



Network Processor Designs for Next-Generation Networking Equipment

White Paper

This white paper provides an overview of network processor technology and explains why RISC-based architectures are insufficient for building next-generation switches and routers. EZchip's TOPcore™ technology will be shown to enable vendors build LAN switches based on network processors optimized for 7-layer packet processing at multi-Gigabit wire speed.

ABSTRACT

The demand for intelligent processing at wire speed has led to the creation of network processors (also called communications processors). Programmable network processors provide system flexibility while delivering the high-performance hardware functions required to process packets at wire speed. While network processors are expected to become the silicon core of next-generation networking equipment, they differ greatly in their architectural design. This white paper examines network processor designs and assesses their ability to build next-generation switches and routers to deliver intelligent processing at Gigabit and Terabit speeds.

NETWORK PROCESSOR DESIGNS

The quest for speed continues to drive the networking industry. Equipment is expected to perform at faster speeds and vendors are racing to bring new products to the market. For this reason, hardware vendors are turning to network processors to support increasingly complex tasks at wire speed, while providing a migration path to the future.

Network processor designs can be divided into three main architecture types: (a) a general RISC-based architecture, (b) an augmented RISC architecture (with hardware accelerators), and (c) network-specific processors. The first two architectures are sufficient for building today's Fast Ethernet products, however they will be unable to provide full 7-layer processing on more than handful of ports at Gigabit speed.

RISC-based Network Processors

Several vendors boast network processor designs based on the integration of multiple "off the shelf" RISC processors into a single chip. By definition, a RISC (Reduced Instruction Set Circuits) is a computer architecture that reduces chip complexity by using simpler instructions than CISC (Complex Instruction Set Circuits) computers.

In a RISC, the microcode layer and associated overhead is eliminated. A RISC maintains a constant instruction size, bans the indirect addressing mode and retains only those instructions that can be overlapped and made to execute in one machine cycle. The RISC machine executes instructions quickly because it does not go through a microcode conversion layer. Yet, the RISC compiler has to generate routines using simple instructions. Performing complex tasks requires many commands, each of which takes a clock.

Drawbacks to RISC-based network processors include their use of numerous commands, the time it takes to perform complex tasks and their inability to modify the data path. For this reason, most RISC-based network processors will be unable to deliver processing performance on more than handful of Gigabit ports.

RISC processors are frequently deployed in parallel to produce high speeds, but this architecture is still constrained by the RISC throughput. Moreover, there is a limit to the number of RISCs that can be incorporated without overly increasing system complexity and the size of the chip.

Augmented RISC-based Network Processors

Tailoring the RISC to networking functions and adding hardware accelerators (boosters) is an alternative approach that speeds processing. Hardware accelerators can copy frames at wire speed to boost performance, but the accelerator itself is neither flexible nor programmable. Again, this approach runs into the limitations previously described for RISC-based network processors.

Some chip vendors employ a combined RISC and ASIC (Application Specific Integrated Circuit) architecture. With this approach, the RISC acts as the core processor and certain tasks are offloaded to ASICs. Hard-wired ASICs provide the speed, but are severely restricted by their inherent inflexibility. ASICs are limited to the functionality embedded in their silicon and cannot be updated to support new features.

Network-Specific Processors

A new wave of network processors, namely network-specific processors, is now being developed to provide the processing performance required for next-generation networking products. Network-specific processors integrate many small, fast processor cores that are each tailored to perform a specific networking task. By optimizing the individual processor cores for packet-processing tasks, network-specific processors overcome the limitations of RISC-based architecture. Network-specific processors can deliver the packet processing performance to handle an appreciable number of ports at Gigabit and Terabit speeds.

With network-specific-processors, exceptionally fast packet processing at high bandwidths is achieved through optimization of both the instruction set and data path. Since each task-oriented core is designed with a specific networking function in mind, it uses a concise instruction set to accomplish the task. It may require as few as 1/10 the number of commands used by a RISC-based processor to accomplish same task.

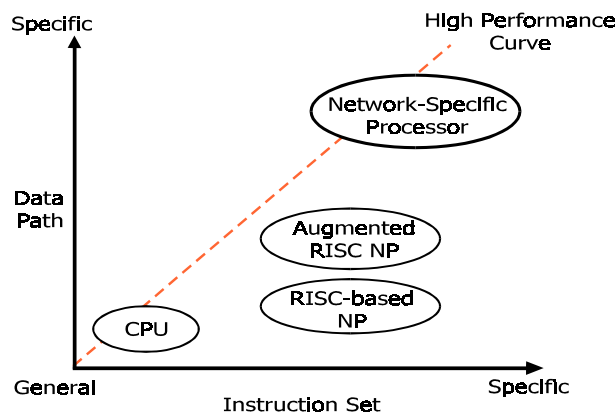


Figure 1. Network processor performance is influenced by a combination of the instruction set and the data path. Higher performance is achieved with fewer and more specific instructions and a shortening of the data path.

PACKET PROCESSING TASKS

Since all packet processing involves four basic tasks – parse, search, resolve and modify – network-specific processors optimize each of these to boost the processing performance.

- Parse – Analyzes and classifies the contents of the packet header and fields.
- Search – Tables are searched for a match between the content that was classified and pre-defined content and rules.
- Resolve – The destination and QoS requirements are resolved and the packet is routed to its destination.
- Modify – Where necessary, the packet is modified (e.g. certain fields within the packet are changed).




Figure 2. The four packet processing tasks in the packet flow.

Significant improvement in performance is obtained by optimization for each task. Digital Signal Processors (DSPs) are a great example of the enhancement achieved through task optimization. DSPs can be viewed as conventional RISC processors with special hardware and instructions for efficiently processing digital signal data. DSPs offer improved performance as compared to a RISC for DSP algorithms that involve fast multiplication of matrixes.

TASK OPTIMIZED PROCESSING CORE (TOPCORE™) TECHNOLOGY

EZchip introduces TOPcore, an innovative technology for application-level switching while maintaining Gigabit and Terabit speeds as a foundation for network-specific processors. Network-specific processors designed with EZchip’s **Task Optimized Processing Core** technology*, abbreviated **TOPcore**, will achieve 10-fold improvements in performance over even advanced RISC-based network processors. TOPcore realizes these performance gains by using a customized instruction set and data path for each packet-processing task. Consequently, the number of clock cycles that are required for complex packet manipulation is minimized; translating into more data processing per clock cycle than RISCs.

* EZchip’s TOPcore™ technology is to be patented.

The TOPcore architecture consists of an array of task-optimized processors (TOPs) each with a customized instruction set and data path, optimized for its specific networking task. Four types of TOPs are featured: TOPparse, TOPsearch, TOPresolve and TOPmodify. Each is tailored to perform its respective function; packet parse and classify, search, resolve and modify.




Figure 3. TOPcore technology customizes the instruction set and data path for each packet-processing task.

TOPparse identifies and extracts the various headers and fields within the packet. It handles all seven layers including fields with dynamic offsets and length.

TOPsearch performs the various table look-ups required for Layer 2 switching, Layer 3 routing, Layer 4 session switching, and Layers 5-7 content switching and policy enforcement. Custom tables are available to search according to any criteria for effective groupings and generalization of conditions. Special support is provided to enable wire speed performance of Layer 5-7 processing such as text strings, which are often very long and of varying sizes (e.g. URLs).

To enable these unprecedented search capabilities, TOPsearch uses a variety of search algorithms, optimized for various searched objects and properties. These algorithms feature innovative enhancements to hash tables, trees and CAMs. Multiple searches using different search methods can be applied simultaneously to yield wire speed throughput, even when applying highly complex network policies.

TOPresolve assigns the packet to its appropriate output port and queue. It forwards the packet to multiple ports for multicast applications. TOPresolve also gathers traffic accounting information on a per flow basis, for network usage analysis and billing.

TOPmodify modifies packet contents in accordance with the results of the previous stages. It modifies relevant fields, e.g. VLAN assignments, Network Address Translation (NAT), QoS priority setting and more.

Operation of all the TOP processors is controlled by a set of software commands, downloaded from the system's host processor. Any change in network policy, e.g. user access, application priority, URL switching criteria, is deployed simply by downloading updated code to the chip. This programmability offers the flexibility to adjust to new intranet and Internet applications through simple changes in software without necessitating changes to system hardware.

REAL PACKET PROCESSING EXAMPLES

The following examples compare the typical number of clock cycles (where each instruction takes a clock cycle) required for the identical packet-processing task using TOPcore technology and a general RISC processor. Each clearly demonstrates that task-optimized processors offer superior speed and performance.

Table 1. Comparison of packet processing tasks performed with EZchip's TOPcore technology versus a general RISC processor**.

	Packet Processing Task	Clock Cycles	
		TOPcore	RISC
1	Parsing a URL in an HTTP/RTSP packet	60	400
2	Searching URL tables	6	200
3	Resolving a multicast routing decision	8	80

Example 1. Parsing a packet. The table above lists the number of clock cycles that are required to parse a typical HTTP or RTSP packet and determine the URL. Using EZchip's TOPparse it takes 60 clock cycles, regardless of the length of the URL, as compared to 400 instructions with a RISC processor (for a URL of 32 characters). With a RISC processor, the longer the URL, the more clock cycles it takes.

Example 2. Searching URL tables. The table above lists the number of clock cycles that are required to lookup the URL in a typical HTTP or RTSP packet in the URL tables. Using EZchip's TOPsearch it takes up to 6 clock cycles, regardless of the length of the URL, as compared to 200 clock cycles with a RISC processor. Longer URLs take the RISC even longer to search.

Example 3. Resolving a multicast routing decision. The table above lists the number of clock cycles that are required to resolve a multicast routing decision based on PIM-SM (Protocol Independent Multicast - Sparse Mode). This type of packet is frequently used for sending multi-client frames across the Internet, such as Voice over IP (VOIP) and videoconferencing. Using EZchip's TOPresolve it takes no more than 8 clock cycles as compared to 80 clock cycles with a RISC processor.

** The actual instructions used in calculating these examples may be obtained under NDA.

SUPERPIPELINE AND SUPERSCALAR ARCHITECTURE

For increased processing power, the task-optimized cores are employed in both a superpipeline and superscalar architecture. The packet processing tasks are pipelined, passing packets from TOPparse to TOPsearch to TOPresolve to TOPmodify.

Superpipelining provides for the rapid execution of the tasks without delay and has the advantage of being scalable. Multiple instruction pipelines are then implemented in a superscalar architecture to execute several different instructions concurrently during a single cycle. The superpipelining and superscalarity of the TOPcore architecture provide the network-specific processor with massive processing power.




Figure 4. TOPcore technology uses an array of processors in a superpipeline and superscalar architecture for massive processing power.

HIGH-BANDWIDTH MEMORY ARCHITECTURE

The previous sections have outlined how the TOPcore technology, with its optimization for tasks, superpipelining and superscalar architecture, has greatly increased the speed of packet processing. Yet in such a high-speed system, overall system performance would slow without a high-bandwidth memory architecture. For this reason, all memory needs of the network-specific processor are embedded on-chip. This provides superior performance that is unmatched by external memory architectures.

BENEFITS OF TOPCORE TECHNOLOGY

The task of providing 7-layer packet processing at wire speeds for next-generation equipment presents huge challenges.

EZchip's TOPcore technology has the following advantages:

- Ten-fold faster performance through modification of instruction sets and data paths.
- Intelligent 7-layer packet processing at Gigabit and Terabit wire speed.
- Scalable and expandable architecture.
- Flexible and programmable processing.

EZchip Technologies is designing a line of network-specific processors based on its TOPcore technology. The EZchip network-specific processor is a single chip IC that combines an array of numerous TOP processors optimized for 7-layer switching at 10 Gigabit speeds. EZchip's network-specific processors can be applied to high-speed switches and routers at both the backbone and the edge.

SUMMARY

As network quality of service and policy enforcement increase in importance, packet-processing functions must probe deeper into the OSI layers. In order to attain high-performance 7-layer switching, next-generation switches will rely on intelligent network processors. Network-specific processors based on EZchip's breakthrough Task Optimized Processing Core (TOPcore) technology will offer equipment vendors improved performance as a clear advantage over RISC-based architectures. By harnessing the strength and flexibility of EZchip's network-specific processor, networking vendors can build scalable, flexible, intelligent products that perform 7-layer packet processing at Gigabit wire speed and beyond.

- TOPcore technology provides more processing power per clock than other network processor designs.
- Superpipeline and superscalar architecture of the processor array boosts performance.
- Embedded memory speeds the memory access time.

EZchip Technologies, Ramat Gabriel Industrial Park, PO Box 184, 10551 Migdal Haemek, Israel
Tel: +972-6-644-9966, Fax: +972-6-654-0124, Email: info@ezchip.com, Web: www.ezchip.com