

Converged Sensor Network Architecture (CSNA)

Ian Dunn (idunn@mc.com), Michael Desrochers (mdesrochers@mc.com), Robert Cooper (rcooper@mc.com)
Mercury Computer Systems, Inc.

IP Networking in Signal Processing

Internet Protocol (IP) networking is becoming a pivotal technology for interconnecting systems of systems, whether the target platform is an aircraft or an aircraft carrier. How broad a role networking plays on these platforms is based on how much performance and determinism system vendors must give up to use it, and where in the platform they must make this tradeoff decision.

Traditionally, IP network connectivity has been relegated to a single host or single board computer (SBC) slot in high-performance embedded computing systems. The result has been a low-performance, hierarchical network architecture suitable only for application command and control. Sensor processing via IP traditionally has been off the table, as the capability of a sensor platform hinges on how much computing can be brought to bear on the signal processing associated with the sensor. Any unnecessary inefficiencies associated with networking must be eliminated in favor of the sensor.

The growth in command and control functionality, coupled with the effects of Moore's Law on reducing the number of slots that must be devoted to signal processing, is however, enabling vendors to build subsystems with more than a single host. Vendors are struggling with how to transition to these more complex, multi-hosted architectures, and are looking to use IP networking technologies to facilitate more convergence in control and data.

This paper defines an architecture called the converged sensor network architecture (CSNA). It is transparent to IP-based networks for the purposes of configuration and control, availability management, and quality-of-service (QoS), while delivering deterministic, high-performance, low-power processing. In addition, we describe a first implementation of this architecture that utilizes existing off-the-shelf multicomputer components.

Before describing that system, we present a reference model for the communication that is oriented toward high-performance, embedded signal processing systems.

Communications Reference Model

- **Sensor Plane:** Acquires, conditions, and forwards the resulting digitized data to the data plane. This data is time-critical, typically with no opportunity for retransmission by the sensor.
- **Data Plane:** Transforms digitized data into information symbols, and routes them over a network to an exploitation function. The data plane must forward data with deterministic, lossless, low-latency throughput.
- **Co-processing Plane:** Provides on-board communication to special-purpose accelerator

functions such as field programmable gate arrays (FPGAs) and general-purpose graphics processing units (GP GPUs).

- **Communications Plane:** Responsible for external data communication with other systems including both upstream remote sensors and separate downstream processing systems. (Control traffic associated with external interfaces is assumed to be part of the Control Plane in this model.)
- **Control Plane:** Concerned with the initialization, configuration and synchronization of sensor processing components in the data plane, including the mapping and/or routing of data through the sensor, data, and co-processing planes.
- **Management Plane:** Concerned with management of the physical assets of the signal processing system including presence, boot, and availability.

RapidIO-based CSNA Solution

We begin this section with a new concept called the Sensor Gateway. This, coupled with a new wire protocol on top of RapidIO for transporting Ethernet frames, is the key enabler in the convergence of IP-networking and sensor computing. The Sensor Gateway is responsible for forwarding control-, data-, and sensor-plane traffic in and out of the distributed signal processor. It provides 10 gigabit Ethernet/IP transport, switching and termination (see [1] for a good introductory reference).

The emergence of low-power, multi-core SoC (system on chip) devices with Ethernet, RapidIO, and PCIe built-in, has enabled a new architecture for IP-based processing in signal processing systems. These high-performance devices can now be configured in software to selectively process packets for an entire subsystem. IP-based command and control packets pass through the Sensor Gateway unfettered, making the subsystem and all of its processing assets visible to the network and interoperable with all network-based assets inside and outside of the subsystem. While external data traffic also passes through the Sensor Gateway unfettered, internal data traffic, often more critical to sensor capability, is mapped directly onto DMA channels, bypassing all of the header processing associated with IP networking.

The Gateway (Figure 1) consists of a VXS-based system with a RapidIO fabric. IP and Ethernet processing is performed by a combination of Mercury Gateway modules, which contain third-party PMC format PCI-X 10 gigabit Ethernet (10 GbE) cards, and Mercury VPA-200 SBCs.

Packet Flow

We developed an encapsulation protocol, called Ethernet-over-RapidIO (EoRIO), which encapsulates Ethernet frames in RapidIO packets. Similar to other IP

encapsulations standards [2, 3], EoRIO maps RapidIO compute node (CN) IDs onto Ethernet MAC addresses. By performing address translation at this level in the OSI stack, a higher degree of network interoperability is preserved at the IP level, in particular, initialization and management schemes built on top of Ethernet broadcast transaction.

EoRIO implements a credit-based scheme whereby the sending IP node is in charge of managing buffer states at the target node.

The scheme is initialized when the connection is first established. The number of outstanding credits and the size of the buffers are configurable-system parameters that can be adjusted to trade off resources and performance.

Packet handling on the gateway uses the fast-path and slow-path concept typical in network routers. The fast path is optimized for doing one job and doing it really well – forwarding unicasting IP packets from the 10 GbE interface to the RapidIO interface and vice versa.

The slow-path processing of all incoming 10 GbE broadcast and multicast packets is offloaded to an application-determined number of VPA-200 SBCs. In the outgoing direction, slow-path Ethernet frame handling is carried out on the originating VPA-200.

The VPA-200 SBCs also support standard Linux-bridging functionality between 10-gigabit and 1-gigabit Ethernet. Ethernet frames move from the 10 GbE interface on the Gateway, over RapidIO, to a VPA-200, through the Linux bridging implementation and then exit the front panel GbE ports.

Traffic Management with RapidIO

Each RapidIO packet carries with it a transaction priority, which enables higher-priority requests to move ahead of lower-priority requests. RapidIO supports three transaction request flows. These flows make it possible for application architects to structure traffic through the system and guarantee deterministic behavior. In a CSNA system, the following traffic classification system is used:

- 2: Data and Co-Processing Planes, Sensor I/O
- 1: Control Plane
- 0: Communications Plane

Traffic classification occurs in two locations in a CSNA system. For external communication, the gateway classifies inbound traffic as a function of the IP port number. A table is used in the forwarding path to look up the appropriate RapidIO transaction priority. The IP packet is assigned this priority and is then forwarded into the RapidIO fabric for delivery to a signal processing CN. For internal communication, IP traffic is classified by the source's EoRIO driver. When an IP transaction is queued by the stack to this driver, it inspects the IP port number and classifies the traffic appropriately.

Availability or Fault Management

Availability management begins with the fact that the Sensor Gateway has dual-redundant network feeds. Each network feed is connected to a unique Sensor Gateway. Only a single gateway is forwarding Ethernet frames to

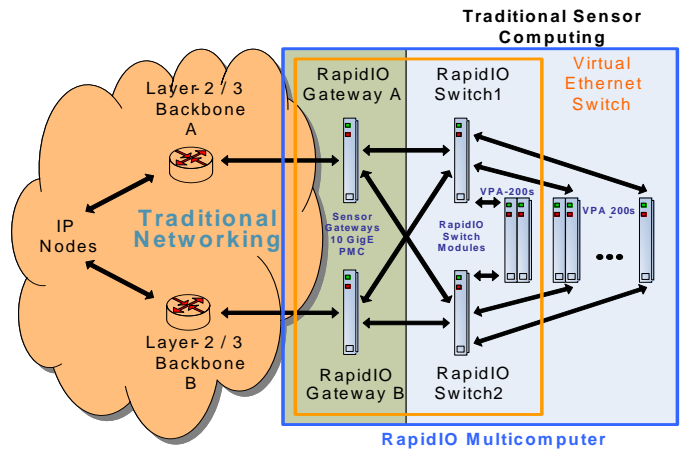


Figure 1: Sensor Gateway Architecture.

and from the network at one time. A network-management function runs inside the CSNA system to handle special control packets involved in computing the cost of redundant routes, enabling the CSNA system to participate and react to a network failure. Signal processing nodes inside a CSNA system track the active gateway through Address Resolution Protocol (ARP) responses. As a consequence, socket-based applications are not aware of network topology changes.

Only a single RapidIO switch module is active for a particular RapidIO source-destination pair. A backup route is made active by updating the immediate crossbar routing table in the path from the source to the destination. Failover to the backup path is fast, and typically occurs before the network application “sees” the fault.

A dedicated error-control channel is provided on each VPA-200. Any runtime RapidIO fabric or DMA errors are reported immediately to a system-management application. The error data includes source- and destination-node IDs, as well as the active RapidIO switch identification. This enables the application to completely coordinate the correct failover strategy.

Performance

Fast path performance on the current prototype is within 75% of the peak bandwidth of the 10 GbE PCI-X based PMC. Complete performance results will be reported in the full paper. Even with the initial results, it is clear that commodity SoC processors are appropriate for high-bandwidth packet forwarding. As eight or more cores become the norm, enough CPU cycles will reside in a single device to terminate 10 Gbps of TCP/IP traffic in software. A renewed focus by processor manufacturers on power consumption will have a continuing positive impact on this issue. The Sensor Gateway and the CSNA architecture provide an attractive path to reach network-centric signal processing.

References

- [1] Rich Seifert, *The Switch Book: The Complete Guide to LAN Switching Technology*, John Wiley & Sons, June 2000.
- [2] K. Fujisawa, A. Onoe, Transmission of IPv6 Packets over IEEE 1394 Networks, IETF RFC 3146, October 2001.
- [3] J. Chu, V. Kashyap, Transmission of IP over InfiniBand (IPoIB), IETF RFC 4391, April 2006.