SLAC - PUB - 4781
November 1988
(T)

IEEE P1596, a Scalable Coherent Interface for GigaByte/sec Multiprocessor Applications*

David B. Gustavson
Stanford Linear Accelerator Center
Stanford, CA 94309

Abstract

IEEE P1596, the Scalable Coherent Interface (formerly known as SuperBus) is based on experience gained during the development of Fastbus (IEEE 960), Futurebus (IEEE 896.1) and other modern 32-bit buses. SCI goals include a minimum bandwidth of 1 GByte/sec per processor; efficient support of a coherent distributed-cache image of shared memory; and support for segmentation, bus repeaters and general switched interconnections like Banyan, Omega, or full crossbar networks.

To achieve these ambitious goals, SCI must sacrifice the immediate handshake characteristic of the present generation of buses in favor of a packet-like split-cycle protocol. Wire-ORs, broadcasts, and even ordinary passive bus structures are to be avoided. However, a lower performance (1 GByte/sec per backplane instead of per processor) implementation using a register insertion ring architecture on a passive "backplane" appears to be possible using the same interface as for the more costly switch networks.

This paper presents a summary of current directions, and reports the status of the work in progress.

Introduction

SuperBus is the working name adopted by a Study Group under the auspices of the Microprocessor Standards Committee of the Technical Committee on Mini and Microcomputers in the IEEE Computer Society. The SuperBus Study Group considered (beginning in November 1987) the need for and feasibility of a very high performance "backplane bus," to be at least an order of magnitude more powerful than the existing standard buses.

Examination of the physical and logical constraints such a system must meet in order to be successful led to a new name, SCI (Scalable Coherent Interface), because it became clear that traditional bus structures would not be able to meet the demands of the next decade: the real goal is to interconnect many powerful processors productively, so that the total power of a system can be increased by merely adding more processors. Our examination of the needs for compute power to handle real engineering problems (e.g. aerodynamic simulation or simulation of large circuit designs) or physics problems (e.g. event reconstruction in the Superconducting SuperCollider) showed that a single bus, even at 1 GByte/sec, would be completely inadequate. Many buses (segmentation for parallelism) joined by selective repeaters would be necessary. Or, more likely, no buses at all, but rather some more general interconnection mechanism.

Many architectures which would be perfectly satisfactory for a single bus become ugly or impractical for assemblages of multiple buses; i.e., they do not scale well. Thus "Scalable" reflects our constraint that the system be smoothly extensible. "Coherent" refers to our requirement for a distributed cache-memory coherence mechanism, similar in concept to that developed for the Futurebus, which can greatly reduce the performance cost of interprocessor communication. "Interface" reflects the generality of our specification, which permits a given module to connect to an unspecified generalized interconnection mechanism, which might be a switching network of any of various kinds, a passive "backplane" forming a register insertion ring, or conceivably even an active bus (i.e. transceivers directly on the backplane).

The SCI standardization project was authorized by the IEEE Standards Board in October 1988, and was assigned the number P1596.

What are the Problems?

Present bus systems are running close to (and in some cases slightly beyond) physics limits; one cannot speed them up much by turning up the clock frequency or increasing transceiver speed or power, unless one shortens them correspondingly. For example, the Next machine uses NuBus (IEEE 1196) protocols at 25 MHz, 2.5 times the 1196 clock rate, but allows only four sockets instead of the 1196's sixteen. If a bus is short enough and is lightly loaded, transceiver and logic speeds do dominate among its various limits, and so its clock rate can be increased.

The fundamental physics limits are the speed of light, which limits the propagation velocity of signals and thus adds delay to handshakes; the capacitance of connectors and transceivers, which so disturbs a bussed signal transmission line that the "ideal transmission line model" is a very poor approximation indeed; and skew, differences in propagation time among a number of parallel signals which threatens to blur the boundary between successive data items.

Other physics problems, such as crosstalk between adjacent signals, are much easier to deal with and have become more economic than fundamental. Distribution of power and ground (nontrivial in the face of very rapid changes of current flow) is also in this category, and so is cooling.

There are also non-physical (logical?) problems. For example, when many processors operate in parallel to solve a given problem, they need to be able to communicate efficiently with one another in order to share resources or to divide the work. This intercommunication can be a significant bottleneck, perhaps using a large fraction of the system bandwidth just accessing one shared semaphore variable over and over. Furthermore, fast processors require fast local storage, so they need their own local copies of data, some of which needs to be shared. These local "cached" copies create logical problems if they are modified, because the various copies can become different or incoherent. Somehow, when one processor modifies data which other processors are using, the other processors have to be notified that their local copies are no longer valid so that they can get a fresh copy.

The cache coherency mechanism developed for Futurebus (and now being adapted to Fastbus) requires each cache controller to observe all other traffic in the system in order to determine whether some of its own data might affect or be affected by the current bus operation. Such a scheme cannot be generalized to highly parallel systems. (In fact it cannot generally work across Fastbus Segment Interconnects, so coherency domains are limited to single segments—though there may be multiple such domains, which may intercommunicate via more explicit mechanisms such as message protocols.)

* Work supported by the Department of Energy, contract DE-AC03-76SF00515.

Poster paper presented at the Nuclear Science Symposium, Orlando, FL, November 9–11, 1988

How will SCI solve these problems?

The shared transmission lines which form buses are extremely difficult to improve beyond what Futurebus and Fastbus have done. The most practical way to do better would be to use an active backplane, which has transceiver chips connected directly to the transmission lines with no connectors or stubs between. This would minimize the capacitance, and would result in uniform and constant loading which would make it possible to compensate for the loading and significantly improve the transmission line behavior. The connectors would be between the modules and the transceivers, so the presence or absence of a particular module would have no effect on the transmission line loading.

This active backplane scheme could also make live insertion and removal feasible, if module power is controlled by the backplane. However, most customers find the active backplane frightening because of the difficulty of replacing it if a failure should occur and thus it has received little support so far. Not all backplane physics problems are solved by this mechanism: the wire-OR glitch would still create delays whenever multiple drivers are permitted to be active on a single line, and bus turn-around (changing from one driver to another, as when changing from read to write or when changing mastership) would require delays for similar reasons.

Furthermore, a bus is inherently a bottleneck because it is shared by too many processors. Processor throughput is so high even today that a few processors can saturate any bus. Heavy loading subjects the users to long waits, slowing the whole system.

Therefore SCI has opted to drop the bus mechanism in favor of high speed unidirectional links. Two models are being supported for using these links. The high performance (and high cost) model uses the links to communicate between the module and a fast switch network, resulting in one GigaByte/sec per module. A lower performance model connects the input and output links of adjacent modules to form a register insertion ring, which can be implemented in printed wiring on a passive backplane structure at low cost but results in only one GigaByte/sec throughput per backplane.

Since even this low-cost version is still much faster than any existing backplane bus system, it seems attractive especially as a transition model over the short term as processors proliferate and costs decrease.

Our hope is to standardize on one module which can operate equally well in either environment, so that processors from many vendors can be developed and used effectively in small quantities initially, and then be moved into a switch environment unchanged when switches become available or necessary or economical for the given application. This provides a nearly unbounded upgrade path for system growth, and should create an attractive market for the manufacturer (high volume) and for the user (low cost due to high volume and competition among manufacturers).

Unidirectional links effectively remove the speed-of-light barrier to system growth: the system size and clock rate are decoupled, and there are no cycle-by-cycle handshakes. Physical signalling problems are greatly simplified because there is always one driver and one receiver, at opposite ends of the link. Signals operate steadily whether there is data flowing or not, which makes it easy to use phase locked loops for data extraction if desired (there is no start-up preamble required). That would make it possible to eliminate skew completely by encoding clock timing with each data bit transmitted, although we do not think this will be necessary yet at our initial 1 GigaByte/sec transmission rate.

We expect to use a central clock as a frequency reference so that only phase errors have to be compensated for during data extraction. Differential signalling, probably ECL but perhaps current steering drivers instead, results in constant current flow between connected modules, enormously simplifying the ground distribution problem compared to normal buses.

We plan to use a narrow 16-bit data path at 2 ns/word (250 MHz clock, both edges active), to control the interface IC pin-count problem and make switch elements more practical. Note that 'differential' implies 2 pins per signal, and 'unidirectional' implies 2 links, one for input and one for output, so we are talking about 64 pins minimum for each SCI interface circuit just on its fast end. A circuit for making switch networks must have at least twice that many, and preferably four or eight times, so the importance of a narrow data path becomes obvious. Actually, the 16 data bits will be accompanied by a clock, a flag bit, and probably a parity bit, so the numbers are somewhat larger than stated above. Modern ECL circuits appear to be able to handle point-to-point transfers at these data rates, but some care will be required with layout and connectors.

We are addressing the logical problems in several ways, trying to keep the system efficient by appropriate choice of protocols and trying to prevent starvation or deadlocks by providing forward-progress mechanisms and doubled queues (which prevent requests from blocking responses).

We are developing a cache coherence mechanism which maintains a distributed directory of users of each data item, so that only those who care have to be notified when shared data is modified. By storing this directory as linked-list pointers in each participating cache, the storage required does not have to be preallocated and there is no intrinsic limit to growth. The proposed mechanism seems simple enough that it should work, but it is not trivial. We must carefully check corner cases, such as what happens if one node decides to remove itself just as another is trying to add itself onto the list. Additional system traffic is required for maintaining coherence, but it is proportional to the information transfer traffic (about double for cached items). This seems a reasonable cost in exchange for the much larger factor of parallelism it makes possible, and for moving spin-waits into caches.

Conclusion

The SCI project is moving rapidly, and has attracted participants from many of the high-performance computer companies. The proposed signalling mechanisms appear to be technically feasible (though not entirely trivial), and there appear to exist logical protocols which are compatible with our goals.

The next phase will be a more careful study of the effects of various compromises and optimizations that could be applied to our logical protocols, and the selection of suitable connectors and packaging mechanisms. There is a lot of work to be done, but the enthusiasm level is high and progress has been rapid, so we are optimistic that we can achieve a workable specification in record time.

If you would like to participate in this work, or if you would like more detailed information, please contact the author:

David B. Gustavson, IEEE P1596 Chairman
Computation Research Group, bin 88
Stanford Linear Accelerator Center
Stanford, CA 94309
U. S. A.
tel: (415) 926-2863
fax: (415) 323-3626
bitnet: DBG@SLACVM