

Network Processing Technology for Terminals Enabling High-quality Services

Yukikuni Nishida[†] and Keiichi Koike

Abstract

This article describes the RENA-SoC, which is a network processor that provides basic functionalities required for a home or SOHO (small office home office) network gateway in the fiber-to-the-home (FTTH) broadband era (RENA: routing engine for Next Generation Network access, SoC: system on a chip). Network processing is performed by a packet processing engine, and the application processing load is distributed over two microprocessors to provide the user with high-quality stable services.

1. Introduction

Managed networks such as the Next Generation Network (NGN) can provide high-quality, secure networking for triple-play services that provide voice and video communications, video delivery, and Internet access. Personal computers (PCs) as well as televisions (TVs) and various sensors will be connected to home networks, enabling operations such as TV program reservation or remote monitoring of sensors. Many other services are likely in the future even from home appliances such as refrigerators and microwave ovens connected to the network.

The gateway is very important for connecting the managed network with home networks and small office home office (SOHO) networks, while also protecting the network and the access points for receiving various services. As shown in **Fig. 1**, the gateway must provide many functions such as security, quality control, and service enhancement, as well as the performance to process packets of Gigabit Ethernet signals.

The main functions of the gateway can be classified into network functions and application functions. Application functions must support voice over Internet protocol (VoIP), a basic service of managed net-

works, and may require the OSGi (Open Services Gateway initiative) framework for managing application software, so that application software can be added and deleted on demand. OSGi runs on the Java VM (virtual machine), which avoids the need to take into consideration differences in platforms, reduces the need to develop application software for each platform (called *bundled* software in Java VM terminology), and makes development inexpensive. In this way, users are expected to be able to get applications at a reasonable cost.

At NTT Laboratories, we are promoting open innovation and have developed the OSGi Service Aggregation Platform (OSAP). This is a base technology that extends OSGi, allowing bundled software from service providers for instance, software that lets the service providers control home appliance, to be downloaded and run on the gateway.

If microprocessor-intensive application software is executed in a multitasking environment, the load on the central processing unit (CPU) will increase and the performance of packet processing, which is a basic gateway function, could decrease and lead to longer packet delays. This could result in a decrease in VoIP quality. To avoid this sort of quality of service (QoS) degradation, we designed an architecture that distributes processing over three processing engines:

- a newly developed gigabit-per-second packet processing engine,

[†] NTT Cyber Space Laboratories
Yokosuka-shi, 239-0847 Japan

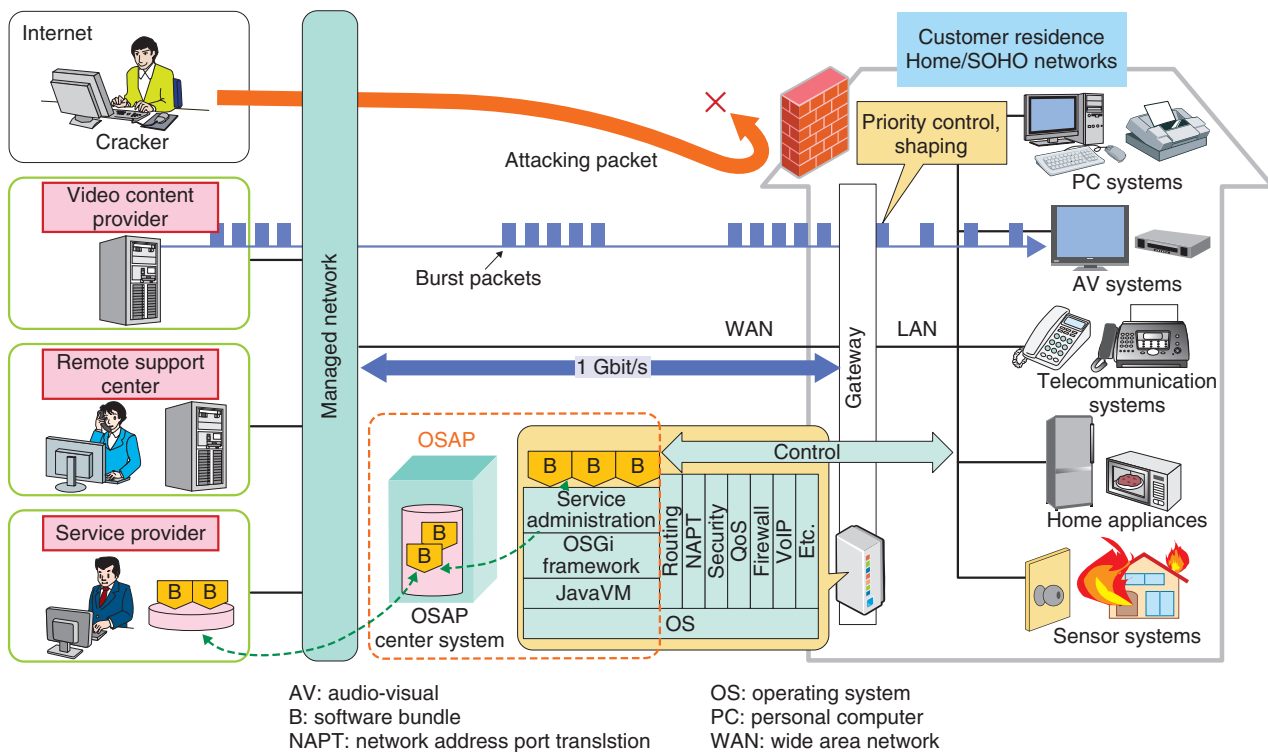


Fig. 1. Home/SOHO networks.

- a CPU for processing primary services such as VoIP,
- a CPU for processing enhanced functions such as OSGi bundles.

Gateways are also required to provide interfaces for devices such as wireless local area network (LAN) cards and USB (universal serial bus) hard disk drives, and a LAN switch (LAN-SW) in order to connect various pieces of network equipment. This can lead to a large number of components and result in increased power consumption and cost, so to reduce power consumption and cost, we developed the RENA-SoC (RENA: routing engine for Next Generation Network access [1], SoC: system on a chip), which is a network processor that incorporates a packet processing engine, two CPUs, PCI (peripheral component interconnect) and USB interfaces, and a LAN-SW on a single chip (Fig. 2).

2. Packet processing engine

The packet processing engine, which contributes greatly to the packet processing performance of the RENA-SoC, is composed of an IPsec block, an IP

switch (IP-SW) block, a LAN switch (LAN-SW) block, and a packet scheduler. The IPsec block handles packet encapsulation for virtual private network communications, high-speed encryption, decryption, and authentication, the IP-SW block handles packet transmission as required for each connected party, including firewall and NAPT (network address port translation) functions, and the packet scheduler controls the QoS. These functional blocks are implemented in hardware, and they achieve processing performance of 1 Gbit/s in each direction (uplink and downlink). To confirm the performance of the packet processing engine, we compared the throughput of an off-the-shelf software-based router and a router implemented using the RENA-SoC for various frame lengths and NAPT table entry numbers. The results are shown in Fig. 3. The software-based router showed throughput dependent on the packet length and the number of simultaneous flows for NAPT's since these affect the software load. However, the RENA-SoC was not affected by these factors, and achieved 1-Gbit/s bidirectional packet processing. This packet processing engine enables low-delay, high-quality communications, even under high-load




Fig. 2. RENA-SoC block diagram.




Fig. 3. Packet transmission performance.

conditions of these sorts, by processing packets at the maximum interface speed of 1-Gbit/s.

The packet processing engine can also support new services through the following added functions.

- Communication quality monitoring function (packet copy function)
- Packet processing support function (improves the flexibility of hardware processing)

These functions are explained in detail below.

3. Communications quality monitoring function

When packets are discarded along the communications path, which causes a QoS decrease in services such as video delivery, it is necessary to detect whether the packets were discarded on the WAN (wide area network) segment or LAN segment. Maintenance costs can be reduced if the detection is not performed by a dispatched worker but by a telemetering function. It is currently done by placing monitoring equipment on the home network segment and dispatching a worker to the location to check the log information, which involves a significant amount of labor. Thus, there is demand for remote realtime monitoring of

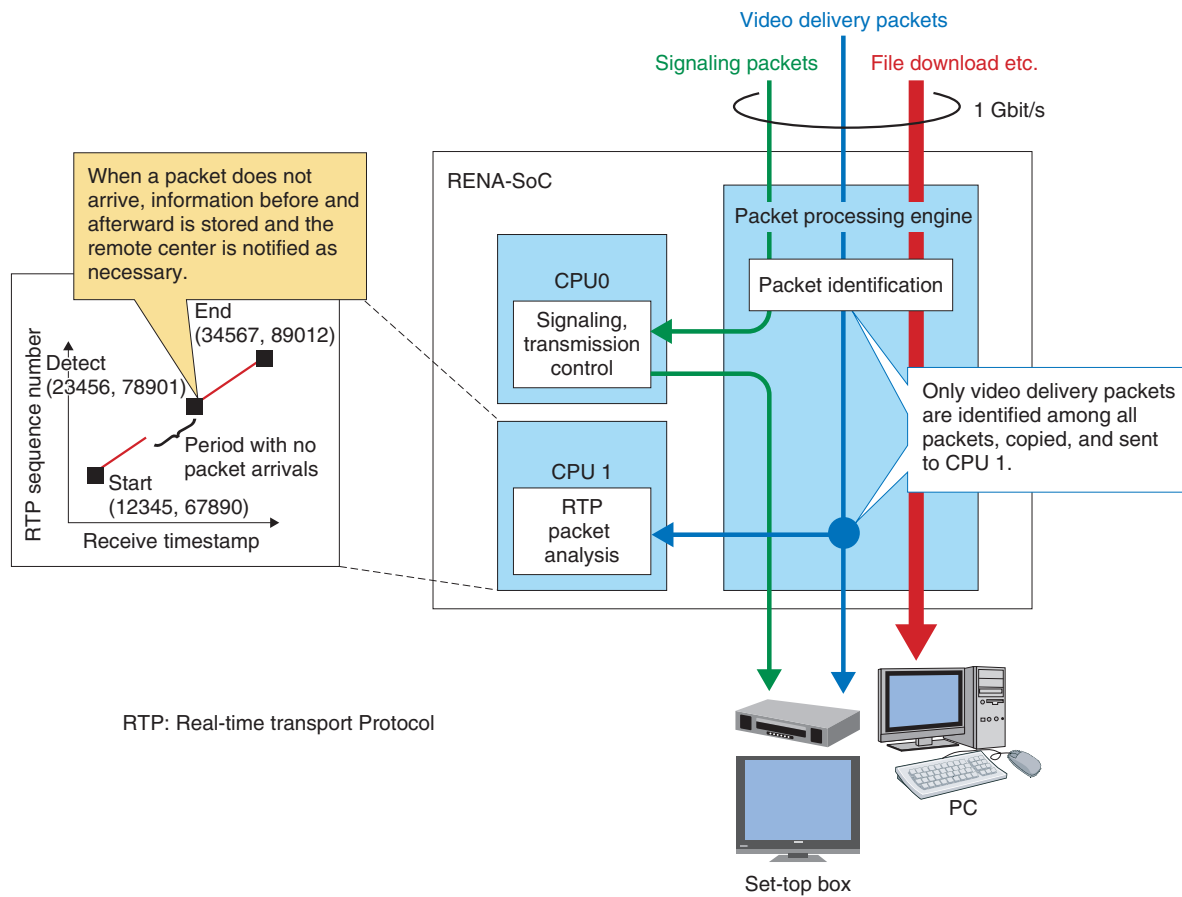


Fig. 4. Quality management for video delivery packets.

communications quality in gateway equipment.

To improve transmission performance when software is used to process packets, often only the headers are read, which avoids the need to read and write entire packets. When packets are captured to monitor quality, entire packets must usually be copied, separated from the packet transmission stream, and passed to the packet capture application. Continuous packet copying makes the processing load heavy and can cause packet transmission performance to decrease significantly, so even if video packets, for example, would have been transmitted correctly, communications quality can be degraded and packet capturing may fail. This can result in faults and errors being detected even when packets could have been received normally.

For these reasons, we have implemented a packet-copy function in hardware that allows communications to continue while copied packets are processed by the CPU and allows realtime packet capturing to

be done without decreasing communications quality. Specifically, the packet processing engine identifies packets to be captured and copies them as shown in **Fig. 4**. Even with packets flowing at 1 Gbit/s, specific packets can be copied to the CPU without affecting the packets being transmitted. The CPU can capture specially selected packets, so quality monitoring can be done without causing further packet loss. Figure 4 shows an example of a function that monitors sequence numbers and timestamps of video delivery packets. It also shows that information before and after the event is saved if a sequence number is not continuous or if packets do not arrive for a period of time. This information is reported to a remote center when necessary, which ensures that the center is notified of decreases in service quality.

The above structures enable immediate remote identification of where a network fault occurred from a remote location.




Fig. 5. Performance of packet processing support functions.

4. Packet processing support functions

Packet processing engines are generally composed mainly of hardware, so they have extremely high-performance packet processing; on the other hand, they are inflexible in terms of support for new functions. If the hardware cannot handle packets, even if it is because of a trivial modification to the packets, all of the packets must be processed by software, which causes a dramatic degradation in performance. To resolve this problem, the RENA-SoC implements packet processing support functions that allow the CPU to directly read and write packet data being processed by the packet processing engine. This allows higher performance than is possible with software-only processing.

When packets are processed using only software, the packet processing engine usually passes the entire packet to the CPU, which performs software-based network processing, and then passes the packet back to the packet processing engine. This uses up a lot of processor resources for the packet processing. By enabling the CPU to directly handle packets being processed by the packet processing engine, we were able to reduce the transfer times between the CPU and packet engine and omit most of the network processing.

As an example, for the case of supporting twice-NAT (RFC 2663), in which IPv4 addresses of both source and destination are rewritten, we compared software-only processing and the use of packet processing support functions (Fig. 5). For software-only processing, the maximum possible transmission rate was 90 Mbit/s; however, with packet processing support functions, we confirmed transmission speeds over 500 Mbit/s.

5. Application protection

By opening up the system to OSGi bundles, as described earlier, we enable various user applications to be executed on the gateway. To prevent these applications from affecting the operation of basic services such as VoIP, the RENA-SoC uses asymmetric multiprocessing, with an independent OS (operating system) operating on each CPU. Having two OSs running on the two CPUs increases the amount of memory required; however, even if, for example, CPU1 stops responding owing to an inadequately tested OSGi bundle, CPU0 can reset only CPU1, which can restart while CPU0 and the packet engine continue to operate independently without being affected. This makes system operation more stable and keeps primary services alive.

6. Summary

We have developed the RENA-SoC, which enables the operation of high-performance systems through its triple-processor architecture, consisting of a packet engine capable of 1-Gbit/s bidirectional packet transmission with support for IPsec, firewalls, and quality control; a CPU for handling basic managed-network services like VoIP; and a CPU for handling enhanced applications such as OSGi bundles or qual-

ity monitoring.

In the future we will continue research and development toward reducing the device's power consumption and providing support for various services.

Reference

- [1] Y. Nishida, K. Kawai, and K. Koike, "A 2Gb/s Network Processor with a 24mW IPsec Offload for residential Gateways," Proc. of ISSCC 2010, Vol. 1, No. 1, pp. 280–281, San Francisco, CA, USA.



Yukikuni Nishida

Senior Research Engineer, Promotion Project 1, NTT Cyber Space Laboratories.

He received the B.E. and M.E. degrees in electrical and computer engineering from Nagoya Institute of Technology, Aichi, in 1992 and 1994, respectively. He joined NTT LSI Laboratories, Kanagawa, in 1994 and has been engaged in research on the design of low-power-consumption DSPs, automatic speech recognition systems, and network appliances. Since 2003, he has been developing network processors performing gigabit-per-second packet transmission for residential gateways. He is currently studying techniques for saving energy in home network equipment. He is a member of the Institute of Electronics, Information and Communication Engineers (IEICE).



Keiichi Koike

Senior Research Engineer, Supervisor, Promotion Project 1, NTT Cyber Space Laboratories.

He received the B.E. and M.E. degrees in electronic engineering from Osaka University in 1985 and 1987, respectively. Since joining NTT LSI Laboratories in 1987, he has been engaged in research on the design of high-speed LSIs. He is currently researching an application-specific LSI for gateway equipment. He is a member of IEICE.