

# Systems Programming and Alternative Operating Systems

## Advanced Operating Systems Tutorial 1

# Goals of Tutorials

- Level M course → assessing critical thinking skills; ability to read research papers, extract key insights
- Tutorials intended to facilitate this:
  - To provide space to discuss the Further Reading highlighted at the end of the lectures in the previous week, to consolidate learning, and emphasise key points of the material
  - You are expected to have read the highlighted papers, and to come to the tutorial prepared to discuss the material
  - Write your own summaries of the papers: what are the key concepts and ideas? what isn't clear? what's unimportant detail?
  - Discuss material that isn't clear in the tutorials → you're not expected to understand everything in the papers

# Reading Research Papers

- Was everyone able to access the papers?
- Experiences with the process of reading a research paper?
  - Critical reading of a research paper is difficult and requires practice; read in a structured manner, not end-to-end, think about the material as you go
  - Advice on paper reading: <http://www.eecs.harvard.edu/~michaelm/postscripts/ReadPaper.pdf>
  - S. Keshav, “How to Read a Paper”, ACM Computer Communication Review, 37(3), July 2007 DOI: [10.1145/1273445.1273458](https://doi.org/10.1145/1273445.1273458)
- Did everyone take notes? Come with questions to discuss?

# Discussion of Papers

- J. Shapiro, “Programming language challenges in systems codes: why systems programmers still use C, and what to do about it”, Proc. PLOS 2006, San Jose, CA, Oct. 2006. DOI:10.1145/1215995.1216004
- Systems programming: constrained memory, I/O performance, data representation, state matters
- Fallacies: factors of 1.5–2 don’t matter; boxed representation can be optimised; the optimiser can fix it; legacy issues insurmountable
- Suggests: annotating code to check application constraints
- Suggests: manual but automatically checked storage management; explicit control over data representation
- The BitC project wasn’t a success, but are the ideas valid?

## Programming Language Challenges in Systems Codes Why Systems Programmers Still Use C, and What to Do About It

Jonathan Shapiro, Ph.D.  
Systems Research Laboratory  
Department of Computer Science  
Johns Hopkins University  
shap@cs.jhu.edu

### Abstract

There have been major advances in programming languages over the last 20 years. Given this, it seems appropriate to ask why systems programmers continue to largely ignore these languages. What are the deficiencies in the eyes of the systems programmers? How have the efforts of the programming language community been misdirected (from their perspective)? What can/should the PL community do address this?

As someone whose research straddles these areas, I was asked to give a talk at this year's PLOS workshop. What follows are my thoughts on this subject, which may or not represent those of other systems programmers.

### 1. Introduction

Modern programming languages such as ML [16] or Haskell [17] provide newer, stronger, and more expressive type systems than systems programming languages such as C [15, 13] or Ada [12]. Why have they been of so little interest to systems developers, and what can/should we do about it?

As the primary author of the EROS system [18] and its successor Coyotos [20], both of which are high-performance microkernels, it seems fair to characterize myself primarily as a hardware systems programmer and security architect. However, there are skeletons in my closet. In the mid-1980s, my group at Bell Labs developed one of the first large commercial C++ applications — perhaps *the* first. My early involvement with C++ includes the first book on reusable C++ programming [21], which is either not well known or has been graciously disregarded by my colleagues.

In this audience I am tempted to plead for mercy on the grounds of youth and ignorance, but having been an active

advocate of C++ for so long this entails a certain degree of *chutzpah*.<sup>1</sup> There is hope. Microkernel developers seem to have abandoned C++ in favor of C. The book is out of print in most countries, and no longer encourages deviant coding practices among susceptible young programmers.

**A Word About BitC** Brewer *et al.*'s cry that *Thirty Years is Long Enough* [6] resonates. It really is a bit disturbing that we are still trying to use a high-level assembly language created in the early 1970s for critical production code 35 years later. But Brewer's lament begs the question: why has no viable replacement for C emerged from the programming languages community? In trying to answer this, my group at Johns Hopkins has started work on a new programming language: BitC. In talking about this work, we have encountered a curious blindness from the PL community.

We are often asked “Why are you building BitC?” The tacit assumption seems to be that if there is nothing fundamentally new in the language it isn't interesting. The BitC goal isn't to invent a new language or any new language concepts. It is to integrate existing concepts with advances in prover technology, and reify them in a language that allows us to build stateful low-level systems codes that we can reason about in varying measure using automated tools. The feeling seems to be that everything we are doing is straightforward (read: uninteresting). Would that it were so.

Systems programming — and BitC — are fundamentally about engineering rather than programming languages. In the 1980s, when compiler writers still struggled with inadequate machine resources, engineering considerations were respected criteria for language and compiler design, and a sense of “transparency” was still regarded as important.<sup>2</sup> By the time I left the PL community in 1990, respect for engineering and pragmatics was fast fading, and today it is all but gone. The concrete syntax of Standard ML [16] and Haskell [17] are every bit as bad as C++. It is a curious measure of the programming language community that nobody cares. In our pursuit of type theory and semantics,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.  
PLOS 2006, Oct. 22, 2006, San Jose, California, United States  
Copyright © 2006 ACM 1-59593-577-0/10/2006...\$5.00

<sup>1</sup> *Chutzpah* is best defined by example. *Chutzpah* is when a person murders both of their parents and then asks the court for mercy on the grounds that they are an orphan.

<sup>2</sup> By “transparent,” I mean implementations in which the programmer has a relatively direct understanding of machine-level behavior.

# Discussion of Papers

- G. Hunt and J. Larus. “Singularity: Rethinking the software stack”, ACM SIGOPS OS Review, 41(2), April 2007. DOI:10.1145/1243418.1243424
- Use of strongly-typed languages to build an operating system; software isolated processes; message passing – is this a sound basis for the system?
- Type-safe message passing through channels; checked state machines for communication protocols (e.g., to control device driver state) – useful tool to help ensure correctness, or over-complex and stifling?
- Small unsafe microkernel, with type-safe system layered above – can the microkernel be written in a safe language?
- Threads and exchange heap; garbage collection – overheads?
- Is the idea of running everything in a virtual machine reasonable?

