

Statement of Research Interests

Eddie Kohler

Many systems today have poor extensibility, flexibility, correctness, and security, all of which are easier to fix when a system is easy to understand. Understanding a real, complex software system should be as easy as reading a page or two of code. This would require high-level comprehension of the system; however, we design systems with architectures that don't facilitate higher-level comprehension, and the tools we use to program, analyze, test, and debug them don't help either. I want to create new architectures and tools that will facilitate higher-level comprehension and test them on real-world systems. I want to build systems that satisfy tough performance requirements, but are still flexible and understandable. This work will combine my interests in systems, programming languages, and software design with a willingness to rethink the way systems should be built.

As a first approach, I want to treat systems as programming languages, an avenue inspired by my graduate school research. In this model, a new programming language is an integral part of each system. This language specifies system-level properties, rather than the data layout and control flow issues controlled by languages like C and C++. Its constructs correspond directly to high-level properties of the system, which may be different even for two systems in the same domain. (For example, an operating system language might specify how file systems, network connections, and other kernel-level objects interact. Different operating systems would probably have different languages.) A program in the language gives us a readable, high-level description of the system, and we can manipulate this program with language processors to optimize or statically check the system as a whole.

The first of two projects that inspired this approach is Prolac, an object-oriented language for network protocol implementation. Prolac was designed to make protocol implementations readable, modular, extensible, and efficient. It includes advanced language features, appropriate syntax, tight integration with unforgiving environments (Unix kernels), and compiler optimizations. Our Prolac TCP implementation is organized in a new, readable and modular way; for example, each TCP protocol extension is readably located in one small set of subclasses. The compiler optimizes most of this organization away, leaving code that performs comparably to commodity TCPs.

The Prolac language is low-level and detailed, however, which complicates adoption. Most people don't want or need to learn another systems programming language. Therefore, in my next project—the Click modular router—we kept the programming language higher level, specifying the interaction between components rather than the instruction-by-instruction behavior of each component.

Click is an architecture for building software routers from modular components called *elements*. Element definitions are written in C++; the Click programming language specifies how elements should be connected in a particular router. A router's high-level behavior is thus exactly determined by its definition in the Click language. Network administrators can create arbitrary feature combinations

by manipulating easy-to-read Click-language files, rather than by hacking kernel source code.

Like formal specifications, the Click language can be used to detect errors statically and to discover global router properties. Unlike formal specifications, the language compiles into working routers, so people can test and modify the way their packets are routed by manipulating Click-language files. Programs analogous to conventional language processors can manipulate these files as well, to optimize them, analyze them, and so forth. We have written several such programs already: a pattern-based optimizer that replaces common element arrangements with faster equivalents; an aligner, which examines a configuration and adds elements to ensure that packet memory is correctly aligned; and a specializer, which creates new, faster elements tailored to how the elements in a configuration are actually being used.

We have implemented the Click architecture on conventional PC hardware with good results. Click routers are flexible—their configurations can be easily read, understood, and changed. Already, performance is good enough for many demanding applications. Furthermore, having the Click language made the core system better by guiding us towards better internal structures and element designs.

Now I want to expand from these lessons into other systems. While my previous work has been in networking, my interest in systems is omnivorous. Networking, operating systems, and distributed applications—particularly over networks of small devices—all seem like fertile ground for this research. Regardless of area, I want to build real systems; I believe you learn from interacting with complexity, and little is more complex than the real world. These new systems probably won't be much like Click or Prolac, because they will be tailored for different problems. Nevertheless, successful system languages may share some basic principles, which I would like to pin down.

This work is fundamentally interdisciplinary in nature. Building a successful system language requires expertise in languages, systems in general, and the particular kind of system being built. I look forward to working in many system areas through close collaboration with peers and, especially, students.