PRODUCT DATASHEET

StarLink II

Switched Fabric Data Interconnect PMC

Features

- Switched Fabric Interconnect PMC
- Up to 440MB/sec sustained uni-directional throughput
- Up to 500MB/sec sustained bi-directional throughput
- 64-bit, 133MHz PCI-X Interface
- PCI to StarFabric bridge (2 ports)
- 6-port StarFabric switch
- 4 ports via PMC Pn4 connectors or
  - 2 ports via front pannel RJ-45's &
  - 2 ports via Pn4 connectors
- Robust LVDS physical layer
- Low power operation
- Supported by Inter-Processor Communications (IPC) Software
- VxWorks driver
- Integrity driver
- Linux diver
- Air-cooled and conduction-cooled versions available

Many signal processing problems demand the use of multiple processors to achieve real-time throughput and response times. For applications of this type, it is invariably necessary to share large amounts of data between the processing nodes. The StarLink II PMC adaptor provides a modern technology solution to this problem with many specific benefits for the military aerospace systems designer.

The StarLink II PMC card provides the user with a flexible, switched fabric board interconnect system that easily scales from a few to many boards. The flexibility of the system is a virtue of the underlying packet switching technology, where data is automatically and transparently routed through a fabric network of switches to its destination.

StarLink II is based on StarFabric technology from Stargen Inc. StarFabric is a high-speed serial, switched-fabric technology. The system is based on two types of device, a PCI to StarFabric bridge, and a StarFabric switch. The elegance of StarFabric is in its simplicity. Memory attached to one node, is made visible in PCI address space to the other nodes in the network. Data movement between nodes is accomplished with the simple act of reading or writing to local PCI memory space which results in a data transfer across the fabric to the destination node.

Mission-Critical Features

StarFabric was designed to meet the many demanding requirements of high-availability systems used in the telecom industry. Many of these requirements have parallels in military real-time computers, such as fault detection and recovery, redundancy, quality of service and low power consumption. The StarLink II interconnect system provides a rich set of these features, combined with a low latency, high throughput data flow capacity.

Features and Performance

The StarLink II PMC card is implemented with two StarFabric devices, a PCI to StarFabric bridge and a six port StarFabric switch. The bridge provides two ports that are connected to the switch. The remaining four ports of the switch are accessible externally. Systems are constructed by simply interconnecting between ports on the cards involved. See Figure 1, StarLink II Block Diagram.

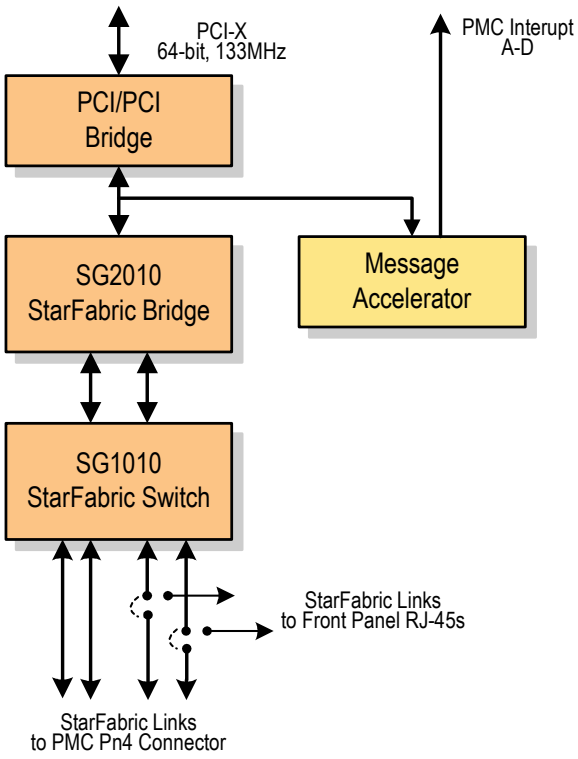

Figure 1: StarLink Block II Diagram

StarFabric links are point-to-point, full duplex, and operate at a link rate of 2.5Gbps. Accounting for the overhead of 8B/10B encoding and packet headers, each link is capable of sustained transfer rates of 220Mbytes/sec, in each direction simultaneously. It is possible to logically bundle two links together to create 440MB/sec connections between nodes. The fabric will automatically recognize the parallel path and permit two links to behave logically as a single, higher bandwidth connection.

The StarLink II PMC is fitted with an FPGA that serves as a hardware assist to accelerate messaging software that communicates over StarFabric. Software, such as CWCECs Inter-Processor Communications can take advantage of hardware controlled memory buffer management that alleviates overhead on the host processors, and reduces message latencies.

Flexibility

The StarLink II interconnect system can be configured in many different topologies. The system supports simple point-to-point connections, basic mesh topologies, and more elaborate topologies with redundant connections for high availability consideration. The system supports routing of packets through up to seven switches, so systems can be scaled to extremely large numbers of nodes. A high-availability system can be constructed by providing for redundant routes between end points. A failure in any one connection will be detected and automatically re-routed over the remaining good connection. Failures are reported so that the application software can take appropriate action.

The StarLink II PMC adaptor provides both the StarFabric bridge and a switch. An interconnected DSP system is constructed solely of these components, with associated interconnecting wiring or backplane. There are no active backplane overlay modules, active hubs, or special cards required. As a result, the logistics costs of maintenance and sparing are the minimum possible. During a development project, reconfiguring the system requires little more than re-arranging standard category-5e cables and re-initializing the software to perform the network discovery process.

Reliability/Availability

StarLink II provides a number of features that permit the construction of high availability systems. The link layer incorporates 8B/10B encoding and CRC checking between nodes, and will retransmit frames as needed. This checking is done on the fly by the hardware, incurring no performance penalty. Transmission failures are reported to the software layer. It is not necessary to add additional "test channel" software to monitor the integrity of the network.

A StarFabric network can be arranged in many different topologies, so as to suit the data flow and redundancy requirements of the system. For high availability systems, it is possible configure a network without any single points of failure. This is done by ensuring that redundant paths exist between nodes in question. During initialization, the fabric will pre-determine alternate routes between nodes, and automatically re-route data to accommodate a failure in wiring, or in a switch.

Physical Layer

The physical layer of StarLink II is based on Low Voltage Differential Signaling (LVDS) connections, operating at 622Mbps. Connections are full-duplex, and consist of four pairs in each direction. (16-pins total). Note the natural fit to PMC and VME standards that provide 64-pins of connectivity. (Pn4 and VME P2). The 8B/10B encoded LVDS link layer is well suited to the electrical environment found in a typical embedded computer system. StarFabric signals can be carried on standard Category-5e unshielded cables up to 12.2M in length. StarLink II can be used between cards in a backplane, and also chassis to chassis. This physical layer has many benefits applicable to the design of rugged deployed systems such as:

- Use of conventional PWB materials and tracking techniques in backplanes
- Use of standard VME and cPCI connectors
- No need for coaxial cabling routed to the backplane connectors
- No need for expensive co-axial 38999 connectors at the chassis
- No extra termination networks
- No TTL signal quality issues, with edge rates affected by temperature
- Use of standard Category 5e cables for development

Software Support

The StarLink II card is supported for use with a variety of real-time operating systems. To afford the highest level of hardware abstraction, CWCEC recommends the Inter-Processor Communication Library (IPC) designed to enable high performance, low latency message passing. IPC allows processors to communicate task-to-task on the same card or within the StarFabric network. It also provides low-overhead block data transfers, segmented block data transfers and signaling between processors to assist in high bandwidth data movement. See the IPC Library datasheet for more details.

For applications that do not need the functionality of IPC, CWCEC provides drivers under VxWorks, INTEGRITY and Linux. This driver provides discovery, initialization and node identification functions. Once the initialization of the network is complete, processing nodes can transfer data between themselves with conventional memory-to-memory mechanisms, including, but not limited to DMA transfers. Contact CWCEC for further information on supported host SBCs and DSP products.

Options

The StarLink II PMC is available in two configurations, one with all four ports connecting to the Pn4 connector for backplane I/O usage, and the other with two ports on the PMC front panel and two ports on the Pn4 connector.

CWCEC supplies rear-panel transition modules (RTM) or cables that plug into the P0 and P2 connectors at the rear of the backplane, providing access to the four StarFabric ports via RJ-45 connectors. The RTM and cabling systems support one or two StarFabric modules per slot.

Application Example

An example application is presented in the following section, with two diagrams representing the example system from different perspectives.

In Figure 2, Application Example, System Block Diagram, an example of a small processing system comprised of a Single Board Computer and two quad PowerPC DSP boards is illustrated. The DSP boards each carry two StarLink II PMC's. Doing so, provides each DSP with the highest possible I/O bandwidth. In many systems, one StarLink II PMC will be sufficient to manage the data I/O requirements of the application.

The network is established with the external cables (or backplane or overlay) that connect between the various cards in the system.

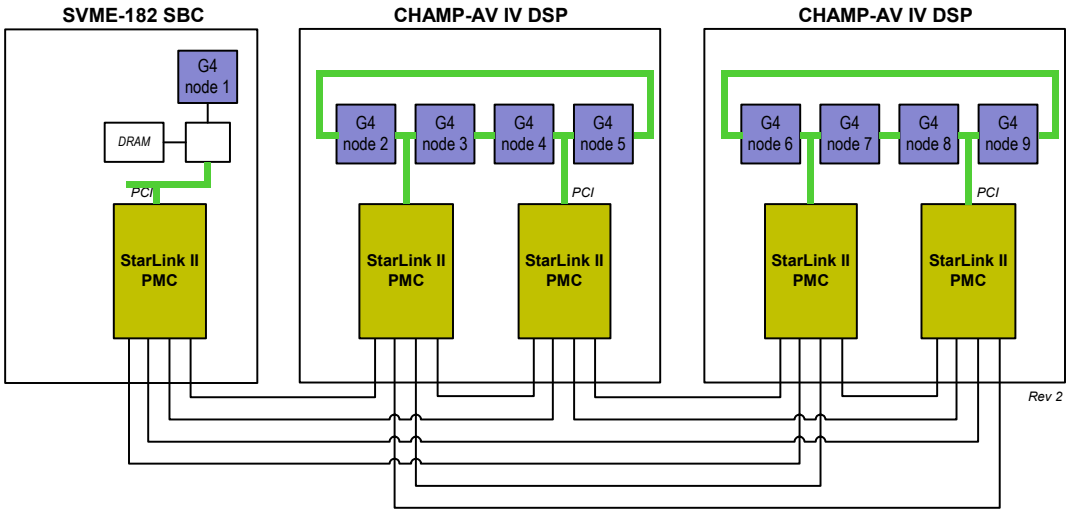

Figure 2: Application Example: System Block Diagram

The same system is depicted from the StarFabric topology perspective in Figure 3, Application Example Topology illustrates the same system from the StarFabric topology perspective. With 5 switches in the system, it is possible to implement a full mesh as depicted where every node has a direct path to each other node. Note that a fully functional StarFabric network could be implemented with just four external links, identified in Figure 3 as the minimum required external cabling. The full mesh implementation will support a greater total bandwidth and redundant connections that would operate in the event of a cable failure.

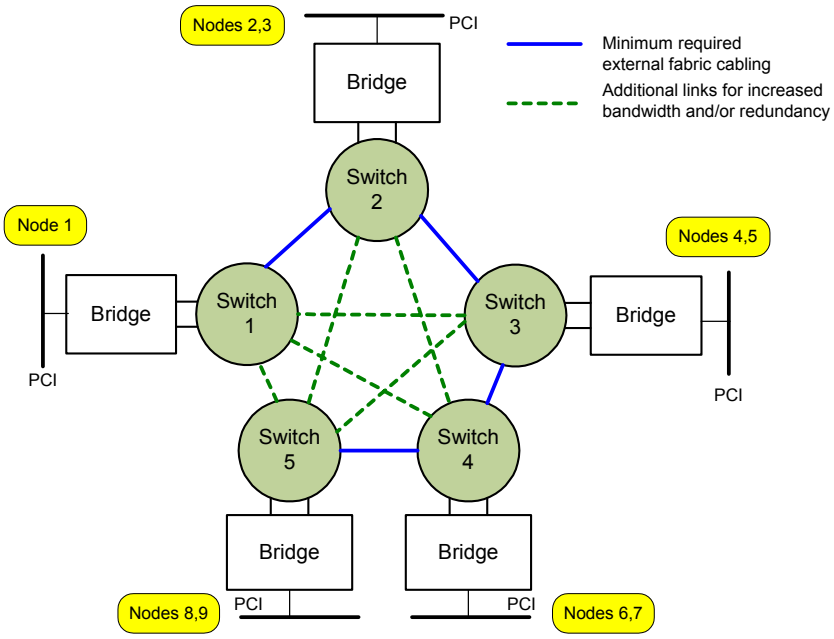

Figure 3: Application Example: Network Topology

Ordering Information

The StarLink II PMC is ordered with the following part numbers. SPMC-233-xyyy denotes air-cooled versions of the product, where "x" defines the ruggedization level, (0,1,2 etc.) and "yyy" identifies a specific configuration. DPMC-233-xyyy denotes conduction-cooled versions of the product, following the same conventions. A formal quote from CWCEC or authorized representative will provide a complete part number and description of the configuration.

Specifications

Table 1: Specifications

RUGGEDIZATION LEVELS*		
SPMC card	Available in level 0 and 100	
DPMC card	Available in level 100 and 200	
POWER REQUIREMENTS		
+3.3V	1.3A Typical	1.8A
DIMENSIONS		
	Size	**Weight**
SPMC card	per IEEE 1386.1	<150 g (<0.33 lb.)
DPMC card	per IEEE 1386.1 (VITA 20-199x)	<160 g (<0.35 lb.)

* See Ruggedization Guideline

Contact Information

To find your appropriate sales representative, please visit:

Website: www.cwcembedded.com/sales or

Email: sales@cwcembedded.com

For technical support, please visit:

Website: www.cwcembedded.com/support1

Email: support1@cwcembedded.com