

RZ 3479 (# 99475) 07/21/03  
Computer Science 11 pages

# Research Report

## Towards High-performance Active Networking

Lukas Ruf<sup>1</sup>, Roman Pletka<sup>2</sup>, Pascal Erni<sup>1</sup>, Patrick Droz<sup>2</sup> and Bernhard Plattner<sup>1</sup>

<sup>1</sup>Computer Engineering and Networks Laboratory  
Swiss Federal Institute of Technology (ETH)  
CH-8092 Zurich  
Switzerland  
{ruf, plattner}@tik.ee.ethz.ch  
pascal@promethos.org

<sup>2</sup>IBM Research  
Zurich Research Laboratory  
8803 Rüschlikon  
Switzerland  
{rap, dro}@zurich.ibm.com

### LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties). Some reports are available at <http://domino.watson.ibm.com/library/Cyberdig.nsf/home>.

**IBM Research**  
Almaden · Austin · Beijing · Delhi · Haifa · T.J. Watson · Tokyo · Zurich

# Towards High-performance Active Networking

Lukas Ruf<sup>1</sup>, Roman Pletka<sup>2</sup>, Pascal Erni<sup>3</sup>, Patrick Droz<sup>2</sup>, Bernhard Plattner<sup>1</sup>

<sup>1</sup> Computer Engineering and Networks Laboratory  
Swiss Federal Institute of Technology (ETH)  
CH-8092 Zürich/Switzerland

{ruf,plattner}@tik.ee.ethz.ch

<sup>2</sup> IBM Zurich Research Laboratory  
Säumerstrasse 4

CH-8803 Rüschlikon/Switzerland

{rap,droz}@zurich.ibm.com

<sup>3</sup> pascal@promethos.org \*

**Abstract.** Network processors have been developed to ease the implementation of new network protocols in high-speed routers. Being embedded in network interface cards, they enable extended packet processing at link speed as is required, for instance, for active network nodes. Active network nodes start using network processors for extended packet processing close to the link. The control and configuration of high-performance active network nodes with network processors such that new services can benefit from the additional processing capacity offered is nontrivial. In this paper, we present PromethOS NP which is a modular and flexible router architecture that provides a framework for dynamic service extension by plugins with integrated support of network processors, namely the IBM PowerNP 4GS3 network processor. We briefly introduce the PowerNP architecture in order to show how our active networking framework maps onto this network processor and provide results from performance measurements. Owing to architectural similarities of network processors, we believe that our considerations are also valid for other network processors.

## 1 Introduction & Motivation

Network processors (NPs) have been developed to ease the implementation of new networking functionalities and services in high-speed routers. The programmable environments provided by processor manufacturers remove the burden of creating application-specific integrated circuits (ASICs) or other hardware components needed for extended or new services. Hence, NPs combine the high performance known from ASICs with the capability to adapt networking functionalities in software, while not requiring expensive modifications in hardware.

---

\* This work is partially sponsored by the Swiss Federal Institute of Technology (ETH) Zürich and the Swiss Federal Office for Education and Science (BBW Grant 99.0533). We would like to acknowledge the great support and funding received by the IBM Zurich Research Laboratory. PromethOS v1 has been developed by ETH as a partner in European Union's Project FAIN (IST-1999-10561).

Modern high-performance active network nodes (hANNs) are built by a set of host CPUs and a set of network interface cards (NICs) populated with on-board NPs. The host CPUs provide extended processing capacity mainly for control-point (CP)<sup>4</sup> functionalities while the NPs carry out extended packet processing for one or more network interfaces. Such an hANN provides three levels of different processing environments (PEs) in which service components can be installed. To benefit from the additional processing capacity, code installation and configuration at run-time require the dynamic creation and configuration of execution environments (EE) on each PE involved.

A framework is essential for the creation and configuration of EEs, code installation, instantiation and configuration as well as the communication among components within a service, as the control and management tasks of an hANN are nontrivial. PromethOS NP provides a framework that copes with the complexity of such an hANN. It is based on an extended version of PromethOS [7], a Linux kernel-space-based NodeOS providing the PromethOS EE. The current implementation of PromethOS NP controls an hANN on an Application Reference Board (ARB) [11], including the IBM PowerNP 4GS3 network processor [4].

In this paper, we present the fundamental design considerations and the architecture of PromethOS NP in Section 2 first. Subsequently, we give a brief introduction to the IBM PowerNP 4GS3 and the ARB in Section 3 and provide further implementation details. Our implementation is then evaluated by performance measurements and the results are presented in Section 4. In Section 5 we review related work, before we summarize and conclude our paper (Section 6).

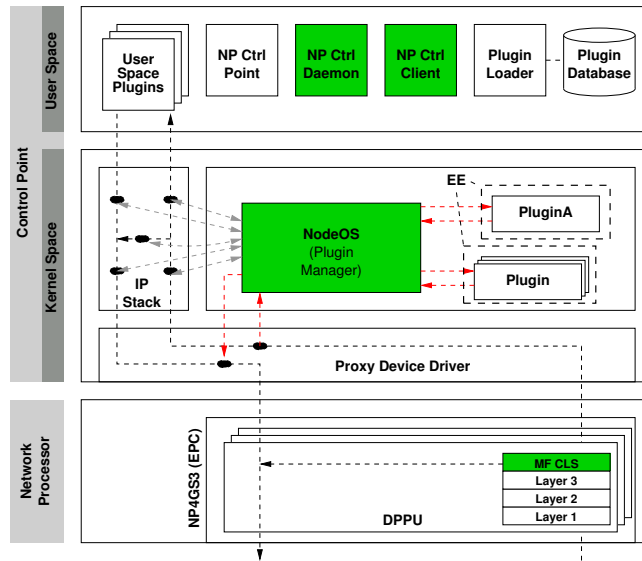
## 2 PromethOS NP

The PromethOS NP framework controls an hANN with NICs that provide NPs for extended packet processing. It is composed of management applications and the PromethOS NodeOS as well as the PromethOS EE. Figure 1 provides an overview of the main components of a PromethOS NP node:

- **Management applications:** The management applications control the NodeOS. Further, they initiate component installation and service configuration. They are implemented by the NP Control Daemon (NP CtrlD), the NP Control Client (NP Ctrl).
- **NodeOS:** The PromethOS NodeOS functionality is provided mainly by the PromethOS plugin manager, which is responsible for the creation, configuration and control of the PromethOS EE. It attaches to the legacy hooks of the IP stack and to the fast-path of the proxy device driver.
- **EE:** The PromethOS EE follows the plugin paradigm, in which plugins are organized as a directed graph of modules.
- **Plugins:** Code components installed in the PromethOS EE are called PromethOS plugins. They provide the service functionality. PromethOS plugins are installed on packet processor engines of the network processor if they provide the packet classification for the current implementation.

---

<sup>4</sup> The CP is responsible for node configuration: it processes management protocols and configures the NP.



**Fig. 1.** PromethOS NP: Architectural Overview

A PromethOS NP node spans all processing environments: by design, PromethOS EEs are located on all three levels, thus providing environments for active service components.

## 2.1 Design Considerations

There are three different approaches to how PromethOS plugins can be implemented on NPs. First, PromethOS plugins can be added in the embedded processor complex (EPC) and run directly on a NP core (specialized processor engine). This has the advantage that no additional copying of the packet is required. As actions are taken directly in the data plane, the overhead of sending the packet to a control point processor is avoided. On the other hand, the instruction memory can hold 32k picocode instructions shared among all NP cores, which suffices for traditional packet-forwarding tasks and advanced networking functions [2] but limits the size, and therefore the functionality of PromethOS plugins. Although theoretically feasible, picocode or parts of it cannot be dynamically reloaded with the current version of the network processor application services (NPAS). This would require all plugins to be downloaded during the initialization phase, thereby losing the benefit of dynamic code loading of the plugin approach. Running plugins on NP cores eliminates bottlenecks due to external interfaces but might add new ones on the code-execution level: Additional limits can arise owing to the scaled-down RISC architecture of the NP cores (e.g., there is no floating-point support). Even though a C-compiler for the NP cores exists, efficient code is closely linked to the hardware and therefore often written directly in picocode, which lacks code portability. A just-in-time compiler which translates an architecturally neutral programming language into

picocode [8] would then be required. A general question is *where* the code will be executed, i.e., on ingress, egress, or both. Active code placed in the data path and executed on NP cores has been evaluated in [8] for a simple active networking language.

Second, the ePPC (embedded PowerPC) in the EPC can be used to run PromethOS plugins. After classification, PromethOS relevant packets are redirected to the CP residing on the ePPC; all other packets in the data path are handled by the NP cores. The former is done by the general PowerPC handler (GPH), an NP core capable of writing the packet into the ePPC's memory and indicate its arrival to the ePPC by means of an interrupt. The packet then traverses the Linux IP stack before being handed to the plugin manager. The plugin identifier found during classification on the NP allows the plugin manager to select the appropriate plugin. Here the advantages are that only PromethOS-relevant packets will be redirected to the ePPC, while the flexibility of the Linux kernel (e.g., Netfilter support) is retained. No additional processor is needed and the approach behaves much like a system-on-a-chip. The approach will eventually encounter performance limitation due to the interface between NP cores and the ePPC. Moreover, the ePPC is clocked at 133 MHz, which might not be enough for extensive plugin processing.

As a third option, the PromethOS plugin manager can run on an Ethernet-attached external CP, usually a general-purpose processor (GPP). This approach is similar to the previous one, but uses a physical interface and the GMII gigabit Ethernet-to-PCI-X bridge to copy packet data into the CP memory. Redirection is done by a guided frame handler (GFH) NP core. Processing of plugins is limited by the clock speed of the attached external GPP CP.

Compared with an approach without NPs the benefits are that packet classification is done by the NP, hence reducing packet handling in the Linux IP stack, while normal data packets are directly forwarded by the NP. In this paper we analyze the latter two approaches, where the plugin manager resides on the ePPC or an external CP. Given its limited functionality, the approach with dynamically (re-)loadable picocode plugins is left for future work.

### **3 The IBM NP4GS3 Network Processor**

#### **3.1 The Power NP4GS3 Architecture**

The IBM PowerNP 4GS3 is composed of an embedded processor complex (EPC), the enqueueer dequeuer scheduler (EDS) blocks, the switch interfaces, the physical MAC multiplexers, embedded SRAM memory, and additional memory interfaces for external memories. The EDS is responsible for hardware flow control and scheduling while the MAC multiplexers transfer packets from/to the physical-layer devices. The main functional blocks of a PowerNP are shown in Figure 2.

The EPC consists of 16 packet processor engines called NP cores each supporting two independent threads, a set of eight specialized coprocessors for each NP core, and an embedded PowerPC 405 microprocessor, all running at 133 MHz. The coprocessors perform asynchronous functions such as longest-prefix lookup, full-match lookup, packet classification, hashing (all performed by two tree search engines (TSE) per NP

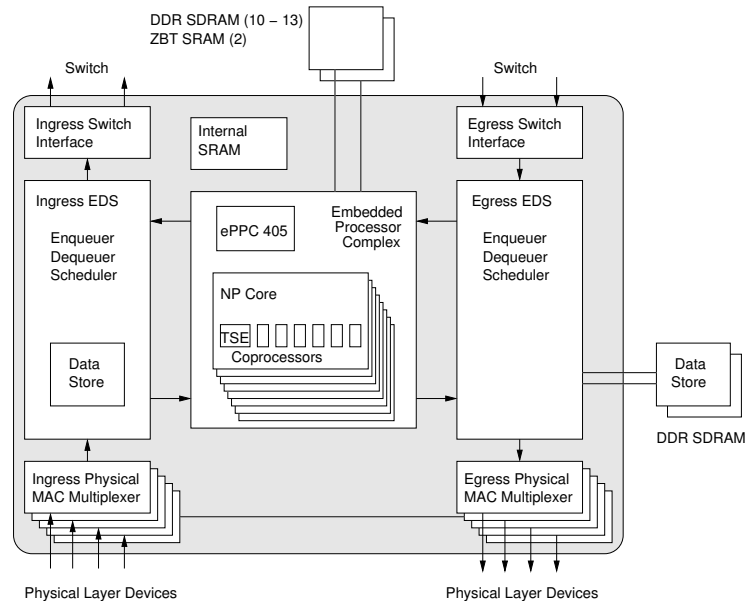


Fig. 2. Main functional blocks of an IBM PowerNP 4GS3.

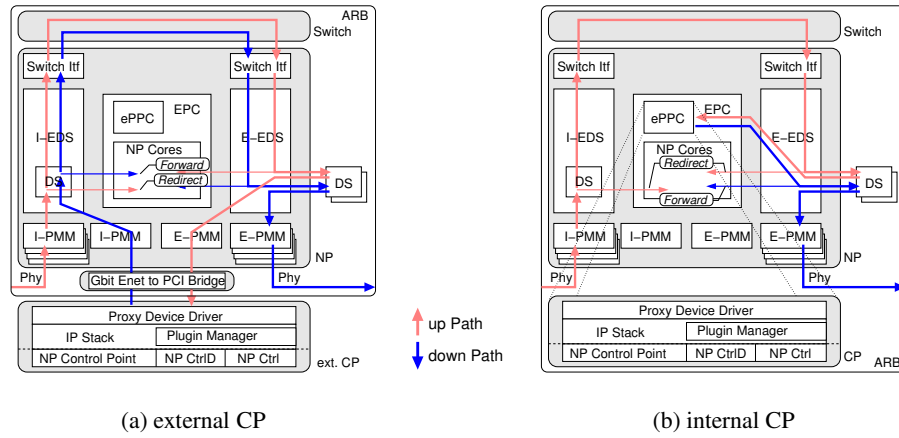
core), data copying, checksum computation, counter management, semaphores, policing, and access to the packet memory. The NP cores are scaled-down RISC processors which execute the so-called picocode. The picocode instruction set is specifically designed for packet processing and forwarding.

Packet processing is divided into two stages: Ingress processing directs packets from the physical interface to the switch, egress processing does the reverse. Every NP core can handle both stages, but usually one is associated virtually at dispatch time for convenience. Threads are dispatched upon packet arrival from the physical interface or the switch, or by an interrupt. Each thread has its own independent set of registers, so there is no overhead in switching threads. When a thread stalls (e.g., when waiting for a coprocessor), multi-threading will switch to the other thread if this one is ready for execution. This dynamic thread execution helps to balance the processor load. A thread entirely processes a stage of a packet, which is called run-to-completion mode. Additional context information (e.g., output interface identifier gained from the IP forwarding lookup) can be transferred from ingress to egress along with the packet.

We based our implementation on the Application Reference Board (ARB) from Silicon Software System [11]. This board provides a BroadCom PCI-X Ethernet controller (BCM5700) [1] for bridging between the application reference board and the host.

### 3.2 PromethOS NP Implementation

Figure 3 gives an overview of the architecture for (a) the external CP and (b) the internal CP on the ePPC. Administration and configuration of classifier rules are handled by



**Fig. 3.** Data Path of packets handled by (a) the external CP and (b) the internal CP on the ePPC.

the NP CtrlID and the NP Ctrl. The client allows a user to manage classification rules and plugin IDs similar to `tc` of the Traffic Control [3] package in Linux. The daemon provides an interface to the client process and talks to the NP using the proxy interface to the NPAS from the NP control point. For this, the daemon performs the necessary translation process and maintains counters of rule hits at the same time. The NP CP uses the Proxy Device Driver to encapsulate control traffic from the CP to the NP.

The implementation of PromethOS on the PowerNP is based on the multi-field classifier from the NPAS which provides a CP API and its corresponding picocode part. Depending on the memory size, up to 5192 multi-field classification rules can be stored. The classifier picocode has been enhanced in order to return the plugin ID (later being used by the plugin manager) if a rule matches. A rule match redirects an incoming packet, including the plugin ID found, to the CP for further processing.

While the redirection *decision* is taken on the ingress (i.e. during packet classification), the redirection *action* occurs at the egress. In the case of an attached external CP, the packet is sent to the physical interface and then traverses the Ethernet-to-PCI bridge to reach the CP. As the plugin ID is already known, the packet will not traverse the full Linux IP stack, but is handed directly to the plugin manager by the proxy device driver (fast-path). After processing, the plugin manager sends the packet back to the NP. It will again traverse the ingress side of the NP, but this time the forwarding decision is taken. Next it traverses the switch and the egress side of the NP as a normal IP packet does. In the case of an internal CP, the GPH sends the packet directly to the ePPC, where it will be handled, and receives it back afterwards for forwarding on the egress.

### 3.3 Performance Characteristics

The following list mentions the performance characteristics of the PowerNP that play a major role for all PromethOS NP configurations, as discussed in Section 2.1.

- **Data Mover Units:** The PowerNP provides five data mover units (DMU). Each DMU moves data at 1 gigabit per second (Gbps) in both ingress and egress directions. Four of them can be configured independently (e.g., as an Ethernet medium access control (MAC)). The fifth pair is directly inter-connected to move data from the egress to the ingress side of the NP4GS3.
- **Ethernet:** Three DMUs are configured as 1000Base-T GMII Ethernet ports. The fourth establishes the connection to the attached external GPP by means of a GMII Gigabit Ethernet-to-PCI-X bridge.
- **Switch interface:** The switch interface consists of two data-aligned synchronous link (DASL) interfaces in each direction. Each of them provides a transfer rate between 3.25 and 4 Gbps surpassing the accumulated bandwidth of the four Gigabit Ethernet interfaces [4]. These interfaces can either connect an NP to a switch fabric, to another network processor, or directly transfer the data from the ingress to the egress interface. Thus, this interface will not cause any performance degradation.
- **Data store coprocessor:** Data are copied into or from the EPC by the data store coprocessor of the NP. The packet throughput depends linearly on the number of bytes copied per packet: Usually only 64 bytes are copied, as this is sufficient for header inspection. The PowerNP achieves 4.80 Gbps of aggregated throughput of Internet-like traffic when doing layer 3 packet forwarding [5]. Depending on the PromethOS configuration, data packets traverse each stage up to two times. Because PromethOS requires additional layer 4 classification we expect that the PowerNP can provide up to 1.5 Gbps throughput.
- **PCI bus:** The ARB can be integrated into an hANN using its Ethernet-to-PCI bridge. The BroadCom PCI-X Ethernet Controller BCM5700 permits bridging at 1 Gbps full duplex. The PCI standard v2.3 defines the following bus transfer rates: 1.1 Gbps for an interface with 32 bits width running at 33 MHz (32b/33MHz), 4.3 Gbps (64b/66MHz), and 8.5 Gbps (62b/133MHz PCI-X 1.0). However, the PCI bus does not provide full duplex. So, if the ARB were placed in a 32b/33MHz PCI system, we could expect a throughput of at most 0.55 Gbps (provided the bus is not used by other devices). Thus, at least 2 Gbps are required from the PCI bus bandwidth to satisfy the ARB.
- **General PowerPC Handler:** The ePPC is connected to the general PowerPC handler (GPH), a NP core with extended capabilities, via shared memory for data transmission. The GPH copies data packets into the external DRAM, and signals this to the ePPC by an interrupt. The reverse process is carried out if the ePPC sends a packet. Passing packets to the ePPC has been designed for the control path, hence we cannot expect high throughput for data-path applications. However, it can be extremely valuable to offload complex data-path processing as encountered in active networks in order to prevent packet redirection to an external CP, as long as the rate is bounded to an acceptable value. As it is difficult to estimate the performance of this interface, we provide empirical results in section 4.1.

From this analysis, we assume that the PowerNP should be powerful enough to carry out packet classification for PromethOS plugins on the one hand, and, on the other hand, to forward packets of other streams at link speed (1 Gbps) simultaneously.



## 4 Evaluation

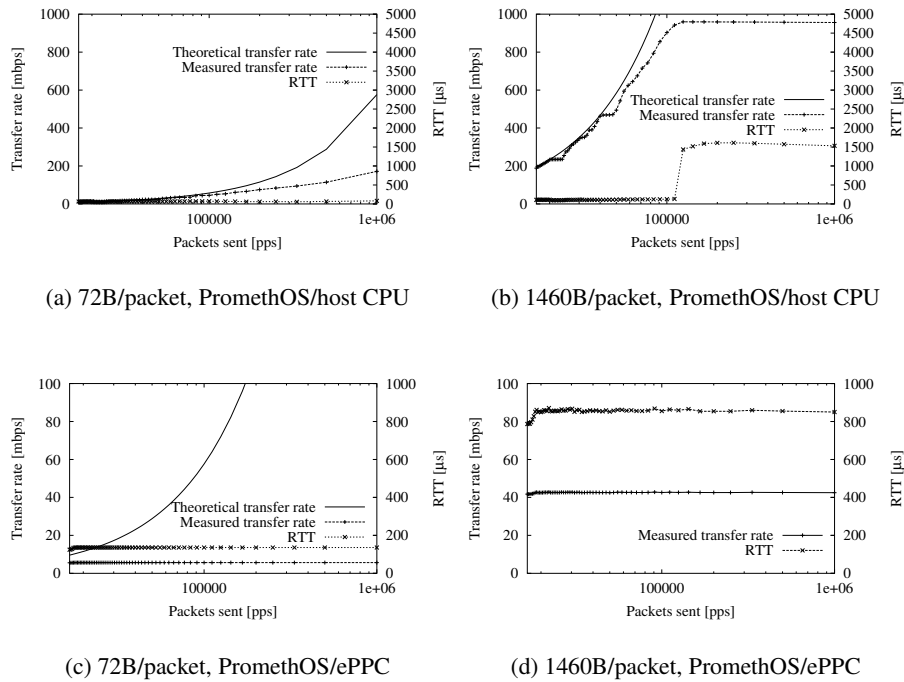
### 4.1 Performance Measurements

Following the analysis of all interfaces involved (cf. section 3), we base our evaluation on an hANN with an Intel Xeon 2.4 GHz processor running Linux 2.4.18 in which the ARB is installed. The ARB operates at 64b/66MHz PCI speed.

We measured the performance of the hANN without real service functionality of the plugins because otherwise throughput would additionally depend on the service complexity rather than on the efficiency of the framework. Packets were sent by a traffic generator (source) to the plugin manager (sink), whereby the plugin manager acts as source and sink at the same time for convenience. Packets were sent out by one Ethernet interface and received on another via crossed cables. Pathes packets were sent along have been visualized in figure 3.

Latency, throughput and packet loss have been measured in two configurations: In the first case PromethOS NP was running on the Ethernet-attached external CP, in the second case it was placed on the CP running on the ePPC. The results are for different packets sizes, namely, 72 and 1460 Bytes. In Figure 4 (a) and (b), we plot the results of

**Fig. 4.** PromethOS NP on the host CPU (a,b) and on the ePPC (c,d) – Transfer Rate and Round Trip Time: (a,c) 72 Bytes per packet; (b,d) 1460 Bytes per packet.



the first configuration in which the NP cores are only used for packet classification. The measurement results achieved for the second configuration are shown in Figure 4 (c) and (d). The packet size corresponds to the number of bytes sent at the Ethernet interface, omitting the internal header (36 Bytes) added by the Linux proxy device driver for signaling. In Figure 4, the transfer rate (TR) is shown in megabits per second (Mbps), the round trip time (RTT) in units of microseconds ( $\mu s$ ), and the packet transfer rate in units of packets per second (pps). For comparison, we also plot the ideal transfer rate, where the number of packets attempted to send corresponds to the number of packets received, assuming all transmission attempts are successful.

**Table 1.** Comparison of transfer rates and round trip times

PromethOS NP on the host CPU:

| 72 Bytes per packet |           |                 | 1460 Bytes per packet |           |                 |
|---------------------|-----------|-----------------|-----------------------|-----------|-----------------|
| TR (pps)            | TR (Mbps) | RTT ( $\mu s$ ) | TR (pps)              | TR (Mbps) | RTT ( $\mu s$ ) |
| 297985              | 171.639   | 81.2            | 81846                 | 955.966   | 1531.4          |
| 20134               | 11.597    | 48.7            | 20110                 | 234.879   | 96.2            |

PromethOS NP on the ePPC:

| 72 Bytes per packet |           |                 | 1460 Bytes per packet |           |                 |
|---------------------|-----------|-----------------|-----------------------|-----------|-----------------|
| TR (pps)            | TR (Mbps) | RTT ( $\mu s$ ) | TR (pps)              | TR (Mbps) | RTT ( $\mu s$ ) |
| 9807                | 5.649     | 135.9           | 3638                  | 42.497    | 849.8           |
| 9640                | 5.553     | 124.3           | 3574                  | 41.471    | 786.6           |

In Table 1, we compare the maximum throughput, the maximum transfer rate, and the minimum round trip time for both configurations. We note the difference in performance between the two configurations.

The increase in latency found in Figure 4 (b) corresponds to the default queue-threshold configuration of the PowerNP. The throughput difference between the two configurations is due to the mailbox communication interface between the EPC and the ePPC. Therefore we conclude that an EE should be dynamically loadable, depending on the expected rate of packets that require processing by a PromethOS plugin.

We further investigated the second configuration in two specially designed evaluation runs: First, we measured the performance of Linux with regard to its capacity of creating, sending and receiving socket buffers without real transmission. Second, we measured the performance of the interface between the NP cores and the ePPC by transferring full-sized packets (1460 Bytes) via the shared memory and interrupt signaling back and forth. As in the configuration mentioned above, two copy operations are performed. In the first measurements, we achieved a transfer rate of 697.39 Mbps. In the second, we were able to measure a transfer rate of 298.04 Mbps. Note that we did not vary the internal socket-buffer limits imposed by Linux which can further improve our results.

## 5 Related Work

VERA [6] implemented a three-level router architecture. VERA itself provides a device driver to the Linux operating system that runs on the host CPU. It interfaces to the IXP1200 Evaluation Board. In [13] resource allocation and scheduling issues are analyzed on a three-level processor hierarchy, and [12] evaluates the performance of the Intel IXP 1200 for vanilla IP packets. In [9], an IXP1200-based network interface card offering four 100T ports was evaluated. On the IXP1200 StrongARM core, Linux is run, but used for initialization and debugging purposes only; processing is carried out in the so-called kernels run on the microEngines of the IXP, while the host CPU is used for extended processing. A very interesting approach to datapath packet processing is provided in [10] where the performance of a Click-based NP software architecture is evaluated.

The Active Packet Editing (APE) approach [14] is a two-level active networking architecture that consists of an active packet processor in software running on a GPP and a packet editor based on an FPGA with content-addressable memory (CAM) for efficiency. The packet processor configures the packet editor, which performs packet classification and simple packet-modification tasks through active packets. Their packet editor prototype achieves slightly less than 1 Gbps of throughput for simple IP header modifications and the packet processor is capable of handling 10 Mbps of small-sized packets.

## 6 Summary, Conclusion and Outlook

In this paper, we introduced PromethOS NP, a framework that eases the use of network processors for high-performance active network nodes. The framework provides extended NodeOS functionality while supporting a PromethOS EE. Our implementation is based on the IBM PowerNP 4GS3 network processor. It is run either on the host CPU (Ethernet-attached external control point) or on the embedded general-purpose processor of the network processor. In both configurations, the NP cores provide packet classification for the fast-path to circumvent legacy packet classification by the network stack of the operating system.

Our performance measurements prove the efficiency of our architecture. PromethOS NP supported by the PowerNP was able to handle Gigabit link speed (~956 Mbps); 297,985 packets per second could be processed without any particular optimization of legacy Linux. In addition, when PromethOS NP was run on the Ethernet-attached external control point (host CPU), the PowerNP provided ample capacity for additional processing. When run on the internal control point (embedded PowerPC), ample room for operating-system optimization was revealed by the performance investigations we carried out.

We are convinced that PromethOS NP in conjunction with the IBM PowerNP 4GS3 provides a flexible and efficient architecture and platform for active services that need to process packets at link-speed. Currently, we are investigating the extended use of the NP cores as well as optimizations of a NodeOS running on the PowerNP creating a multiprocessor high-performance active node.

## References

- [1] Broadcom Corporation. BCM5700 PCI-X 10/100/1000BaseT controller, 2002.
- [2] R. Haas, C. Jeffries, L. Kencl, A. Kind, B. Metzler, R. Pletka, M. Waldvogel, L. Freléchoux, and P. Droz. Creating advanced functions on network processors: Experience and perspectives. *IEEE Network*, 17(4), July 2003.
- [3] B. Hubert et al. Linux Advanced Routing & Traffic Control. <http://lartc.org>, 2003.
- [4] IBM Corp. IBM PowerNP NP4GS3 databook. <http://www.ibm.com>, 2002.
- [5] IBM Corp. LinleyBench 2002 test results, IBM PowerNP NP4GS3. <http://www.chips.ibm.com/techlib>, 2002.
- [6] S. Karlin and L. Peterson. VERA: An extensible router architecture. In *Proceedings of the 4th International Conference on Open Architectures and Network Programming (OPENARCH)*, pages 3–14, April 2001.
- [7] R. Keller, L. Ruf, A. Guindehi, and B. Plattner. PromethOS: A dynamically extensible router architecture supporting explicit routing. In *Proceedings of the Fourth Annual International Working Conference on Active Networks IWAN*, volume 2546 of *Lecture Notes in Computer Science*, Berlin, Heidelberg, December 2002. Springer Verlag.
- [8] A. Kind, R. Pletka, and B. Stiller. The potential of just-in-time compilation in active networks based on network processors. In *Proceedings of IEEE OPENARCH '02*, pages 79–90, June 2002.
- [9] K. Mackenzie, W. Shi, A. McDonald, and I. Ganey. An Intel IXP1200-based network interface. In *Proceedings of the Workshop on Novel Uses of System Area Networks at HPCA (SAN-2 2003)*, 2003.
- [10] N. Shah, W. Plishker, and K. Keutzer. NP-Click: A programming model for the Intel IXP1200. In *Proceedings of 9th International Symposium on High Performance Computer Architectures (HPCA), 2nd Workshop on Network Processors*, February 2003.
- [11] Silicon Software System. Application reference board for the IBM PowerNP NP4GS3 network processor user manual.
- [12] T. Spalink, S. Karlin, and L. Peterson. Evaluating network processors in IP forwarding. Technical Report TR-626-00, Department of Computer Science, Princeton University, November 2000.
- [13] T. Spalink, S. Karlin, L. Peterson, and Y. Gottlieb. Building a robust software-based router using network processors. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, pages 216–229, October 2001.
- [14] N. Takahashi, T. Miyazaki, and T. Murooka. APE: Fast and secure active networking architecture for active packet editing. In *Proceedings of IEEE OPENARCH '02*, pages 104–113, June 2002.