interrupts delivered to driver as IPC from kernel

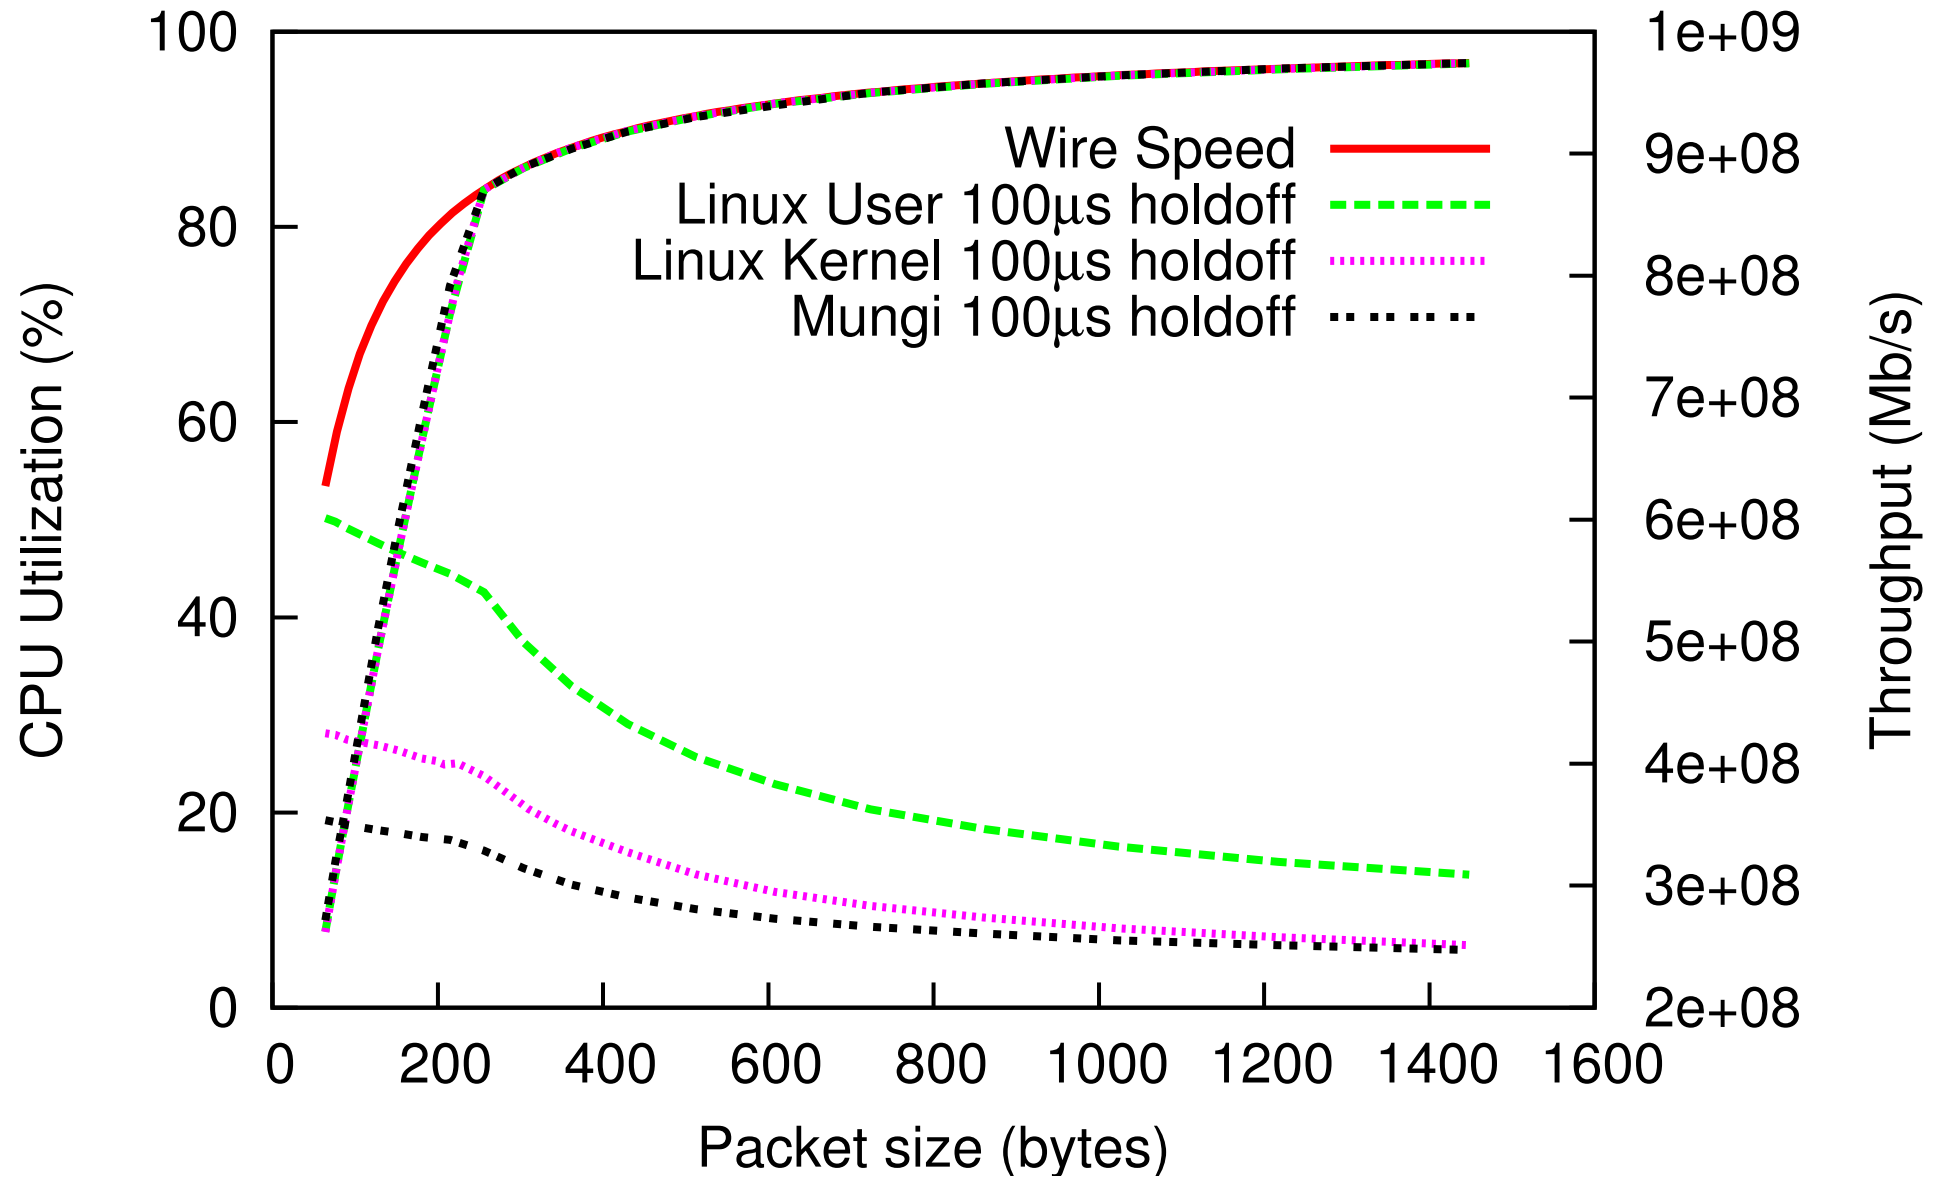
HIGH-PERFORMANCE USER-LEVEL DRIVERS

- L4 device drivers are *always* outside the kernel (at user level)
 - ★ Interrupts delivered to driver as IPC from kernel
- Potentially higher communication overhead
 - ★ past experience with user-level drivers:
≥50% performance degradation

HIGH-PERFORMANCE USER-LEVEL DRIVERS

- L4 device drivers are *always* outside the kernel (at user level)
 - ★ Interrupts delivered to driver as IPC from kernel
- Potentially higher communication overhead
 - ★ past experience with user-level drivers:
 $\geq 50\%$ performance degradation
- L4 IPC performance is very high
 - ★ with well-designed driver interfaces can achieve good performance

USER-LEVEL DEVICE DRIVER PERFORMANCE



Gigabit Ethernet echo on 900MHz Itanium-2 with 66MHz 64-bit PCI

USER-LEVEL DRIVERS: ONGOING WORK

- Complete driver framework and methodology
 - ease development of high-performance drivers
 - reduce driver complexity
 - drivers portable between systems (L4 and Linux)

USER-LEVEL DRIVERS: ONGOING WORK

- Complete driver framework and methodology
 - ease development of high-performance drivers
 - reduce driver complexity
 - drivers portable between systems (L4 and Linux)
- Integration with I/O system
 - Linux VFS layer integration
 - user-level network protocol stacks
 - componentised protocol stacks

USER-LEVEL DRIVERS: ONGOING WORK

- Complete driver framework and methodology
 - ease development of high-performance drivers
 - reduce driver complexity
 - drivers portable between systems (L4 and Linux)
- Integration with I/O system
 - Linux VFS layer integration
 - user-level network protocol stacks
 - componentised protocol stacks
- Driver encapsulation
 - use hardware mechanisms to limit DMA
 - use software mechanisms to limit trust in drivers
 - goal: untrusted device drivers
- Collaboration between NICTA and UNSW Gelato project
 - 1 PhD student

FORMAL VERIFICATION OF MICROKERNEL

- First (and most significant) step to making the TCB *trustworthy!*

FORMAL VERIFICATION OF MICROKERNEL

- First (and most significant) step to making the TCB *trustworthy!*
- Involves three aspects:
 1. develop a formal specification of the kernel API
 2. prove safety properties about the API
 - note link to API revision project
 3. prove that implementation is correct
 - note link to L4 slimming project

FORMAL VERIFICATION OF MICROKERNEL

- First (and most significant) step to making the TCB *trustworthy!*
- Involves three aspects:
 1. develop a formal specification of the kernel API
 2. prove safety properties about the API
 - note link to API revision project
 3. prove that implementation is correct
 - note link to L4 slimming project
 - never before been done for any protected kernel!
 - L4 ($\approx 10,000$ loc) is at limit of tractability

FORMAL VERIFICATION OF MICROKERNEL

- First (and most significant) step to making the TCB *trustworthy*!
- Involves three aspects:
 1. develop a formal specification of the kernel API
 2. prove safety properties about the API
 - note link to API revision project
 3. prove that implementation is correct
 - note link to L4 slimming project
 - never before been done for any protected kernel!
 - L4 ($\approx 10,000$ loc) is at limit of tractability
- Ultimate goal: Verify a *real* kernel (L4) that is used in production

FORMAL VERIFICATION OF MICROKERNEL

- First (and most significant) step to making the TCB *trustworthy*!
- Involves three aspects:
 1. develop a formal specification of the kernel API
 2. prove safety properties about the API
 - note link to API revision project
 3. prove that implementation is correct
 - note link to L4 slimming project
 - never before been done for any protected kernel!
 - L4 ($\approx 10,000$ loc) is at limit of tractability
- Ultimate goal: Verify a *real* kernel (L4) that is used in production
- Reduce risk: pilot project
 - collaboration with NICTA Formal Methods Program

L4 VERIFICATION PILOT PROJECT

- Purpose:
 - ★ test tools and get experience with them
 - ★ feasibility study
 - ★ bridge cultural gap between theorists and kernel hackers
 - ★ develop project plan for full-scale verification project

L4 VERIFICATION PILOT PROJECT

- Purpose:
 - ★ test tools and get experience with them
 - ★ feasibility study
 - ★ bridge cultural gap between theorists and kernel hackers
 - ★ develop project plan for full-scale verification project
- Operates on a “slice” of the API (address-space management)
 - ★ done formal model of (simplified) API slice
 - ★ proved some properties (absence of loops in mappings)
 - ★ refined specification down to source code
 - page table operations and simplified mapping database

L4 VERIFICATION PILOT PROJECT

- Purpose:
 - ★ test tools and get experience with them
 - ★ feasibility study
 - ★ bridge cultural gap between theorists and kernel hackers
 - ★ develop project plan for full-scale verification project
- Operates on a “slice” of the API (address-space management)
 - ★ done formal model of (simplified) API slice
 - ★ proved some properties (absence of loops in mappings)
 - ★ refined specification down to source code
 - page table operations and simplified mapping database

L4 VERIFICATION PROJECT

- Timeline:
 - ★ pilot project concludes January 2005
 - ★ main project: 3 years from February 2005
 - will be based on new L4 API
- Resources: \approx 20 person years

L4 VERIFICATION PROJECT

- Timeline:
 - ★ pilot project concludes January 2005
 - ★ main project: 3 years from February 2005
 - will be based on new L4 API
- Resources: \approx 20 person years
 - ★ presently:
 - 2 FTE researchers
 - 1 PhD student
 - 1 honours
 - ★ definitely looking for students!

L4 VERIFICATION PROJECT

- Timeline:
 - ★ pilot project concludes January 2005
 - ★ main project: 3 years from February 2005
 - will be based on new L4 API
- Resources: \approx 20 person years
 - ★ presently:
 - 2 FTE researchers
 - 1 PhD student
 - 1 honours
 - ★ definitely looking for students!
 - ★ desired background:
 - good systems background and aptitude for maths, or
 - maths background and aptitude for systems

PROJECT: L4 WCET ANALYSIS

- Hard real-time systems need guarantees of worst-case latencies
 - ★ Need *worst-case execution time* (WCET) of kernel operations

PROJECT: L4 WCET ANALYSIS

- Hard real-time systems need guarantees of worst-case latencies
 - ★ Need *worst-case execution time* (WCET) of kernel operations
- Usually done:
 - ★ by code inspection
 - unfeasible for non-trivial kernel
 - ★ experimentally
 - unreliable: cannot guarantee 100% coverage!

PROJECT: L4 WCET ANALYSIS

- Hard real-time systems need guarantees of worst-case latencies
 - ★ Need *worst-case execution time* (WCET) of kernel operations
- Usually done:
 - ★ by code inspection
 - unfeasible for non-trivial kernel
 - ★ experimentally
 - unreliable: cannot guarantee 100% coverage!
 - ★ by educated guess
 - highly unreliable (RTLinux got it wrong!)
 - typically leads to vastly pessimistic latencies

PROJECT: L4 WCET ANALYSIS

- Hard real-time systems need guarantees of worst-case latencies
 - ★ Need *worst-case execution time* (WCET) of kernel operations
- Usually done:
 - ★ by code inspection
 - unfeasible for non-trivial kernel
 - ★ experimentally
 - unreliable: cannot guarantee 100% coverage!
 - ★ by educated guess
 - highly unreliable (RTLinux got it wrong!)
 - typically leads to vastly pessimistic latencies
- Need systematic execution-time analysis of complete kernel
 - never been done for non-trivial kernel!

PROJECT: L4 WCET ANALYSIS

- Hard real-time systems need guarantees of worst-case latencies
 - ★ Need *worst-case execution time* (WCET) of kernel operations
- Usually done:
 - ★ by code inspection
 - unfeasible for non-trivial kernel
 - ★ experimentally
 - unreliable: cannot guarantee 100% coverage!
 - ★ by educated guess
 - highly unreliable (RTLinux got it wrong!)
 - typically leads to vastly pessimistic latencies
- Need systematic execution-time analysis of complete kernel
 - never been done for non-trivial kernel!
 - hired an expert (Stefan Petters) to do the project
 - looking for students!

IGUANA: OS PERSONALITY FOR EMBEDDED SYSTEMS

- Remember, L4 is a “strict” microkernel:
 - does not provide any services
 - does not provide policies (or only very few)
 - provides mechanisms
- L4 aspires to be generic kernel, suitable for all kinds of systems

IGUANA: OS PERSONALITY FOR EMBEDDED SYSTEMS

- Remember, L4 is a “strict” microkernel:
 - does not provide any services
 - does not provide policies (or only very few)
 - provides mechanisms
- L4 aspires to be generic kernel, suitable for all kinds of systems
- Almost any system requires a set of core services:
 - process management
 - memory management
 - security management

... based on some system-wide policies

IGUANA: OS PERSONALITY FOR EMBEDDED SYSTEMS

- Remember, L4 is a “strict” microkernel:
 - does not provide any services
 - does not provide policies (or only very few)
 - provides mechanisms
- L4 aspires to be generic kernel, suitable for all kinds of systems
- Almost any system requires a set of core services:
 - process management
 - memory management
 - security management

... based on some system-wide policies
- Iguana provides these (or at least more tools for providing them)
 - designed for use in embedded systems
 - designed to minimise trusted computing base

WHAT DOES IGUANA PROVIDE?

- ★ Convenient way of using L4 primitives
 - OO-style method invocations instead of explicit IPC calls
 - IDL compiler for automatic generation of stubs

WHAT DOES IGUANA PROVIDE?

- ★ Convenient way of using L4 primitives
 - OO-style method invocations instead of explicit IPC calls
 - IDL compiler for automatic generation of stubs
- ★ Protection framework for access rights management
 - capability based, flexible
 - able to model most standard security models

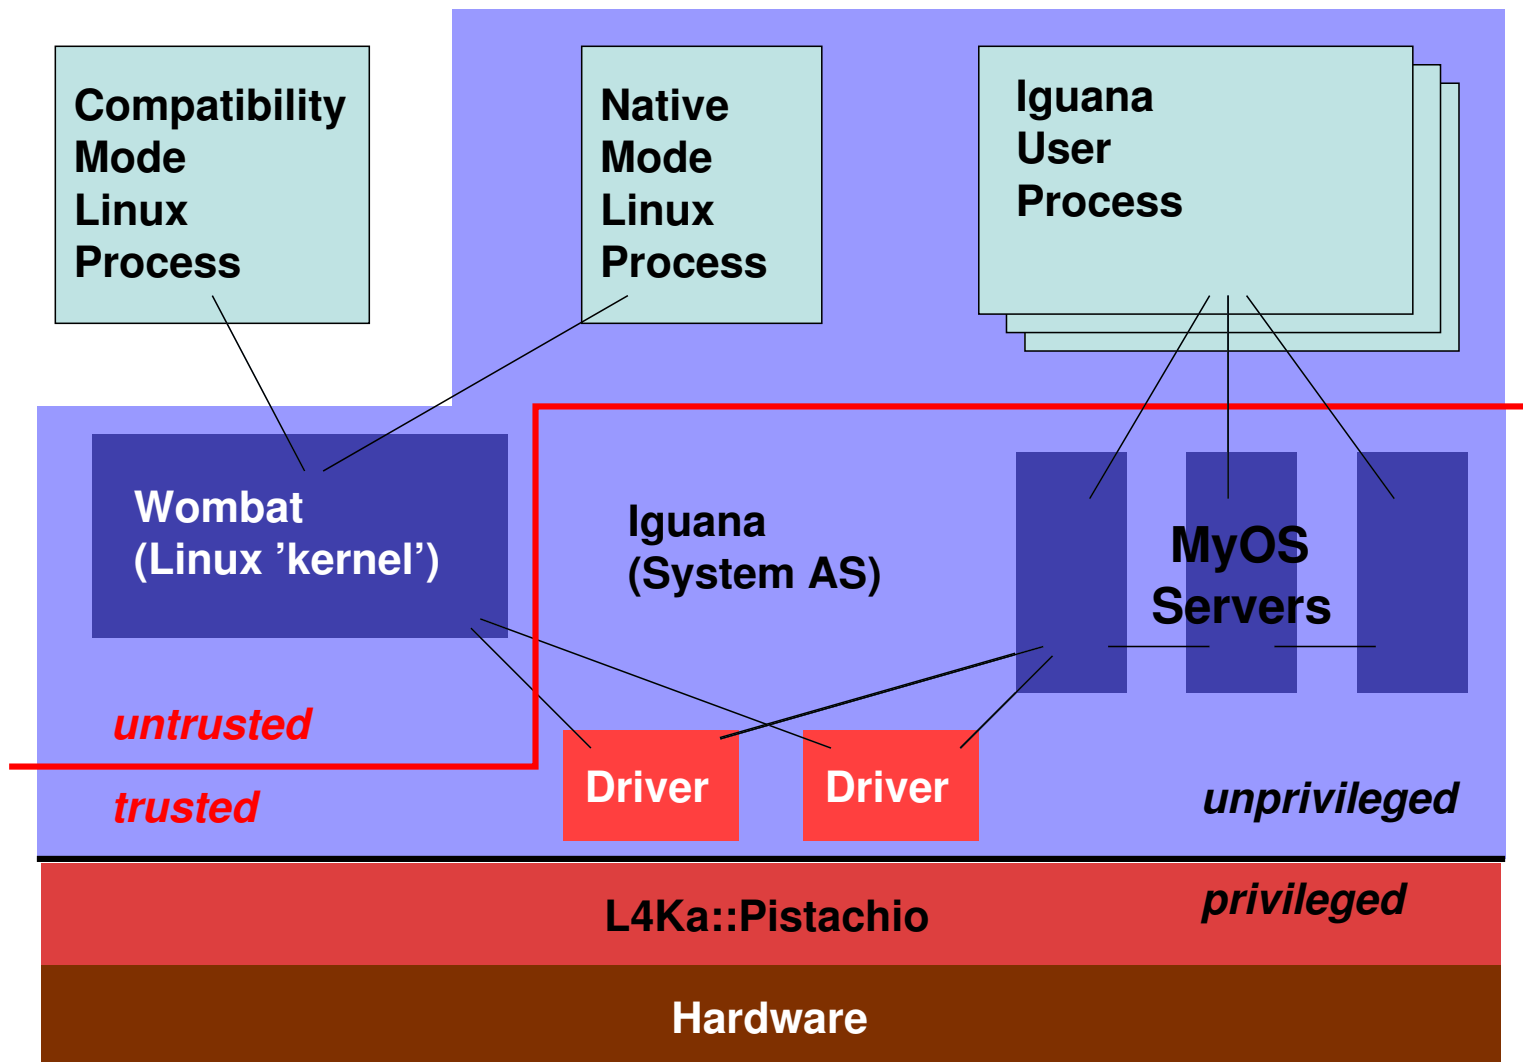
WHAT DOES IGUANA PROVIDE?

- ★ Convenient way of using L4 primitives
 - OO-style method invocations instead of explicit IPC calls
 - IDL compiler for automatic generation of stubs
- ★ Protection framework for access rights management
 - capability based, flexible
 - able to model most standard security models
- ★ Virtual memory management
 - allocation, deallocation, sharing, ...
 - single-address-space view, supporting FASS on ARM

WHAT DOES IGUANA PROVIDE?

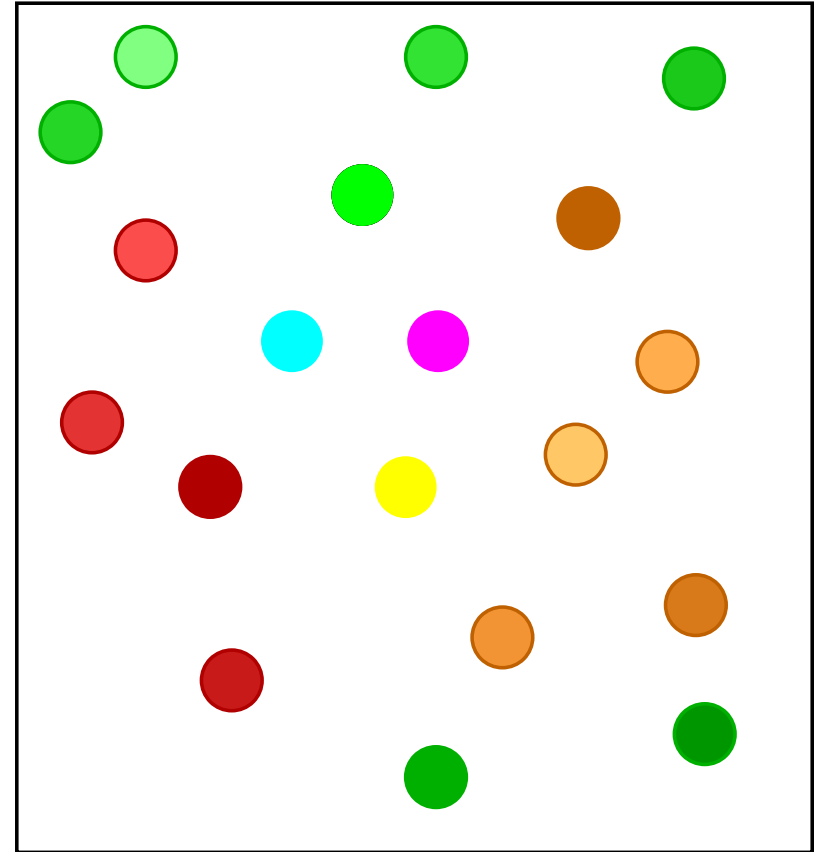
- ★ Convenient way of using L4 primitives
 - OO-style method invocations instead of explicit IPC calls
 - IDL compiler for automatic generation of stubs
- ★ Protection framework for access rights management
 - capability based, flexible
 - able to model most standard security models
- ★ Virtual memory management
 - allocation, deallocation, sharing, ...
 - single-address-space view, supporting FASS on ARM
- ★ Protection-domain (process) management
- ★ Thread management

SAMPLE IGUANA SYSTEM



IGUANA : BASIC APPROACH

- Basic idea: single address space (SAS)



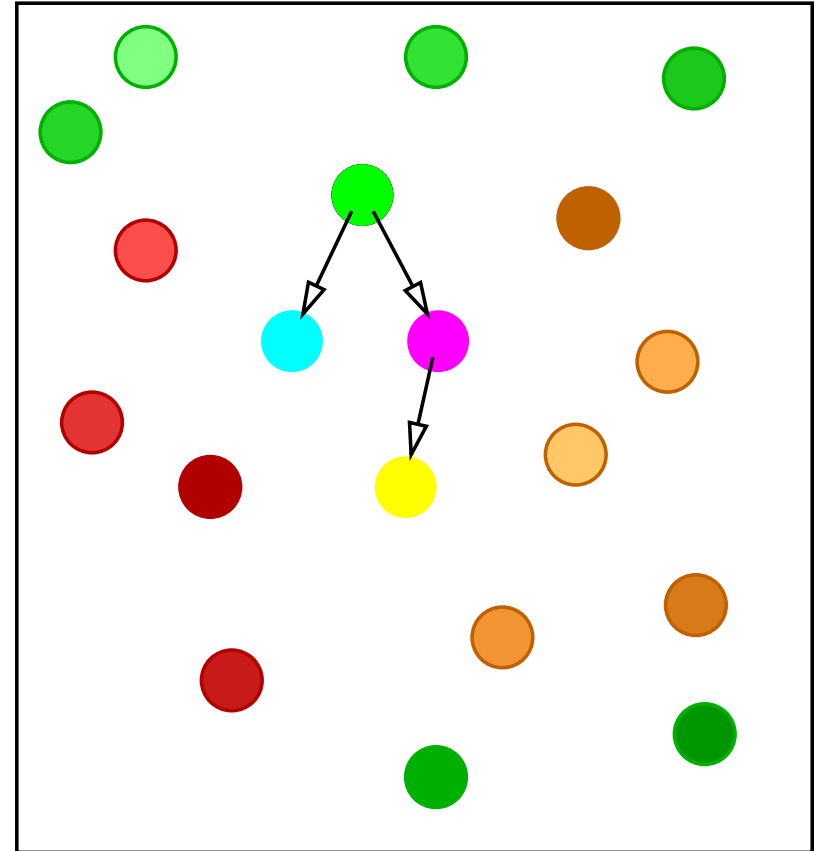
IGUANA : BASIC APPROACH

- Basic idea: single address space (SAS)

★ eases sharing of data

→ minimises copying

→ no problems with pointers



IGUANA : BASIC APPROACH

- Basic idea: single address space (SAS)

- ★ eases sharing of data

- minimises copying

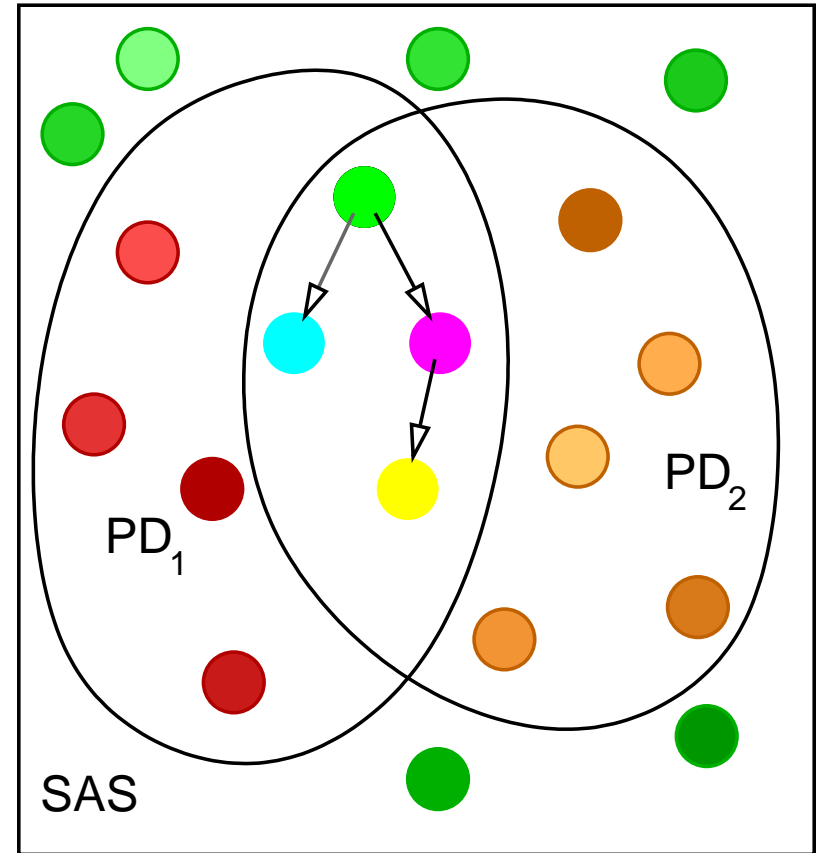
- no problems with pointers

- Per-process *protection domains*

- ★ enforce security policy

- any access is subject to access control

- ★ do not interfere with sharing



IGUANA : BASIC APPROACH

- Basic idea: single address space (SAS)

- ★ eases sharing of data

- minimises copying

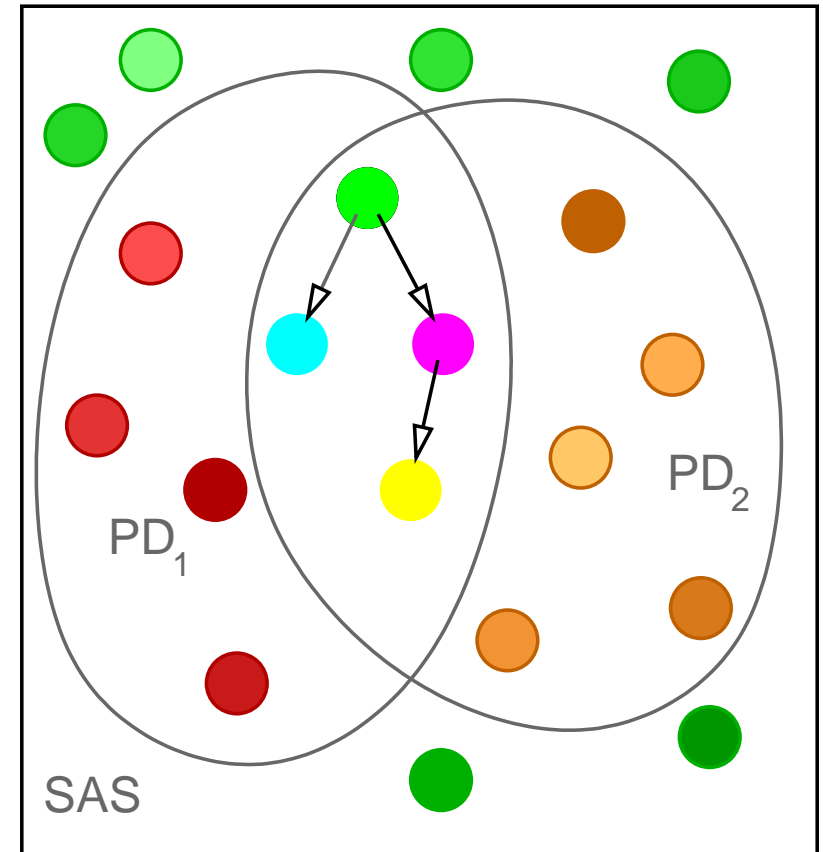
- no problems with pointers

- Per-process *protection domains*

- ★ enforce security policy

- any access is subject to access control

- ★ do not interfere with sharing



- SAS layout supports fast-address-space switching on ARM

- avoids AS overlaps for non-shared data without use of PID relocation

- advantage: 1MB domain granularity instead of 32MB for PID relocation

- less internal fragmentation

IGUANA CONCEPTS

- Memory section
 - unit of VM allocation and protection
 - can be an encapsulated object with methods and data

IGUANA CONCEPTS

- Memory section

- unit of VM allocation and protection

- can be an encapsulated object with methods and data

- Thread

- execution abstraction, as in L4

IGUANA CONCEPTS

- Memory section

- unit of VM allocation and protection
- can be an encapsulated object with methods and data

- Thread

- execution abstraction, as in L4

- Server

- thread associated with memory section
- invoked through methods with well-defined interfaces

IGUANA CONCEPTS

- Memory section
 - unit of VM allocation and protection
 - can be an encapsulated object with methods and data
- Thread
 - execution abstraction, as in L4
- Server
 - thread associated with memory section
 - invoked through methods with well-defined interfaces
- Protection domain
 - defines access and resource rights of a thread
 - corresponds to a process in traditional OS

IGUANA CONCEPTS

- Session
 - client-server (or peer-to-peer) communication channel
 - amortises authentication cost over many invocations

IGUANA CONCEPTS

- Session

- client-server (or peer-to-peer) communication channel
- amortises authentication cost over many invocations

- Capability

- represents access rights
- basis of protection

IGUANA CONCEPTS

- Session

- client-server (or peer-to-peer) communication channel
- amortises authentication cost over many invocations

- Capability

- represents access rights
- basis of protection

- Resource token

- represents resource usage right
- basis of resource management

IGUANA CONCEPTS

- Session

- client-server (or peer-to-peer) communication channel
- amortises authentication cost over many invocations

- Capability

- represents access rights
- basis of protection

- Resource token

- represents resource usage right
- basis of resource management

- External Space

- address space extern to Iguana's SAS
- for legacy support and large processes

IGUANA PHILOSOPHY

- Small and lightweight
- Strong yet unintrusive protection
- Support for resource management
- Legacy support
- Code and concept re-use

IGUANA PHILOSOPHY

- Small and lightweight
 - geared towards embedded systems
 - allow optimal utilisation of hardware
- Strong yet unintrusive protection
- Support for resource management
- Legacy support
- Code and concept re-use

IGUANA PHILOSOPHY

- Small and lightweight
 - geared towards embedded systems
 - allow optimal utilisation of hardware
- Strong yet unintrusive protection
 - hide protection machinery from most apps
 - able to emulate most standard protection models
- Support for resource management
- Legacy support
- Code and concept re-use

IGUANA PHILOSOPHY

- Small and lightweight
 - geared towards embedded systems
 - allow optimal utilisation of hardware
- Strong yet unintrusive protection
 - hide protection machinery from most apps
 - able to emulate most standard protection models
- Support for resource management
 - in principle, although it isn't implemented yet!
- Legacy support
- Code and concept re-use

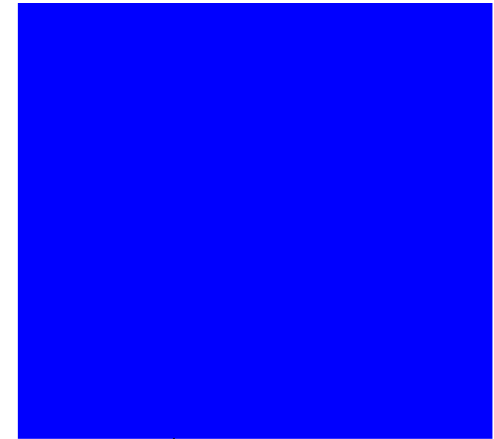
IGUANA PHILOSOPHY

- Small and lightweight
 - geared towards embedded systems
 - allow optimal utilisation of hardware
- Strong yet unintrusive protection
 - hide protection machinery from most apps
 - able to emulate most standard protection models
- Support for resource management
 - in principle, although it isn't implemented yet!
- Legacy support
 - designed to run Linux server
- Code and concept re-use

IGUANA PHILOSOPHY

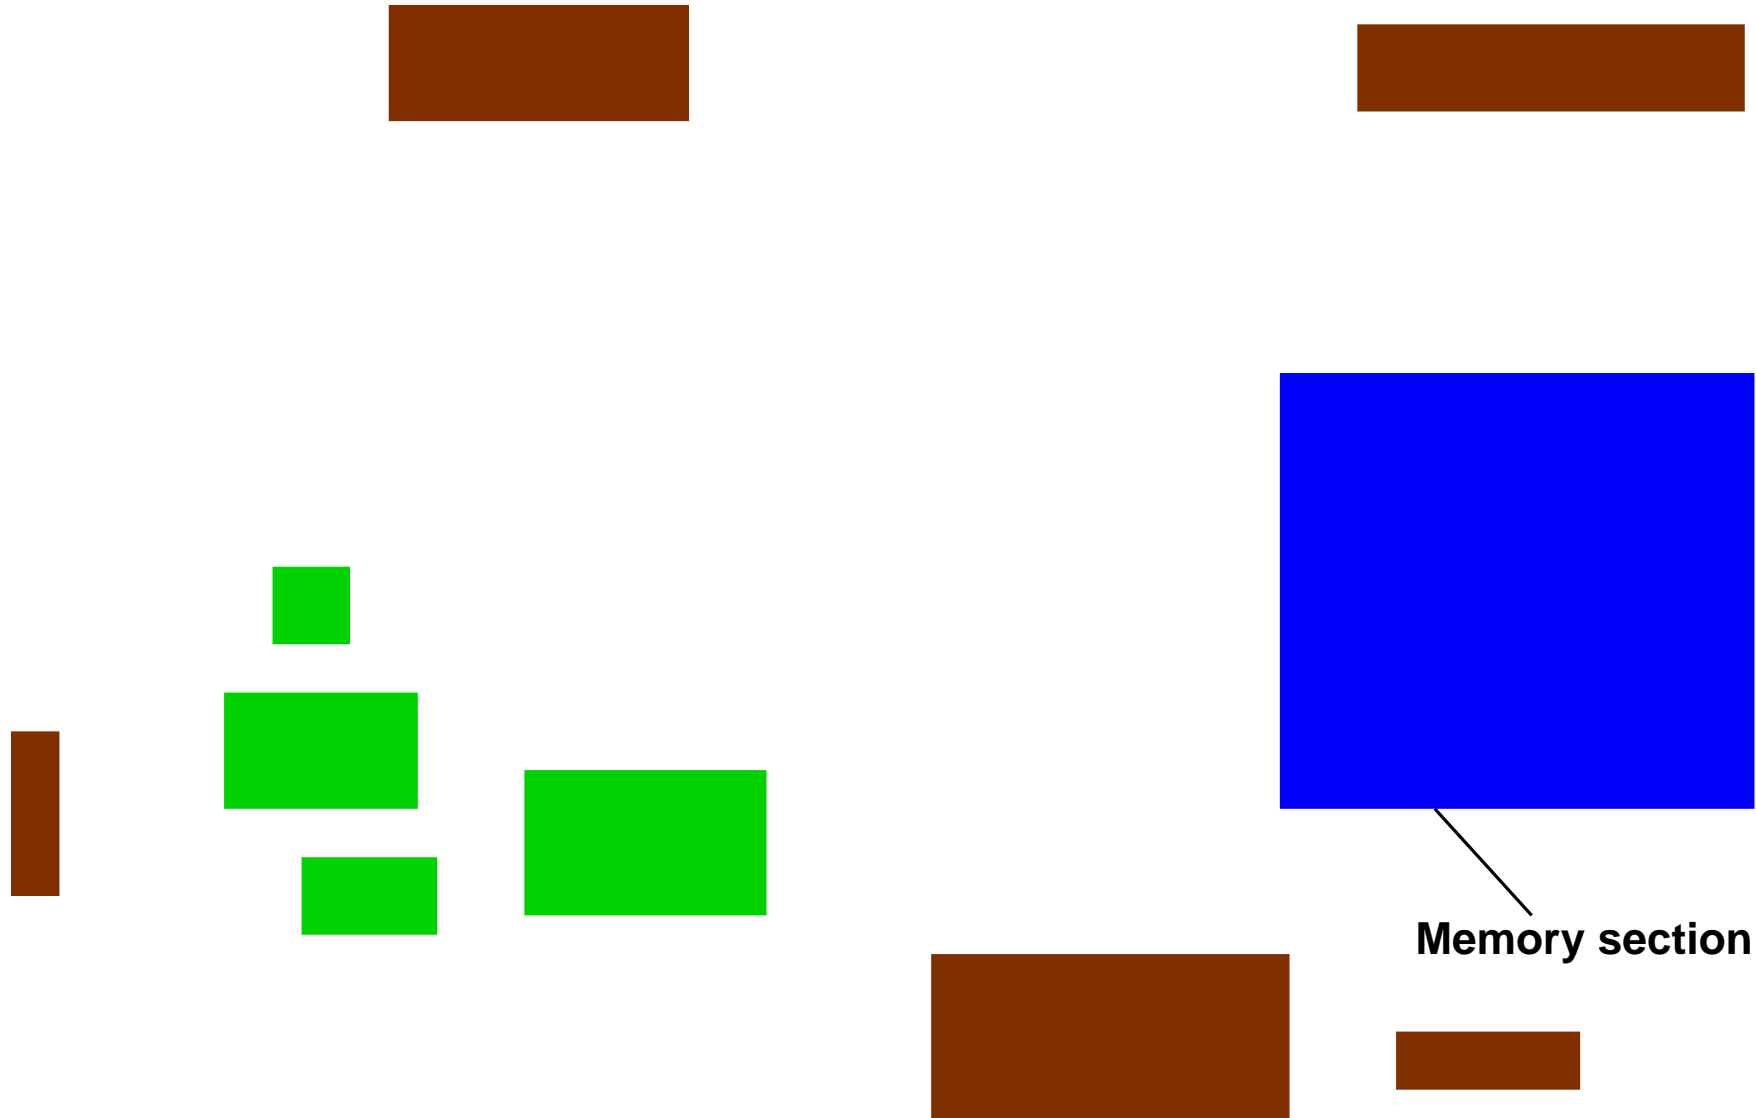
- Small and lightweight
 - geared towards embedded systems
 - allow optimal utilisation of hardware
- Strong yet unintrusive protection
 - hide protection machinery from most apps
 - able to emulate most standard protection models
- Support for resource management
 - in principle, although it isn't implemented yet!
- Legacy support
 - designed to run Linux server
- Code and concept re-use
 - utilise what we've learned in and developed for Mungi project

IGUANA CONCEPTS

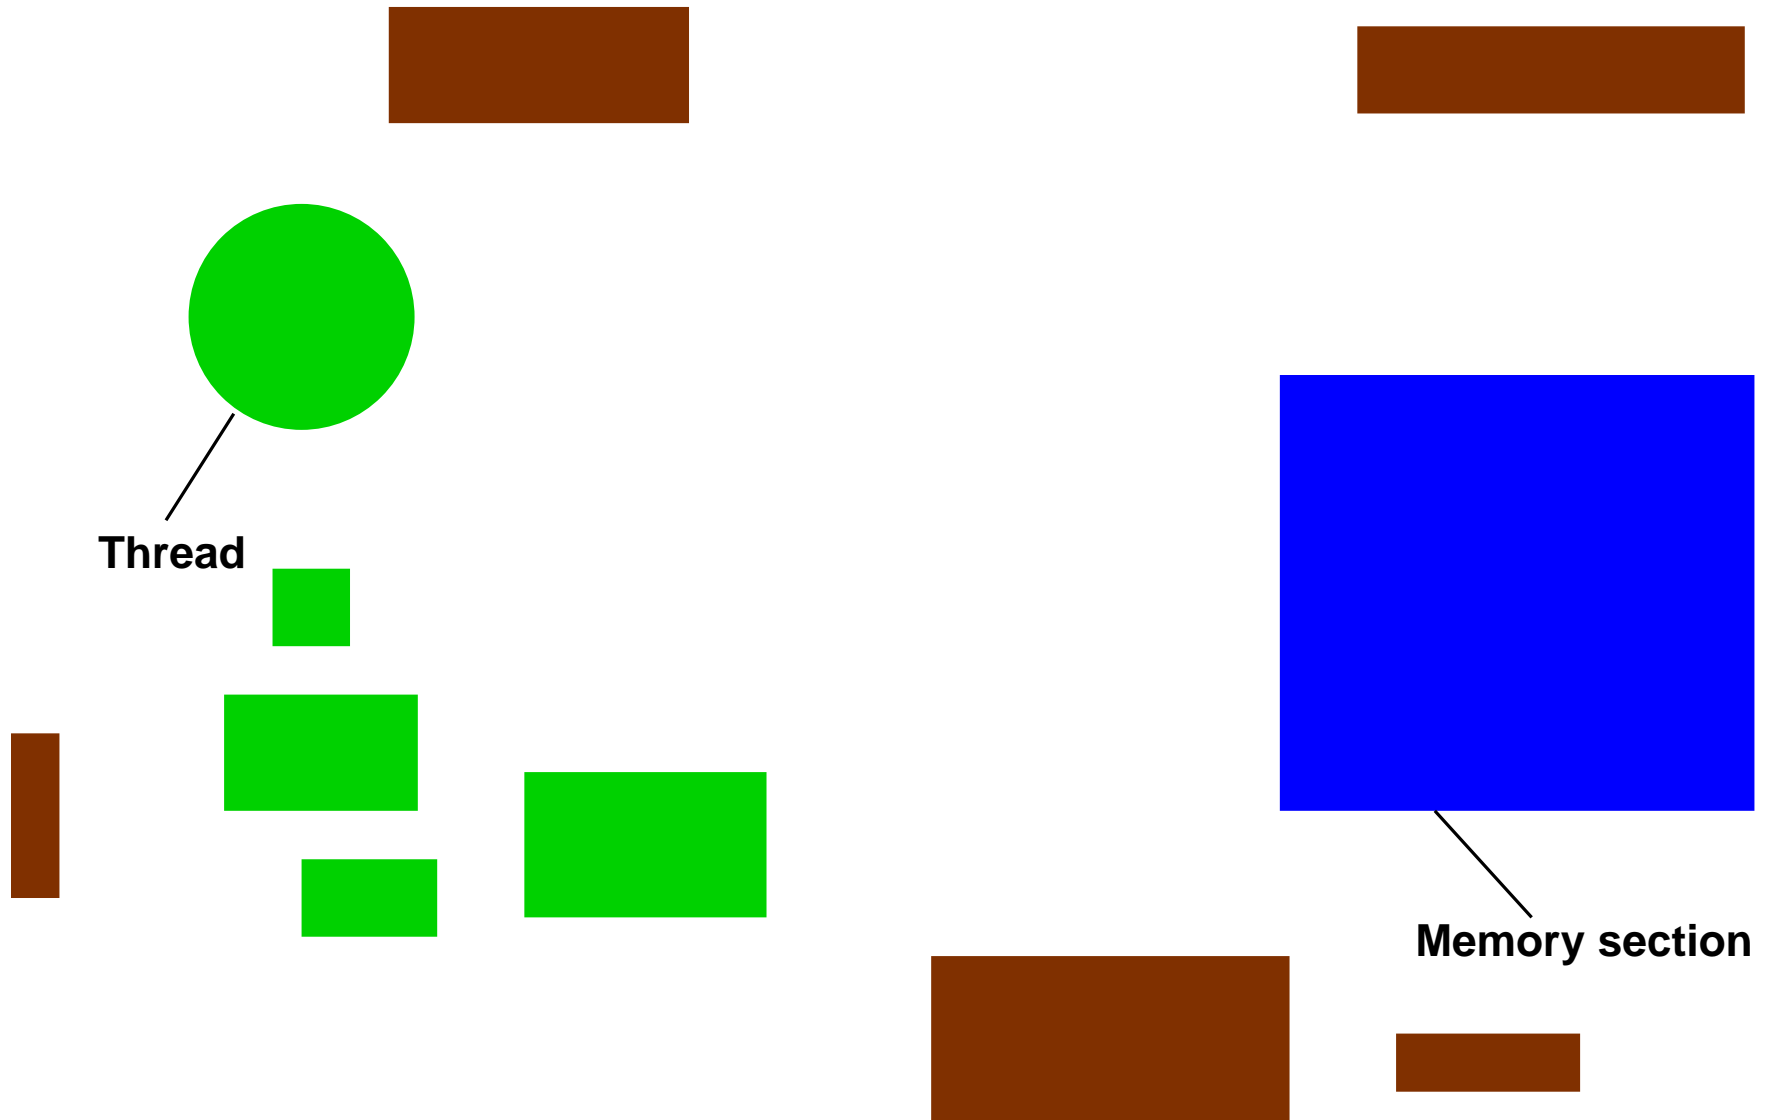


Memory section

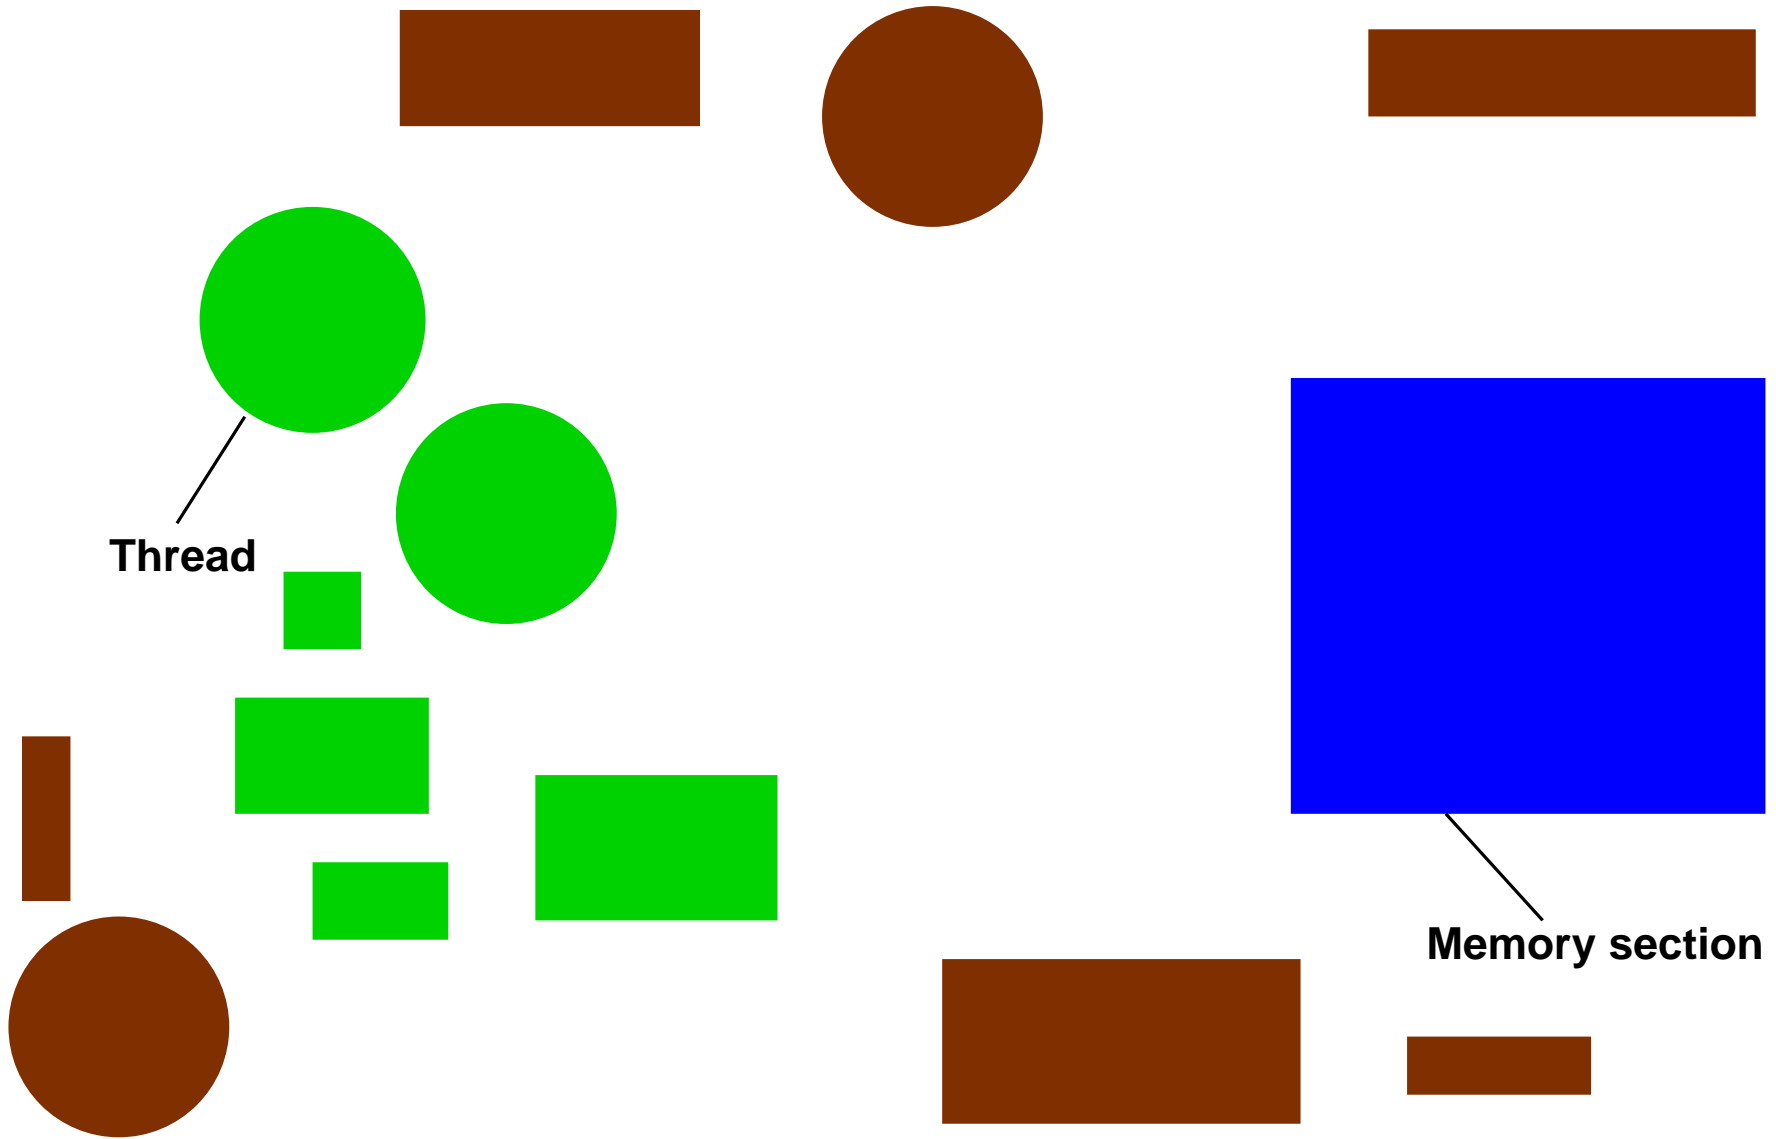
IGUANA CONCEPTS



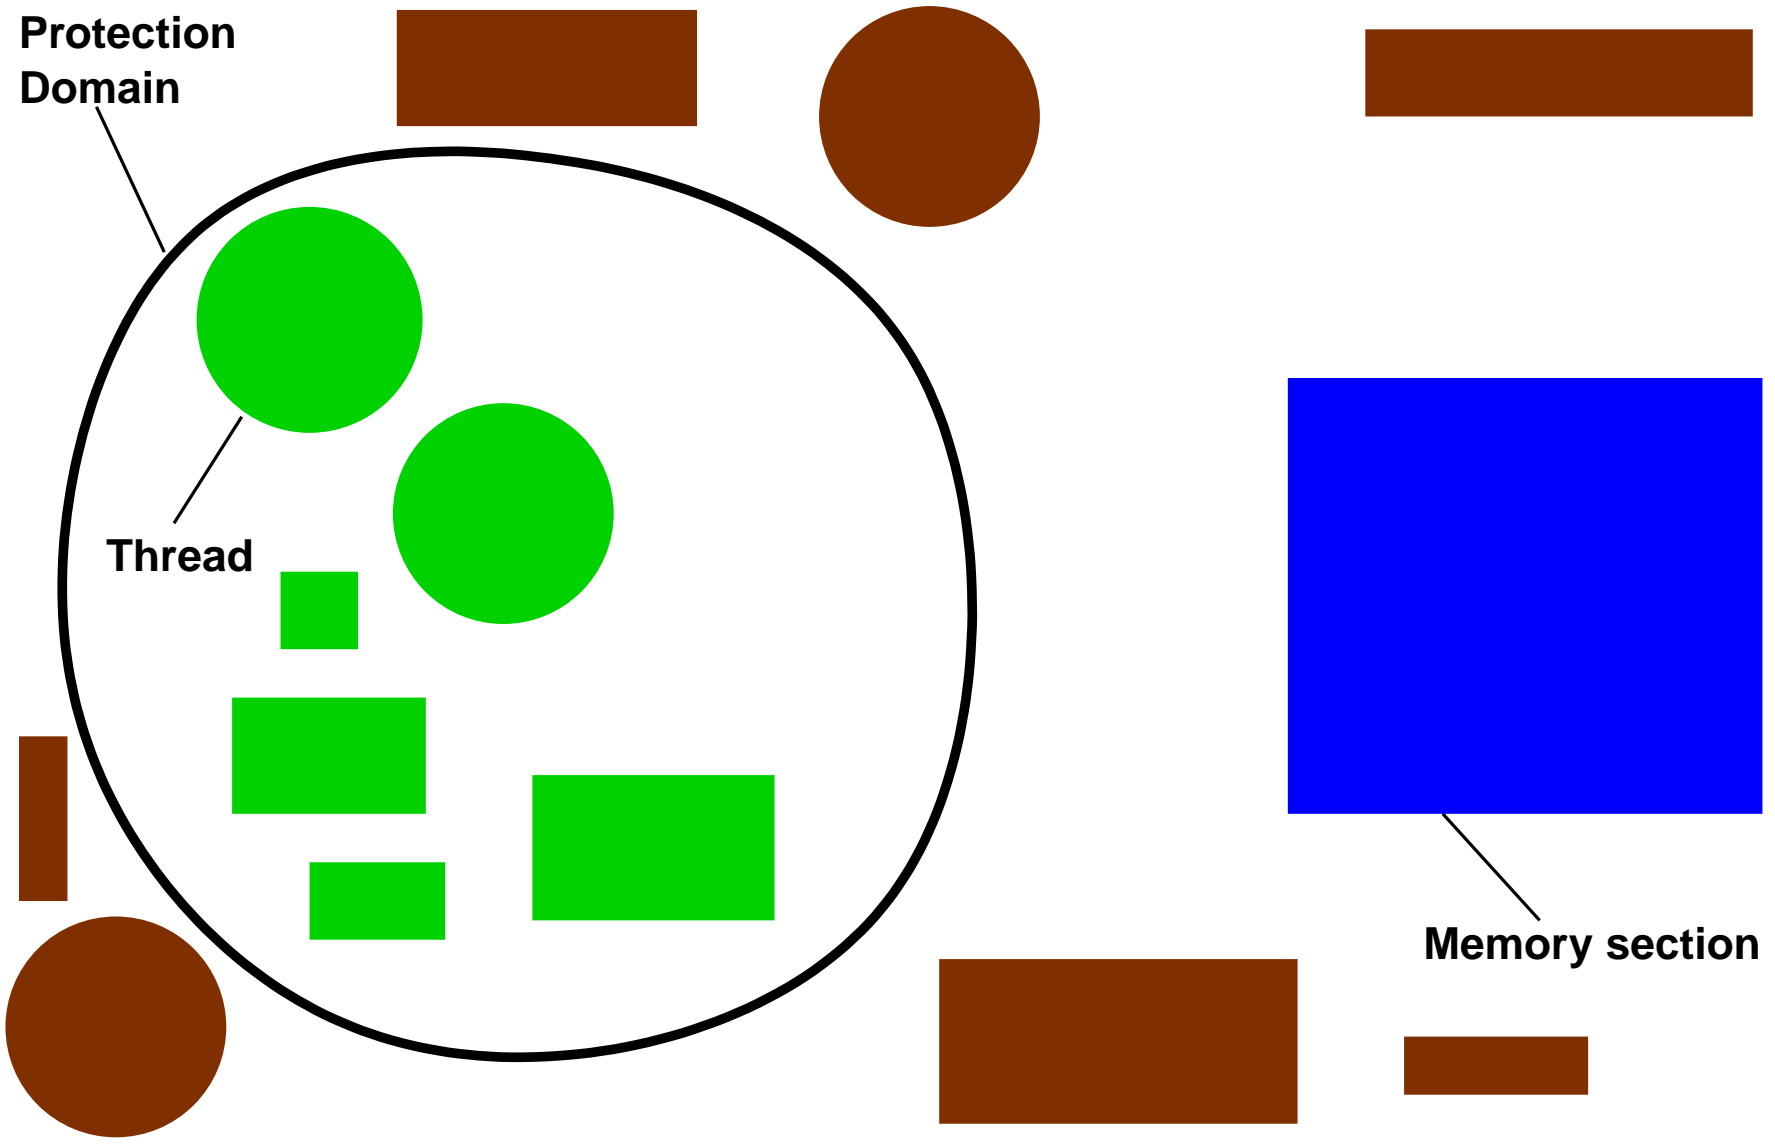
IGUANA CONCEPTS



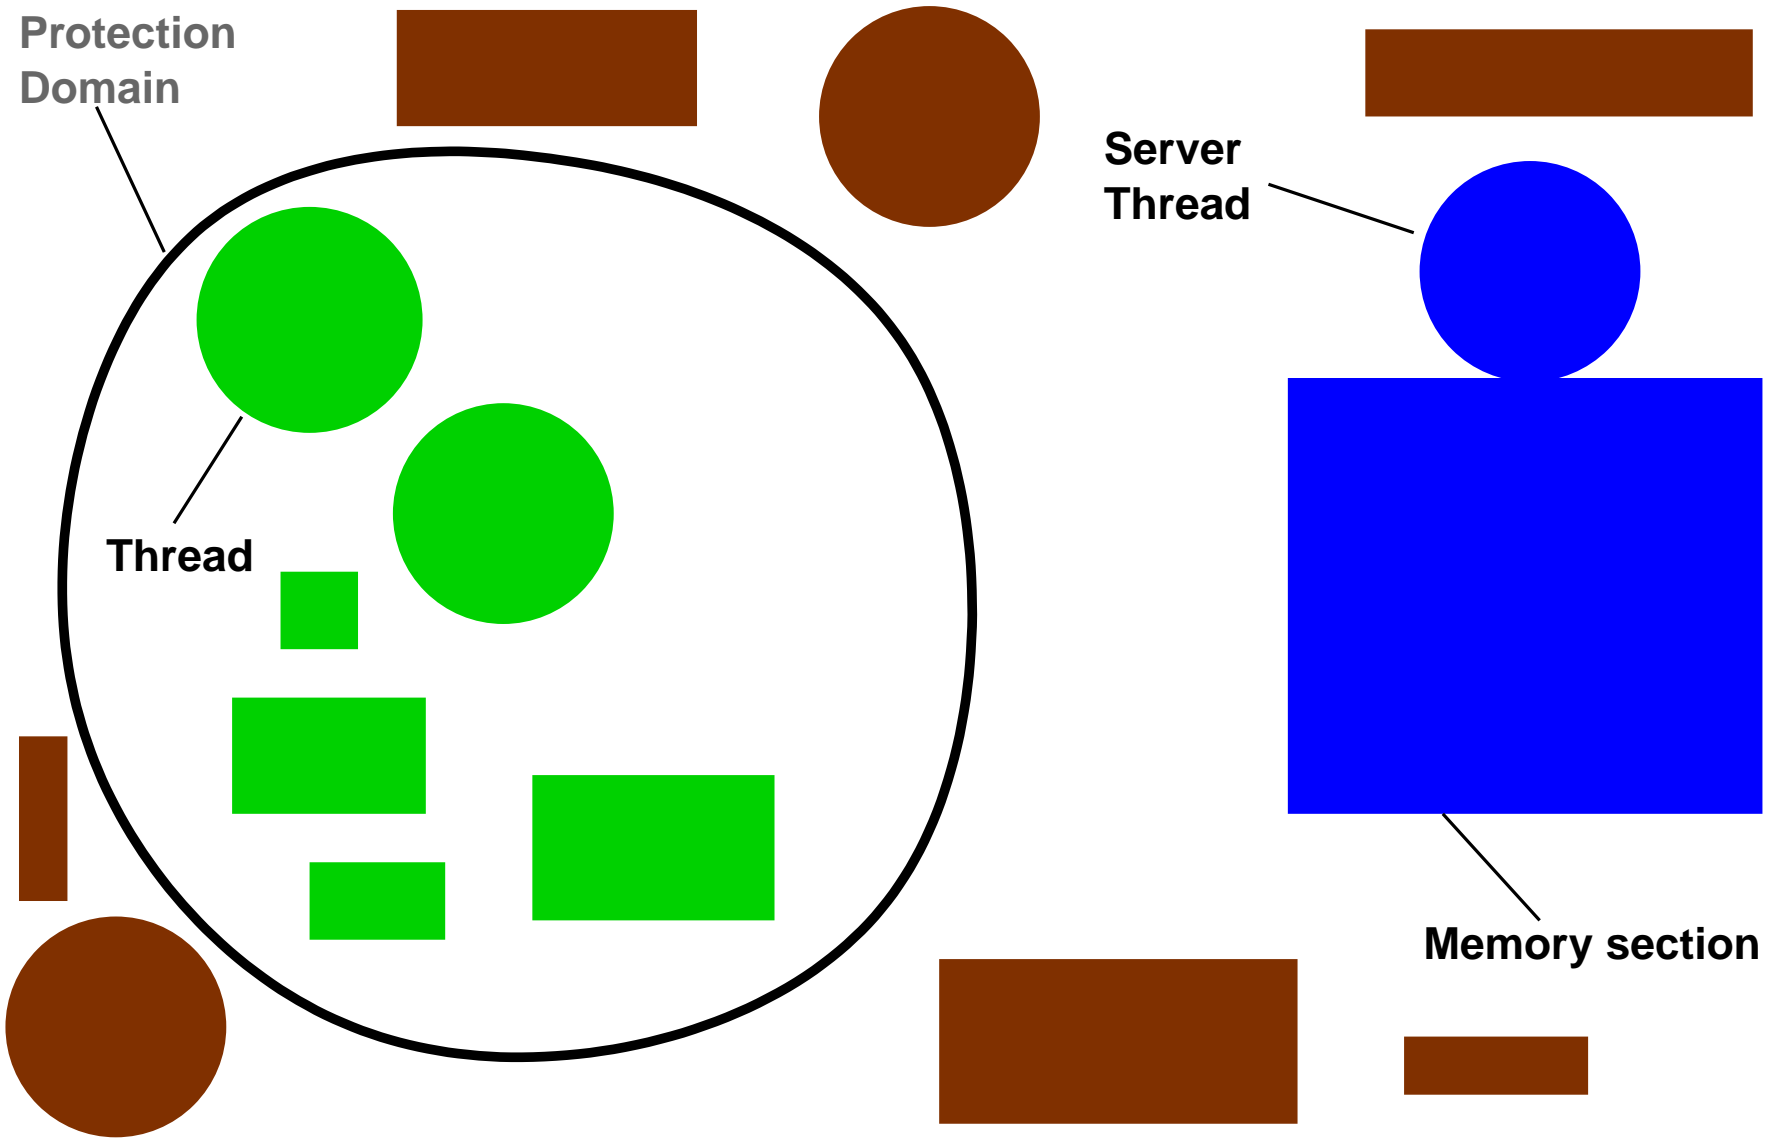
IGUANA CONCEPTS



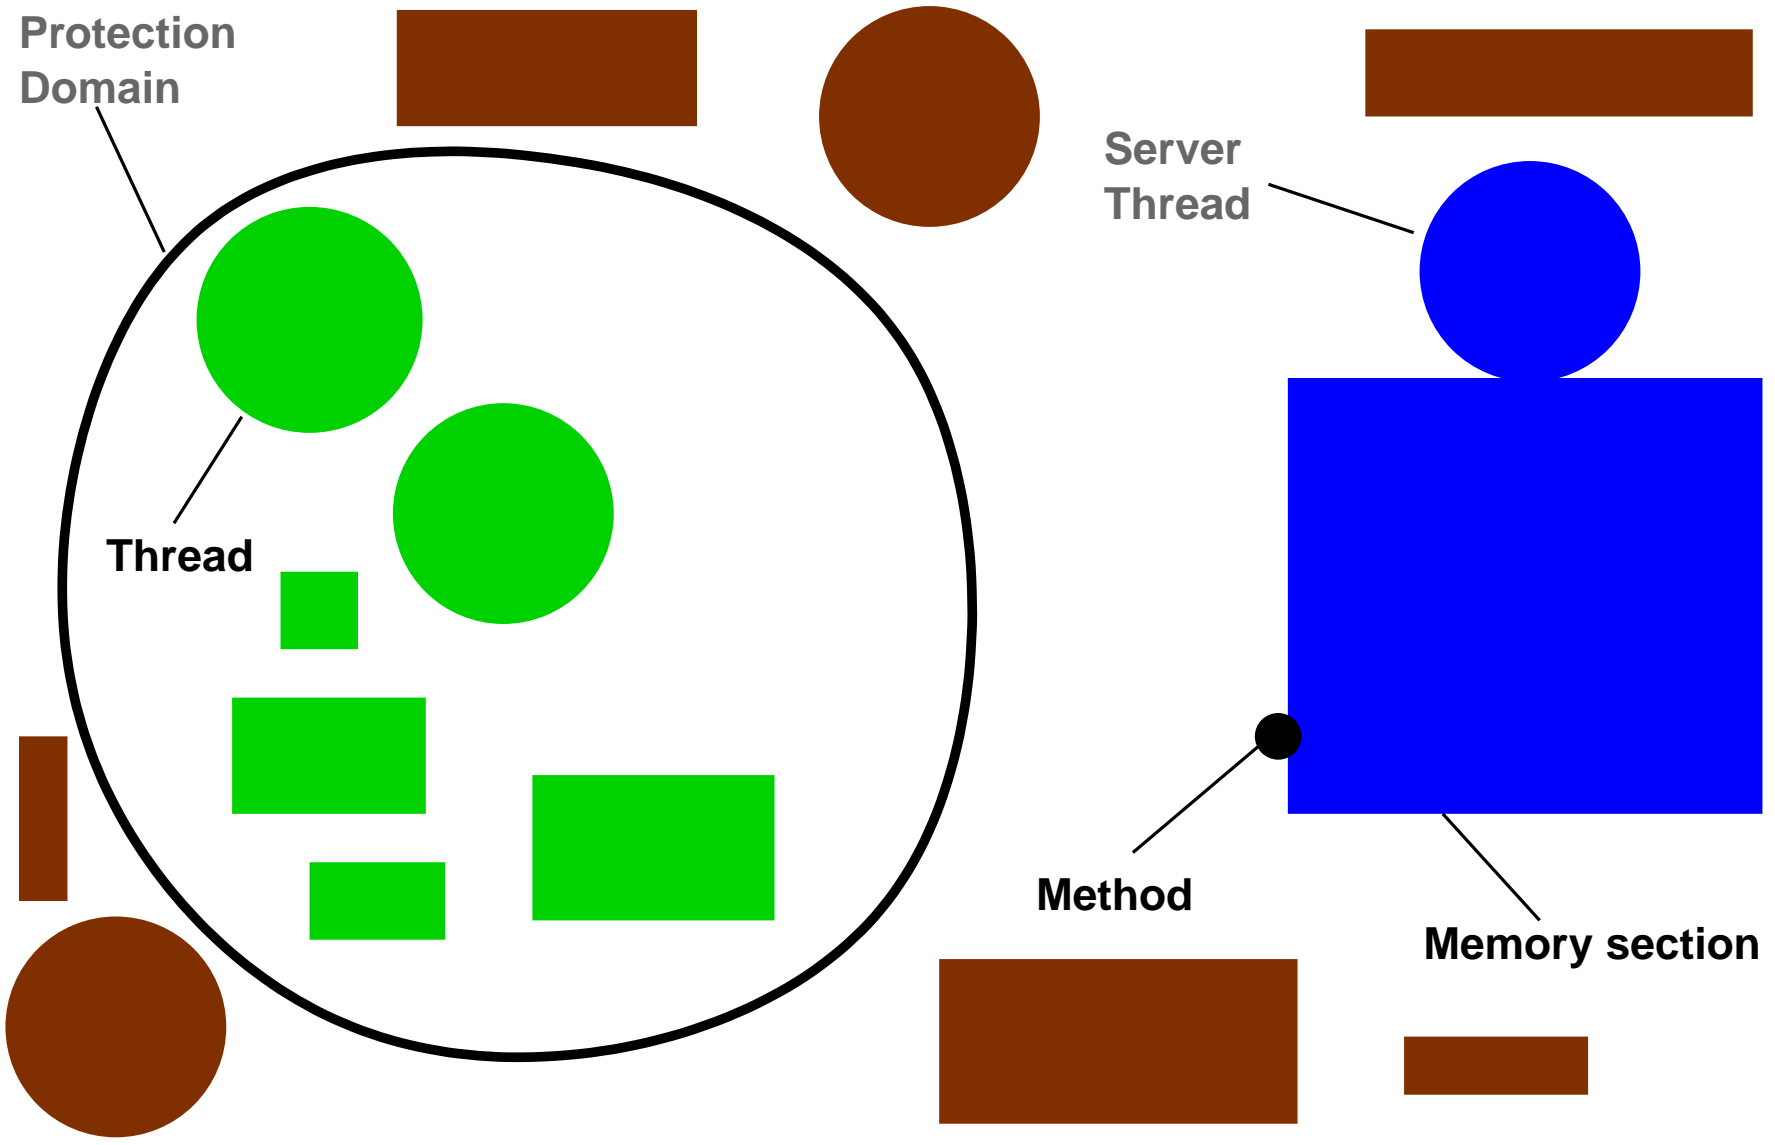
IGUANA CONCEPTS



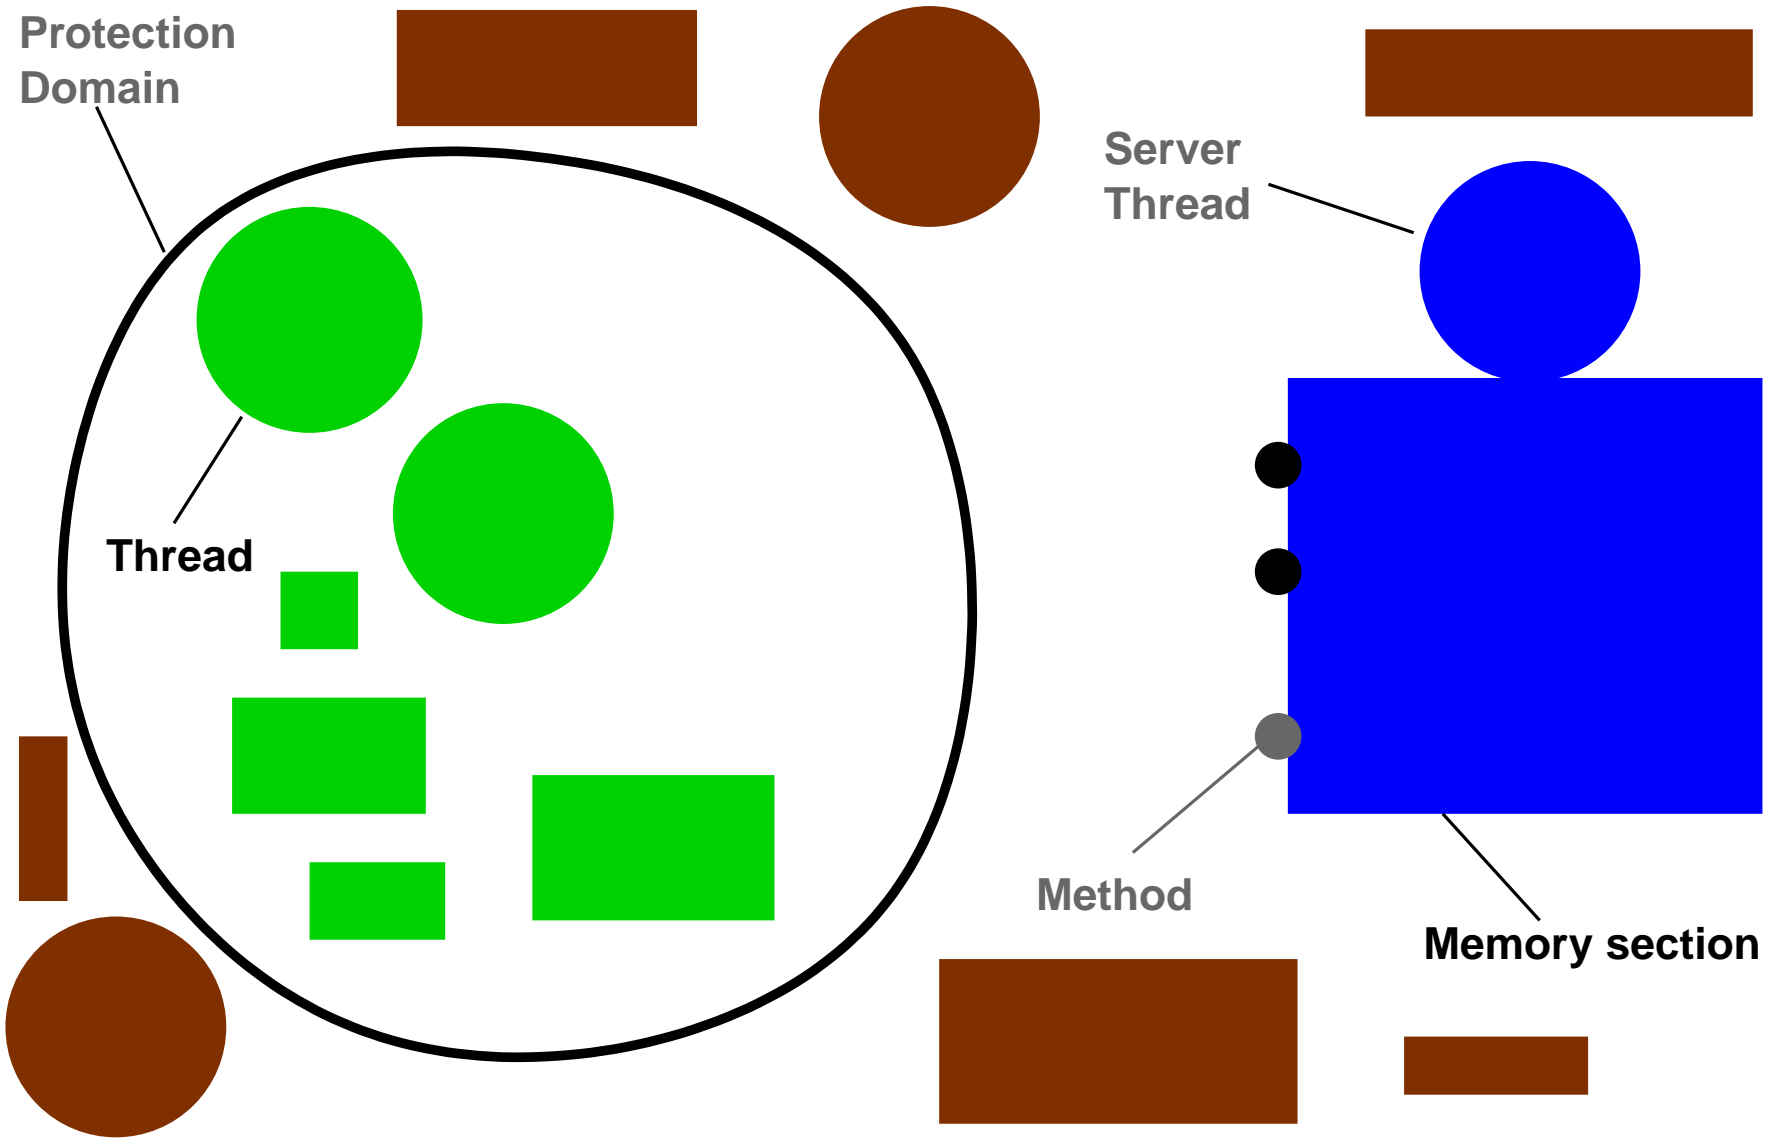
IGUANA CONCEPTS



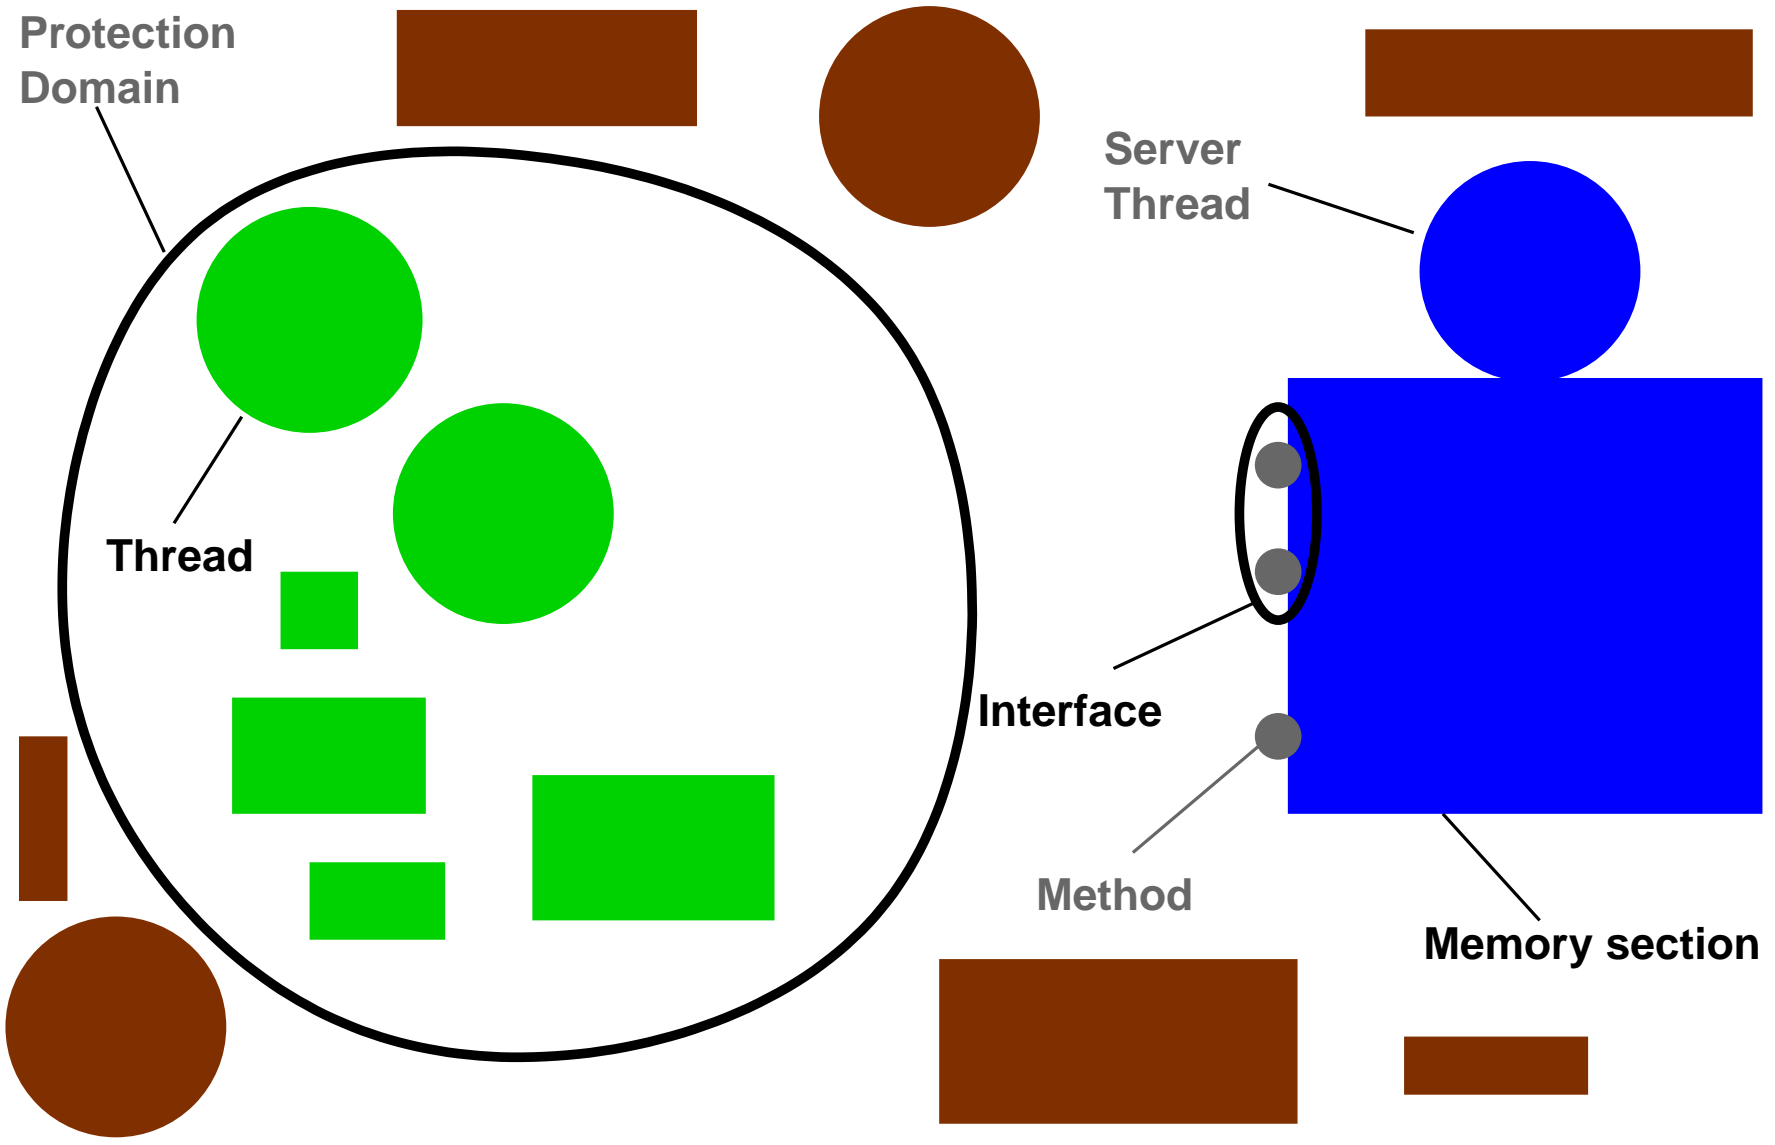
IGUANA CONCEPTS



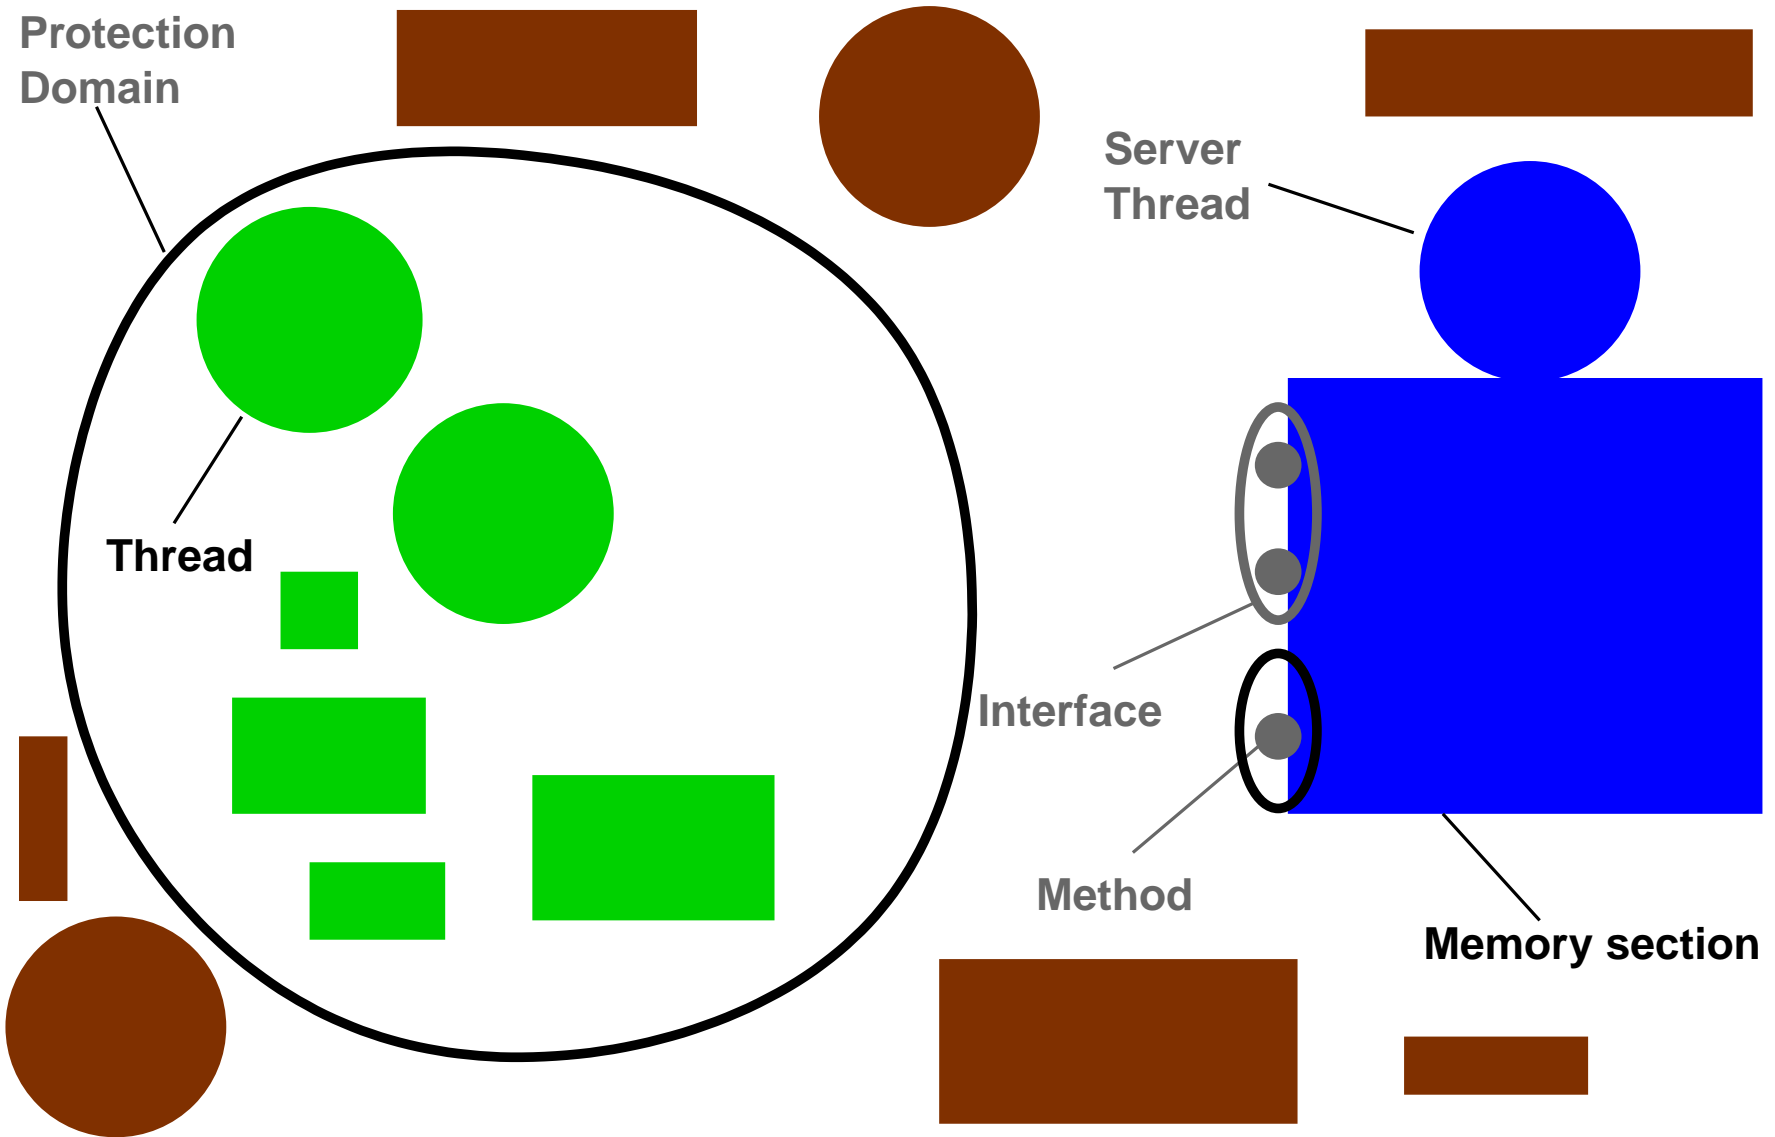
IGUANA CONCEPTS



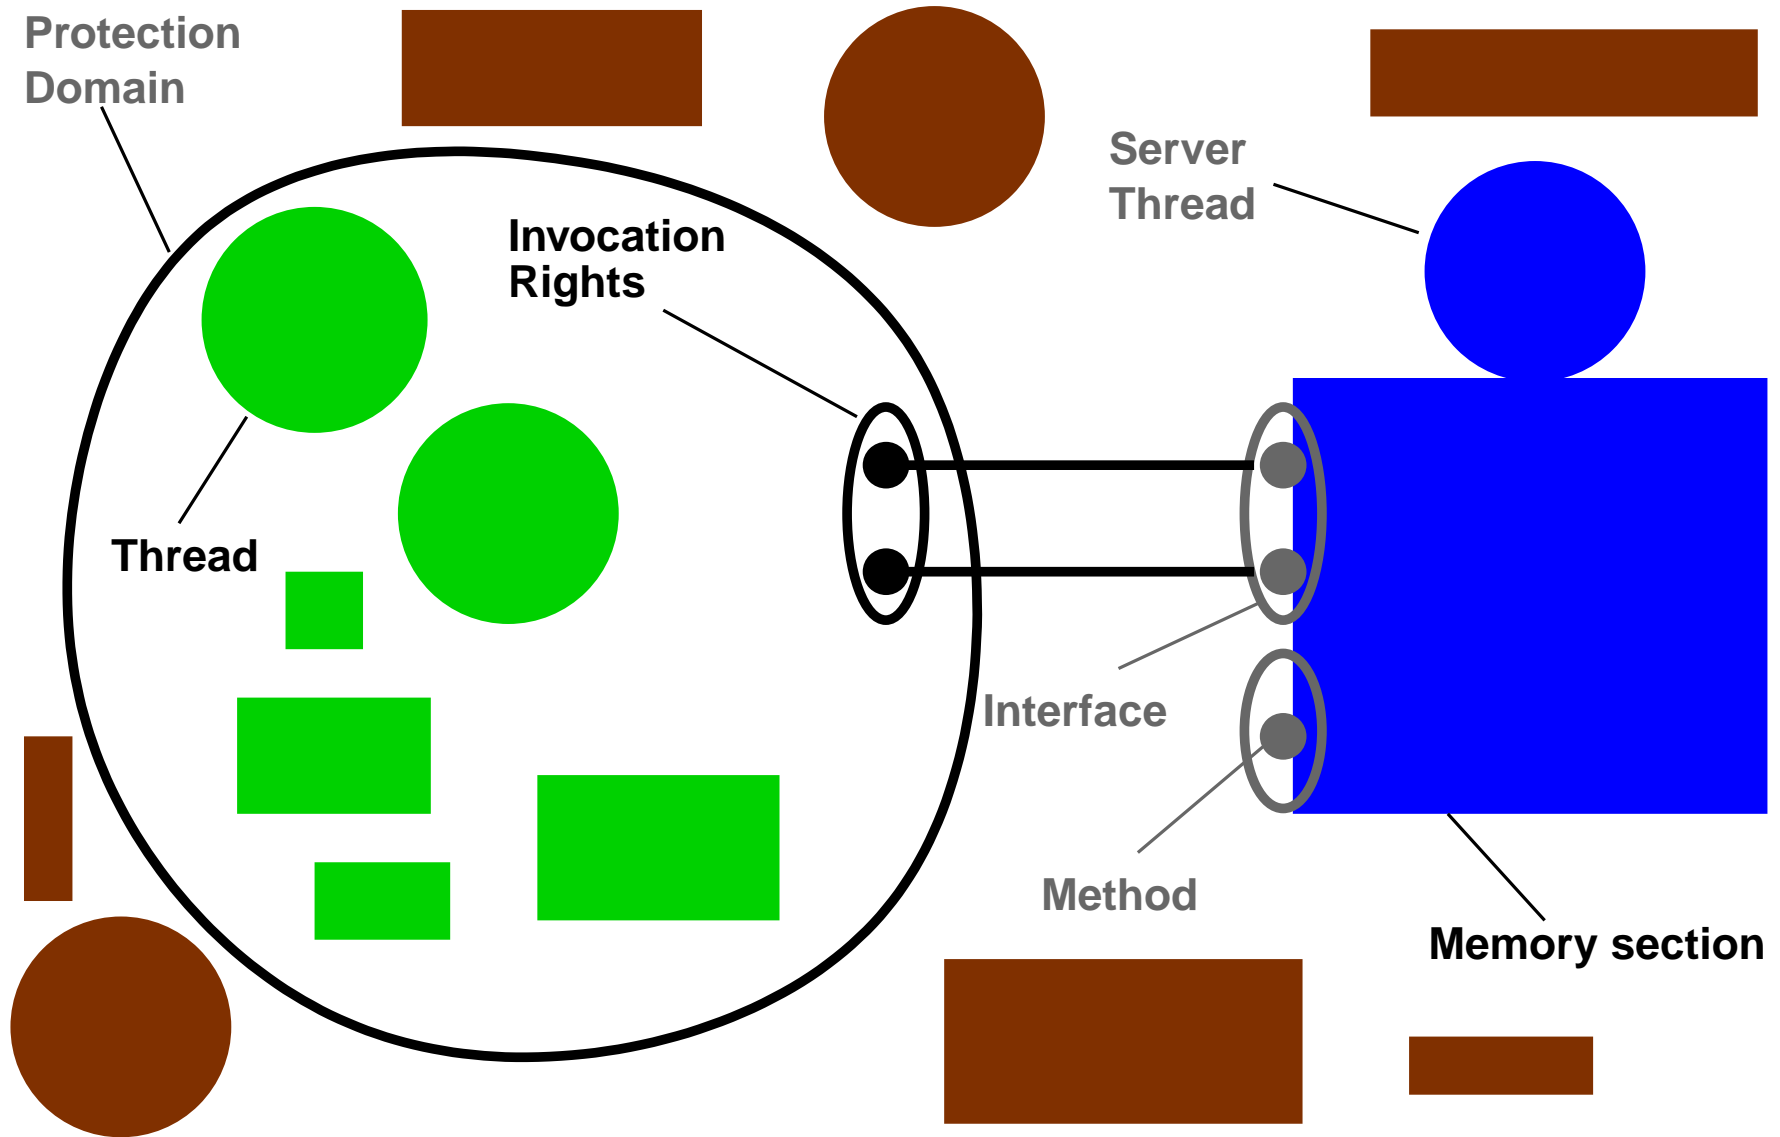
IGUANA CONCEPTS



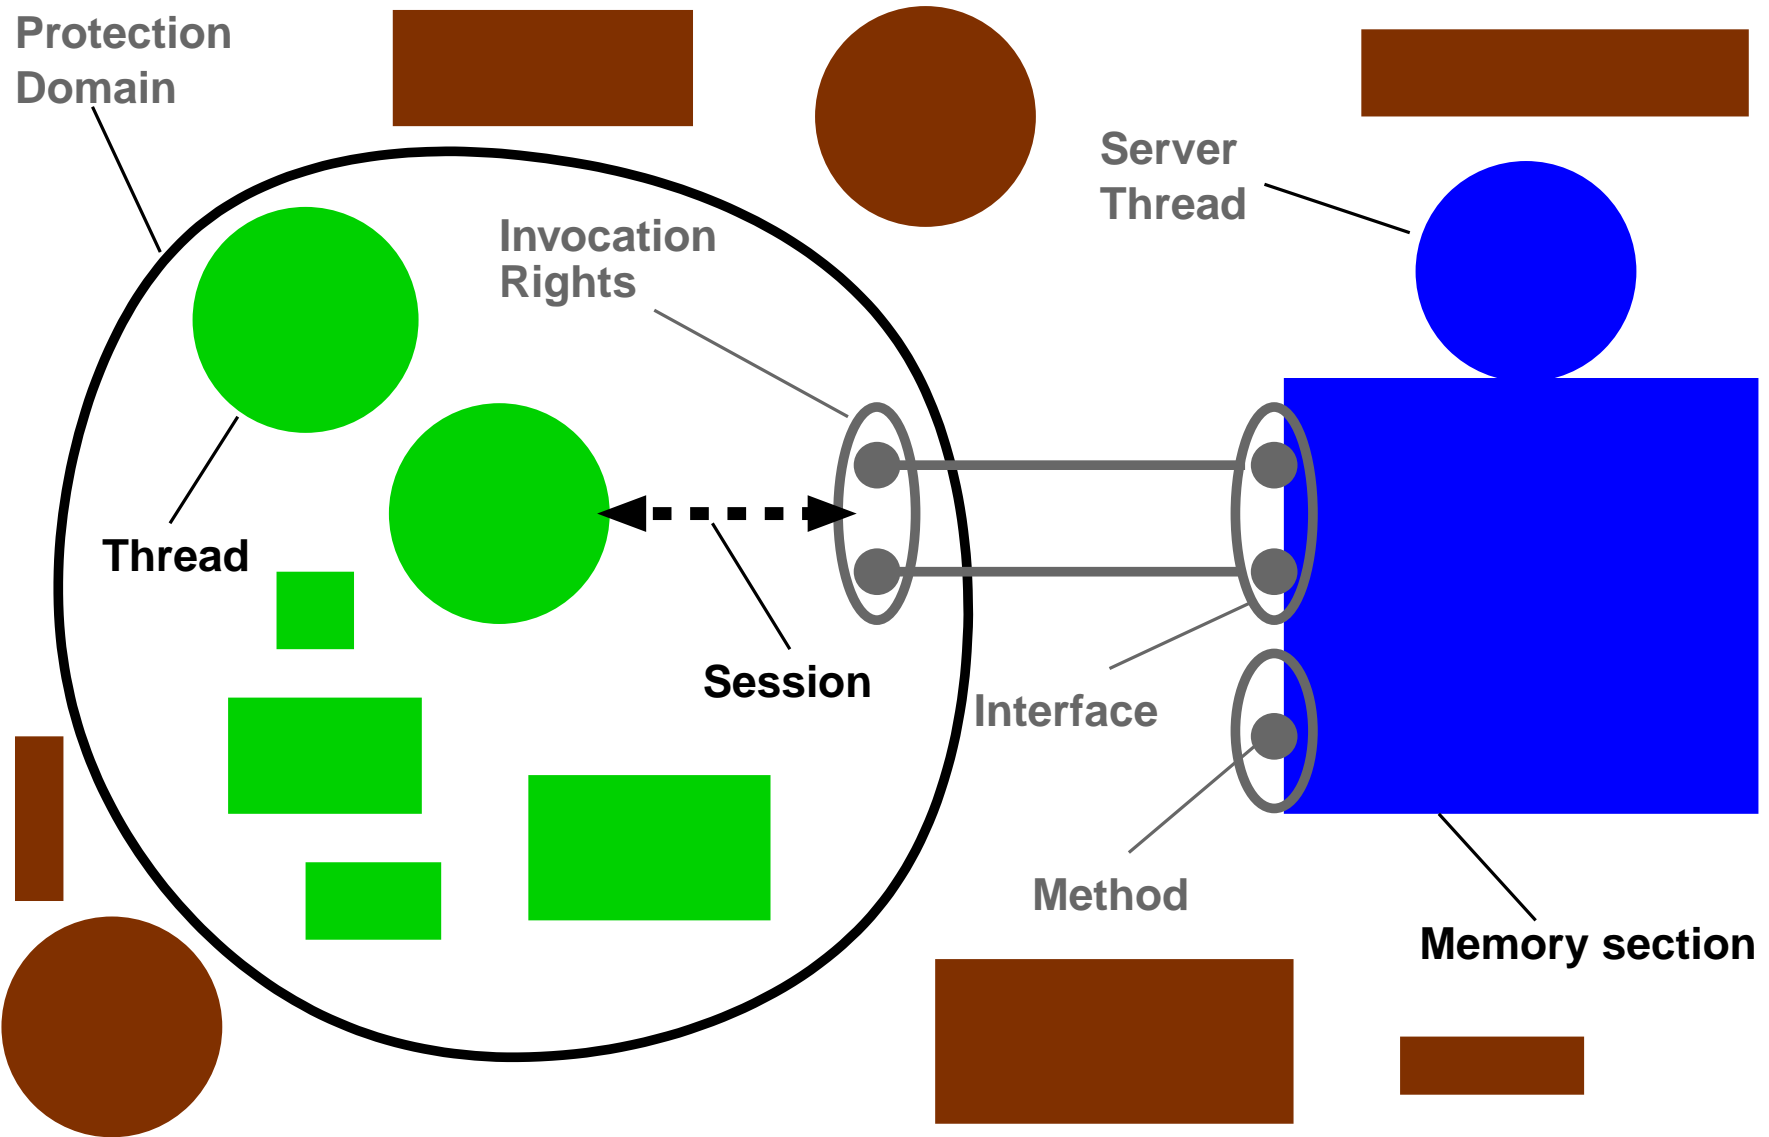
IGUANA CONCEPTS



IGUANA CONCEPTS



IGUANA CONCEPTS



IGUANA OBJECTS

- Six kinds of *objects*
 1. memory sections
 2. threads
 3. protection domains (PDs)
 4. sessions
 5. resource tokens (restoks)
 - not yet implemented, not covered here
 6. external spaces
 - not full Iguana objects
 - serve as proxies for non-Iguana objects

IGUANA OBJECTS

- Six kinds of *objects*
 1. memory sections
 2. threads
 3. protection domains (PDs)
 4. sessions
 5. resource tokens (restoks)
 - not yet implemented, not covered here
 6. external spaces
 - not full Iguana objects
 - serve as proxies for non-Iguana objects
- Access controlled by *capabilities*

OBJECTS: COMMONALITIES

- Objects have a unique name — *object ID* (OID)
 - OIDs are addresses in Iguana's SAS
 - only for memory sections does this address correspond to actual memory

OBJECTS: COMMONALITIES

- Objects have a unique name — *object ID* (OID)
 - OIDs are addresses in Iguana's SAS
 - only for memory sections does this address correspond to actual memory
- Objects have *methods* that can be invoked
 - ★ one method that exists for all objects: *destroy*
 - ★ each kind of object has a set of pre-defined methods

OBJECTS: COMMONALITIES

- Objects have a unique name — *object ID* (OID)
 - OIDs are addresses in Iguana's SAS
 - only for memory sections does this address correspond to actual memory
- Objects have *methods* that can be invoked
 - ★ one method that exists for all objects: *destroy*
 - ★ each kind of object has a set of pre-defined methods
- Objects are created by invoking constructor on a PD:
 - `kind_cap = pd->new_kind(args) ;`

OBJECTS: COMMONALITIES

- Objects have a unique name — *object ID* (OID)
 - OIDs are addresses in Iguana's SAS
 - only for memory sections does this address correspond to actual memory
- Objects have *methods* that can be invoked
 - ★ one method that exists for all objects: *destroy*
 - ★ each kind of object has a set of pre-defined methods
- Objects are created by invoking constructor on a PD:
 - `kind_cap = pd->new_kind(args) ;`
- Methods are grouped into *interfaces*
 - ★ interfaces also have unique IDs (IIDs) that are OID + interface number
 - ★ interfaces have capabilities
 - grant rights to invoke an interface's methods
 - ★ all pre-defined methods belong to separate interfaces
 - i.e., access is individually protected

DIFFERENCES TO MUNGI

- SAS name space extended to *all* objects
 - not just memory
 - supports unified access control and resource management

DIFFERENCES TO MUNGI

- SAS name space extended to *all* objects
 - not just memory
 - supports unified access control and resource management
- Secondary storage not part of SAS
 - meant to run on 32-bit hardware
 - meant to run on embedded systems without disk

DIFFERENCES TO MUNGI

- SAS name space extended to *all* objects
 - not just memory
 - supports unified access control and resource management
- Secondary storage not part of SAS
 - meant to run on 32-bit hardware
 - meant to run on embedded systems without disk
- Provision for management of all resources (i.e. objects)
 - only half worked out yet
 - one honours student working on details

DIFFERENCES TO MUNGI

- SAS name space extended to *all* objects
 - not just memory
 - supports unified access control and resource management
- Secondary storage not part of SAS
 - meant to run on 32-bit hardware
 - meant to run on embedded systems without disk
- Provision for management of all resources (i.e. objects)
 - only half worked out yet
 - one honours student working on details
- Client-server model instead of migrating-threads model
 - more familiar to potential users
 - avoids some (as yet) unresolved resource management issues

DIFFERENCES TO MUNGI

- SAS name space extended to *all* objects
 - not just memory
 - supports unified access control and resource management
- Secondary storage not part of SAS
 - meant to run on 32-bit hardware
 - meant to run on embedded systems without disk
- Provision for management of all resources (i.e. objects)
 - only half worked out yet
 - one honours student working on details
- Client-server model instead of migrating-threads model
 - more familiar to potential users
 - avoids some (as yet) unresolved resource management issues
- Aimed at commercial deployment ⇒ more conventional
 - Mungi is a pure research system
 - useful for trying out things

IGUANA CAPABILITIES

- A capability is a token that confers some access right(s)

IGUANA CAPABILITIES

- A capability is a token that confers some access right(s)
- Two kinds of capabilities in Iguana:
 - ★ *master capability*
 - created when an object is created
 - confers rights on all methods of object
 - allows creation of further capabilities
 - ★ *invocation capability*
 - created when an interface is created
 - confers right to invoke methods of a single interface

IGUANA CAPABILITIES

- A capability is a token that confers some access right(s)
- Two kinds of capabilities in Iguana:
 - ★ *master capability*
 - created when an object is created
 - confers rights on all methods of object
 - allows creation of further capabilities
 - ★ *invocation capability*
 - created when an interface is created
 - confers right to invoke methods of a single interface
- Capabilities are only active if stored in PD's *capability lists*
 - as in Mungi

MEMORY SECTIONS

- Memory sections represent virtual memory
 - ★ allocation of a certain amount of virtual memory:

```
mem_cap = pd->new_mem( size ) ;
```

MEMORY SECTIONS

- Memory sections represent virtual memory
 - ★ allocation of a certain amount of virtual memory:
$$mem_cap = pd \rightarrow new_mem(size) ;$$
- Memory sections are the only objects that support user-defined methods
 - others have pre-defined (standard) methods only
- Used to provide encapsulated services:
 - ★ service = memory (data) + server (thread) + methods

MEMORY SECTIONS...

- To create a service:

- ★ register a server thread on memory section

base → `new_server(thread_id) ;`

- *base* is the base address (OID) of the memory section

- ★ register interfaces (user-defined methods)

base = `iid` → `new_cap() ;`

- *iid* refers to number of new interface

MEMORY SECTIONS...

- To create a service:

- ★ register a server thread on memory section

 - `base->new_server(thread_id);`

 - `base` is the base address (OID) of the memory section

- ★ register interfaces (user-defined methods)

 - `base = iid->new_cap();`

 - `iid` refers to number of new interface

- Registering interfaces supports user-defined methods

- ★ remember: each interface can have one or more methods

 - interface number only interpreted by server

 - similarly, the method number is an opcode delivered to the interface

- ★ IIDs and method numbers allocated by system implementor

 - part of the service's interface protocol

MEMORY SECTIONS: PSEUDO METHODS

- Read (R), write (W), execute (X) are logically considered methods
 - ★ subjects them to same protection mechanisms as other methods
 - ★ no actual methods exist corresponding to those operations

MEMORY SECTIONS: PSEUDO METHODS

- Read (R), write (W), execute (X) are logically considered methods
 - ★ subjects them to same protection mechanisms as other methods
 - ★ no actual methods exist corresponding to those operations
- Further pseudo-method is *clist* (C)
 - ★ allows holder to add an object to the Clist array

THREADS

- Iguana threads are essentially L4 threads:
 - ★ threads within same PD operated on by plain L4 syscalls
 - correspond to local L4 threads (i.e., same L4 AS)
 - `ExchangeRegisters`, `IPC`

THREADS

- Iguana threads are essentially L4 threads:
 - ★ threads within same PD operated on by plain L4 syscalls
 - correspond to local L4 threads (i.e., same L4 AS)
 - `ExchangeRegisters`, `IPC`
 - ★ direct IPC to non-local threads is not allowed
 - use method invocations (corresponding to server thread)
 - presently not enforced by Iguana
 - requires enhancements to L4 (forthcoming API) to do efficiently
 - will provide attribute to ensure enforcement (at a cost)

THREADS

- Iguana threads are essentially L4 threads:
 - ★ threads within same PD operated on by plain L4 syscalls
 - correspond to local L4 threads (i.e., same L4 AS)
 - `ExchangeRegisters`, `IPC`
 - ★ direct IPC to non-local threads is not allowed
 - use method invocations (corresponding to server thread)
 - presently not enforced by Iguana
 - requires enhancements to L4 (forthcoming API) to do efficiently
 - will provide attribute to ensure enforcement (at a cost)
- Certain operations require privileges
 - ★ e.g. thread creation and deletion done by privileged L4 `ThreadControl()` call
- Done by Iguana on invocation of appropriate methods

THREAD OPERATIONS

- Thread creation:

```
thread_cap = pd->new_thread(&l4_tid) ;
```

- ★ returns two kinds of thread IDs

- * Iguana thread ID (*tid*), part of the *thread_cap*

- used for protection and other Iguana-specific purposes

- * L4 thread ID (*l4_tid*)

- used for L4 syscalls

THREAD OPERATIONS

- Thread creation:

```
thread_cap = pd->new_thread(&l4_tid) ;
```

- ★ returns two kinds of thread IDs

- * Iguana thread ID (*tid*), part of the *thread_cap*

- used for protection and other Iguana-specific purposes

- * L4 thread ID (*l4_tid*)

- used for L4 syscalls

- New thread created *inactive*

- ★ can be activated by:

- L4 syscall `ExchangeRegisters()` (local threads only)

- Iguana method `tid->start(ip, sp)`

THREAD OPERATIONS...

- Obtain L4 thread ID

 - `l4tid = tid->l4_tid();`

- Obtain own thread ID

 - `tid = myself();`

- Obtain protection domain of thread

 - `pd = tid->domain();`

- Obtain and modify scheduling parameters

 - `tid->schedule_info(&info);`

SESSIONS

- Sessions reduce authentication overheads of repeated calls

SESSIONS

- Sessions reduce authentication overheads of repeated calls
- Prior to invoking methods on a service, must establish session

```
session = pd->new_session(server);
```

- ★ establishes session between target PD and server

SESSIONS

- Sessions reduce authentication overheads of repeated calls
- Prior to invoking methods on a service, must establish session

```
session = pd->new_session(server);
```

★ establishes session between target PD and server

★ `server` is a PD ID

→ **Note:** This is likely to change

SESSIONS

- Sessions reduce authentication overheads of repeated calls

- Prior to invoking methods on a service, must establish session

```
session = pd->new_session(server);
```

- ★ establishes session between target PD and server

- ★ server is a PD ID

→ **Note:** This is likely to change

- ★ Iguana informs the server by invoking its notification method

```
server->session_created(pd);
```

- ★ Iguana notifies remaining partners if the session is destroyed

```
pd_or_server->session_destroyed(session);
```

IGUANA CAPABILITIES

- Iguana capabilities are user-level objects

→ *password capabilities*, consisting of OID and password



→ Length of password is configurable (normally \geq 64 bits)

IGUANA CAPABILITIES

- Iguana capabilities are user-level objects

→ *password capabilities*, consisting of OID and password



→ Length of password is configurable (normally ≥ 64 bits)

- Same model as in Mungi

→ implicit presentation

→ two-level structure (Clists caps array, Clists)

→ same confinement approach

EXTERNAL SPACES

- External spaces are “raw” L4 address spaces
 - not part of Iguana SAS
- Provided to deal with restrictions of Iguana model
 - 32-bit address space may not be large enough to share between all protection domains
 - legacy support (e.g. strict `fork()` semantics) may require separate address spaces

EXTERNAL SPACES

- External spaces are “raw” L4 address spaces
 - not part of Iguana SAS
- Provided to deal with restrictions of Iguana model
 - 32-bit address space may not be large enough to share between all protection domains
 - legacy support (e.g. strict `fork()` semantics) may require separate address spaces
- External spaces come at a cost
 - ★ unable to make full use of fast address-space switching on ARM
 - ★ not well integrated with Iguana world
 - no fine-grained access control provided by Iguana capabilities
 - not allowed to communicate with any PD other than creator
 - not even with Iguana — cannot invoke methods
 - enforced via L4 redirectors

EXTERNAL SPACES — OPERATIONS

- Creation requires explicit specification of KIP and UTCB address

```
es = pd->new_es (kip, utcb_area);
```

- Thread creation also requires arguments similar to L4

```
l4tid = es->new_thread(pager, scheduler, starter, utcb);
```

HARDWARE ACCESS

- Device drivers need to access raw hardware features
- Iguana provides a (static) `hardware` object for this

- ★ physical memory access:

```
hardware->back_mem(adr, p_adr, caching);
```

- maps the memory section (`adr`) to the specified physical address with specified caching attributes

- ★ interrupt association:

```
hardware->register_interrupt(tid, irq);
```

- registers the specified thread as the handler of the specified interrupt

RESOURCE TOKENS

- Iguana's resource management mechanism
- Note: presently this only exists conceptually
 - details of the model still need to be worked out
 - however, model is based on our experience with a similar model in Mungi

RESOURCE TOKENS

- Iguana's resource management mechanism
- Note: presently this only exists conceptually
 - details of the model still need to be worked out
 - however, model is based on our experience with a similar model in Mungi
- Basic idea: all resources have a price that must be paid by the user
- Model provides great flexibility for defining charging details
- Generalisation of Mungi bank accounts

WOMBAT: A PORTABLE LINUX SERVER ON L4

MOTIVATION

- Provide Linux API
- Support architectures other than x86
- Be easily maintainable
- Integrate with Iguana

LEGACY APIS ON L4

- Microkernel, like L4, provides mechanisms, not services
 - services implemented by user-level server
 - microkernel is OS agnostic
 - can, in principle, provide any OS API (OS “personality”)
 - can provide multiple OS APIs

LEGACY APIS ON L4

- Microkernel, like L4, provides mechanisms, not services
 - services implemented by user-level server
 - microkernel is OS agnostic
 - can, in principle, provide any OS API (OS “personality”)
 - can provide multiple OS APIs
- Multiple OS APIs are a classical *motivation* for microkernels

LEGACY APIS ON L4

- Microkernel, like L4, provides mechanisms, not services
 - services implemented by user-level server
 - microkernel is OS agnostic
 - can, in principle, provide any OS API (OS “personality”)
 - can provide multiple OS APIs
- Multiple OS APIs are a classical *motivation* for microkernels
- Simplest approach is to port an existing monolithic kernel
 - e.g., Unix, Linux
 - even several running concurrently

LEGACY APIS ON L4

- Microkernel, like L4, provides mechanisms, not services
 - services implemented by user-level server
 - microkernel is OS agnostic
 - can, in principle, provide any OS API (OS “personality”)
 - can provide multiple OS APIs
- Multiple OS APIs are a classical *motivation* for microkernels
- Simplest approach is to port an existing monolithic kernel
 - e.g., Unix, Linux
 - even several running concurrently
- Past experience poor:
 - Mach Unix server
 - IBM Workplace OS (Mach-based)

LEGACY APIS ON L4

- Microkernel, like L4, provides mechanisms, not services
 - services implemented by user-level server
 - microkernel is OS agnostic
 - can, in principle, provide any OS API (OS “personality”)
 - can provide multiple OS APIs
- Multiple OS APIs are a classical *motivation* for microkernels
- Simplest approach is to port an existing monolithic kernel
 - e.g., Unix, Linux
 - even several running concurrently
- Past experience poor:
 - Mach Unix server
 - IBM Workplace OS (Mach-based)
- **Attributed to poor Mach IPC performance, large cache footprint**

PORTING MONOLITHIC KERNELS: PROS

- Code reuse: Existing services
- Code reuse: Legacy support
- Coder reuse
- Buzzword reuse

PORTING MONOLITHIC KERNELS: PROS

- Code reuse: Existing services
 - can use existing code unchanged (file systems, network stacks, drivers, ...)
 - dramatic reduction of cost of building new system
 - potential reduction in maintenance cost (but...)
- Code reuse: Legacy support
- Coder reuse
- Buzzword reuse

PORTING MONOLITHIC KERNELS: PROS

- Code reuse: Existing services
 - can use existing code unchanged (file systems, network stacks, drivers, ...)
 - dramatic reduction of cost of building new system
 - potential reduction in maintenance cost (but...)
- Code reuse: Legacy support
 - can support huge number of existing apps (if sufficiently compatible)
 - dramatic reduction of cost of deploying new system
- Coder reuse
- Buzzword reuse

PORTING MONOLITHIC KERNELS: PROS

- Code reuse: Existing services
 - can use existing code unchanged (file systems, network stacks, drivers, ...)
 - dramatic reduction of cost of building new system
 - potential reduction in maintenance cost (but...)
- Code reuse: Legacy support
 - can support huge number of existing apps (if sufficiently compatible)
 - dramatic reduction of cost of deploying new system
- Coder reuse
 - can provide a familiar target for programmers
 - reduction of cost of maintaining system and new app development
- Buzzword reuse

PORTING MONOLITHIC KERNELS: PROS

- Code reuse: Existing services
 - can use existing code unchanged (file systems, network stacks, drivers, ...)
 - dramatic reduction of cost of building new system
 - potential reduction in maintenance cost (but...)
- Code reuse: Legacy support
 - can support huge number of existing apps (if sufficiently compatible)
 - dramatic reduction of cost of deploying new system
- Coder reuse
 - can provide a familiar target for programmers
 - reduction of cost of maintaining system and new app development
- Buzzword reuse
 - can jump on Linux bandwagon while really doing something more sensible...
 - good PR value ;-)

PORTING MONOLITHIC KERNELS: CONS

- Performance

- ★ Microkernel inherently adds overhead

- microkernel-mediated rather than direct access to hardware

- IPC to server (4 mode switches) rather than syscall (2 mode switches)

⇒ Legacy system on L4 will **always** be slower than native!

PORTING MONOLITHIC KERNELS: CONS

- Performance

- ★ Microkernel inherently adds overhead

- microkernel-mediated rather than direct access to hardware

- IPC to server (4 mode switches) rather than syscall (2 mode switches)

⇒ Legacy system on L4 will **always** be slower than native!

- ★ Question is: *How much slower???*

PORTING MONOLITHIC KERNELS: CONS

- Miss out on most advantages of microkernels

PORTING MONOLITHIC KERNELS: CONS

- Miss out on most advantages of microkernels
 - ★ software-engineering advantages
 - microkernels can provide strongly encapsulated components with small interfaces
 - Linux approach is unmaintainable in the long run [Schach *et al*, Maintainability of the Linux Kernel, IEE Proceedings Software, **149** 18–23 (2002)]

PORTING MONOLITHIC KERNELS: CONS

- Miss out on most advantages of microkernels
 - ★ software-engineering advantages
 - microkernels can provide strongly encapsulated components with small interfaces
 - Linux approach is unmaintainable in the long run [Schach *et al*, Maintainability of the Linux Kernel, IEE Proceedings Software, **149** 18–23 (2002)]
 - ★ reduction of *trusted computing base*
 - trusting 2–3Mloc vs 10kloc kernel plus core services and drivers

PORTING MONOLITHIC KERNELS: CONS

- Miss out on most advantages of microkernels
 - ★ software-engineering advantages
 - microkernels can provide strongly encapsulated components with small interfaces
 - Linux approach is unmaintainable in the long run [Schach *et al*, Maintainability of the Linux Kernel, IEE Proceedings Software, **149** 18–23 (2002)]
 - ★ reduction of *trusted computing base*
 - trusting 2–3Mloc vs 10kloc kernel plus core services and drivers
 - ★ potential for formal verification of TCB?
 - might be feasible if it's small
 - definitely infeasible with something the size of Linux

PORTING MONOLITHIC KERNELS: CONS

- Miss out on most advantages of microkernels
 - ★ software-engineering advantages
 - microkernels can provide strongly encapsulated components with small interfaces
 - Linux approach is unmaintainable in the long run [Schach *et al*, Maintainability of the Linux Kernel, IEE Proceedings Software, **149** 18–23 (2002)]
 - ★ reduction of *trusted computing base*
 - trusting 2–3Mloc vs 10kloc kernel plus core services and drivers
 - ★ potential for formal verification of TCB?
 - might be feasible if it's small
 - definitely infeasible with something the size of Linux
- Main-stream systems have notoriously poor real-time support
 - typically unsuitable for hard real time

HYBRID SYSTEMS: BEST OF TWO WORLDS?

- Legacy server and native apps running side-by-side
 - code/coder/buzzword reuse
 - software-engineering advantages for critical/new parts
 - hard real-time apps not affected by legacy system

HYBRID SYSTEMS: BEST OF TWO WORLDS?

- Legacy server and native apps running side-by-side
 - code/coder/buzzword reuse
 - software-engineering advantages for critical/new parts
 - hard real-time apps not affected by legacy system
- How about performance?
 - L4 primitives much faster than Mach and others
 - fast enough?

L4 ON LINUX: HISTORY

L⁴LINUX

- First port of Linux kernel to L4
- Done for ix86 at Dresden [SOSP 97]
 - ★ modified architecture-specific part of Linux to use L4 syscalls
 - ★ binary compatibility: syscall redirection to Linux server
 - syscall exception delivered to L⁴Linux via exception IPC
 - ★ better performance using modified libc
 - syscall trap replaced by L4 IPC to Linux server
 - ★ re-done for 2.2, 2.4, 2.6
 - L⁴Linux implementation highly dependent on Linux internals
 - significant effort for forward porting (start from scratch each time)

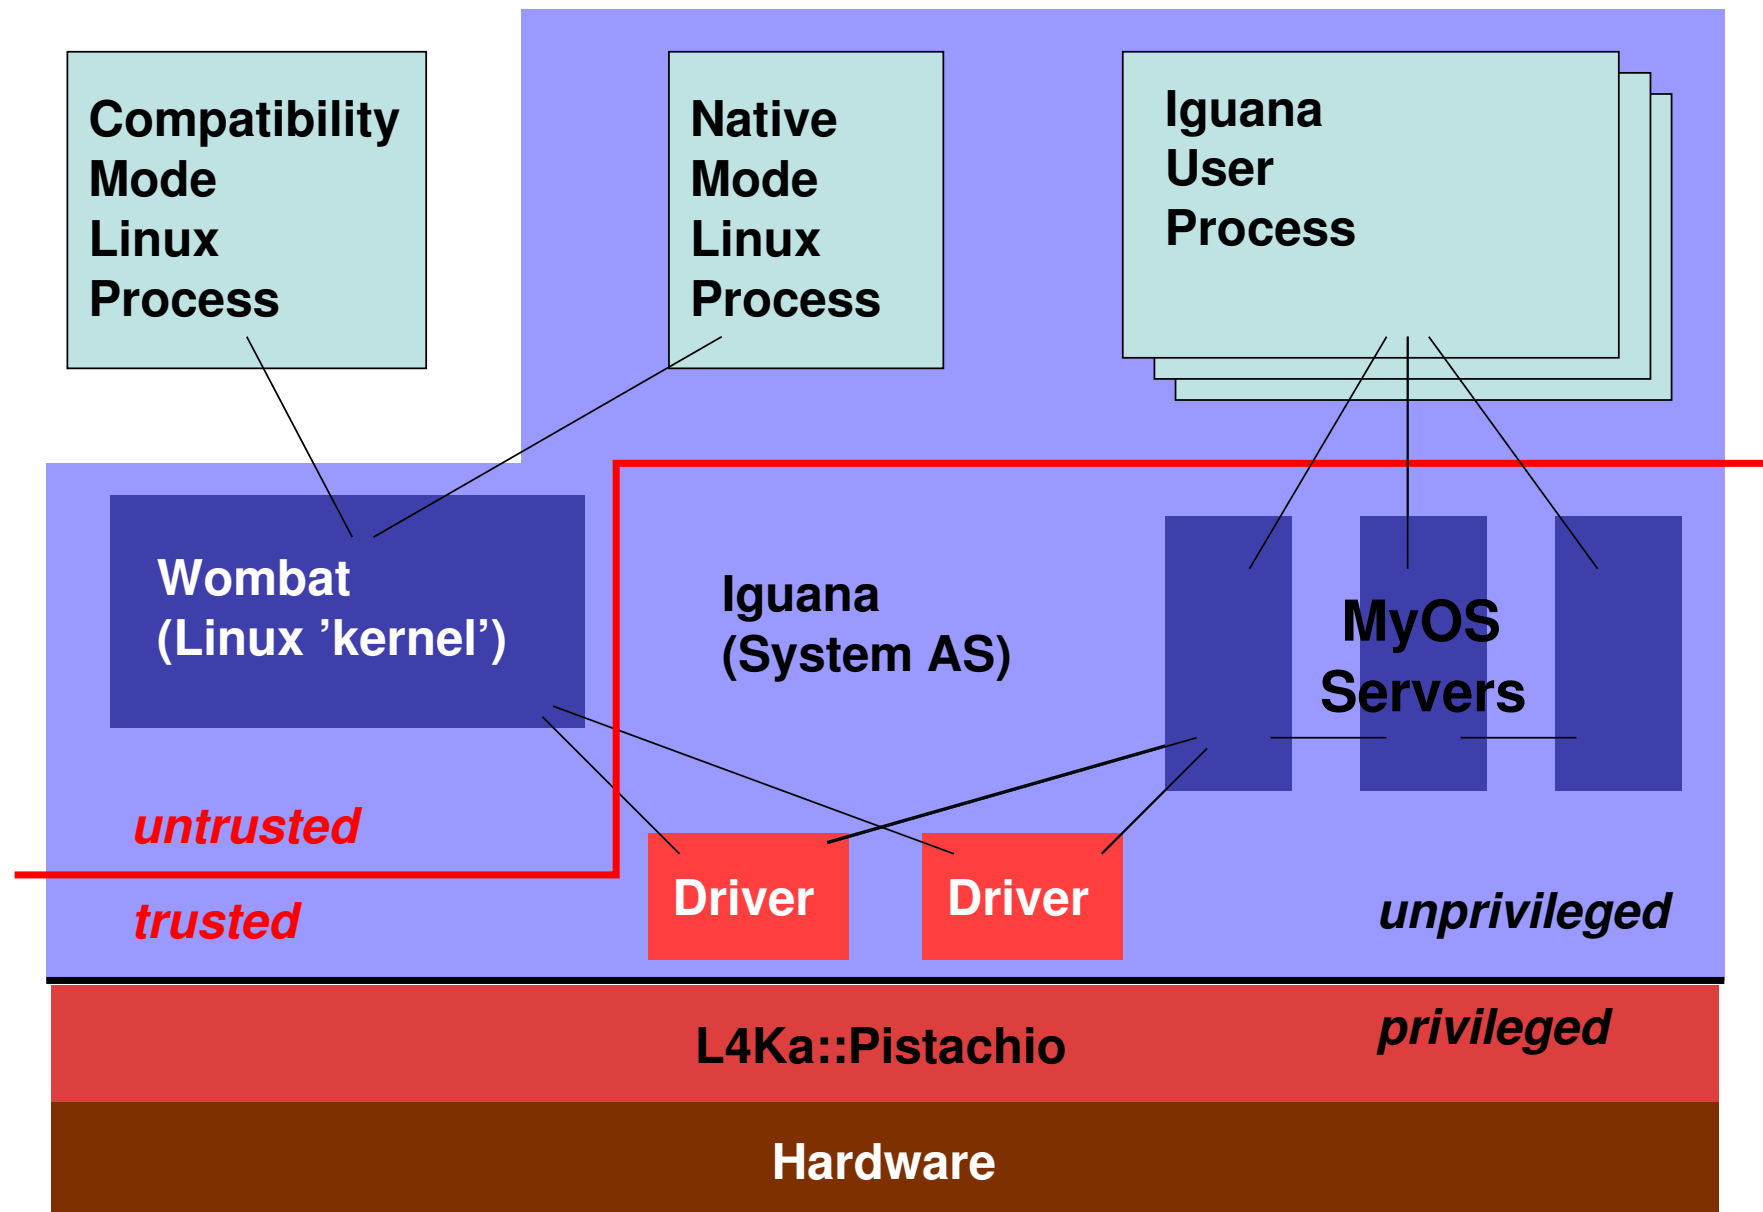
WOMBAT LINUX SERVER

- Done at NICTA 03–04 (not yet released)
 - done from scratch, independent of Dresden work
- Portable between architectures
 - based on Iguana
 - presently runs on x86, ARM, MIPS64
 - MIPS64 supports 32-bit and 64-bit executables

WOMBAT LINUX SERVER

- Done at NICTA 03–04 (not yet released)
 - done from scratch, independent of Dresden work
- Portable between architectures
 - based on Iguana
 - presently runs on x86, ARM, MIPS64
 - MIPS64 supports 32-bit and 64-bit executables
- Much reduced dependence on Linux internals
- Binary compatibility via syscall redirection
- Proper Linux scheduling of Linux apps
- Potential to achieve good performance on ARM

A SYSTEM RUNNING WOMBAT



WOMBAT IMPLEMENTATION

- Dresden L⁴Linux approach:
 - ★ applied modifications to architecture-dependent part of kernel
 - very x86-specific

WOMBAT IMPLEMENTATION

- Dresden L⁴Linux approach:
 - ★ applied modifications to architecture-dependent part of kernel
 - very x86-specific
- Wombat approach:
 - ★ introduced new `arch/14/` in Linux source
 - kept as architecture-neutral as possible
 - architecture-dependent code in `arch/14/sys-arm` etc

WOMBAT IMPLEMENTATION

- Dresden L⁴Linux approach:
 - ★ applied modifications to architecture-dependent part of kernel
 - very x86-specific
- Wombat approach:
 - ★ introduced new `arch/14/` in Linux source
 - kept as architecture-neutral as possible
 - architecture-dependent code in `arch/14/sys-arm` etc
 - ★ otherwise only modified 3 Linux source files
 - two are one-line bug fixes (patches submitted)
 - third is `printk` (additional early debug output)

WOMBAT IMPLEMENTATION

- Dresden L⁴Linux approach:
 - ★ applied modifications to architecture-dependent part of kernel
 - very x86-specific
- Wombat approach:
 - ★ introduced new `arch/14/` in Linux source
 - kept as architecture-neutral as possible
 - architecture-dependent code in `arch/14/sys-arm` etc
 - ★ otherwise only modified 3 Linux source files
 - two are one-line bug fixes (patches submitted)
 - third is `printk` (additional early debug output)
 - ★ tracked head revision of Linux from 2.5.x to 2.6.4
 - now at 2.6.6, will do 2.9 within a week

PAGE FAULTS AND EXCEPTIONS

- Wombat has a single server thread per CPU, like L⁴Linux
 - ★ is L4 pager and exception handler for Linux processes

PAGE FAULTS AND EXCEPTIONS

- Wombat has a single server thread per CPU, like L⁴Linux
 - ★ is L4 pager and exception handler for Linux processes
 - ★ Linux “syscalls” are L4 exceptions
 - use `swi` instruction
 - Pistachio uses an invalid jump causing a *prefetch abort* for its system calls
 - easy to distinguish L4 and Linux syscalls

PAGE FAULTS AND EXCEPTIONS

- Wombat has a single server thread per CPU, like L⁴Linux
 - ★ is L4 pager and exception handler for Linux processes
 - ★ Linux “syscalls” are L4 exceptions
 - use `swi` instruction
 - Pistachio uses an invalid jump causing a *prefetch abort* for its system calls
 - easy to distinguish L4 and Linux syscalls
 - ★ other exceptions are reflected back as Linux signals

PAGE FAULTS AND EXCEPTIONS

- Wombat has a single server thread per CPU, like L⁴Linux
 - ★ is L4 pager and exception handler for Linux processes
 - ★ Linux “syscalls” are L4 exceptions
 - use `swi` instruction
 - Pistachio uses an invalid jump causing a *prefetch abort* for its system calls
 - easy to distinguish L4 and Linux syscalls
 - ★ other exceptions are reflected back as Linux signals
- Have single runnable Linux user process per CPU, unlike L⁴Linux
 - ★ differs from Dresden/Karlsruhe L⁴Linux approaches
 - ★ helps to maintain Linux scheduling behaviour

SYSTEM-CALL PROCESSING

- Access to application memory by Wombat:
 - ★ provide `copy_to/copy_from` functions
 - ★ look up page table and find page(s) in server's memory pool
 - ★ on MIPS could instead run Wombat in supervisor mode
 - would give direct access to user's address space

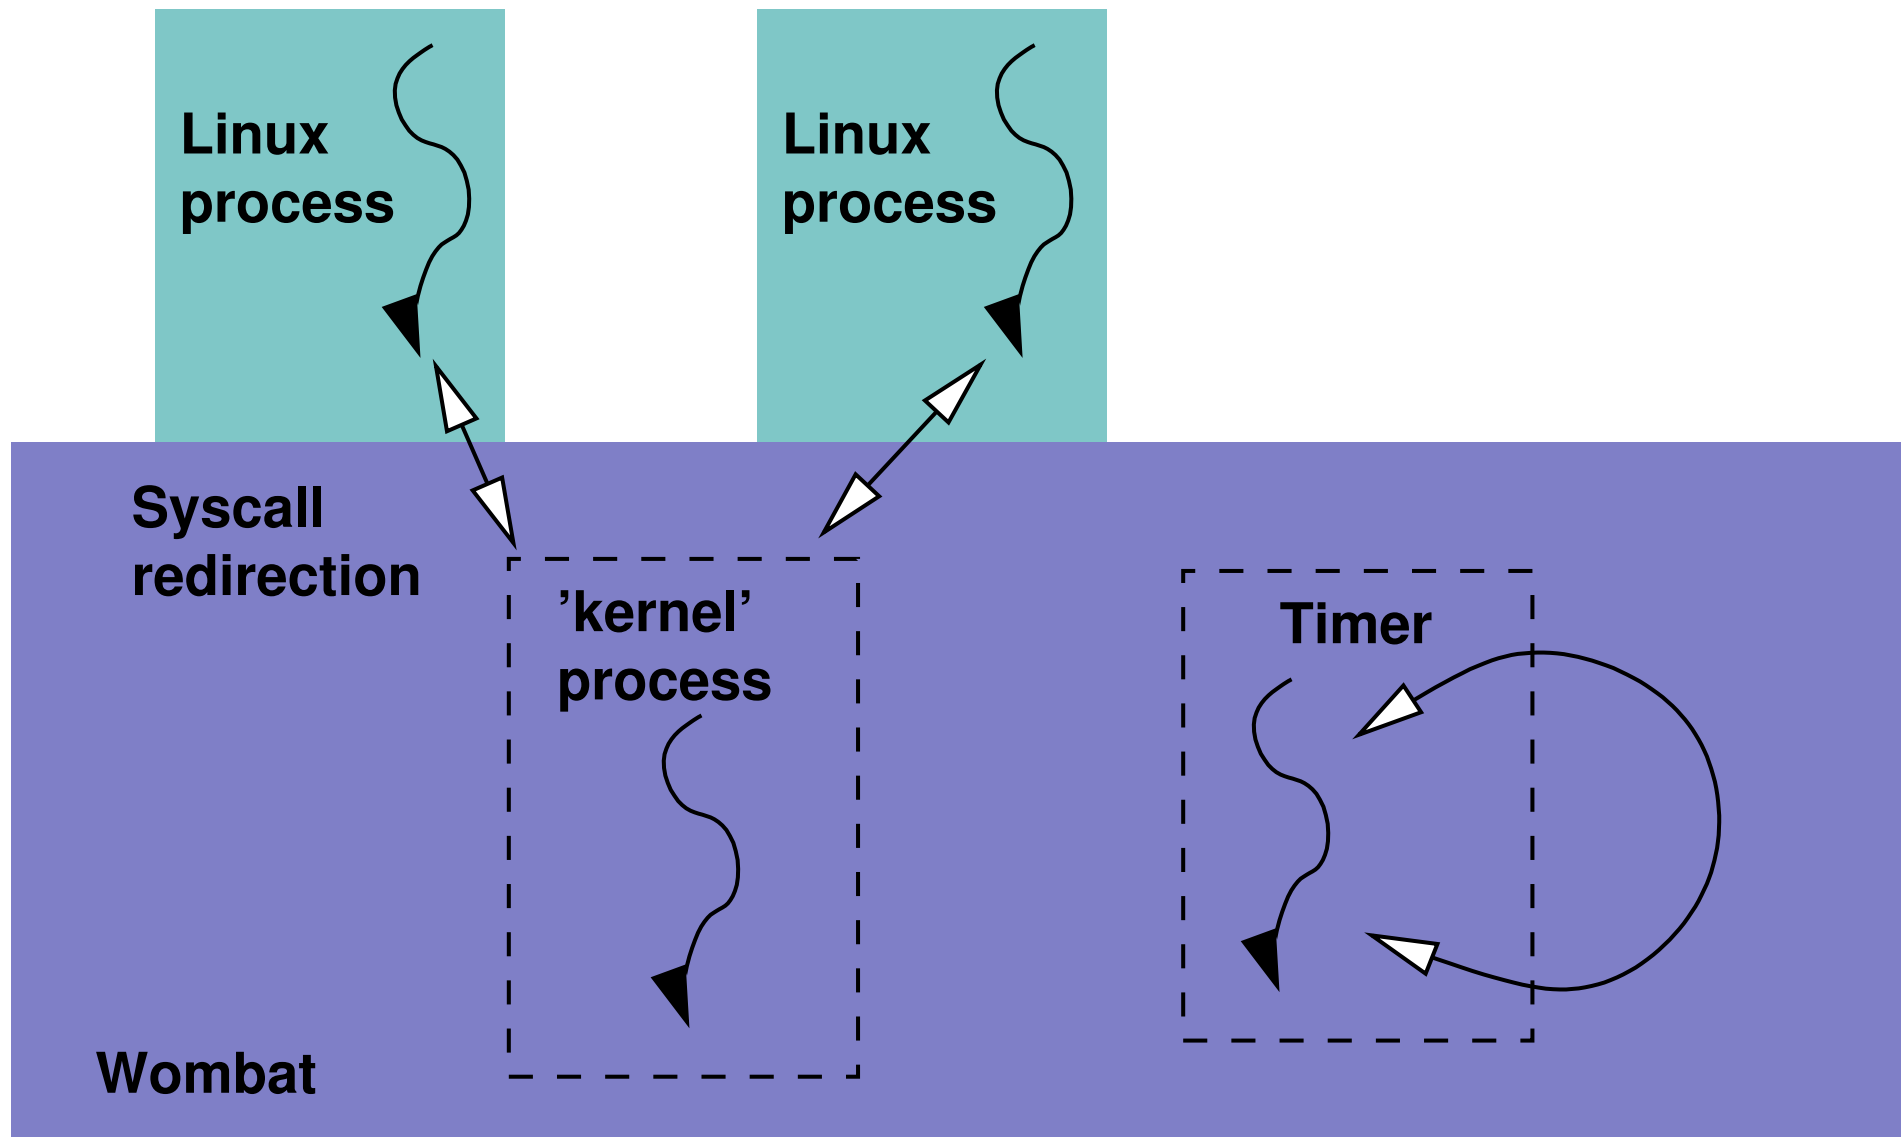
SYSTEM-CALL PROCESSING

- Access to application memory by Wombat:
 - ★ provide `copy_to/copy_from` functions
 - ★ look up page table and find page(s) in server's memory pool
 - ★ on MIPS could instead run Wombat in supervisor mode
 - would give direct access to user's address space
- Context switching:
 - ★ Wombat maintains pointer to state of current Linux process
 - ★ On a process switch, switches state and updates pointer
 - on ARM, state is banked user registers

SYSTEM-CALL PROCESSING

- Access to application memory by Wombat:
 - ★ provide `copy_to/copy_from` functions
 - ★ look up page table and find page(s) in server's memory pool
 - ★ on MIPS could instead run Wombat in supervisor mode
 - would give direct access to user's address space
- Context switching:
 - ★ Wombat maintains pointer to state of current Linux process
 - ★ On a process switch, switches state and updates pointer
 - on ARM, state is banked user registers
- Signal processing is architecture dependent
 - emulates Linux signal handling

LINUX PROCESSES



WOMBAT SCHEDULING

- Wombat has second thread: timer
 - ★ high-priority thread normally in timed wait
 - ★ maintains Linux time slice

WOMBAT SCHEDULING

- Wombat has second thread: timer
 - ★ high-priority thread normally in timed wait
 - ★ maintains Linux time slice
 - ★ when woken:
 - checks whether same Linux process is still running
 - if so, sets *reschedule* flag
 - if flag already set then sets Linux process' *total quantum* to zero
 - forces preemption, and an IPC to the scheduler (the “kernel” thread)

WOMBAT SCHEDULING

- Wombat has second thread: timer
 - ★ high-priority thread normally in timed wait
 - ★ maintains Linux time slice
 - ★ when woken:
 - checks whether same Linux process is still running
 - if so, sets *reschedule* flag
 - if flag already set then sets Linux process' *total quantum* to zero
 - forces preemption, and an IPC to the scheduler (the “kernel” thread)
- Scheduling happens when Linux process enters kernel
 - ★ checks *reschedule* flag, if set:
 - invokes Linux scheduler
 - perform context switch, reset *reschedule* flag

WOMBAT SCHEDULING

- Actual scheduling decision made by *unmodified* Linux scheduler
 - unlike L⁴Linux

WOMBAT SCHEDULING

- Actual scheduling decision made by *unmodified* Linux scheduler
 - unlike L⁴Linux
- Advantages of this approach:
 - ★ no changes required to (architecture-independent) scheduler
 - ★ Linux scheduling policy maintained (sort-of)
 - ★ no locking required other than what is already in Linux

WOMBAT SCHEDULING

- Actual scheduling decision made by *unmodified* Linux scheduler
 - unlike L⁴Linux
- Advantages of this approach:
 - ★ no changes required to (architecture-independent) scheduler
 - ★ Linux scheduling policy maintained (sort-of)
 - ★ no locking required other than what is already in Linux
- Disadvantages:
 - ★ Linux scheduling only approximately maintained
 - actual scheduling happens next time process enters kernel *after* time slice expired
 - processes run normally slightly longer than Linux time slice
 - CPU-bound process runs up to twice as long

WOMBAT SCHEDULING

- Actual scheduling decision made by *unmodified* Linux scheduler
 - unlike L⁴Linux
- Advantages of this approach:
 - ★ no changes required to (architecture-independent) scheduler
 - ★ Linux scheduling policy maintained (sort-of)
 - ★ no locking required other than what is already in Linux
- Disadvantages:
 - ★ Linux scheduling only approximately maintained
 - actual scheduling happens next time process enters kernel *after* time slice expired
 - processes run normally slightly longer than Linux time slice
 - CPU-bound process runs up to twice as long
 - Should be able to get exact Linux scheduling (but at a cost!)

LINUX IDLE TASK

- Linux's idle task executes `cpu_idle()`
 - ★ architecture-specific function
 - ★ in Wombat this IPCs to the "kernel" thread
 - ★ that one just sleeps until the next time tick or interrupt
 - ★ could just as well be made to go into a low-power mode

WOMBAT PERFORMANCE (PRELIMINARY)

Benchmark	Server	Linux	trampoline		IPC	
			L4	ratio	L4	ratio
getpid()	L ⁴ Linux	1.68	9.66	3.7	3.95	2.4
	Wombat	0.41	1.94	4.7		
2-proc ctxsw	L ⁴ Linux	7.0	18.2	2.6	16.2	2.3
	Wombat	4.03	9.66	2.4		
pipe	L ⁴ Linux	29.0	69.4	2.4	22.2	1.8
	Wombat	6.3	16.5	2.6		

- L⁴Linux 2.0 on 133MHz Pentium [SOSP 97], some improvements since
- Wombat 2.6 on Pentium-4 Xeon 2.66Ghz, *unoptimised*
 - overhead should be smaller on RISC processors

IGUANA DEVICE DRIVERS

- Drivers are Iguana services (data + server thread + methods)

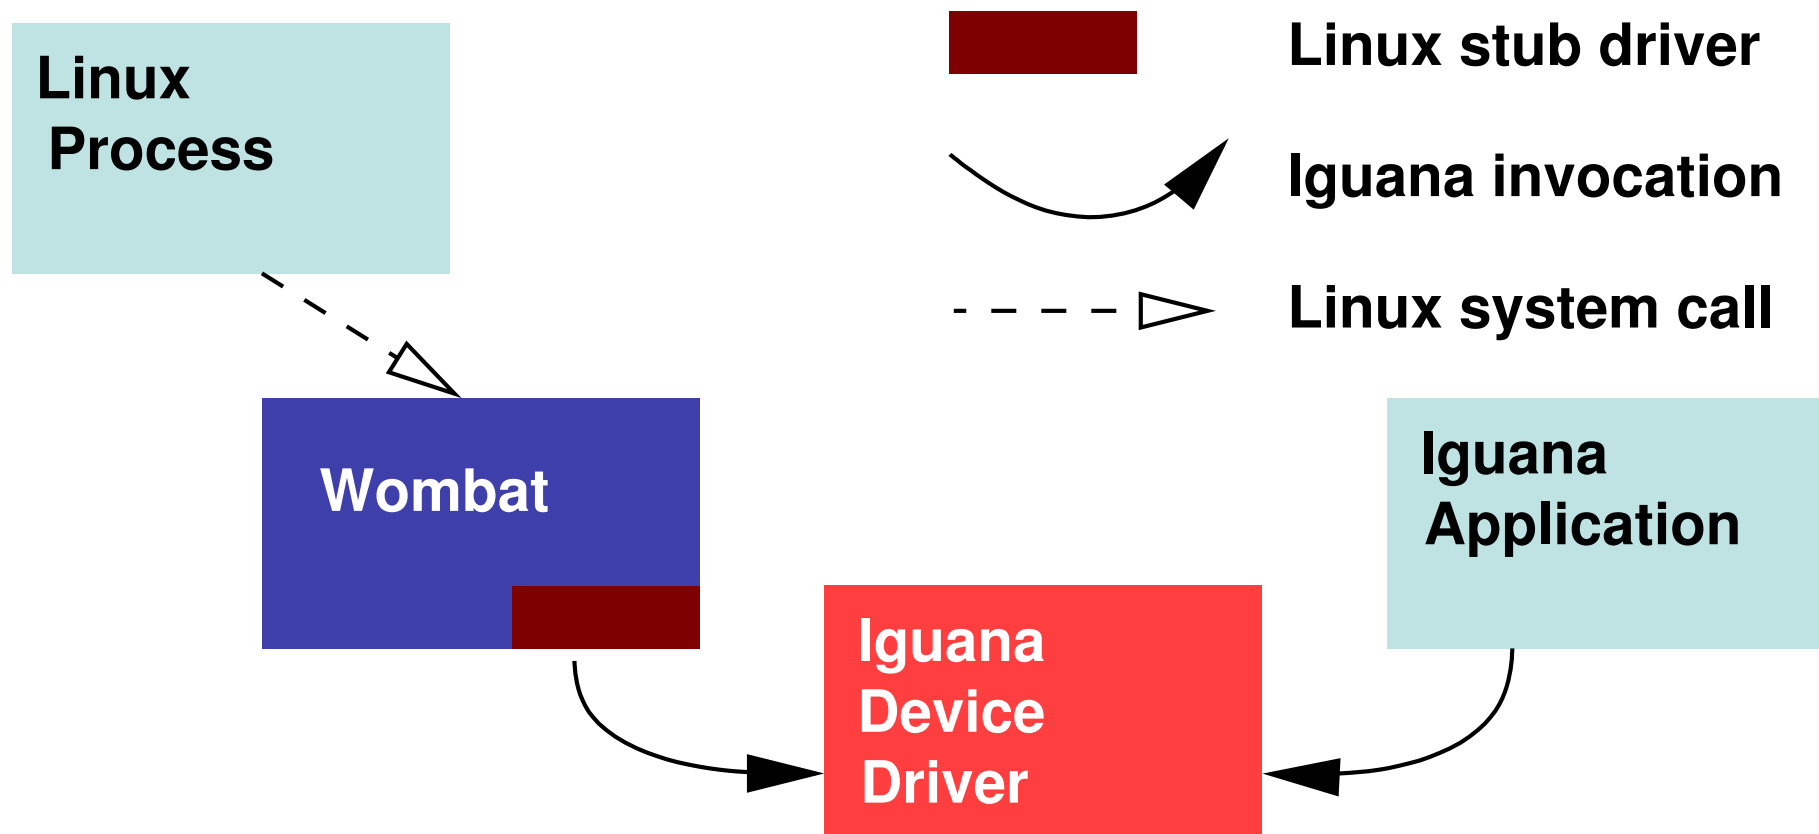
IGUANA DEVICE DRIVERS

- Drivers are Iguana services (data + server thread + methods)
- At startup obtain access to hardware resources
 - ★ can be done first-come, first-served
 - ★ better: restrict allocation to privileged service
 - ... by restricting capabilities to the `hardware` object
 - privileged service can allocate device memory (using `hardware->back_mem`)
 - can then hand capability for device-backed memory section to driver
- Driver is then able to operate the hardware

IGUANA DEVICE DRIVERS

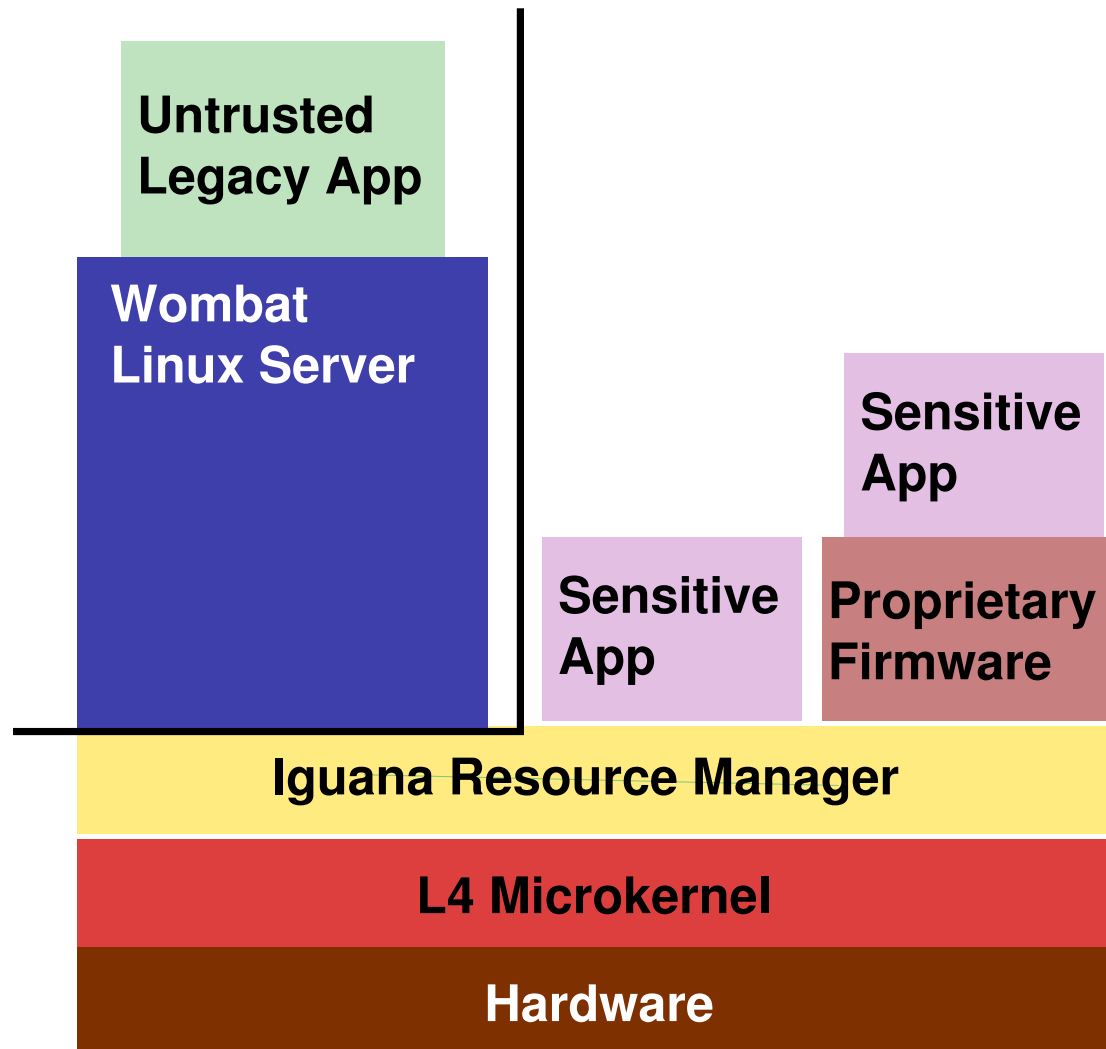
- Drivers are Iguana services (data + server thread + methods)
- At startup obtain access to hardware resources
 - ★ can be done first-come, first-served
 - ★ better: restrict allocation to privileged service
 - ... by restricting capabilities to the `hardware` object
 - privileged service can allocate device memory (using `hardware->back_mem`)
 - can then hand capability for device-backed memory section to driver
- Driver is then able to operate the hardware
- Provides methods as per Iguana driver protocol

SHARING IGUANA DEVICE DRIVERS WITH WOMBAT



- Wombat has stub driver module
 - invokes Iguana driver
 - sample exists for console

TARGET SYSTEM ARCHITECTURE



Similar to Dresden DROPS system

STATUS: IGUANA, WOMBAT AND DRIVER FRAMEWORK

- Alpha versions in use at two multinational companies
- Beta release planned within a week
 - including modified Pistachio to support it
- Production release planned for January
- Deployment in products planned for late next year

STATUS: IGUANA, WOMBAT AND DRIVER FRAMEWORK

- Alpha versions in use at two multinational companies
- Beta release planned within a week
 - including modified Pistachio to support it
- Production release planned for January
- Deployment in products planned for late next year
- Future work:
 - proper SMP implementation
 - support for heterogeneous multiprocessors (rSoC)
 - design and implementation of proper resource management
- Personnel
 - 1 honours student at present
 - *hackers needed!*

COMPONENT-BASED SOFTWARE ARCHITECTURE

- Goal: component architecture that is appropriate for microkernel
 - based on Iguana

COMPONENT-BASED SOFTWARE ARCHITECTURE

- Goal: component architecture that is appropriate for microkernel
 - based on Iguana
- To support:
 - ★ effective software-engineering techniques
 - ★ robustness by encapsulation without destroying performance
 - ★ hot-swapping / hot upgrading of components
 - ★ real-time (WCET) analysis of system
 - ★ ultimately, formal verification of whole system

COMPONENT-BASED SOFTWARE ARCHITECTURE

- Goal: component architecture that is appropriate for microkernel
 - based on Iguana
- To support:
 - ★ effective software-engineering techniques
 - ★ robustness by encapsulation without destroying performance
 - ★ hot-swapping / hot upgrading of components
 - ★ real-time (WCET) analysis of system
 - ★ ultimately, formal verification of whole system
- Project just commencing
 - collaboration with NICTA Empirical Software Engineering Program
 - *students/hackers needed!*

RECONFIGURABLE SoC

- Aims:
 - ★ get experience with reconfigurable computing
 - ★ understand requirements for kernel support
 - ★ provide OS services required

RECONFIGURABLE SoC

- Aims:
 - ★ get experience with reconfigurable computing
 - ★ understand requirements for kernel support
 - ★ provide OS services required
- Example: Embeddable GPS receiver
 - ★ aim: general satellite navigation research platform
 - ★ to be adaptable to different systems (GPS, Galileo, Glonass, Japanese, Chinese, Indian systems)
 - ★ to be usable for developing systems before protocols finalised
 - ★ will be open platform, available to researchers
 - ★ will form basis of follow-on projects

SECURE MOBILE CODE

- Combines language and OS techniques for secure execution
- Code may be:
 - accompanied by a proof certificate establishing its trustworthiness
 - otherwise sandboxed execution
- Supported by ARC
- Lead by Manuel Chakravarty
 - ★ presently working on theoretical foundations
 - ★ later to be integrated with Iguana
 - ★ requires *configurable protection domains*

SYSTEMATIC TESTING

- Software used commercially \Rightarrow needs to work!
- Present approach: ad-hoc
 - unit testing
 - regression testing

SYSTEMATIC TESTING

- Software used commercially \Rightarrow needs to work!
- Present approach: ad-hoc
 - unit testing
 - regression testing
- Better approach: Automatic generation of tests
 - based on formal description of API
 - based on static analysis of implementation (source code)

SYSTEMATIC TESTING

- Software used commercially \Rightarrow needs to work!
- Present approach: ad-hoc
 - unit testing
 - regression testing
- Better approach: Automatic generation of tests
 - based on formal description of API
 - based on static analysis of implementation (source code)
- Collaboration with NICTA Formal Methods Program
 - ★ in planning phase
 - ★ *students needed!*

OTHER PROJECTS

- Distributed mobile file system
 - access you files on your PDA/phone/... wherever you are

OTHER PROJECTS

- Distributed mobile file system
 - access you files on your PDA/phone/... wherever you are
- Reflective real-time systems architecture
 - better approach to managing soft real-time & best-effort

OTHER PROJECTS

- Distributed mobile file system
 - access you files on your PDA/phone/... wherever you are
- Reflective real-time systems architecture
 - better approach to managing soft real-time & best-effort
- Hot swapping/upgrading of system components
 - collaboration with IBM research

OTHER PROJECTS

- Distributed mobile file system
 - access you files on your PDA/phone/... wherever you are
- Reflective real-time systems architecture
 - better approach to managing soft real-time & best-effort
- Hot swapping/upgrading of system components
 - collaboration with IBM research
- Distributed embedded systems
 - e.g. cars
 - pilot project SUNswift

OTHER PROJECTS...

- Secure GUI

- window system that ensures isolation of clients

OTHER PROJECTS...

- Secure GUI
 - window system that ensures isolation of clients
- JVM on L4

OTHER PROJECTS...

- Secure GUI
 - window system that ensures isolation of clients
- JVM on L4
- API emulation
 - support for μ ITRON, Symbian OS, QNX, ...

OTHER PROJECTS...

- Secure GUI
 - window system that ensures isolation of clients
- JVM on L4
- API emulation
 - support for μ ITRON, Symbian OS, QNX, ...
- Kernel ports
 - SPARC, Hitachi, Motorola

OTHER PROJECTS...

- Secure GUI
 - window system that ensures isolation of clients
- JVM on L4
- API emulation
 - support for μ ITRON, Symbian OS, QNX, ...
- Kernel ports
 - SPARC, Hitachi, Motorola
- Simulators
 - tools for OS development and architecture research
 - ARM, PowerPC, ...

SUMMARY: OS RESEARCH AT NICTA

- Improve reliability and trustworthiness of embedded systems
 - ★ make a *real* impact on the practice of embedded systems R&D

SUMMARY: OS RESEARCH AT NICTA

- Improve reliability and trustworthiness of embedded systems
 - ★ make a *real* impact on the practice of embedded systems R&D
- Comprehensive approach used on using microkernel technology
- Combination of formal and practical approaches
- Leverage relevant competencies available within NICTA

SUMMARY: OS RESEARCH AT NICTA

- Improve reliability and trustworthiness of embedded systems
 - ★ make a *real* impact on the practice of embedded systems R&D
- Comprehensive approach used on using microkernel technology
- Combination of formal and practical approaches
- Leverage relevant competencies available within NICTA
- Ample opportunities to participate in cutting-edge research
 - summer projects
 - honours theses
 - PhD theses
 - employment