



University  
of Glasgow

# Advanced Topics in Systems Programming

Advanced Systems Programming (M)

Lecture 1

# Introduction and Course Administration

# Contact Details and Website

- Lecturer and course coordinator:
  - Dr Colin Perkins / S101b Lilybank Gardens / [colin.perkins@glasgow.ac.uk](mailto:colin.perkins@glasgow.ac.uk)
  - No scheduled office hours – make appointments by email if necessary
- Lecture notes and other materials online:
  - <https://csperkins.org/teaching/2018-2019/adv-systems-programming/> (or on the School's Moodle site)
  - Paper handouts and lecture recordings will not be provided
    - The act of taking notes helps learning
    - The act of revising by reviewing your notes and other textbooks gives a useful additional perspective on the material, helps you understand the topic in depth
    - Lecture recordings discourage questions – it's already hard enough to ask a question in a large lecture theatre, but still harder when you know it's being recorded

# Rationale

- Technology shift: desktop PC → laptop, tablet, smartphone, cloud
  - Mobile, power-aware, concurrent, real-time, connected
  - But still programmed in C, running some variant of Unix – technology stack that's becoming increasingly limiting
- This course will explore new techniques for safer, more effective, systems programming
  - Programming in an unmanaged environment, where data layout and performance matter
  - Operating systems kernels, device drivers, low-level networking code

# Aims and Objectives

- The course aims to explore the features of modern programming languages and operating systems that can ease the challenges of systems programming, considering type systems and run-time support.
- It will review the research literature on systems programming and operating system interfaces, discuss the limitations of deployed systems, and consider how systems programming might evolve to address the challenges of supporting modern computing systems.
- Particular emphasis will be placed on system correctness and secure programming, to ensure the resulting systems are safe to use in an adversarial environment.

# Intended Learning Outcomes (1/2)

- By the end of the course, students should be able to:
  - To discuss the advantages and disadvantages of C as a systems programming language, and to compare and contrast this with a modern systems programming language, for example Rust; to discuss the role of the type system, static analysis, and verification tools in systems programming, and show awareness of how to model system properties using the type system to avoid errors;
  - To discuss the challenges of secure low-level programming and write secure code in a modern systems programming language to perform systems programming tasks such as parsing hostile network input; show awareness of security problems in programs written in C;
  - To discuss the advantages and disadvantages of integrating automatic memory management with the operating system/runtime, to understand the operation of popular garbage collection algorithms and alternative techniques for memory management, and know when it might be appropriate to apply such techniques and managed run-times to real-time systems and/or operating systems;
  - ...

# Intended Learning Outcomes (2/2)

- ...
- To understand the impact of heterogeneous multicore systems on operating systems, compare and evaluate different programming models for concurrent systems, their implementation, and their impact on operating systems; and
- To construct and/or analyse simple programming to demonstrate understanding of novel techniques for memory management and/or concurrent programming, to understand the trade-offs and implementation decisions.

# Pre-requisites

- The following courses are pre-requisites:
  - Systems Programming (H)
  - Operating Systems (H)
  - Networked Systems (H)
  - Functional Programming (H)
- You are expected to be familiar with the C programming language, and to understand the basics of operating systems and networking
- A conceptual understanding of functional programming is assumed, but Haskell programming is not required



# Course Outline and Timetable

Week	Lecture	Laboratory
1	#1: Introduction	#1: Introducing Rust (1/2)
2	#2: Types and Systems Programming	#2: Introducing Rust (2/2)
3	#3: Type-based Modelling and Design	#3: Types and Traits
4	#4: Resource Ownership and Memory Management	#4: Ownership, Pointers, and Memory
5	#5: Garbage Collection	
6	#6: Concurrency	#5: Concurrency
7	#7: Coroutines and Asynchronous Programming	#6: Coroutines and Asynchronous Code
8	#8: Security Considerations	
9	#9: Open Issues and Future Directions	
10		

# Laboratory Sessions


- Lab sessions are in Boyd Orr 1028, 13:00-14:00 on Mondays
- Self-study handouts will be available online before each lab – work through the handouts at own pace
- I'll be available in the labs to answer questions
- Lab handouts have been tested using Rust 1.31.1 on OS X 10.12.6, but are expected to run on the Linux machines in the lab

# Assessment


- This is a level M course, worth 10 credits
- Assessment is by examination (80%) and coursework (20%)
  - Sample exam paper, with worked answers, will be made available on the website/Moodle before the end of the semester
  - 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
- Coursework:

Coursework	Date Set:	Date Due:	Weighting:	Topic:
Exercise 1	Lecture 4	Lecture 6	10%	Memory Management (essay)
Exercise 2	Lecture 6	Lecture 9	10%	Concurrent Programming (code)

# Recommended Reading (1/2)



Steve Klabnik and Carol Nichols, "The Rust Programming Language", 2nd Edition, 2018, ISBN 978-1-59327-828-1 ([Amazon](#), [free online edition](#)).

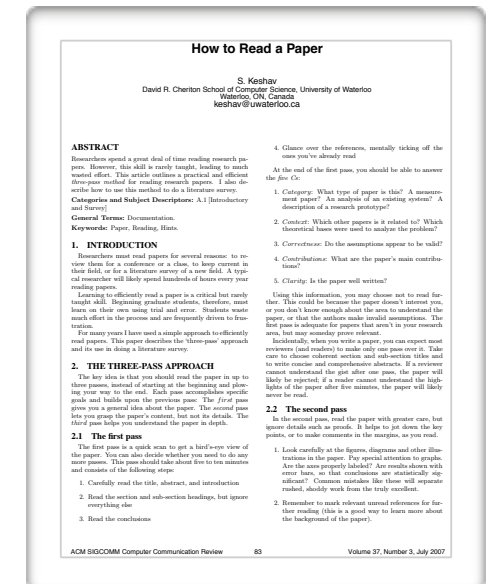


Jim Blandy and Jason Orendorff, "Programming Rust", O'Reilly, 2018, ISBN 978-1-491-92728-1 ([Amazon](#)).

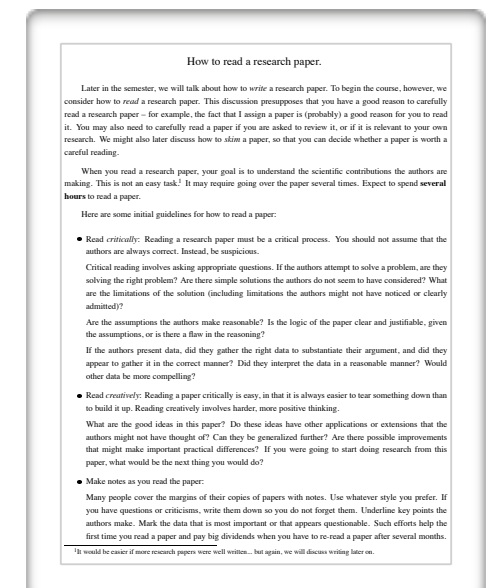
- The Rust Programming language (<https://rust-lang.org/>) will be used to illustrate principles of ownership, memory management, and type-driven programming – read one of the above books
- You are expected to learn the basics of programming in Rust

# Recommended Reading (2/2)

- Research papers will be cited to illustrate some concepts
  - Citations with URLs and/or DOIs are provided
  - Resolve DOIs via <http://dx.doi.org/>
  - All papers are accessible at no cost from the campus networks
- You are expected to read these research papers
  - Critical reading of a research paper requires practice
  - Read in a structured manner, not end-to-end, thinking about the material as you go
    - Focus on the concepts, not the details
  - Realise that research papers are written to explore new ideas
    - Some will be good ideas, some less so
    - Some will be interesting but impractical
    - What's impractical today might be important tomorrow – changes in technology and/or society can change what's feasible/desirable
  - Think and judge for yourself!



S. Keshav, "How to Read a Paper", ACM Computer Communication Review, 37(3), July 2007  
DOI: 10.1145/1273445.1273458



<http://www.eecs.harvard.edu/~michaelm/postscripts/ReadPaper.pdf>

# Systems Programming


- What is systems programming?
- The state of the art
- Challenges and limitations

# 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

# Memory Management and Data Representation

- Predictability
  - Timing of allocations must be bounded for real-time applications
  - Bounds on memory usage
- Data locality
  - Cache line sharing impacts performance
- Data representation
  - Device drivers that must control via fixed layout control registers
  - Network protocol implementations must conform to specified packet layout
  - Ensuring data is aligned and packed into cache lines for high performance
- Systems programming languages offer control of memory management and data representation – others languages lack such controls
  - e.g., stack vs. heap allocation and explicit pointers in C, compared to Java where all objects are allocated on the heap and accessed via implicit references




Smaller values of **stride** indicate data with better spatial locality; **size** is the total amount of data accessed

Source: Bryant & O'Hallaron, "Computer Systems: A Programmer's Perspective", 3rd Edition, Pearson, 2016, Fig. 6.41. <http://csapp.cs.cmu.edu/3e/figures.html> (Permission granted for lecture use with attribution)



# I/O Operations

- Network performance increasingly a bottleneck:
  - Chart shows Ethernet bit rate over time – wireless links follow a similar curve
  - Closely tracking exponential growth over time – unlike CPU speed, which stopped growing significantly mid-2000s
  - MTU remains constant but packet rate increases → fewer cycles to process each packet
- SSD performance on a similar trend for file system access
- I/O performance of systems software critical to overall system performance
  - Device drivers, network protocol stack, file system




# Management of Shared State

- Systems programs responsible for coordinating shared mutable state:
  - State shared across layers/between kernel and applications
    - Data structures for zero-copy networking
    - Header processing and state for the TCP stack
  - State shared between threads
    - Internal state of the kernel
    - File systems
    - Network code
  - Highly performance critical
- Systems programming languages allow sharing data between layers and/or threads – other languages disallow/discourage such sharing

# Performance

- Systems infrastructure performance fundamentally affects overall system and application performance
  - Mobile devices have limited battery life
  - Data centre efficiency and power consumption
- Systems components often the bottleneck in terms of overall performance and power efficiency
  - Simply because they're the basis on which the higher-level systems depend



Source: N. Jones, "The Information Factories", Nature, v.561, p.163–166, Sep. 2018. DOI: [10.1038/d41586-018-06610-y](https://doi.org/10.1038/d41586-018-06610-y)

# Systems Programming

- Systems programming languages offer low-level control of data representation, memory management, I/O, and sharing.
- They are high-performance – concrete rather than high abstraction



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 State of the Art

- Most devices run some variant of Unix as their operating system, programmed in C
  - Original version of Unix written in assembly for PDP-7 minicomputer in 1969
  - Ported to the PDP-11/40 in early 1970s, re-writing into C at that language was developed
    - “The PDP-11/40 was designed to fit a broad range of applications, from small stand alone situations where the computer consists of only 8K of memory and a processor, to large multi-user, multi-task applications requiring up to 124K of addressable memory space. Among its major features are a fast central processor with a choice of floating point and sophisticated memory management, both of which are hardware options.” <https://pdos.csail.mit.edu/6.828/2005/readings/pdp11-40.pdf>
- macOS, iOS, Linux and Android are moderns variants and reimplementations of Unix
- This has proven surprisingly resilient and portable – but is it still the right model?



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

# Unix and C: Strengths

- 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 was due to the C programming language
  - Simple, easy to understand, easy to port to new architectures
  - Explicit pointers, memory allocation, and control over data representation
  - Uniform treatment of memory, device registers, and data structures – easy to write device drivers, network protocols, and interface with external formats
  - Weak type system allows aliasing and sharing

```
struct {
    short errors      : 4;
    short busy        : 1;
    short unit_sel     : 3;
    short done         : 1;
    short irq_enable   : 1;
    short reserved     : 3;
    short dev_func     : 2;
    short dev_enable   : 1;
} ctrl_reg;

int enable_irq(void)
{
    ctrl_reg *r = 0x80000024;
    ctrl_reg tmp;

    tmp = *r;
    if (tmp.busy == 0) {
        tmp.irq_enable = 1;
        *r = tmp;
        return 1;
    }
    return 0;
}
```

Example: hardware access in C

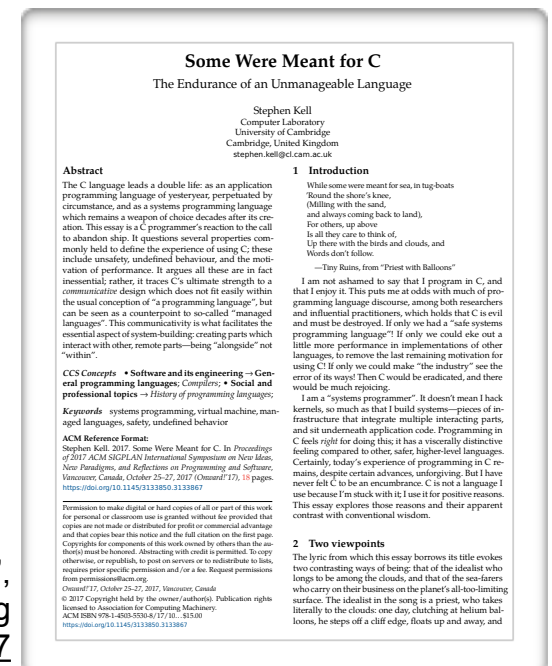
# Unix and C: Weaknesses

- Unix APIs reflect 1970s/1980s minicomputer architectures
  - Sockets and file system APIs significant performance bottlenecks
  - Security architecture insufficiently flexible
  - No portable APIs for mobility, power management, etc.
  - Assumes professional, interactive, systems administration
- C programming language
  - Limited concurrency support → memory model for pthreads poor supported
  - Undefined behaviour, buffer overflows → security risks
  - Weak type system → difficult to reason about correctness, effectively model problem domain

# Unix and C

- Unix has proven surprisingly resilient and portable – but is it still the right model?
  - No – numerous work-arounds for its limitations exist:
    - Kernel bypass networking
    - Increasingly baroque package management
    - Containers and sandboxing
- The C programming language is increasingly a liability
  - Too easy to introduce security vulnerabilities
  - Too easy to trip over undefined behaviour
  - Insufficient abstractions

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)






# Challenges and Limitations


- What changes in the environment are affecting systems programs?
  - The end of Moore's law
  - Increasing concurrency – imposed due to hardware changes
  - Increasing need for security – the Internet
  - Increasing mobility and connectivity

# The End of Moore's Law (1/2): Physical Limits




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/>




- Moore's law: advances in manufacturing double transistor count every two years
- But, rapidly approaching physical limits:
  - 10nm process  $\rightarrow$  features  $\sim 40$  atoms across
  - Transistors *will* stop shrinking soon

# The End of Moore's Law (2/2): Dennard Scaling



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/>




- Dennard scaling: smaller transistors reduce capacitance and voltage → frequency increase without corresponding increase in power consumption
- Scaling relation breaks down eventually, due to leakage → clock frequency increase stalls
- System performance constraints become more acute

Power consumption  $\propto C \cdot F \cdot V^2$


- C = capacitance (transistor size)
- F = frequency (clock rate)
- V = voltage

# Increasing Concurrency




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 limits clock frequency, but Moore's law gives more transistors → used to increase number of cores
- Concurrency related problems become more severe:
  - Ensuring correctness, avoiding race conditions and deadlocks
  - Ensuring good performance

# Increasing Need for Security



- Weaponisation of the Internet
- The combination of C and Unix has not proven easy to secure

# Increasing Mobility and Connectivity

- Computing increasing implies mobile devices
  - Always on – but constrained by limitations of battery power
  - Always connected – by increasingly heterogenous networks
  - Personal – but vulnerable
- Do we have the APIs, tools, and programming models to make effective use of such devices?

# Next Steps in Systems Programming

# 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



# What is a Modern Type System?

- A modern type system is expressive enough to:
  - Provide useful guarantees about program behaviour
    - Prevent buffer overflows, use-after-free bugs, race conditions, iterator invalidation, ...
  - Provide a model of the problem that prevents inconsistencies in the solution, while avoiding run-time overheads
    - No cost abstractions – compile-time checking that has no run-time cost
    - Describe constraints on program behaviour in the types – the compiler as a debugger

```
ACCEPT(2)                                BSD System Calls Manual                                ACCEPT(2)

NAME
    accept -- accept a connection on a socket

SYNOPSIS
    #include <sys/socket.h>

    int accept(int socket, struct sockaddr *restrict address,
               socklen_t *restrict address_len);

DESCRIPTION
    The argument socket is a socket that has been created with socket(2),
    bound to an address with bind(2), and is listening for connections after
    a listen(2). accept() extracts the first connection request on the queue
    of pending connections, creates a new socket with the same properties of
    socket, and allocates a new file descriptor for the socket. If no pend-
```

Common bug in networking code: call `read()` on listening socket, not socket returned from `accept()` that represents the connection

Both file descriptors are represented as type `int` – compiler can't check for such misuse

If *listening socket* and *connected socket* were separate types, and the `read()` call took a connected socket parameter, this bug would be detected at compile-time

Trivial example of important principle – try to describe behaviour in types so compiler can detect logic errors

# What is Functional Programming?

- A programming style that highlights:
  - Pure, referentially transparent, functions
  - No side effects
  - No shared mutable state
  - Control over I/O

with language support for functions as first class types

- Pure functional languages constrain programs
  - Haskell is a testbed for exploring pure functional programming – what are its benefits and costs?
    - Principled, but perhaps impractical for large-scale systems programs
    - Not intended as an industrial-strength platform for deployment
    - The pure functional style is unsuited to some programs, but natural for others

- But concepts are widely applicable
  - Pure functional code is easy to test and debug – no hidden state
  - Pure functional code is thread safe – no side effects or mutable state
  - Eliminating shared mutable state and controlling I/O avoids race conditions
- Use functional programming ideas where they make sense – prevent certain classes of bugs

# Improving Memory Management and Safety

- C has manual memory management
  - Pointers, `malloc()` and `free()`
  - Arrays represented as pointers to their first element, and don't store length
  - Access outside allocated memory “undefined” but no checks to prevent
  - Good reasons for this at the time:
    - Machines were slow and had limited memory
    - Bounds checks and garbage collection too expensive
    - Not all of these are still valid
- These design decisions lead bugs:
  - Use after free, memory leaks, buffer overflows
  - (C++ also suffers from) iterator invalidation
- Modern type systems can eliminate these *classes* of bug
  - Enforce bounds checks
  - Enforce ownership of data – code that tries to use data after it's been freed won't compile; similar for iterator invalidation

```
int foo[5];

// The following are equivalent:
foo[3] = 42;
*(foo+3) = 42;
*(3+foo) = 42;
3[foo] = 42;
```

# Improving Security

**Vulnerabilities By Type**

Year	# of Vulnerabilities	DoS	Code Execution	Overflow	Memory Corruption	Sql Injection	XSS	Directory Traversal	Http Response Splitting	Bypass something	Gain Information	Gain Privileges	CSRF	File Inclusion	# of exploits
<a href="#">1999</a>	894	<a href="#">177</a>	<a href="#">112</a>	<a href="#">172</a>			<a href="#">2</a>	<a href="#">7</a>		<a href="#">25</a>	<a href="#">16</a>	<a href="#">103</a>			<a href="#">2</a>
<a href="#">2000</a>	1020	<a href="#">257</a>	<a href="#">208</a>	<a href="#">206</a>		<a href="#">2</a>	<a href="#">4</a>	<a href="#">20</a>		<a href="#">48</a>	<a href="#">19</a>	<a href="#">139</a>			
<a href="#">2001</a>	1677	<a href="#">403</a>	<a href="#">403</a>	<a href="#">297</a>		<a href="#">7</a>	<a href="#">34</a>	<a href="#">123</a>		<a href="#">83</a>	<a href="#">36</a>	<a href="#">220</a>		<a href="#">2</a>	<a href="#">2</a>
<a href="#">2002</a>	2156	<a href="#">498</a>	<a href="#">553</a>	<a href="#">435</a>	<a href="#">2</a>	<a href="#">41</a>	<a href="#">200</a>	<a href="#">103</a>		<a href="#">127</a>	<a href="#">74</a>	<a href="#">199</a>	<a href="#">2</a>	<a href="#">14</a>	<a href="#">1</a>
<a href="#">2003</a>	1527	<a href="#">381</a>	<a href="#">477</a>	<a href="#">371</a>	<a href="#">2</a>	<a href="#">49</a>	<a href="#">129</a>								
<a href="#">2004</a>	2451	<a href="#">580</a>	<a href="#">614</a>	<a href="#">410</a>	<a href="#">3</a>	<a href="#">148</a>	<a href="#">291</a>								
<a href="#">2005</a>	4935	<a href="#">838</a>	<a href="#">1627</a>	<a href="#">657</a>	<a href="#">21</a>	<a href="#">604</a>	<a href="#">786</a>								
<a href="#">2006</a>	6610	<a href="#">893</a>	<a href="#">2719</a>	<a href="#">663</a>	<a href="#">91</a>	<a href="#">967</a>	<a href="#">1302</a>								
<a href="#">2007</a>	6520	<a href="#">1101</a>	<a href="#">2601</a>	<a href="#">954</a>	<a href="#">95</a>	<a href="#">706</a>	<a href="#">884</a>								
<a href="#">2008</a>	5632	<a href="#">894</a>	<a href="#">2310</a>	<a href="#">699</a>	<a href="#">128</a>	<a href="#">1101</a>	<a href="#">807</a>								
<a href="#">2009</a>	5736	<a href="#">1035</a>	<a href="#">2185</a>	<a href="#">700</a>	<a href="#">188</a>	<a href="#">963</a>	<a href="#">851</a>								
<a href="#">2010</a>	4652	<a href="#">1102</a>	<a href="#">1714</a>	<a href="#">680</a>	<a href="#">342</a>	<a href="#">520</a>	<a href="#">605</a>	<a href="#">275</a>	<a href="#">8</a>	<a href="#">234</a>	<a href="#">282</a>	<a href="#">238</a>	<a href="#">86</a>	<a href="#">73</a>	<a href="#">1493</a>
<a href="#">2011</a>	4155	<a href="#">1221</a>	<a href="#">1334</a>	<a href="#">770</a>	<a href="#">351</a>	<a href="#">294</a>	<a href="#">467</a>	<a href="#">108</a>	<a href="#">7</a>	<a href="#">197</a>	<a href="#">409</a>	<a href="#">206</a>	<a href="#">58</a>	<a href="#">17</a>	<a href="#">557</a>
<a href="#">2012</a>	5297	<a href="#">1425</a>	<a href="#">1459</a>	<a href="#">843</a>	<a href="#">423</a>	<a href="#">243</a>	<a href="#">758</a>	<a href="#">122</a>	<a href="#">13</a>	<a href="#">343</a>	<a href="#">389</a>	<a href="#">250</a>	<a href="#">166</a>	<a href="#">14</a>	<a href="#">624</a>
<a href="#">2013</a>	5191	<a href="#">1455</a>	<a href="#">1186</a>	<a href="#">859</a>	<a href="#">366</a>	<a href="#">156</a>	<a href="#">650</a>	<a href="#">110</a>	<a href="#">7</a>	<a href="#">352</a>	<a href="#">511</a>	<a href="#">274</a>	<a href="#">123</a>	<a href="#">1</a>	<a href="#">205</a>
<a href="#">2014</a>	7946	<a href="#">1598</a>	<a href="#">1574</a>	<a href="#">850</a>	<a href="#">420</a>	<a href="#">305</a>	<a href="#">1105</a>	<a href="#">204</a>	<a href="#">12</a>	<a href="#">457</a>	<a href="#">2104</a>	<a href="#">239</a>	<a href="#">264</a>	<a href="#">2</a>	<a href="#">401</a>
<a href="#">2015</a>	6484	<a href="#">1791</a>	<a href="#">1826</a>	<a href="#">1079</a>	<a href="#">749</a>	<a href="#">218</a>	<a href="#">778</a>	<a href="#">150</a>	<a href="#">12</a>	<a href="#">577</a>	<a href="#">748</a>	<a href="#">367</a>	<a href="#">248</a>	<a href="#">5</a>	<a href="#">127</a>
<a href="#">2016</a>	6447	<a href="#">2028</a>	<a href="#">1494</a>	<a href="#">1325</a>	<a href="#">717</a>	<a href="#">94</a>	<a href="#">497</a>	<a href="#">99</a>	<a href="#">15</a>	<a href="#">444</a>	<a href="#">843</a>	<a href="#">600</a>	<a href="#">87</a>	<a href="#">7</a>	<a href="#">1</a>
<a href="#">2017</a>	14714	<a href="#">3154</a>	<a href="#">3004</a>	<a href="#">2805</a>	<a href="#">745</a>	<a href="#">503</a>	<a href="#">1516</a>	<a href="#">274</a>	<a href="#">11</a>	<a href="#">629</a>	<a href="#">1706</a>	<a href="#">459</a>	<a href="#">327</a>	<a href="#">18</a>	<a href="#">6</a>
<a href="#">2018</a>	16555	<a href="#">1852</a>	<a href="#">3035</a>	<a href="#">2451</a>	<a href="#">400</a>	<a href="#">516</a>	<a href="#">2001</a>	<a href="#">509</a>	<a href="#">11</a>	<a href="#">709</a>	<a href="#">1374</a>	<a href="#">247</a>	<a href="#">461</a>	<a href="#">31</a>	<a href="#">4</a>
<a href="#">2019</a>	4	<a href="#">2</a>				<a href="#">1</a>									
<b>Total</b>	110603	<a href="#">22685</a>	<a href="#">30435</a>	<a href="#">17226</a>	<a href="#">5043</a>	<a href="#">7438</a>	<a href="#">13667</a>	<a href="#">3823</a>	<a href="#">162</a>	<a href="#">5880</a>	<a href="#">10104</a>	<a href="#">4877</a>	<a href="#">2123</a>	<a href="#">2195</a>	<a href="#">4333</a>
<b>% Of All</b>		20.5	27.5	15.6	4.6	6.7	12.4	3.5	0.1	5.3	9.1	4.4	1.9	2.0	


Half of all reported security vulnerabilities are memory safety violations that should be caught by a modern type system – buffer overflows, use-after-free, memory corruption, treating data as executable code

Use of type based modelling of the problem domain can help address others – by more rigorous checking of assumptions

Source: <https://www.cvedetails.com/vulnerabilities-by-types.php>

# Improving Support for Concurrency

- Software increasingly concurrent to make effective use of multicore hardware
- Common abstraction: threads, locks, shared mutable state
- Prone to race conditions:
  - Too many or too few locks held
  - Locks held at the wrong time
  - Locks don't compose
- Program in a function style: avoid races by avoiding shared mutable state
- Avoid races by enforcing single ownership of data – message passing, rather than sharing



# Improving Correctness

- Modern systems programming languages can eliminate certain classes of bug that are common in C
  - Use-after-free, memory leaks, buffer overflows, iterator invalidation
  - Data races in multi-threaded code
  - Don't fix the bug – eliminate the class of bugs
- Modern type systems allow for better modelling of the problem domain, and so more checking of code for consistency
  - Define types representing the problem domain, rather than using generic types – e.g., if you pass around `PersonId` and `VehicleId` rather than `int`, the compiler will warn if you pass the wrong type of identifier to a function
  - Represent program states in the types – e.g., `ListeningSocket` vs. `ConnectedSocket`
  - Modern languages allow you to define types and abstractions easily and without run-time cost – type-first design allows code that is correct by construction
- Use the compiler to debug your design

# Next Steps in Systems Programming

- People can't manage the complexity – need better tooling to help
  - C and Unix solve many systems programming problems
  - Control over data representation, memory management, sharing of state
  - Emerging languages and systems give the same degree of control – with added safety
- Strongly typed languages
- Types help model the problem domain, structure code
- Types and associated tooling help detect logic errors early – correct by construction
- This course will explore these ideas using the Rust programming language

# Summary

- What is systems programming?
- The state of the art
- Challenges and limitations
- Next steps in systems programming