



University  
of Glasgow

# Review and Future Directions

Advanced Systems Programming (M)

Lecture 9

# Review

- Systems Programming
- Memory Management
- Concurrency
- Security

# Systems Programming

# What is Systems Programming?

- Infrastructure components, operating systems, device drivers, network protocols, services
- Systems programs tend to be constrained by:
  - Memory management and data representation
  - I/O operations
  - Management of shared state
  - Performance



J. Shapiro, "Programming language challenges in systems codes: why systems programmers still use C, and what to do about it", Workshop on Programming Languages and Operating Systems, San Jose, CA, October 2006. DOI:[10.1145/1215995.1216004](https://doi.org/10.1145/1215995.1216004)

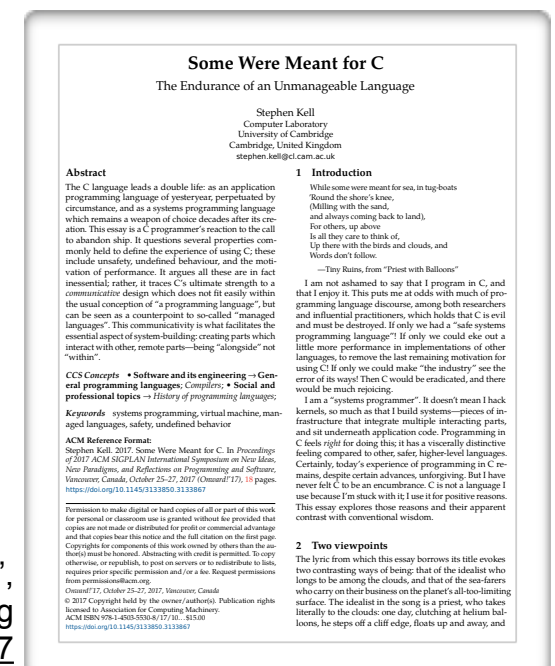


# State of the Art: Unix and C

- Unix gained popularity due to portability and ease of source code access, but also:
  - Small, relatively consistent set of API calls
  - Low-level control
  - Robust and high performance
  - Easy to understand and extend
- Portability due to the C programming language
  - Simple, easy to understand, easy to port
  - Explicit pointers, memory allocation, and control over data representation; uniform treatment of memory, devices and data structures
  - Weak type system allows aliasing and sharing




<https://dave.cheney.net/2017/12/04/what-have-we-learned-from-the-pdp-11>  
Image credit: Dennis Ritchie




S. Kell, “Some were meant for C: The endurance of an unmanageable language”, International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software, Vancouver, BC, Canada, October 2017. ACM. DOI:[10.1145/3133850.3133867](https://doi.org/10.1145/3133850.3133867)

# No More Moore?



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten  
New plot and data collected for 2010-2015 by K. Rupp

Source: <https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/>



- Breakdown of Dennard scaling, impending end of Moore's law → concurrency
- Ubiquitous networks and cyber-physical systems pushing security to the fore
- Existing approaches to systems programming no longer sufficient

Do you agree with this claim?

# Next Steps in Systems Programming

- Need a programming language and environment that:
  - Improves memory management and safety – while maintaining control over allocation and data representation
  - Improves security – eliminates common classes of vulnerability
  - Improves support for concurrency – eliminates race conditions
  - Improves correctness – eliminates common classes of bug
- Advances in programming language design are starting to provide the necessary tools – and beginning to be applied to systems languages
  - Modern type systems
  - Functional programming techniques

# Why is Strong Typing Desirable?

- “Well-typed programs don’t go wrong” – Robin Milner
- The result is well-defined – although not necessarily correct
  - The type system ensures results are consistent with the rules of the language, but cannot check if you calculated the right result
  - A strongly-typed system will only ever perform operations on a type that are legal – there is no undefined behaviour
- Types help model the problem, check for consistency, and eliminate common classes of bug

# Segmentation fault (core dumped)

Segmentation faults should never happen:

- Compiler and runtime should strongly enforce type rules
- If program violates them, it should be terminated cleanly
- Security vulnerabilities – e.g., buffer overflow attacks – come from undefined behaviour after type violations

# Challenges in Strongly-typed Systems Programming

- Four fallacies:
  - Factors of 1.5x to 2x in performance don't matter
  - Boxed representation can be optimised away
  - The optimiser can fix it
  - The legacy problem is insurmountable
- Four challenges:
  - Application constraint checking
  - Idiomatic manual storage management
  - Control over data representation
  - Managing shared state

- Many good ideas in research languages and operating systems – only recently that these issues have been considered to make *practical* tools



J. Shapiro, "Programming language challenges in systems codes: why systems programmers still use C, and what to do about it", Workshop on Programming Languages and Operating Systems, San Jose, CA, October 2006. DOI:[10.1145/1215995.1216004](https://doi.org/10.1145/1215995.1216004)

# The Rust Programming Language


- A practical, strongly typed, systems programming language
  - Enumerated types and pattern matching: **Option** and **Result**
  - Structure types and traits as alternative to object oriented design
  - Ownership, borrowing, and multiple reference types
  - Little is novel – adopts ideas from numerous other languages:
    - Syntax → C, Standard ML, and Pascal
    - Basic data types → C and C++
    - Enumerated types and pattern matching → Standard ML
    - Traits → Haskell type classes
    - The ownership and borrowing rules → Cyclone
- Strongly typed and permits no-cost abstractions





# Type-driven Development

- **Define the types first**
  - Define concrete numeric types, identifiers
  - Define enum types to represent alternatives
  - Indicate optional values, results, error types
- Using the types as a guide, **write the functions**
  - Write the input and output types
  - Write the function, using the structure of the types as a guide
  - Make state machines explicit
  - Consider ownership of data
- **Refine** and edit types and functions as necessary
  - Use the compiler as a tool to help you debug your design
- Don't think of the types as checking the code, think of them as a plan, a model, for the solution – and as machine checkable documentation



Type-driven development approach  
adapted from: E. Brady, "[Type-Driven Development with Idris](#)",  
Manning, March 2017.

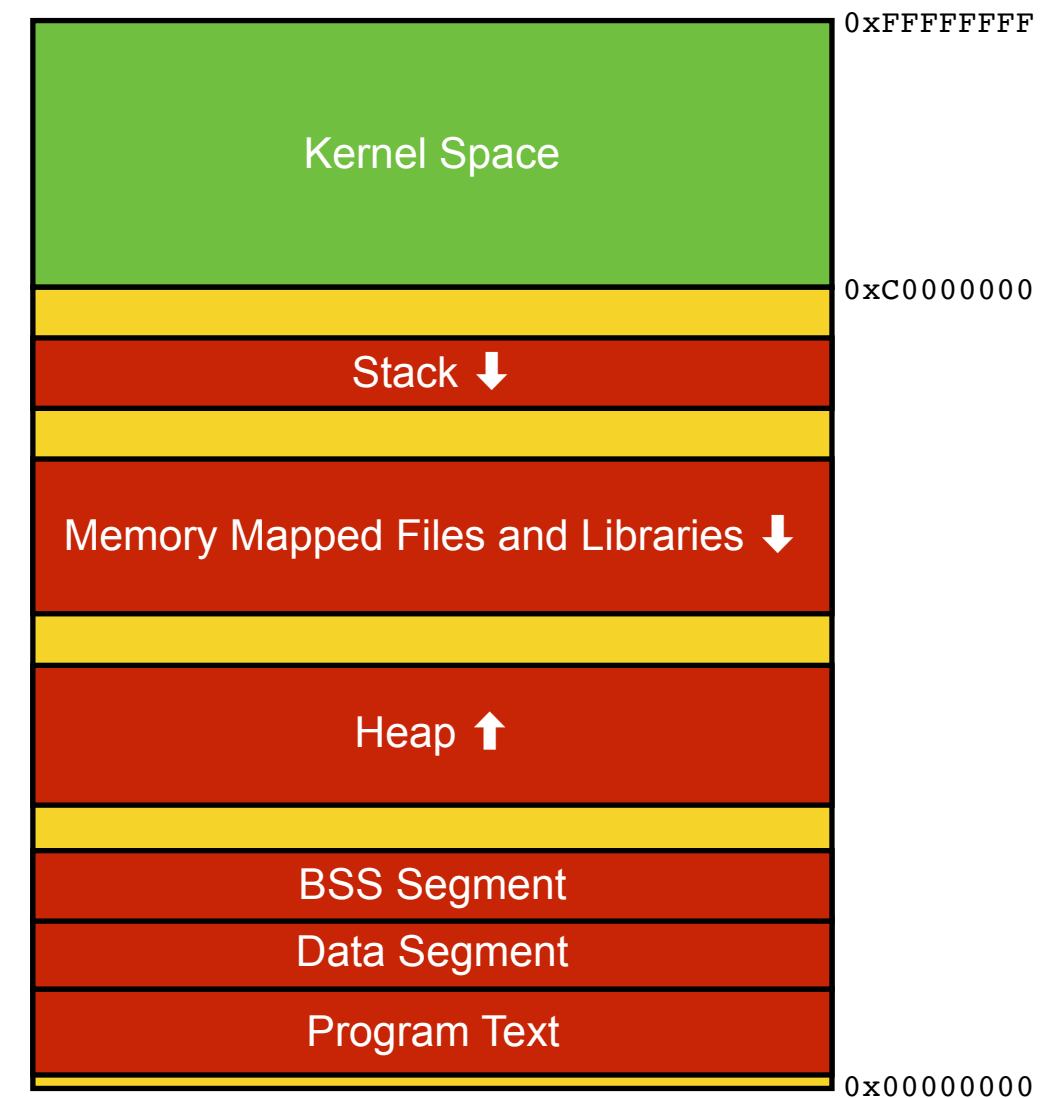
Is this a good approach to systems development?



# Memory Management

# Layout of a Processes in Memory

- Layout of process address space:
  - Kernel at top of address space
  - Program text, data, and global variables at bottom of virtual address space
  - Heap allocated upwards, above BSS
  - Stack grows downwards, below kernel
  - Memory mapped files and shared libraries between these



Typical addresses on 32 bit machines


See also <http://duartes.org/gustavo/blog/post/anatomy-of-a-program-in-memory/>

# Automatic Memory Management

- Automatic memory management distrusted by systems programmers
  - Perceived high processor and memory overheads, unpredictable timing
  - But, memory management problems are common:
    - Unpredictable performance
      - Calls to `malloc()`/`free()` can vary in execution time by several orders of magnitude
    - Memory leaks
    - Memory corruption and buffer overflows
    - Use-after-free
    - Iterator invalidation
- New automatic memory management schemes solve many problems
  - Garbage collectors → lower overhead, more predictable
    - Also system performance improvements made overhead more acceptable
  - Region-based memory management → predictability, compile time guarantees

# Reference Counting

- Simplest automatic heap management
- Each allocation also allocates space for an additional **reference count**
  - An extra **int** is allocated along with every object
  - Counts number of references to the object
    - Increased when new reference to the object is created
    - Decremented when a reference is removed
  - When reference count reaches zero, there are no references to the object, and it may be reclaimed
    - Reclaiming object removes references to other objects
    - May reduce their reference count to zero, so triggering further reclamation
- Incremental, predictable, and understandable
- Potentially high cost; memory leaks with cyclic data structures



Source: P. Wilson, "Uniprocessor garbage collection techniques", Proc IWMM'92, DOI:10.1007/BFb0017182

# Region-based Memory Management

- Allocate objects with lifetimes corresponding to regions matching program scope
- Track object ownership, and changes of ownership:
  - What region owns each object at any time
  - Ownership of objects can move between regions
- Deallocate objects at the end of the lifetime of their owning region
  - Use scoping rules to ensure objects are not referenced after deallocation
- Efficient and predictable; correctness guarantees prevent common bugs
- Constrains the type of programs that can be written
- Forces programmer to consider resource ownership early in the design – but eliminates run-time misbehaviours and costs



# Garbage Collection

- Garbage collection algorithms:
  - Mark-sweep
  - Mark-compact
  - Copying collectors
  - Generational algorithms
- Incremental approaches
- Real-time collection

## Uniprocessor Garbage Collection Techniques

Paul R. Wilson

University of Texas  
Austin, Texas 78712-1188 USA  
(wilson@cs.utexas.edu)

**Abstract.** We survey basic garbage collection algorithms, and variations such as incremental and generational collection. The basic algorithms include reference counting, mark-sweep, mark-compact, copying, and treadmill collection. *Incremental* techniques can keep garbage collection pause times short, by interleaving small amounts of collection work with program execution. *Generational* schemes improve efficiency and locality by garbage collecting a smaller area more often, while exploiting typical lifetime characteristics to avoid undue overhead from long-lived objects.

### 1 Automatic Storage Reclamation

*Garbage collection* is the automatic reclamation of computer storage [Knu69, Coh81, App91]. While in many systems programmers must explicitly reclaim heap memory at some point in the program, by using a “free” or “dispose” statement, garbage collected systems free the programmer from this burden. The garbage collector’s function is to find data objects<sup>1</sup> that are no longer in use and make their space available for reuse by the the running program. An object is considered *garbage* (and subject to reclamation) if it is not reachable by the running program via any path of pointer traversals. *Live* (potentially reachable) objects are preserved by the collector, ensuring that the program can never traverse a “dangling pointer” into a deallocated object.

This paper is intended to be an introductory survey of garbage collectors for uniprocessors, especially those developed in the last decade. For a more thorough treatment of older techniques, see [Knu69, Coh81].


#### 1.1 Motivation

Garbage collection is necessary for fully modular programming, to avoid introducing unnecessary inter-module dependencies. A routine operating on a data structure should not have to know what other routines may be operating on the same structure, unless there is some good reason to coordinate their activities. If objects must be deallocated explicitly, some module must be responsible for knowing when *other* modules are not interested in a particular object.

<sup>1</sup> We use the term object loosely, to include any kind of structured data record, such as Pascal records or C structs, as well as full-fledged objects with encapsulation and inheritance, in the sense of object-oriented programming.

P. R. Wilson, “Uniprocessor garbage collection techniques”, Proceedings of the International Workshop on Memory Management, St. Malo, France, September 1992. DOI: [10.1007/BFb0017182](https://doi.org/10.1007/BFb0017182)

# Memory Management Trade-offs



- Rust pushes memory management complexity onto the programmer
  - Predictable run-time performance, low run-time overheads
  - Uniform resource management framework, including memory
  - Limits the programs that may be expressed – matches common patterns in good C code
- Garbage collection imposes run-time costs and complexity, but simpler for the programmer


Does Rust make the right trade-off here?

# Concurrency



# Concurrency, Threads, and Locks

- Operating systems expose concurrency via *processes* and *threads*
  - Processes are isolated with separate memory areas
  - Threads share access to a common pool of memory
- The processor/language memory models specify how concurrent access to shared memory works
  - e.g., synchronise by explicitly locking critical sections
    - **synchronized** methods and statements in Java
    - **pthread\_mutex\_lock()/pthread\_mutex\_unlock()**
  - Limited guarantees about unlocked concurrent access to shared memory



# Limitations of Lock-based Concurrency

- Major problems with lock-based concurrency:
  - Difficult to define a memory model that enables good performance, while allowing programmers to reason about the code
  - Difficult to ensure correctness when composing code
    - Difficult to enforce correct locking
    - Difficult to guarantee freedom from deadlocks
  - Failures are silent – errors tend to manifest only under heavy load
  - Balancing performance and correctness difficult – easy to over- or under-lock systems

# Composition of Lock-based Code

- Correctness of small-scale code using locks can be ensured by careful coding (at least in theory)
- A more fundamental issue: lock-based code does not compose to larger scale
  - Assume a correctly locked bank account class, with methods to credit and debit money from an account
  - Want to take money from **a1** and move it to **a2**, without exposing an intermediate state where the money is in neither account
  - Can't be done without locking all other access to **a1** and **a2** while the transfer is in progress
  - The individual operations are correct, but the combined operation is not
- This is lack of abstraction a limitation of the lock-based concurrency model, and cannot be fixed by careful coding
- Locking requirements form part of the API of an object

```
a1.debit(v)  
a2.credit(v)
```

Preemption exposes  
intermediate state



# Alternative Concurrency Models

- Concurrency increasingly important
  - Multicore systems now ubiquitous
  - Asynchronous interactions between software and hardware devices
- Threads and synchronisation primitives problematic
- Are there alternatives that avoid these issues?
  - Transactions
  - Message passing

# Transactions

- Are transactions a reasonable programming model?
- Is transactional memory a realistic technique?
  - Assumption: shared memory system, doesn't work with distributed and networked systems – is this true?
- Concurrent Haskell:
  - Monadic IO; do notation; IORefs; spawning threads
  - Type system separates state and stateless computation
  - The STM interface
- Do its requirements for a purely functional language, with controlled I/O, restrict it to being a research toy?
- How much benefit can be gained from transactional memory in more traditional languages?



T. Harris, S. Marlow, S. Peyton Jones, and M. Herlihy, "Composable Memory Transactions", Communications of the ACM, 51(8), August 2008. DOI:10.1145/1378704.1378725.

<http://www.cmi.ac.in/~madhavan/courses/pl2009/reading-material/harris-et-al-cacm-2008.pdf>

Is this a realistic systems programming model?

# Message Passing Systems


- System is a set of communicating processes – no shared mutable state
- Communication via exchange of messages
  - Immutable or linearly typed to ensure safety

```
use std::sync::mpsc::channel;
use std::thread;

fn main() {
    let (tx, rx) = channel();

    thread::spawn(move || {
        let _ = tx.send(42);
    });

    match rx.recv() {
        Ok(value) => {
            println!("Got {}", value);
        }
        Err(error) => {
            // An error occurred...
        }
    }
}
```



Does this offer sufficient benefit over shared state concurrency?

# Coroutines and Asynchronous Code

- Provide language and run-time support for I/O multiplexing on a single thread, in a more natural style

```
fn read_exact<T: Read>(input: &mut T, buf: &mut [u8]) -> Result<(), std::io::Error> {  
    let mut cursor = 0;  
    while cursor < buf.len() {  
        cursor += input.read(&mut buf[cursor..])?;  
    }  
}
```



```
async fn read_exact<T: AsyncRead>(input: &mut T, buf: &mut [u8]) -> Result<(), std::io::Error> {  
    let mut cursor = 0;  
    while cursor < buf.len() {  
        cursor += await!(input.read(&mut buf[cursor..]))?;  
    }  
}
```

- Runtime schedules **async** functions on a thread pool, yielding to other code on **await!()** calls → low-overhead concurrent I/O

Do the benefits outweigh the cooperative multitasking concern?

# Security



# Memory Safety in Programming Languages


- A **memory safe** language is one that ensures that the only memory accessed is that owned by the program, and that such access is done in a manner consistent of the declared type of the data
  - The program can access memory through its global and local variables, or through explicit references to them
  - The program can access heap memory it allocated via an explicit reference to that memory
  - All accesses obey the type rules of the language
    - Array bounds are respected
    - Memory cannot be accessed after it's freed
    - References can only refer to objects of matching type
    - ...

Memory Safe	Memory Unsafe
Java, Scala, C#, Rust, Go, Python, Ruby, Tcl, FORTRAN, COBOL, Modula-2, Occam, Erlang, Ada, Pascal, Haskell, ...	C, C++, Objective-C

# Impact of Memory Unsafe Languages

- Lack of memory safety breaks the machine abstraction
  - With luck, program crashes – segmentation violation
  - If unlucky, memory unsafe behaviour corrupts other data owned by program
    - Undefined behaviour occurs
    - Cannot predict without knowing precise layout of program in memory
    - Difficult to debug
    - **Potential security risk** – corrupt program state to force arbitrary code execution

# Impact of Memory Unsafe Languages – Security (2/2)



Source: Matt Miller, Microsoft, presentation at BlueHat IL conference, February 2019  
[https://github.com/Microsoft/MSRC-Security-Research/tree/master/presentations/2019\\_02\\_BlueHatIL](https://github.com/Microsoft/MSRC-Security-Research/tree/master/presentations/2019_02_BlueHatIL)

- ~70% of Microsoft security updates fix bugs relating to unsafe memory usage
  - Interestingly, this has not significantly changed in >10 years
  - We're not getting better at writing secure code in memory unsafe languages

# Mitigations for Memory Unsafe Languages (1/2)

- Use modern tooling for C and C++ development:
  - Compile C code with, at least, **clang -W -Wall -Werror**
  - Use **clang** static analysis tools during debugging:
- Use modern C++ language features:
  - I advise against C++ programming – language is too complex to understand
  - But if you have to use C++, modernise the code base, to use newer – safer – language features and idioms, where possible
- Consider rewriting the most critical sections of the code in a memory safe language:
  - Rust can call C functions directly, and can be compiled into a library that can be directly linked with C or C++
  - Objective-C and Swift can similarly be combined in the Apple ecosystem

# Mitigations for Memory Unsafe Languages (2/2)

- Writing safe code in unsafe languages is difficult
  - Easy to rationalise each individual bug – “How could anyone write that?”

```
if ((err = SSLHashSHA1.update(&hashCtx, &signedParams)) != 0)
    goto fail;
goto fail;
... other checks ...
fail:
    ... buffer frees (cleanups) ...
    return err;
```

<https://dwheeler.com/essays/apple-goto-fail.html>

- But, people will continue to make mistakes
- Memory safe languages eliminate common classes of vulnerability; strong type systems help detect mistakes early

# Parsing and LangSec

- Can we combine better parsing with modern, strongly typed, languages to improve network security?
- Is performance good enough?
- Can we improve the way we design protocols?



S. Bratus, M. L. Patterson, and A. Shubina. The bugs we have to kill. ;login:, 40(4):4–10, August 2015. <http://langsec.org/papers/the-bugs-we-have-to-kill.pdf>

# Causes of Security Vulnerabilities

- Security vulnerabilities generally caused by persuading a program to do something that the programmer did not expect
  - Write past the end of a buffer
  - Treat user input as executable
  - Confuse a permission check
  - Violate an assumption in the code
- Strong typing makes assumptions explicit
  - Use explicit types rather than generic types
  - Define safe conversion functions

# No Silver Bullet

Memory safety and strong typing won't eliminate security vulnerabilities



But, used carefully, they eliminate certain classes of vulnerability, and make others less likely by making hidden assumptions – and their consequences – visible

**Do you agree with this characterisation?**



# Future Directions

- Scaling systems programming

# Concurrency

- Still no good solutions for concurrent programming
  - Transactional memory doesn't sit well with impure imperative languages
  - Message passing has issues with back pressure, race conditions, deadlocks, large-scale orchestration
  - Asynchronous code is a work-around for heavyweight threads
- None scale to GPU programming → implicit parallelism?



# Data Centre Scale Computing

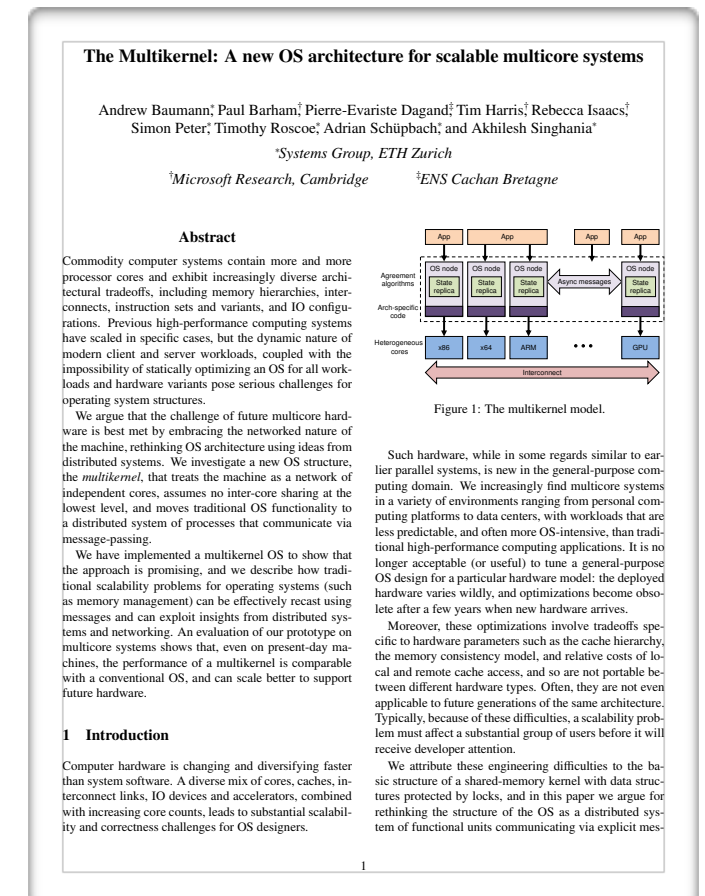
How do we  
scale systems  
programming to the  
scale of a data centre?

- Distributed memory computers with millions of cores
- Hard to configure nodes and communications at this scale – automatic configuration, scaling, tuning, and fault tolerance essential
- Unclear that the future should be large-scale Linux clusters – what comes next?



# Distributed Operating Systems

- Barrelfish – a message passing kernel for multicore systems
  - Each core runs a separate operating systems instance
  - **All** communication between cores is via message passing; no “main” CPU to act as central orchestrator
- A research prototype to test an idea, not a product
  - Where is the boundary for a Barrelfish-like system?
  - Distinction between a distributed multi-kernel and a distributed system of networked computers? Should there be such a distinction?
  - How does it relate to concurrent programming, data centres, distributed edge computing, etc?




A. Baumann *et al.*, “The Multikernel: A new OS architecture for scalable multicore systems”, Proc. ACM Symposium on Operating Systems Principles, 2009. DOI:10.1145/1629575.1629579

# Future Systems Programming Languages

- Increasingly seeing research ideas incorporated into mainstream programming languages
  - Rust, Swift – abstraction, resource management, strong and expressive type systems, low-level control
  - Erlang, Go – concurrency, fault tolerance, communication
- Choose the language based on problem domain, available tooling, and expertise
  - Good solutions for programming single systems
  - We still have a lot to learn about scaling to large clusters



# Scaling Systems Programming



- Unix and C made sense when programming simple single-core systems
  - And form the basis of most existing systems – **but what comes next?**
- Modern massively concurrent and distributed applications need new approaches:
  - Low-level performance still crucial, but also need type systems and abstractions to hide the complexity – *no-one* can hold all the details in their head
  - Tooling is essential – to check invariants, document assumptions, ensure consistency
  - Deployment, configuration, and management must become increasingly autonomic – DevOps, infrastructure as code, self-managing

# What are we missing?

# Assessment



# Assessment

- This is a level M course, worth 10 credits
- Coursework (20%)
  - Essay (10%) and programming exercise (10%)
  - Exercise 1 marks available from teaching office
  - Exercise 2 marks will be returned by 22 March 2019
- Examination (80%):
  - Material from the lectures, labs, and cited papers is examinable
  - Aim is to test your understanding of the material, not to test your memory of all the details; explain why – don't just recite what
  - Two hours duration; answer three out of four questions
  - Sample paper will be provided by 22 March 2019

# The End



J. Mickens. The night watch. ;login: logout, pages 5–8, November 2013.  
[https://www.usenix.org/system/files/1311\\_05-08\\_mickens.pdf](https://www.usenix.org/system/files/1311_05-08_mickens.pdf)