PMC-StarLite

DSP SIGNAL PROCESSING SOLUTIONS

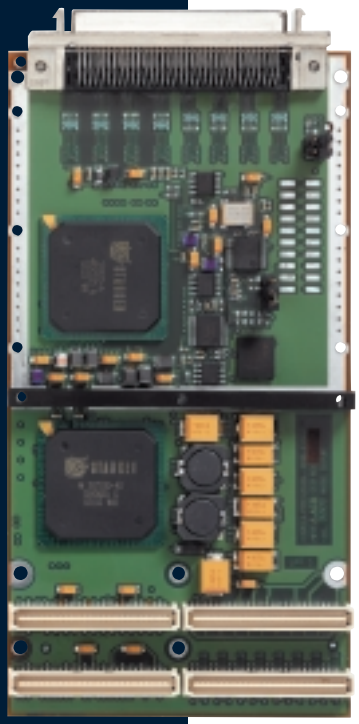

- StarFabric inter-board Interconnect PMC
- StarFabric-to-PCI bridge
  -32/64 bit, 33/66 MHz PCI Interface & 2 StarFabric Ports
- StarFabric 6 port switch
  - 5 Gbps peak (duplex) aggregate bandwidth on each port
  - 30 Gbps peak Non Blocking Switching for High availability
- 4 StarFabric ports to the front or rear
- 2 StarFabric ports to StarFabric-to-PCI bridge
- Additional passive any-to-any front/back plane, low cost & maintenance
  - Scalable over limitless board sets in the VME environment
  - Additional front or rear routing adapter for development
- VxWorks drivers for Radstone's G4DSP and PPC SBC cards
- Air and Conduction-Cooled

Product Overview

The PMC-StarLite is a Switched-Fabric interconnect PMC, designed for linking Radstone's VME boards together, particularly Radstone's flagship digital signal processing product, the G4DSP.

As signal processing problems become increasingly demanding of bandwidth, using many multiples of processors and boards, there needs to be a corresponding method of transferring large amounts of data between processors-both on-board and off-board. The PMC-StarLite is designed to dramatically increase inter-board communication bandwidth, over existing shared bus technologies.

The underlying technology of the PMC-StarLite is StarFabric from StarGen Inc. StarFabric is a high-speed serial, switched-fabric technology which exist as a switch and PCI endpoints-termed edge nodes. StarFabric maps directly into PCI address space. The Fabric appears to a software application as multiple PCI domains separated by non transparent bridges.

Each PMC-StarLite provides a 6 port StarFabric switch with each port capable of 5 Gbps peak communications, (in the form of 2.5 Gbps peak transmit + 2.5 Gbps peak receive duplex). The six-port switch is connected to a StarFabric-to-PCI bridge (64 bit/66 MHz) to provide a software seamless connection between multiple compute nodes.

Four of the six ports are routed to the PMC front panel for front I/O or to the PMC P4 connector for rear I/O. The remaining two ports are connected on-board to the PCI-to-StarFabric bridge. For development, a front or rear panel adapter can be fitted so that the user's fabric network routing can be configured using standard 'Ethernet' (CAT-5 RJ45 terminated) cables. This way, unlimited options are available for connecting heterogeneous board sets.

The PMC-StarLite is available in air and conduction-cooled formats and a VxWorks driver is available for use with Radstone's G4DSP and PPC SBC products.

RADSTONE TECHNOLOGY

Why settle for less

StarFarbric Introduction :

StarFabric from StarGen Inc. is a networking solution based on non-blocking packet switching to provide 99.999% reliable high-bandwidth data communications between compute nodes in a flexible network. For more information on StarFabric visit the StarGen website at:
(http://www.stargen.com)

As a brief outline, there are two main StarFabric components used on the PMC-StarLite, these are:-

i. The StarFabric to PCI bridge-the SG2010.
ii. The 6 port StarFabric switch, the SG1010.

StarGen's PCI-to-StarFabric bridge device, the SG2010, interfaces 64 bit or 32 bit PCI buses operating at 66 MHz or 33 MHz to StarFabric. The bridge translates PCI traffic into serial framed format for transmission across the packet switched fabric. The bridge function supports legacy address routed traffic which provides 100% compatibility with existing PCI software including configuration, BIOS, OS and drivers.

The SG1010 StarFabric 6-port non-blocking switch device routes the frames and operates credit based flow control in concert with the SG2010 that prevents the loss of any packets and reserves bandwidth for high priority traffic. The switch has 30 Gbps peak aggregate throughput.

Many StarFabric features are common to both devices:

At the physical layer, StarFabric uses popular modern high-speed serial communication technology. Each port consists of 4 pairs transmit and 4 pairs receive, and operates IEEE 1596.3/ TIA/EIA-644 2.5V LVDS at 622 Mbps. Each port therefore operates at a peak of 2.5 Gbps full duplex. The clock is encoded in the data so that there are no sideband signals to worry about. Coupled with this, the signaling employs 8B/10B encoding and CRC on all frames, which combines to give a robust, open-standard physical layer:

At layer 1, StarFabric has very high fail-over capabilities to maintain five nines Quality-of-Service (QoS). StarFabric links continue to work if up to 3 out of 4 pairs are broken. The data is re-stripped and resent without any packet loss, which adds robustness to cable connections. To increase bandwidth on a link, ports can simply be bundled together (by register settings) to operate as one link. The links are Hot Pluggable, though the software in driver code is needed to detect the new configuration.

One of StarFabric's chief advantages over many of the new switched fabrics is that it supports legacy PCI addressing as well as source-destination path (channel) based addressing.

The inclusion of PCI-addressing means that it can it can be 100% compatible with existing PCI, BIOS, drivers, application software and operating systems. This minimizes impact and hence time to market.

However, in connecting multiple computing nodes together, the concept of path addressing where packets contain source and destination addresses is advantageous. With path based routing, multiple traffic classes, including Multicast and Isochronous, can be supported. This allows StarFabric to carry heterogeneous traffic, such as voice, video and asynchronous data simultaneously.

The use of credit-based flow control for each traffic class also means that bandwidth for different classes can be reserved such that the network can be optimized for a particular application.

PMC-StarLite Architecture:

Simplicity with balanced bandwidth.

The PMC-StarLite's architecture is designed for optimally matched bandwidth and maximum user flexibility. The two StarFabric ports from the PCI-to-StarFabric bridge are routed to two of the ports on the StarFabric switch. This means that the StarFabric I/O data rate is matched to the 66 MHz 64 bit PCI bus device interfacing to the PMC-StarLite. The four remaining ports from the StarFabric switch to be routed to the outside world. (See Figure 1)

Two ports are necessary between the bridge and the switch to support the DMA operations on PCI. Each 4 bit StarFabric port is approximately 2.5 Gbps full duplex, (i.e. 2.5 Gbps peak in each direction simultaneously). Due to the use of 8B/10B encoding which is commonly used in high speed serial data communications for line balancing, only 8 out of 10 bits of StarFabric data contains useful information. Therefore, each StarFabric port carries approximately 2 Gbps or 250 Mbytes/s full duplex. The PCI bus however, is not duplex and can only operate at full speed in one direction at a given instant. 64 bit, 66 MHz PCI has a theoretical peak bandwidth of 528 Mbytes/sec. Therefore, two StarFabric ports are required to support full speed single direction PCI transfers which occur during long DMA operations. Working from this, during normal bi-directional traffic flow, it can be shown that two StarFabric ports have almost double the capacity of a 66 MHz 64 bit PCI bus and therefore will not create any I/O bottleneck.

The overall result is a fully balanced bandwidth architecture that allows the full potential of any attached PCI Master to be realized.

Other features of the architecture include serial EEPROMs fitted to both the bridge and switch device for true PC-Like plug-and-play compatibility. The bridge also has 2 Mbyte of FLASH, sufficient to contain a lightweight embedded application.

The board is fitted with 8 LEDs which indicate the synchronized operation of both internal and external links. Also fitted is a programmable device which can by configured by Radstone for the user to utilize all of the 8 general purpose I/O lines provided by each StarFabric chip. By default, the programmable device is used to interface the PMC-StarLite's switch diagnostic port to a PC host. This allows access to view and update the switch's internal registers, which aids optimization in demanding applications.

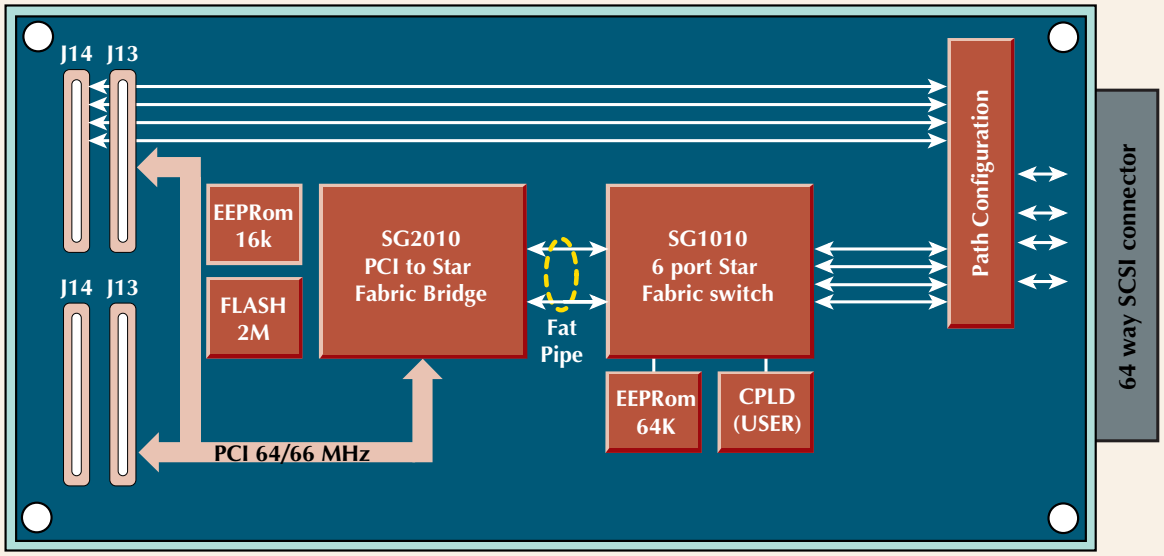

Figure 1: PMC-StarLite Electrical Architecture

Front-I/O

The PMC-StarLite provides 4 ports at the front or rear – for maximum flexibility, scalability and ease of maintenance.

For front I/O the StarFabric ports are routed to a 68-way SCSI connector for inter-board connectivity. From that connector there are 4 options:

- CAT-5 Ethernet wiring using an RJ-45 adapter
- SCSI-II cabling, SCSI to SCSI, or SCSI to RJ-45s
- A custom designed flexi-rigid front plane
- Twisted pair ribbon cable via a SCSI to IDC adapter

For instance, in development environments, the front panel StaRJ45 adapter is simply plugged on to the PMC-StarLite's front panel connector and can be securely mounted via a bezel mounting. This provides 8x RJ-45 connectors, (2 per port) which allows simple hot pluggable reconfiguration of the network. Once development is complete, the adapter is removed and an alternative can be used for the finalized deployed system – such as a PCB front plane.

The front plane takes the form of simple passive PCB wire board and may be fabricated on flexible film. Such boards, though requiring careful tracking and custom routing for a particular system, are relatively inexpensive to develop and maintain.

The fact that there are only half as many connectors in the signal path, makes a front plane a better electrical channel. Front planes with SCSI connectors are also much more robust mechanically than RJ45's in harsh environments. Radstone have already built front planes which a flush fit with the metal work of conduction-cooled VME cards and therefore require no alteration to standard ATR chassis. This passive front plane methodology has been used by Radstone to connect multiple SHARC based DSP cards together for a number of years.

Rear-I/O

For Rear-I/O, through PMC P4, StarFabric can be taken directly to a host board and then off the host board via P2 or P0, (providing that the host board provides adequate tracking to P2 or P0) to a passive rear-I/O backplane. The host board or carrier card ideally needs to support differential tracking from PMC P4 to either a StarFabric edge node or through to P0 (P2).

PMC-StarLite Software Drivers

The PMC-StarLite is supported by a VxWorks driver, which has been developed to run on Radstone's G4DSP and PowerPC SBC motherboards running a Radstone VxWorks BSP.

The driver configures the SG2010 Bridge on the PMC-StarLite to work in its Gateway-Only mode so that it will be the only device in the StarFabric to be mapped into the local PCI memory map at power-on.

The same Vxworks driver code is also run on all other processors that have a PCI-to-Fabric bridge on their PCI-bus.

The driver will determine whether or not a particular PMC-StarLite is the root-node for the network, by reading back register settings affected by a jumper on the PMC-StarLite. Any StarFabric network only has one root-node. (Parallel fabrics only have one root-node per fabric.)

The enumeration of the StarFabric network happens automatically at power-on, as part of the StarFabric protocol, and needs no software control.

The driver performs a number of separate tasks:-

- **Discovery**–the driver determines structure of the network connected to the local PCI-to-Fabric bridge (the SG2010). The network just appears to software as PCI domains separated by non-transparent bridges
- **Event Detection**–all fabric events are directed to the root-node where the interrupts that are generated are handled and cleared. The user will be able to install an event specific interrupt service routine for each of the events
- **Doorbell Interrupt Handling**–the driver will handle doorbell interrupts generated by the local PCI-to-Fabric bridge (SG2010). The user will be able to install an interrupt service routine for this interrupt
- **Memory Mapping**–the driver will map blocks of local PCI memory space to access physical memory on other compute nodes in the fabric network
- **Multicast Traffic Provision**–the driver will map blocks of local PCI memory space to write to memory on other boards. It will also allocate/ reserve bandwidth for multicast groups
- **Semaphores**–the driver will allow easy use of Fabric semaphores
- **Remote Configuration cycles**–the driver will allow type 0 PCI configuration cycles to be performed on another board
- **SROM programming**–the driver allows access to program the PMC-StarLite's SROMs so that the power-on configuration of the fabric can be changed

In addition to the drivers developed by Radstone, The PMC-StarLite comes with an extensive Fabric Primitives Library (FPL) developed by StarGen. A generic driver which uses these primitives that are accessible via IOCTL calls, is also available for customers who want to use the full extent of StarFabric's capabilities. This means that the PMC-StarLite is suitable for both DSP and telecomm customers.

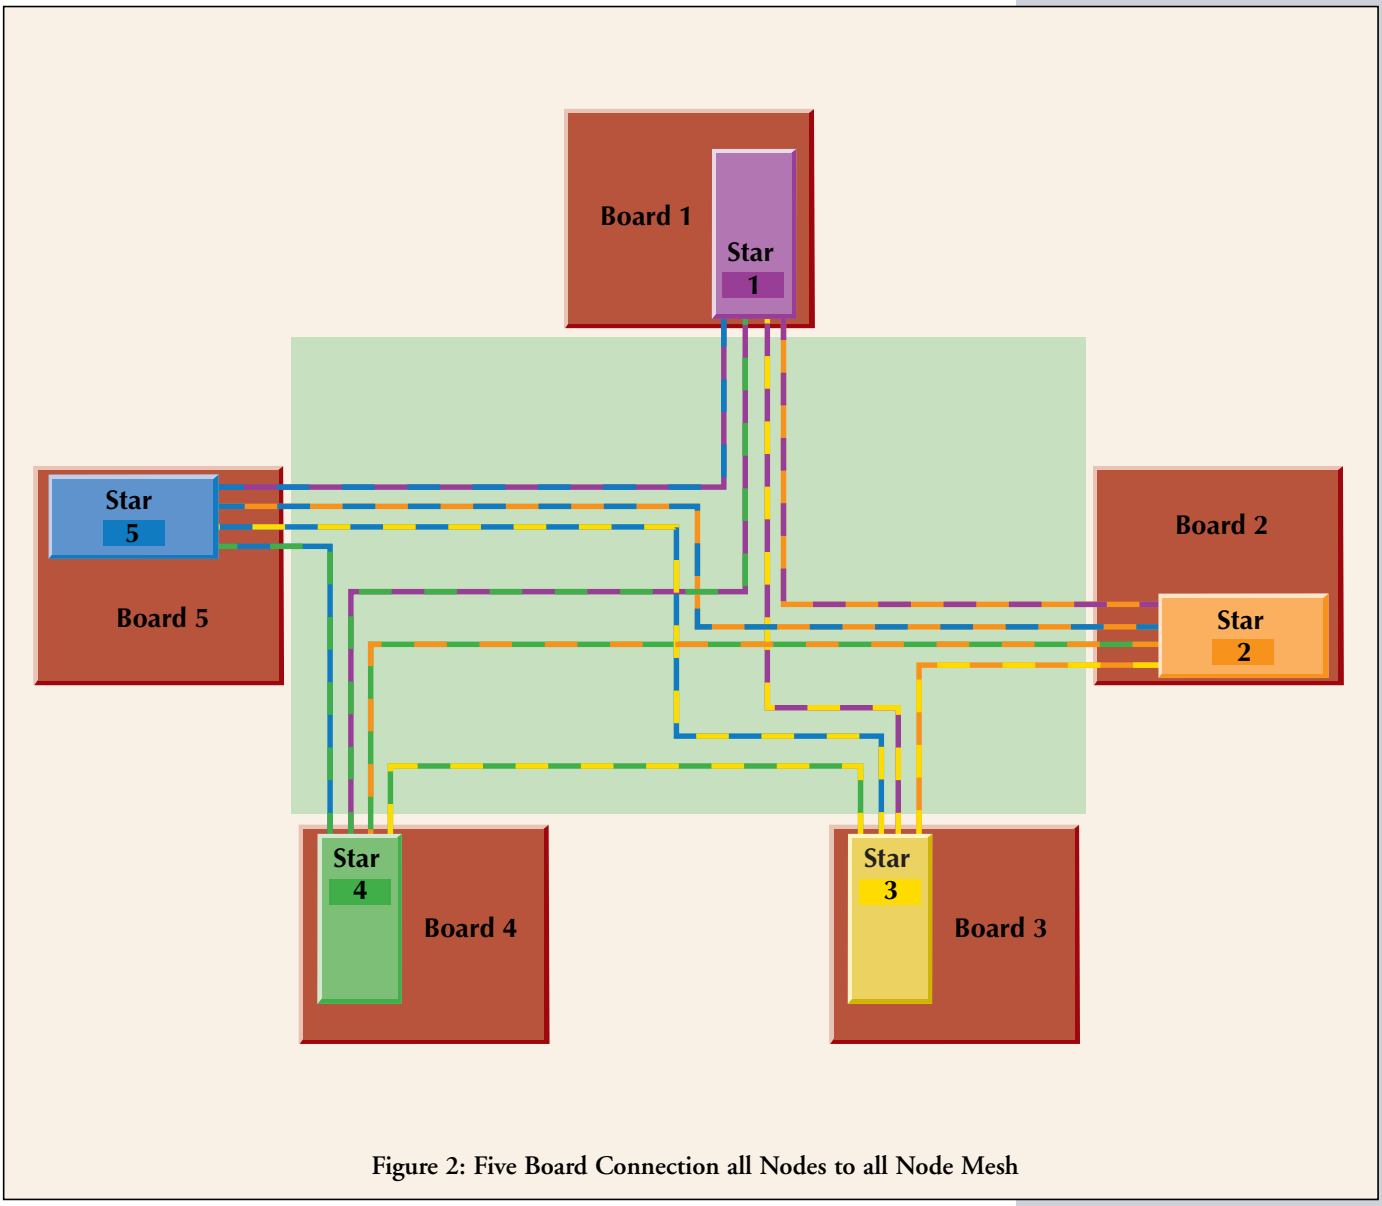

Figure 2: Five Board Connection all Nodes to all Node Mesh

NOTE: Nine boards can be connected in an any-to-any manner with addition of a second PMC-StarLite on each board

Data transfers explained

Figure 3 shows how data transfers can be carried out using a PMC-StarLite. Processors 1-3 can be processors on different boards, or processors with independent memory maps on the same board. The PMC-StarLite resides in PCI space of the local processor. The local processor can access the PMC-StarLite and the PMC-StarLite can access local memory of the processor.

In this example, the PMC-StarLite on processor 1 is configured to map an area of PCI memory that will multicast to processors 2 and 3. This means that data written to this area PCI memory by processor 1 will be written to local memory of both processors 2 and 3 - the data is simulataneously multicast to two locations.

The figure also shows that the PMC-StarLite on processors 2 and 3 are configured so that both of these processors can read from or write to the same area of local memory on processor 1 using local PCI accesses.

Built -In-Test (BIT)

As with all Radstone Boards, the board is bundled with BIT code. For the PMC-StarLite, the initial release of BIT is designed to run on the host board under VxWorks. BIT coverage is best performed in this manner, since the primary aim is to test each of the ports.

The board has low level functional testing and JTAG testing completed at the factory. Failure is indicated early in the set-up procedure, via reporting of PCI configuration cycles, which are sufficient to detect correct operation of both the SG2010 and SG1010.

Applications

StarFabric designs can span from chip-to-chip and board-to-board to room size networks. (See Figure 2) Inexpensive unshielded twisted pair copper cable can yield distances of approximately 5 meters. The fragile link fail-over capability of StarFabric makes it ideal as a chassis to chassis interconnect especially in harsh environments.

Under the StarFabric protocol, if a particular path fails, traffic can be re-routed over an alternative path automatically with no loss of data. The ability to implement parallel (redundant) fabrics makes StarFabric particularly attractive in remote applications where maintenance is difficult.

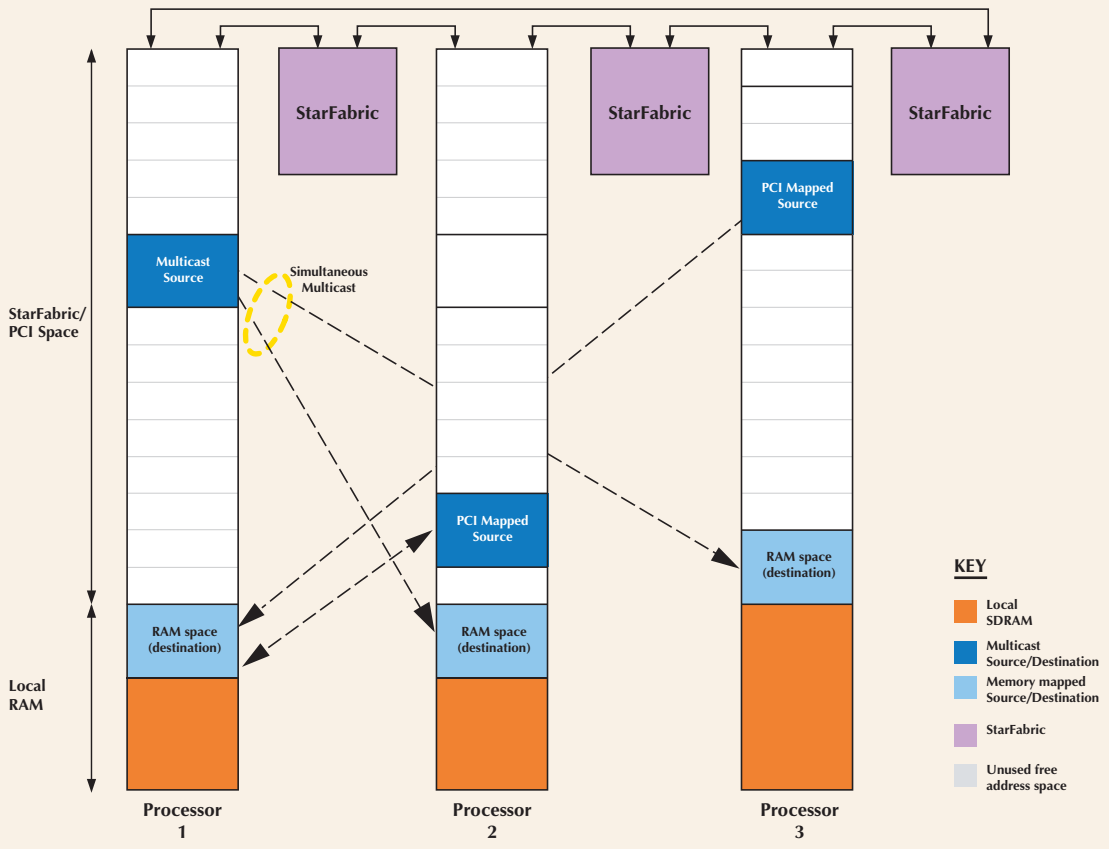

Figure 3: Data Transfer Possibilities using StarFabric between Processor Nodes

Specification

Parameter	Component	Quantity	Comments
Size	single PMC	1	Conduction-Cooled Heatsink layers) 74 x 149 mm
Power	5.5 Watts (Maximum)		
Power Supply	5V		
Weight	87 grams	1	Approximately
V I/O	3v3 or 5v PCI 2.5V LVDS		
PCI interface	32 bit/33 MHz-64 bit/66 MHz		3v3 or 5v signalling
User rearI/O	StarFabric ports	4	
FrontI/O	StarFabric ports	4	
Feature			
PCI end point	SG2010A	1	StarFabric to PCI interface : 64 bit 66 MHz
StarFabric Switch	SG1010A	1	6 Port StarFabric Switch
FLASH	AM29LV160	1	16 Mbytes-(2M x 8) on the SG2010
EEPROMs	ATMEL 25640N	1	64 Kbytes on SG1010
	Microchip 25AAL160	1	16 Kbytes on SG2010
Control CPLD	Xilinx XC9572XL	1	Reset, synchronisation and StarView interface
3v3 Power Supply	Elantec 4562 2 Amp switching	1	Efficient for good thermal ability
1v5 Power Supply	Elantec 4562 4 Amp switching	1	Efficient for good thermal ability
Clocks	VCC1-3BE-62M208	1	62.208 MHz crystal : 45/55 sym. 20ppm. Ind temp
Front I/O connector	Amp 787394-7	1	68 way SCSI-II Sheilded

Note: For ordering information, further options and accessories please contact your local Radstone sales office.

UK
Radstone Technology
Water Lane
Towcester
Northants NN12 6JN
Telephone: +44 (0) 1327 359444
Facsimile: +44 (0) 1327 359662
E-mail: sales@radstone.co.uk
Videoconferencing: +44 (0) 1327 359495
(by arrangement)

USA
Radstone Technology Corporation
50 Tice Boulevard
Woodcliff Lake, NJ 07677-7645
Telephone: +1 800 368-2738
Facsimile: +1 (201) 391-2899
E-mail: sales@radstone.com

Eastern Area: +1 (201) 391-2700
Mid-Atlantic Region: +1 (201) 391-2700
Southeast Region: +1 (321) 308-0240
Western Region: +1 (909) 974-1141
Central Region: +1 (480) 964 5407

FRANCE
Radstone SARL
10 Av. du Québec
BP 116
91 944 Courtaboeuf Cedex
Telephone: +33 (0) 1 64 46 04 03
Facsimile: +33 (0) 1 69 28 03 40
E-mail: radstonefr@radstone.com

RADSTONE
TECHNOLOGY

Meet us on the Web at www.radstone.com

Printed in UK Publication RT259 4K09/2002