Distributed Protocol Offload using FPGA SmartNICs

(PI) Shuai Mu[1], Mike Ferdman[1], Peter Milder[2]
[1]Department of Computer Science, [2]Department of Electrical and Computer Engineering

Summary

Distributed systems provide high performance and high availability in the presence of system failures, scale to large datasets that exceed the capacity of a single node, and often simultaneously achieve both of these goals. However, achieving high performance with pure software implementations of distributed systems presents significant challenges, as many small network packets must hop around between the nodes for each request, resulting in low throughput and high latency, especially at the tail of the latency distribution.

This proposal will explore Distributed Protocol Offload (DPO) for SmartNICs. DPO is our novel approach to hardware-acceleration of distributed systems, a mechanism that will retain the relative simplicity, flexibility, and maintainability of software implementations, while achieving extremely high throughput and bare-minimum operation latency. Rather than offloading the entire network stack and distributed protocol to the SmartNIC in a way that mimics existing software, DPO will partition the distributed protocol handling into two parts, separating out the simple common cases with minimal functionality and no error handling, while leaving the full network stack, complex protocol implementation, packet re-transmission, and error handling to the tried-and-true robust software. Our goal for DPO is to serve as a demonstration of a simple, generic, flexible, and highly performant approach to hardware acceleration of distributed protocols.