

Speeding Up SS7 With UltraSPARC Ili

Wayne Fischer, Director of Strategic Programs, Force Computers, a Solectron Subsidiary

Historically, telecommunication systems have been implemented using closed and proprietary hardware and software in an effort to ensure high availability and reliability, and because standard computer telecom technology and hardware has been lacking. Today, telecommunications is being transformed to a distributed, loosely coupled model driven by market demand for new, customized services, and by progress in hardware and software technology. Service provider networks are evolving to a client/server model that will require fail-safe, redundant platforms running standard open operating systems that support multiple telecommunication-centric application software environments.

Signaling System 7 (SS7) is one of the key technologies used to link the telephone system with open computing platforms. SS7 switching hosts add intelligence to the telecom system by utilizing databases that can be accessed to provide value-added services throughout the network. Advanced Intelligent Network (AIN) is another telecom system innovation that has enabled many new custom calling services. For example, calls from pre-determined numbers can be blocked, coded with a distinctive ring, or forwarded to another number in compliance with instructions in a computer database accessed through the SS7 network.

One particularly important open systems telecom application employs network elements called Intelligent Peripherals (IPs) and Service Control Points (SCPs), which are the central office platforms upon which new AIN-based services can be provided. IPs interface to actual telephone lines within the exchange and intercept or create calls using SS7. Local exchange carriers are basing IPs at the central office, making their services available to all subscribers no matter where they are located. Typical services provided by IPs include messages that play to inform the caller that the number of a subscriber has changed, and speech recognition services that provide hands-free dialing for mobile phone users.

As the second leading provider of VMEbus embedded computing products,



Wayne Fischer

Force Computers is dedicated to meeting the needs of telecom vendors for high performance, flexible, rugged and reliable open computing systems. Force Computers builds a rich portfolio of single board VMEbus computers and systems based on SPARC, 68K and PowerPC technologies for embedded UNIX and real-time applications. The company stays abreast of new computing standards and technologies through its participation in the VITA Standards Organization and its chairing of the committees on VME64x, ETL and IndustryPacks on VME64x.

Fastest SPARC embedded platform

Force Computers' newest embedded computing system, the teraforce-50 (see figure 1), provides the ideal platform for IP switching hosts and SS7 processing hosts. The teraforce-50 sets a new performance standard for the Solaris environment, providing a level of computing horsepower that previously could be achieved only with a multiprocessor system at a considerably higher cost.

The speed improvements in the teraforce-50 come from its use of the new UltraSPARC Ili processor, the fastest SPARC processor ever designed for embedded computing. The new processor combines tighter design rules, increased memory and peripheral bandwidth and a new architecture to provide estimated SPECint95 and SPECfp95 benchmarks of 10 and 12 respectively (see figure 2). The UltraSPARC Ili has an integrated I/O, DRAM and UPA64S graphics interface which provide dramatic performance improvements relative to the previous generation SPARC processor. For example, maximum E-Cache read bandwidth has been increased to 1.2 GB/s, maximum DRAM random read bandwidth to 350 MB/s and memory from DRAM to UPA64S has been raised to 550 MB/s.



Figure 1.

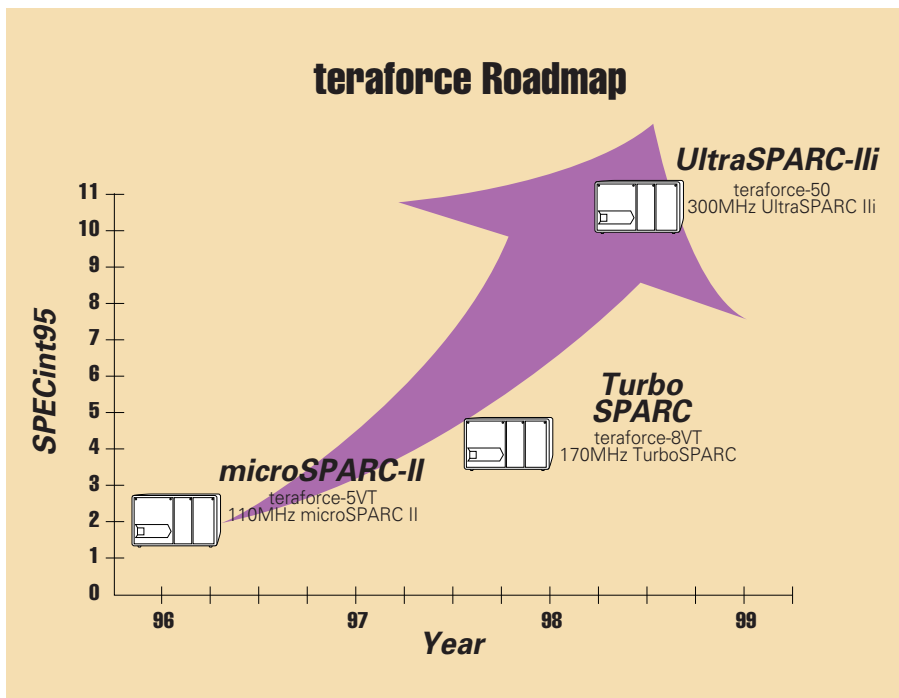


Figure 2.

The new SPARC V9 architecture provides additional performance gains. Each functional area on the UltraSPARC IIi maintains decentralized control, allowing many activities to overlap. Sustained performance of up to four instructions per cycle is supported, even in the presence of conditional branches and cache misses, by a decoupled fetch and dispatch unit. Load buffers on the input side of the execution unit, together with store buffers on the output side, decouple pipeline execution from data cache misses. Instructions predicted to be executed are issued in program order to multiple functional units. These instructions are executed in parallel and may be completed out of order.

Increased Flexibility

The teraforce-50 is designed to be extremely flexible, allowing customers to configure precisely targeted solutions and to rapidly bring those solutions to market. For example, the system provides 18 slots for peripherals—exceptionally high in embedded computers. This was accomplished by taking advantage of the high integration of the UltraSPARC IIi processor to limit processing elements to only two slots. The processor integrates a PCI bus controller to interface directly with the PCI bus; an I/O memory management unit to manage virtual to physical memo-

ry address mapping using translation lookaside buffers for improved performance; an external cache unit to handle instruction and data cache misses efficiently and to provide high transfer bandwidth; and a memory and UPA64S control unit to manage all transactions to DRAM and the UPA64S subsystem typically used for high performance graphics.

Expansion to three slots accommodates the addition of Sun Creator Graphics for applications requiring advanced

2D and 3D graphics, or the addition of up to one gigabyte of memory. Up to four storage modules can also be accommodated, for mixing and matching hard disk, CD-ROM, tape or floppy drives, while dual Ethernet and dual fast SCSI-2 interfaces are standard. For further flexibility, the teraforce-50 provides a choice of back-panel loadable power supplies—110/220 volt AC today for most industrial applications, and a future 48 volt DC version will extend the teraforce family into telecommunications applications which require a DC power source.

Rugged and Reliable

The teraforce-50 meets telephone central office ruggedness and reliability requirements. The VME boards used in the teraforce-50 use dependable Eurocard standard pin and socket connectors (see figure 3). The 2.54mm connectors are designed for use in telecom equipment (IEC-603-2) and are now widely used in industrial control applications as well. Eurocard boards feature card guides for solid rear backplane connectors alignment, front panel retaining mechanisms, injector/extractor handles, and EMC protection features to minimize electromagnetic interference.

The teraforce-50 ships with a rugged, rack-mountable system chassis with a 20-slot VMEbus subsystem, multiple access doors with thumb screws, vertical forced air cooling, user-configurable, back-panel I/O connectors with EMI protection and strain relief. Many competitive systems




Figure 3.

avoid a front door on the chassis by placing card mount handles outside the metal frame of the system. By providing a front door, the teraforce-50 provides additional shock resistance and EMI protection. Power supplies can be swapped out through the rear panel, while all boards plug in easily from the front.

At the same time, the teraforce-50 uses 30 percent less power than UltraSPARC I implementations, so systems stay cooler and embedded applications run more reliably. VME provides tremendous thermal management advantages over the conventional desktop PCs fre-

quently used in embedded applications. The standard desktop design provides the inherent reliability problems of irregular airflow paths that can result in overheating. VME overcomes these problems by providing a clear, even passage of airflow over all active, heat producing cards.

The new teraforce-50 combines the high performance and low power consumption of the UltraSPARC Iii processor with Force Computers' embedded computing expertise to deliver a flexible, reliable, space-saving solution for high-end embedded control applications, and to move the teraforce platform into the

telecommunications applications arena. The flexibility of the teraforce-50 opens up a wealth of possibilities for both systems designers and for Force Computers. From an inexpensive platform for drafting or running telecommunications applications to a highly reliable UltraSPARC Iii workstation for industrial applications to a powerful Solaris-based graphics workstation—the teraforce-50 covers a lot of computing territory.

Visit Force Computers at <http://www.forcecomputers.com>

VMEbus FAQ

Frequently Asked Questions for the Beginning VMEbus User

John Rynearson, Technical Director, VITA

Question: What are the issues surrounding live insertion for VME64x?

First we need to understand what live insertion means and in what context it is being defined. Live insertion usually means that an electronic module can be removed and then reinserted into a system while the system remains under power. Another popular term for "live insertion" is "hot swap". The assumption is that removal of the module and reinsertion will cause no electrical harm to the system. Why is this capability desirable? Systems are usually powered down for either repair or reconfiguration. However some systems are configured in such a manner that it is not acceptable to power down all or part of the system. Telecommunications system fall into this category. Other systems may control industrial equipment such as robots, sorters, etc. that must remain powered up. Whatever the reason live insertion is becoming a desirable feature for many systems.

A Design Issue

Live insertion is really a system design issue. What are the application requirements that make live insertion a desired

feature? Including live insertion support in a system costs money. If your application doesn't require live insertion, then save money and don't put it in. If your application does require live insertion, then a careful analysis must follow.

System Operation

First, must your system continue to function during live insertion? If the only requirement is that the system remain powered, but that its functions may cease temporarily, then a simple method for "safeing" the software and the hardware that it controls is all that is required.

On the other hand if the system must continue to function then a careful analysis of the system at a modular level is required. Each module must be examined to determine how its possible failure and replacement can be handled.

Conventional, High Availability, and Fault Tolerance

A conventional computer system will have many single points of failure. That is, a failure in a specific component will cause the entire system to cease functioning. Systems that fall into this category cannot really maintain functionality when a component fails. At the other end of the spectrum are fault tolerant systems. A system is fault tolerant if it is

defined to have no single point of failure. Fault tolerant systems are constructed with redundant components so that if one component fails another is available to take its place. However, fault tolerant systems are expensive and may not always be necessary. A less expensive approach is a "high availability" system. High availability systems don't meet the definition of fault tolerant systems in that they may have one or more single points of failure. However, they can be used in applications where certain functions need to be redundant and others don't.

Repair, Reconfigure, or Both

Is the purpose of supporting live insertion for repair, reconfiguration, or both? Supporting live insertion for repair assumes that when a module goes bad the system can detect the failure, take the affected module off-line, and notify the system operator. When the module is replaced the system must detect the replacement, qualify the new module for system operation, and then bring the module smoothly back on-line. Reconfiguration, on the other hand, can be done as either replacement, enhancement, or both. For example, a trivial case would be to install additional memory into a system. When the module is installed, the system must be notified that a new resource has been added. The system must then qualify the