

Enabling Efficient Packet Processing



by Ivo Bolsens
CTO
Xilinx, Inc.
ivo.bolsens@xilinx.com

I recently attended an intriguing Clean Slate workshop at Stanford University. The topic focused on the idea of reinventing the Internet infrastructure in order to solve some of its shortcomings and enable new applications and services (<http://cleanslate.stanford.edu>).

Driven by an increase in rich media applications such as video, service convergence, and broadband access deployment, global IP traffic is showing explosive growth. Furthermore, IP data network traffic is highly dynamic, with constantly changing demands and substantially different characteristics when compared to telephony voice.

Today, IP backbones are designed to have core routers directly connected through static circuits. This assumption of static links means that links must be grossly over-designed to deal with peak traffic uncertainties and link failures. As IP traffic grows, such a network design implies bigger links and bigger routers – neither of which can be scaled cost-effectively. Moreover, the reliability and security of the networking infrastructure to support a growing demand of dependable applications becomes a serious challenge.

Network Virtualization

The Clean Slate researchers agreed that “network virtualization” would become an essential technology to provide future value-added services and robustness.

Network virtualization is a capability that allows you to create multiple yet very different concurrent networks on top of the same physical infrastructure of diverse resources, and would imply a functionality

A clean-slate approach to reinventing the Internet.



that extends beyond current packet routing and forwarding.

Today, traditional networking features nodes with fixed architectures. Flexibility in networks comes from packet switching itself – guided by look-up tables – plus the relatively slow evolution of protocol development and node software updates. To support desirable concepts like virtualization and technological convergence, it will be necessary to be much more flexible in terms of node architectures and protocol definition. Therefore, data processing (whether packetized or something else) may be carried out using very different styles, depending on widely differing applications. The protocols providing the rules for such processing will also vary widely and require much more dynamic definition.

From a Xilinx® perspective, FPGA technology is experiencing a parallel shift that complements the future direction of networking. Many people see FPGAs as ASIC substitutes; once programmed, the devices have the characteristics of fixed hardware. Today, however, FPGAs are providing a basis for delivering flexible, “soft” processing architectures that can be matched to particular problem instances.

For example, you can essentially build your own network processing unit

(NPU) every time rather than forcing problems onto fixed architectural solutions such as an ASSP or NPU (be it a connection of multi-threaded engines or deeply pipelined parallel data paths). Moreover, the soft architecture can evolve over time in response to changes in demand and usage.

Today’s FPGA architectures enable the creation of soft packet processing architectures that, depending on application functionality, can be configured as highly pipelined parallel data paths or networks of multi-threaded micro-engines.

These soft processing architectures allow the memory architecture, the connection of the different packet processing functions, and the packet operational units to be tailored to the application at hand. This results in the implementation of packet processing functions in FPGAs that are outperforming NPUs and ASSPs in throughput, power consumption, and cost by at least one order of magnitude.

An important component is the availability of programming tools that harness the highly concurrent capabilities of the FPGA and move away from traditional hardware design paradigms toward new, flexible protocol implementation paradigms.