

UK Systems Research Challenges 2020 Workshop:

JIT-as-a-Service

Tom Spink

tspink@inf.ed.ac.uk

Senior Researcher at the *University of Edinburgh*

It is widely observed that dynamic programming languages are gaining ever more popularity. For example, the TIOBE programming community index identifies that 6 out of 10 of the most popular programming languages in use today are dynamic languages. Typically, these languages rely on just-in-time (JIT) compilation for their efficient implementation.

While programmers enjoy the ease-of-use and productivity gains offered by dynamic programming languages, their widespread use has created a new, serious problem: the development and performance of fast and optimising JIT compilers has not kept up with the pace and amount of performance-critical code developed in dynamic languages. Many users are frustrated with the poor runtime performance of their code, and dynamic programming languages like Java and Python have gained a reputation for being slow. This phenomenon is aggravated by the fact that there is not just a single JIT compiler in use, but typically each programming language has its own JIT engine, which needs to be continually maintained, improved, and ported to the many processor architectures (Intel x86, ARM, RISC-V, ...) in use today, further stretching scarce and valuable compiler developer resources.

In this talk, I will describe how we are *rethinking the foundations* of how we provide JIT compilation to clients.

We want to liberate JIT compilation from the shackles of device constraints, e.g. limited processing speed, memory or energy in mobile devices. Instead, JIT compilation can be done anywhere in the cloud—from edge to data center. This enables entirely new opportunities for JIT compilers, for example, by offloading JIT compilation, code caching and sharing already compiled code between users and devices, or through the use of hardware accelerators deployed in edge or data centre servers and specifically designed to deliver highest code quality at low latency/power consumption.

Central to this, we firmly believe that high-bandwidth, low-latency 5G communication is a *game changer* not only for the mobile phone market, but its impact will directly affect most areas of computing. By 2020, 50 billion smart devices are expected to be in use. 5G will help support the massive growth in the Internet-of-Things and enable devices to communicate with each other seamlessly through the convergence of mobile communications and computing.

The key to *unlock* this 5G JIT compilation revolution is to investigate *JIT compilation as a service*, i.e. a JIT compiler is no longer a deeply integrated module in a particular programming language runtime, but it will become a transparent cloud hosted microservice that can be shared with and accessed by many clients. Cost and effort of this JIT compilation service is amortised across users, programming languages and architectures, making JIT-as-a-Service a cost-effective proposition in response to a changing IT landscape.