



NPA-1/NPA-2/NPA-3

Access Network Processors with Integrated Traffic Management

Product Brief

Highlights

- Single-chip, programmable, wire-speed network processor with 10-Gigabit aggregate throughput
- Scaled-down version of EZchip's NP-3 network processor targeting Ethernet access applications
 - Software compatible with EZchip's NP-2 and NP-3
- Flexible processing with programmable packet parsing, classifying, modifying and forwarding enabled through integrated Task Optimized Processors (TOPs)
- Integrated 10-Gigabit traffic management with hierarchical scheduling, supporting services defined by the Metro Ethernet Forum, e.g. MEF9 and MEF15
- On-chip Fabric Interface Controller (FIC) functionality for direct interfacing to Ethernet fabrics enabling system-wide traffic management
- Integrated hardware implemented search engines
- Integrated memory for lookup tables and statistics counters
 - Extension of lookup tables to external DRAM
 - Extension of statistics counters to external DRAM
- Integrated TCAM for on-chip ACL processing
- On-chip OAM protocol processing offload
- On-chip IEEE1588v2 clock sync processing offload
- 1-lane PCI-Express 2.5Gbps external host CPU interface for configuration and control
- RGMII data-only interface to host CPU
- Supports oversubscription beyond 10-Gigabit processing throughput by smart classification

Package / Process / Power

- Package: BGA 676 pins, 1.00 mm pitch, 27x27 mm
- Process: TSMC 90nm
- Est. power dissipation: 6-10W typical
- RoHS compliant
- Industrial operating temperature range -40C° to 85C° ambient

Target Applications

- Line card, service card and pizza box applications
 - Optical access (GPON/EPON OLTs and ONUs)
 - Wireless backhaul and base station aggregation (3G/4G and WiMax)
 - Access routers
 - Fiber and copper Ethernet access switches
 - Ethernet demarcation devices
 - Copper access (DSLAMs)
- Programming delivers a variety of applications such as L2 switching, Q-in-Q, PBT, T-MPLS, VPLS, MPLS and IPv4/IPv6 routing



Models

- Three footprint and software compatible models:
 - **NPA-1**: Eight 1-Gigabit Ethernet ports (SGMII/SERDES)
 - **NPA-2**: 16x 1-Gigabit Ethernet ports (SGMII/SERDES)
 - **NPA-3**: Configurable to eight 1-Gigabit Ethernet ports (SGMII/SERDES) and two 10-Gigabit Ethernet ports (XAUI); or twelve 1GE and one 10GE ports; or four 10GE ports
- SGMII interfaces support 2.5Gbps SERDES for GPON applications

Ordering Information

Device	Part No.	Device	Part No.
NPA-3	20777901	NPA-1	20777701
NPA-3 RoHS	20775801	NPA-1 RoHS	20775601
NPA-2	20777801		
NPA-2 RoHS	20775701		

Detailed Feature List

Integrated Traffic Management

- 10-Gigabit traffic manager providing queuing and scheduling on all transmitted traffic on all ports
- Per Flow Queuing (PFQ) with 4 level hierarchical scheduling:
 - 32 ports/channels
 - 256 sub-ports
 - 4K classes/subscribers
 - 16K flow queues
- Policing: Per-flow metering, marking and policing
- Configurable WRED profiles
- Shaping: Single and Dual leaky bucket controlling committed/peak rate/bursts (CIR, CIB, PIR, PIB) with IFG (Inter Frame Gap) emulation for accurate rate control
- Scheduling: WFQ and priority scheduling at each hierarchical level
- Work conserving and non-work conserving schedulers
- Frame size from 1 byte to 16K bytes
- Up to 0.5 Gbyte total frame buffering in external DRAM
- Per-frame timestamp and timeout drop
- Dynamic hitless reconfiguration and resource allocation
- LAG shaping
- Out-of-band flow control per physical port (SGMII or XAUI) or logical channel (total of 32 channels)

Programming

- Single-image programming model with no parallel programming or multi-threading
- Automatic allocation of frames to processing engines (TOPs) with passing of messages among TOPs
- Automatic ordering of frames
- In-service software updates
- Large code space memory for multiple and complex applications
- Microcode compatible with EZchip's NP-2 and NP-3 network processors

Integrated Search Engines

- Performs flexibly defined lookups in switching, routing, classification and policy tables
- Programmable size and contents of search keys and results (associated information) per table
- Support for long keys and long results per table entry
- Table entries stored in integrated memory for fastest lookup time
- Tables may be stored in external DRAM memory

- On-chip state learning and updates of millions of addresses, sessions and flows per second

Integrated TCAM

- Enables fast lookups through tables with wildcards, such as Access Control Lists (ACL)

Statistics and Counters

- Per-flow statistics for programmable events, traffic metering, policing and shaping
- Programmable threshold settings and threshold exceeded notification
- Dynamic allocation and auto association between counters and flows. Counters are automatically recycled when a flow is deleted or aged
- Hardware implementation of token bucket per flow (srTCM, trTCM or MEF5)

OAM Offload

- Per OAM session state tracking and reporting
- 802.1ag and 802.3ah compliant OAM offload
- Dedicated timer hardware blocks
 - KeepAlive frame generation for precise and accurate session maintenance operations
 - KeepAlive watchdog timers for fastest detection time
- Flexible statistics collection on a per session basis

Integrated FIC Functionality

- For architectures that adapt standard Ethernet switches as the SF solution, the NPA integrates the FIC functionality
- Allows use of standard low-cost Ethernet switches as the backplane switch fabric
- Direct connection from NPA on the line card to the backplane Ethernet switch
- NPA provides for system-wide QoS with per COS and per-flow congestion management

Sync Ethernet

- Enables on-board clock generation schemes using an external or recovered clock reference
- Provides output clock selection from each Serdes lane recovered clock

IEEE1588v2

- On-chip IEEE1588v2 clock sync processing offload for precise time synchronization among remote nodes and switches
- Can operate as the clock master, boundary clock, transparent node or a combination thereof
- Provides an accurate RTC, adjustable from the control CPU or an external source, and provides input and output timestamping for time and delay measurement

NPA Architecture

EZchip's NPA is a highly-flexible network processor with integrated traffic management targeting Ethernet network access platforms (ONT/OLT GPON/EPON), copper access platforms (DSLAM) and demarcation devices and 3G/4G WiMAX base stations aggregation and backhaul. The NPA provides high integration, programmable packet processing and advanced flow-based bandwidth control at 10-Gigabit aggregate throughput.

Through programming the NPA delivers a variety of applications such as L2 switching, Q-in-Q, PBT, T-MPLS, VPLS, MPLS and IPv4/IPv6 routing. The integrated traffic management provides advanced QoS for flow-based service level agreements (SLA) and for enabling triple-play services (voice, video, data).

The NPA integrates several functions normally found in several chips:

- Programmable packet processing at 10-Gigabit throughput
- Traffic manager
- Classification search engines
- OAM processing offload
- Synchronous Ethernet and IEEE 1588v2 clock sync offload
- On-chip memory and external DRAM memory for lookup tables and statistics
- On-chip TCAM for ACLs
- Interface to external TCAM or lookup RAM
- On-chip Fabric Interface Controller (FIC)
- Combinations of up to sixteen 1-Gigabit or up to four 10-Gigabit Ethernet MACs, with serial interfaces

NPA provides exceptionally **flexible packet processing** enabling system designers to future proof their designs to support new protocols and features through s/w updates. Packet parsing is supported for any field anywhere in the packet. Various table lookup options are provided with support for long lookup keys and results. Flows are classified based on any combination of extracted packet information. Any packet header and content can be edited and packets can easily be replicated to support multicast applications. A 'run to completion' processing model guarantees support for processing scenarios of any complexity. Large code space is provided to support complex applications as well as true hitless code updates.

Each TOP processor type employs a unique architecture with a customized, function-specific data path and instruction set. This minimizes the number of clock cycles required for complex packet manipulation and provides exceptionally fast packet processing. TOP performance is boosted by a super-scalar architecture in which multiple instances of the TOPs operate in parallel within each pipeline stage.




Figure 1. NPA Block Diagram

NPA uses a simple **single-image programming model** with no parallel programming or multi-threading. Allocation of the TOPs processing engines to incoming frames, passing messages between the TOPs as well as maintaining the ordering of frames is completely transparent to the programmer and performed in hardware. Large code space memory is available to support multiple and complex applications while providing headroom for adding new features. Full support for hitless code updates is provided to reduce system downtime for maintenance.

NPA offers extensive **traffic management** capabilities for traffic transmitted on all NPA interfaces. Flows transmitted to the network links and the control CPU interface can be assigned with specific QoS settings, queued and aggregated to enforce SLAs for services, subscribers, virtual ports and ports.

NPA supports DiffServ and IntServ services and a wide variety of QoS mechanisms. These include:

- Classification – assigning frames to specific flows with applicable QoS parameters.
- Metering – measuring per-flow traffic and determining compliance with traffic parameters. Single and Two Rate Three Color Metering (srTCM and trTCM) are used, compliant with the Metro Ethernet Forum specifications.
- Marking – individual frame's compliance.
- Congestion Avoidance – profile-based WRED early packet discards based on priority, available memory and metering results.
- Traffic Conditioning – enforcing rules on traffic flows, e.g. policing in which packet dropping is applied to non-compliant packets, or shaping which schedules a flow to conform to its assigned parameters (no packet dropping). Single and dual leaky bucket shaping is applied.
- Congestion Management – hierarchical scheduling of flows to the various interfaces using a priority scheme and Weighted Fair Queuing (WFQ).

- Fabric Interface Chip functionality – inter device messaging to avoid switch fabric congestion and target output queue congestion.

When using the NPA as a line card device in a chassis, the NPA enables **use of standard low-cost Ethernet switches as the chassis backplane switches**. In this configuration NPA connects directly to the backplane through 1-Gigabit, 2.5-Gigabit or 10-Gigabit serial interfaces, and provides the required QoS features on top of the Ethernet switch features. The NPA's integrated TM queues and buffers manage the traffic transmitted to the backplane Ethernet switch and provide system-wide congestion management per COS and per flow.

NPA contains **integrated memory** for search tables and statistics counters. In addition, standard DDR-II SDRAM and RDRAM-II offer additional external memory. The Data and Search External DRAM is used for Traffic Manager frame buffering and for additional search tables while the Control and Statistics External DRAM is used as the TM control memory and for additional statistics counters. All external memories are ECC protected when using RDRAM-II.

NPA features **integrated search engines** that perform lookups for implementing diverse applications in layer 2-4 switching/routing and layer 5-7 deep packet processing. These search engines deliver programmable lookups in a combination of tables. NPA can store the **lookup tables in integrated memory** to reduce board complexity, power dissipation and cost. Applications can also utilize the Data and Search External DRAM.

NPA supports several types of **lookup tables** – direct access tables, hash tables, trees and FastIP – each is flexibly defined and used for various applications. Tables may be used for forwarding and routing, flow classification, access control, etc. Numerous tables of each type can be defined, stored in internal memory and/or external memory and searched through per packet. For maximum flexibility, the key size, result (i.e. associated data) size, and number of entries are all user-programmed per table.

Patented **search algorithms** enable high-speed lookups in trees and hash tables, which are stored in the embedded memory or external memory. All search algorithms for hash, trees and FastIP are implemented in hardware for maximum performance and simplicity. Hash lookup performance is nearly deterministic regardless of the hash table size. FastIP is an innovative implementation of the routing table as a series of direct access tables as opposed to a tree. The FastIP data structure improves lookup performance and reduces the number of memory accesses required, making it ideally suited for IPv4 routing tables. FastIP can be used to

perform best match IP address longest-prefix-match lookups. Large result (associated information) can be stored per each entry and retrieved upon a lookup match.

In addition, an **integrated TCAM** is available for performing fast lookups through tables with multiple wildcards such as Access Control Lists (ACL). The TCAM lookups are performed in parallel to algorithmic lookups performed by the integrated TOPsearch engines of NPA.

Stateful auto learning and updating of table entries and the session's state is performed entirely by NPA's TOPs with no intervention required by the host CPU. This provides for an extremely high rate of millions per second, for the addition of new flows or deletion of old flows to/from the flow table. Result information, statistics counters, and per-flow rate limiters can be automatically associated with new flows and recycled when deleting old flows. Auto learning requires external Search Memory.

On-chip **OAM support** tracks individual sessions and offloads the host CPU from the task of generating and monitoring OAM messages. Dedicated configurable timers generate KeepAlive frames per thousands of sessions. These frames are processed by the TOPs for flexible formatting and accounting on a per session basis. Dedicated h/w monitors thousands of sessions and verifies that KeepAlive messages arrive within configurable intervals and according to a specified rate. The host CPU is alerted to any session failing to meet the minimum messaging rate.

NPA provides **support for Synchronous Ethernet and the IEEE1588v2** protocol for precise time synchronization among remote nodes and switches. For Synchronous Ethernet NPA enables on-board clock generation schemes using an external or recovered clock reference. It provides clock recovery per Serdes lane and configurable output clocks for each Serdes lane. NPA can operate as the clock master, boundary clock, transparent node or combination of these. NPA provides accurate RTC, adjustable from the control CPU or an external source, and provides input and output time stamping for time and delay measurements. Time synchronization is assisted with the on chip OAM block for periodical synchronization transactions without requiring CPU intervention, to enable data plane implementation of IEEE1588 protocol on-chip.

For application scenarios that load the NPA with traffic bursts that exceed 10Gbps, NPA provides **smart oversubscription** that assigns preferential processing for high priority traffic. Packets can be classified according to VLAN, priority bits, selected list of SA or DA, MPLS tag and more. The NPA compares the status of its input queues and if a threshold is exceeded, a drop decision can be taken for packets classified as lower priority.

System Configurations

NPA's flexibility and integration allows system vendors to deliver cost effective solutions that can easily adapt to changing market requirements. Illustrated below are several sample solutions.

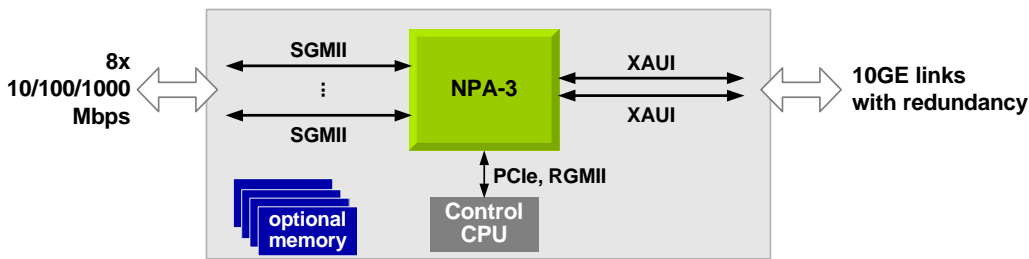


Figure 2. 10GE access switch or 10GE inter-carrier demarcation box (NNI) with redundancy

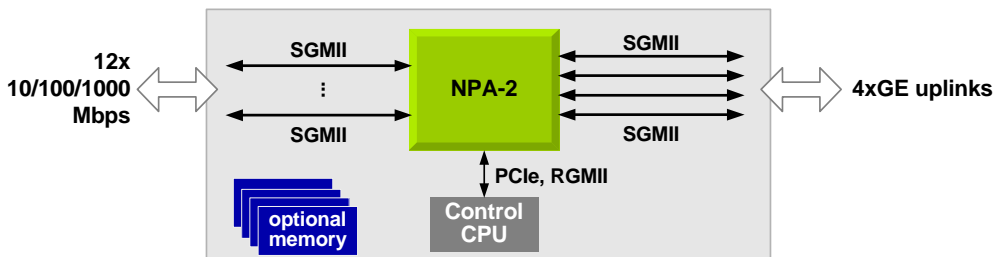


Figure 3. GE access switch or GE demarcation box (UNI)

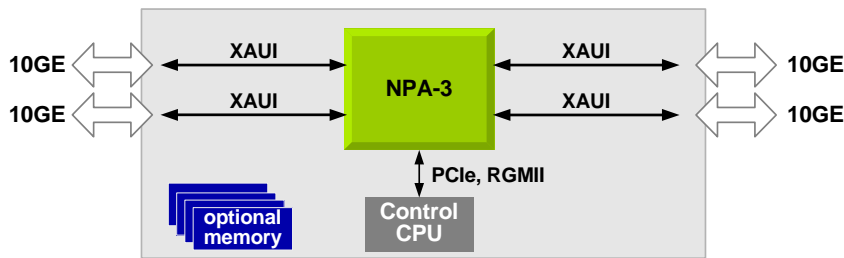


Figure 4. 10GE inter-carrier demarcation box (NNI)

In the following applications, NPA provides switching and QoS to Radio and DSL line cards, and uplinks to aggregation networks:

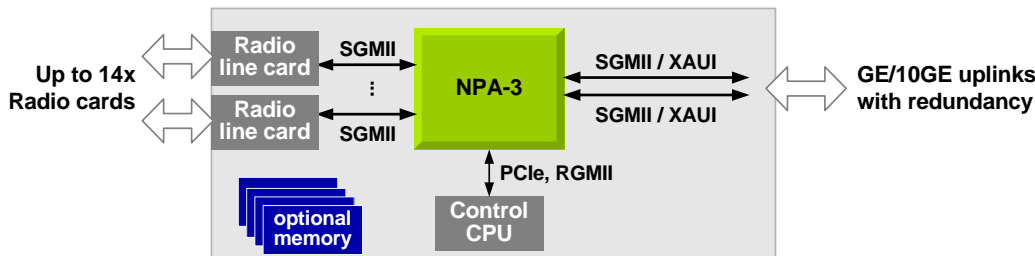


Figure 5. Wireless backhaul aggregation switch

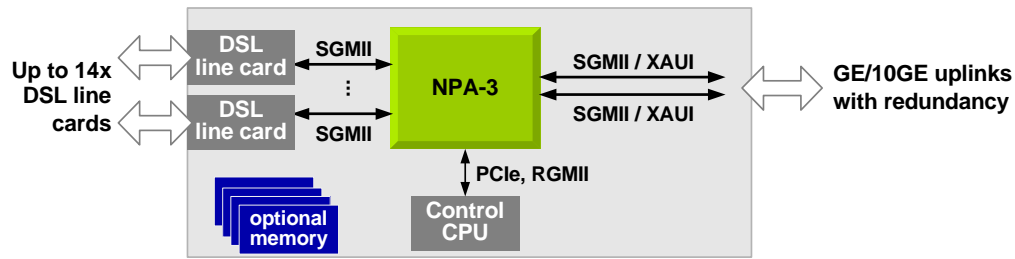


Figure 6. DSLAM ADSL/VDSL

In the following application, NPA provides switching and QoS to GPON/EPON links, and uplinks to aggregation networks:




Figure 7. GPON OLT

Direct Connection to Ethernet Switches

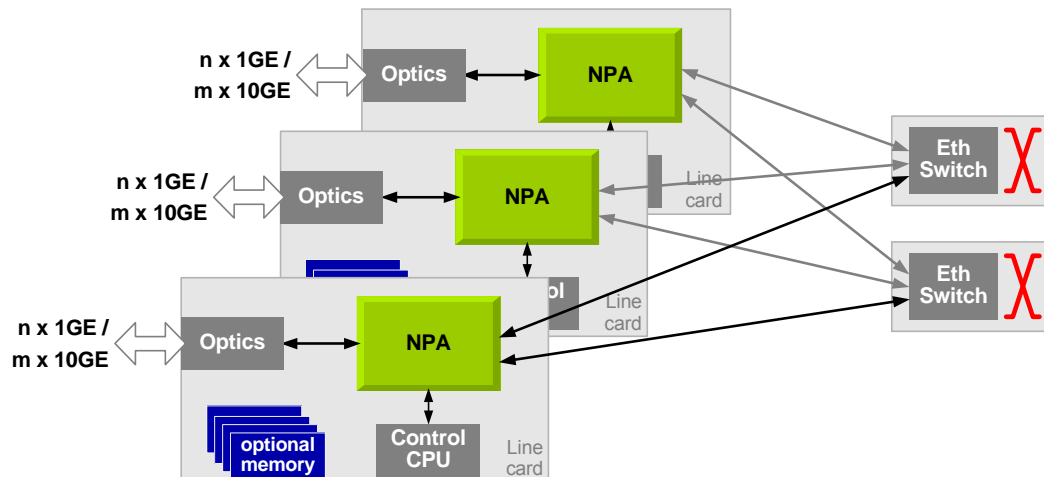


Figure 8. Direct connection to Ethernet switches with end-to-end system-wide QoS

About EZchip

EZchip Technologies is a fabless semiconductor company that provides Ethernet network processors. EZchip provides its customers with solutions that scale from 1-Gigabit to 100-Gigabits per second with a common architecture and software across all products. EZchip's network processors provide the flexibility and integration that enable triple-play data, voice and video services in systems that make up the new Carrier Ethernet networks. Flexibility and integration make EZchip's solutions ideal for building systems for a wide range of applications in telecom networks, enterprise backbones and data centers. Visit our web site at www.ezchip.com.



Email: ezsupport@ezchip.com • Web: www.ezchip.com

EZchip Technologies Inc. • 900 E Hamilton Ave, Suite 100, Campbell, CA 95008, USA • Tel: (408) 879-7355, Fax: (408) 879-7357
 EZchip Technologies Ltd. • 1 Hatamar Street, PO Box 527, Yokneam 20692, Israel • Tel: +972-4-959-6666, Fax: +972-4-959-4166

©2010 EZchip Technologies. All rights reserved. Information is subject to change without notice. Revised; December 31, 2010.