

SCRAMNet+

Shared Memory—Speed, Determinism, Reliability, and Flexibility For Distributed Real-Time Systems

Abstract

Using multiple computers in a single real-time system requires that accurate outputs be generated on time, every time. Any network that connects advanced real-time systems must have the low-latency and determinism to support highly dynamic applications, the throughput to support vast amounts of data in highly complex applications, and guaranteed data delivery to ensure that critical data cannot be lost.

The Shared Common Random Access Memory Network (SCRAMNet+) from Curtiss-Wright Controls Electronic Systems meets the unique demands of dedicated real-time systems such as advanced aircraft flight simulators where multiple computers are used to solve an intricate real-time problem. SCRAMNet+ is a Commercial Off The Shelf (COTS) real-time network that provides low data latency, system determinism, high throughput and guaranteed data delivery while providing flexibility to the real-time designer.

Introduction

In SCRAMNet+'s replicated shared memory network, each card stores its own copy of the network's shared memory set. The SCRAMNet+ card's host computer accesses this memory as if it were its own. Any time data is changed in one SCRAMNet+ node's memory, the network protocol automatically broadcasts the changes to all other nodes via the SCRAMNet+ serial data link. The SCRAMNet+ Network's transmission speed and unidirectional ring topology

cut data latency to a minimum and provide network determinism.

Much of the speed advantage of the replicated shared memory network comes from the fact that it is solely a hardware communications link. Once the network node is properly configured, no additional software is required to make it operate as a real-time link. No real-time drivers are needed, and there are no time consuming software routines needed to pack, queue, transmit, de-queue, and unpack messages. The ring transmission methodology automatically sends all updates to every node on the network, and all SCRAMNet+ nodes can write to the network simultaneously. This eliminates overhead such as node addressing and network arbitration protocols, and reduces system latency.

If a value in SCRAMNet+ memory is updated by one computer of a typical ten-node ring, this new value will be available in the memory of all other computers on the ring within 9.0 microseconds! This is the total application-to-application communications time (latency) from when a value is changed on one



computer to when the value is available in memory to the farthest computer on the ring. Interrupts are communicated over the ring with the same speed and simplicity of data. For the 10-node network example above, a master CPU could interrupt all nine other nodes within the same 9.0 microsecond elapsed time.

When using a distributed architecture in highly dynamic real-time systems such as aircraft flight simulators, this low-latency communication is an absolute necessity. For example, effective fighter jet simulations must subject pilots, aircraft systems, and weapons systems to the most realistic representation of air combat possible. Any loss of realism severely degrades the simulation's value as a training and evaluation tool.

In this atmosphere, where simulated objects travel at speeds of Mach 2 and exert wildly variable course changes through three dimensions, updates to the system must be made as frequently as possible to accurately plot the positions of moving aircraft. A node latency of even a few hundred microseconds can destroy the realism of the whole simulation, especially when dealing with sensitive electronic components like missile radar systems. In the 10-node SCRAMNet+ ring in the previous example a 9.0 microsecond worst-case update latency would ensure that a simulated aircraft traveling at Mach 2 could travel less than 0.6 centimeters between position updates.

Variable Bandwidth

The speed of SCRAMNet+ provides high throughput, and SCRAMNet+ has the capability to automatically increase its network data bandwidth through the use of variable-length data packets. As a node's throughput demands increase, the hardware automatically packs the data into larger packets minimizing network overhead, while increasing throughput to meet the demands of the system.

High Effective Bandwidth

By eliminating unneeded network traffic, data filtering can increase effective bandwidth exponentially. In many cyclical real-time systems, algorithms and tasks execute at fixed, clock-driven frame rates (for example, 100 Hz, 50 Hz, 25 Hz, etc.). This synchronous

architecture scheme leads to deterministic system performance, but results in new data values being computed more frequently than is always necessary. There may be many model iterations where model outputs are actually unchanged, and without some method of detecting unchanged output data, they can needlessly consume network bandwidth. The SCRAMNet+ design incorporates a programmable data filter to eliminate these unnecessary updates. SCRAMNet+ data is only transmitted to the other nodes on the network if those "writes" to the shared-memory produce a data value change or if the location is set to cause an interrupt on the network. This data filtering technique has been shown to significantly increase the effective throughput of the network. Actual measurements on an aircraft flight simulation application have shown that this technique filtered out approximately 75% of the network traffic. This alone increased the "effective" bandwidth of the network by 400% for this application.

System Determinism

The particular ring design used by SCRAMNet+ is very deterministic. The following characteristics apply when a four-byte data packet is employed.

- ◆ Messages are fixed in length; 82 bits long.
- ◆ The ring access time is between 100 nanoseconds and 800 nanoseconds.
- ◆ The node latency time is between 250 nanoseconds and 800 nanoseconds per node, depending on whether the node is "active" at the time of data receipt.
- ◆ Nominal delays are around 250 nanoseconds.
- ◆ Propagation delay incurred in transmission media is 5 nanoseconds/meter of cabling.

SCRAMNet+ guarantees that data will always pass through a network node within a 550 nanosecond "window", and propagation delay is only dependent upon cable length. Designers can narrowly define the time when a packet will reach any given node on the SCRAMNet+ network. This deterministic ability means that in the 10-node network example above, designers could pinpoint the arrival of a packet at the furthest node on the ring to within 2.5 microseconds.



System Reliability

In any real-time system, accurate data must arrive at the proper location in the correct order or data latency, system configuration, and all other considerations are irrelevant. A typical SCRAMNet+ network has a bit error rate of approximately 10⁻¹⁵, meaning that an error could occur roughly once every 76 days of continuous, 24-hour, 100% bandwidth-saturated network use. Few applications require this level of network saturation, so most users will experience a much longer time between potential bit-errors. In the unlikely event that a network error does occur, the SCRAMNet+ automatic error correction mode detects and corrects errors at the hardware level.

Automatic Error Correction

When operated in the PLATINUM mode, SCRAMNet+ provides automatic error detection and correction. This mode of operation insures that data and interrupts arrive error free and in the proper order with minimal infringement upon network bandwidth. It is fast, efficient and user transparent.

The source or originating node (as identified by the ID field) is responsible for removing returning packets and comparing them with the packet retained in the insertion buffer. If no errors are detected, then the packet is removed from the network and the insertion buffer. However, in the event of errors, all or a portion of the data packet is re-transmitted from the insertion buffer. If the packet does not return in a pre-determined time frame, a "watch dog timer" is triggered which also causes re-transmission of the entire packet.

While there are other techniques of ensuring data delivery such as redundant transmission of all data packets, these techniques merely reduce error rates to a point where it is highly unlikely that an error will affect data delivery, and redundant transmission cuts network bandwidth in half. The SCRAMNet+ automatic error detection/correction scheme guarantees data delivery and ensures that precious network bandwidth isn't squandered on multiple transmissions of data.

Flexibility

Once the real-time requirements for high-speed, low-latency, system determinism and reliable data delivery have been met, designers can focus on implementing their system. By supporting many types of buses and computer operating systems, flexible topologies, and long-distance connections, SCRAMNet+ gives designers alternatives to find their best solution.

The SC150 Series

The proven SCRAMNet+ SC150 series is designed to provide the speed and deterministic performance of shared-memory while avoiding the data collisions and computer-to-computer latencies often caused by the complex, time-consuming layers of software and hardware protocols used in typical networking systems.

Memory writes to SCRAMNet's replicated shared memory at one computer are instantly sent to all other replicated shared memories on the network at 1.50 Mb/s via high-speed fiber optic or coaxial cable. There are no bus contentions, time-consuming protocols, or non-deterministic message crashes. Data latency is only 250-800 nanoseconds per node. This automatic data transfer requires no software intervention or backplane loading and overall data latency at each computer is measured in mere microseconds!

Up to 256 nodes (each separated by up to 3500 meters) can be contained on a single fiber optic ring, with both data transmission and computation occurring simultaneously. The newest addition to the SCRAMNet+ Network – the SC150e – takes full advantage of the speed and efficiency of today's computers to dramatically improve data throughput between SCRAMNet+ memory and the host computer.

The SC150e boards perform memory-to-host transfers at 70 MBytes/second – over four times faster than other shared-memory networks.

Standard networks may offer sufficient data transfer rates between boards, but they are unable to match the simplicity and superior host-to-board communication achieved by the SCRAMNet+ Network's shared memory architecture.

Multiple Buses

The SCRAMNet+ modular design provides support for a variety of buses. Our custom SCRAMNet+ ASIC provides functions such as message reception and transmission arbitration, network control, memory management, interrupt handling, error detection and correction, and node configuration. This core logic provides easy porting to different host CPUs and enhances system-level network maintenance and troubleshooting because the same core logic is used as the basis of all SCRAMNet+ boards.

This modular approach allows SCRAMNet+ SC150 boards to provide direct support for popular bus systems including PMC, PCI and VME6U form factors. Hardware support for other devices is provided by a "rehostable" SCRAMNet+ board, which can be implemented in non-supported or custom systems. Support for CompactPCI® platforms can be implemented using hardware bridges or PMC carrier boards.

In advanced real-time systems, different processors must perform a variety of highly specialized tasks. Specific computing platforms may be particularly well suited to some of these tasks, while relatively unsuitable for others. The variety of SCRAMNet+ boards available provides designers with the flexibility to use many different computers in their real-time system.

When updating existing SC150 systems, designers can mix advanced Single Board Computers with Silicon Graphics®, DEC™, and Sun™ computers; Legacy ISA/EISA-based PCs; PCI-based Pentium® PCs; and existing, custom-built platforms. When developing systems from scratch, designers are free to choose from a wide variety of workstations and rack-mounted

systems to find the processor which most-closely matches the specific price/performance needs of each function in their design. This flexible approach allows SCRAMNet+ users to pick and choose the most appropriate system for each individual portion of their real-time system.

Long Distance Connections

Designers must also contend with physical distance considerations, and the SCRAMNet+ network provides the capability for long-distance connections between nodes. The SC150 network uses a fiber-optic transmission approach as an integral part of the basic design. The standard 62.5/125 mm wave guide provides a maximum node separation of 300 meters when using 850 nm LEDs. For applications requiring distances of greater than 300 meters (up to 3,500 meters), optional 1300 nm LEDs are available for SC150 series boards. SCRAMNet+ SC150 boards can also employ low-cost coaxial cable as a transmission medium for connecting nodes separated by 30 meters or less.

It is important to note that the different transmission mediums available on SCRAMNet+ SC150 are linked to the network card via a modular media access card (MAC), which plugs onto the main board. This allows most SCRAMNet+ cards to be upgraded to fiber optic or converted to coaxial cable by changing only the MAC.

LinkXchange Switch

For users with more complex switching needs, the SCRAMNet+ implementation of our LinkXchange switch offers up to 16 SC150 ports and can be cascaded to provide additional ports. The LinkXchange switch permits implementation of SCRAMNet+ networks with various subdivisions and topologies, while retaining the simplicity and transparency of the SCRAMNet+ architecture. Because LinkXchange is transparent to the network, using LinkXchange has a negligible effect upon system latency.

The LinkXchange can isolate different SCRAMNet+ nodes or sub-rings. LinkXchange is a full 16x16 crossbar matrix switch and is controlled via RS-232. LinkXchange allows the individual nodes or sub-rings

connected to its ports to be combined into any number of independent rings, and provides total flexibility in the configuration and use of scarce real-time resources.

SCRAMalyzer

The SCRAMalyzer is a portable system used to monitor, record, and analyze the performance of a SCRAMNet+ Network. The SCRAMalyzer collects data traffic in real-time and reports the network throughput and latency through easy-to-interpret graphs and other on-screen displays. Its easy-to-use interface helps you to visually determine the impact of simulation events and system design changes on your real-time network.

The SCRAMalyzer is designed to be minimally intrusive to a network system. Network activity can be automatically logged at all times or only for a duration of time that you specify. The outputs from the monitor are real-time graphical displays of throughput and latency, including various statistical data reports used in post-test analysis. You can select specific areas of SCRAMNet+ memory to record and view, and results can be saved and printed for later analysis.

Potential Applications

The SCRAMNet+ Network has many different commercial and military uses. Shared memory technology is ideally suited for applications with critical control loop timing requirements, such as simulation systems. Each SCRAMNet+ node provides the high-speed transfer of data between the control and feedback systems that may be running on many dissimilar computers in the system. Typical simulation systems include:

- ◆ Aircraft Simulators.
- ◆ Missile Simulators.
- ◆ Power Generation Plant Simulators.
- ◆ Ship Simulators.
- ◆ Land Vehicle Simulators.
- ◆ Mission Planning Simulators.
- ◆ Force-Level Training Simulators.
- ◆ Virtual Reality Systems.
- ◆ Location-Based Entertainment Systems.

Tight control-feedback loops required by many instrumentation and control systems demand the low-latency performance of the SCRAMNet+ Network for accurate, real-time data acquisition and control. SCRAMNet+ provides the nearly-instantaneous availability of critical data that enables system evaluation and adjustment in real-time. Typical instrumentation and control applications include:

- ◆ Flight Test.
- ◆ Wind Tunnels.
- ◆ Test Stands.
- ◆ Robotics.
- ◆ Particle Accelerators.
- ◆ Industrial Process Control.

Summary

To connect multiple computers in real-time, a designer needs a network that combines high-speed, low-latency, system determinism, high bandwidth and guaranteed data delivery. The SCRAMNet+ Network provides these attributes in a mature COTS product that offers flexible configurations, long-distance connections, and support for various buses. Thousands of SCRAMNet+ nodes are currently in use, meeting the speed, determinism, reliability and flexibility requirements of real-time designers worldwide.

Product specifications mentioned herein are subject to change without notice. SCRAMNet and LinkXchange are registered trademarks of Curtiss-Wright Controls Electronic Systems. CompactPCI is a registered trademark of the PCI Industrial Computer Manufacturers Group. Silicon Graphics is a registered trademark of Silicon Graphics Inc. DEC is a trademark of Digital Equipment Corporation. Sun is a trademark of Sun Microsystems. SelBUS is a trademark of Encore Computer Corporation. All other trademarks or registered trademarks mentioned herein are the sole property of their respective owners. © 2004, Curtiss-Wright Controls Electronic Systems, All Rights Reserved.