BUS AND RING SYSTEM INTERCONNECTIONS FOR DATA ACQUISITION AND CONTROL

V.I. Vinogradov, INR RAN, RF.

Abstract

The traditional bus systems based on 8/16 bit microcomputer- and register-oriented interfaces for data acquisition and control are analyzed. Bus- and nonbus-type system and network interconnections are compared. The possible transition from traditional bus to indirect point-to-point interconnections for control systems is discussed. The low interaction-rate nodes of message-passing systems can be exchanged for high-rate ones in operating control networks, with direct control access to distributed equipment based on shared distributed memory subsystems and a new generation microprocessors (µP).

Advanced distributed-memory data acquisition and control systems based on a new generation 16/32-bit general purpose microprocessors and Digital Signal Processors (DSP) with RAMLINK-type (SCI-like) distributed memory interconnections are proposed and discussed. Advanced distributed µP systems are proposed and analyzed on the basis of a new generation of DSP and distributed memory.

The new international standards for distributed data processing and memory interconnections are discussed. But for control systems an international standard for fieldbuses has not been selected after more than 20 years, though there are some national standards developments. In this case register-oriented bus interfaces on traditional message-passing networks can be changed to distributed shared-memory direct-access interconnections.

Some new decisions for fieldbus and instrumentation input-output subsystems can be based on bus and ring interconnections, which must be developed according to measurement, acquisition and control (MAC) requirements, using new 32-bit address fields with direct access to distributed memory, data block transfer, a simple data transfer protocol and connecting of different types of existing and new devices. The new generation of 16/32-bit µP- and DSP-based systems, with high speed links, can be a new basis for development of distributed memory systems with symmetrical or non-symmetrical structures. The ways of developing such a system are presented and discussed.

1. MICROCONTROLLERS, MICROPROCESSORS AND DSPs

Digital control and data acquisition and processing real-time systems address a need for monitoring and control of a large number of parameters. Any simple control system includes a controller which generates actuator commands, received from an operator, and readbacks, provided by sensors. The controller processes the signals to achieve a desired response and can thus modify the frequency of response of the system. The computation element can be realized with analog or digital components. Analog controllers process a signal, can be used for very high bandwidth systems and can be realized with inexpensive components. However they suffer from component aging and temperature drift and are limited to very simple algorithms.

DIGITAL CONTROLLERS approximate them, but are not affected by component aging and temperature drift and they provide stable performance with sophisticated techniques and algorithms. Programmable microcontrollers make it easy to upgrade, maintain and design systems for optimal and adaptive control. The choice of µP is critical in the performance and behavior of a digital controller. Available choices are microcontrollers, DSPs and general purpose µPs, including high cost RISC architectures. A control system requires real-time signal processing. The signal and gain coefficients must be represented accurately without loss of resolution for the smallest and largest magnitudes. Floating-point arithmetic may be necessary if gain coefficients and signals are time-varying or with large dynamic ranges [1-3].

High performance is required if the sampling of signals at discrete time intervals requires high speed processing. The µP should finish the processing as soon as possible. Thus very fast instruction cycle and multiplication times are needed. Peripheral integration is important from system cost, simplicity of design and board-size points of view. Typical peripherals on such chips are I/O pins, parallel and serial interfaces, DACs and ADCs. Digital microcontrollers are primarily designed to replace hardwired logic for data acquisition and decision making in control systems. The special architecture and high performance of Digital Signal Processors (DSP) make it possible to implement a wide variety of modern Digital Control algorithms and Data Acquisition systems. The most popular DSP's are from Texas Instruments, Motorola and Analog Devices.

MOTOROLA has created compatible families of µPs since 1974, beginning with the 6800 family of microcontrollers, then next the 68000 line and now PowerPC RISC µPs. Motorola's DSPs, based on the DSP56000 architecture, have compatibility across an entire product line. All family members benefit from the same extended triple-bus Harvard architecture. On-Chip Emulation (OnCE) was first proposed by Motorola and is useful for field analysis and repair. Wait and stop modes reduce power consumption. The Do loop instruction directs a µP to repeat a series of operations a specified number of times without any branching time.

Motorola offers 3 closely related families of DSP: DSP5610X, DSP5600X and DSP960X, along with development tools and peripherals. DSP56100 (16-bit) is optimized for the standards of digital cellular communications. The DSP56000 (24 bits) has become the standard for digital audio applications. The new member, DSP56004, is used in automobile audio systems. The DSP96000 (32 bit) family combines color graphics, sound and communications for medical, industrial and other applications. Previously limited to military, aerospace and scientific systems, DSP- technology is soon be used far more widely. The DSP56000 family is designed as a 16-bit programmable core µP for DSPs, supplemented with configurations of memory and serial communications interfaces, timers, a host interface and in-chip coder-decoders for analog signals [1].

TEXAS INSTRUMENTS produces several on-chip DSP families, including field-programmable microcontrollers and fixed- and floating-point µPs. Field Programmed Modules (FPM) are Multi-Time Programmable (MTP) µPs with EPROM and EEPROM all on a single chip. Ten year data retention, 10,000 write cycles and an internal charge pump that generate a 12V programming voltage means that the end product can be adapted. The TMS370 family has more than 10 FPM members. CMOS technology gives low current consumption. The 256 bytes of EEPROM on a TMS370C710 have the advantage of being modifiable in the field. Coming from a microcontroller family these devices are well suited to industrial applications. The sensor signal µP family TMS400 was developed to serve sensor signal conditioning applications with high accuracy and precision, where average system power must be in the microwatt range [2].

The FIXED POINT 16-BIT DSPs from TI give the lowest cost per MIPS in their class. The TMS320C5X generation of DSPs gives 50 MIPS of performance at 5V and 40 MIPS at 3V. A parallel logic unit provides high speed without modifying the ALU status or registers. Software wait states provide an external interface for use of slower off-chip memory and I/O.

The FLOATING POINT 32-bit DSP TMS320C30 generation integrates system control and intensive math functions for fast data movement and high-speed data processing. Powerful instructions, parallelism and buses provide flexibility in performance up to 50 MFLOPS. A high degree of parallelism allows one to perform up to 10 operations in a single instruction cycle (for example 2 data accesses, a multiply, an ALU operation and a DMA). The TMS320C32 is an enhanced low cost 32-bit DSP manufactured in 0.72 µm EPICS CMOS technology. The C32 includes all the features associated with a general purpose controller (similar to a RISC/CISC µP), but also provides additional ones. The TMS320C40 generation of DSP is effective for intensive parallel data processing in real time and has on-chip communication ports and support for shared global memory and its own pair of extended buses which give direct processor-to-processor accesses. The device provides scalability, fault tolerance and reduced bus saturation.

ANALOG DEVICES INCORPORATED (ADI) has digital signal processing experience, including Mixed-Signal Peripherals (MSPeripherals) and Mixed-Signal Processors (MSProcessors). In 1982 ADI introduced the first DSP in CMOS technology with fixed- and floating-point building blocks, program sequencers, data address generators and register files. The first single chip CMOS DSP ADSP-2100 was introduced in 1986. The company defines a digitally integrated approach to analog systems of signal processing as a first step in a natural evolution in the area of Microcomputers (µC). There are two compatible ADI signal processor families: the 16-bit fixed point ADSP-2100 and 32-bit floating point ADSP-21000s. High performance serial ports (SPORTs) are bi-directional, double buffered interfaces and feature user configurable clocking, framing and word length. A Parallel Host Interface Port serves as an asynchronous interface between the signal and host µP, acting as a block of 8 dual ported registers (8/16 bit interface). An ADSP-2181 can respond to 11 interrupts (up to 6 external). Two Serial ports, SPORT0 and SPORT1, allow multiprocessor communications with word length from 3 - 16 bits (bi-directional) [3].

THhe FLOATING POINT ADSP-21000 family has a complete set of arithmetic operations, including min/max, y/x etc. The parallel 9 port general purpose data register file transfers up to 9 operands to and from the computation units and memory. The program sequencer supports zero overhead looping, single cycle setup and exit for multiple program loops and delayed and nondelayed branching. The DSP-21020 handles 32-bit IEEE floating point, 40-bit IEEE floating point and 32-bit fixed point data formats. System tests and on-chip emulations are provided by an IEEE JTAG boundary scan serial port. The ADSP-21010 is a low cost 32-bit version of the ADSP-21020.

The SUPER HARVARD ARCHITECTURE COMPUTER ADSP-21060/62 family (SHARC) consists of 32-bit DSP microcomputers optimized for high per-formance applications, built on the same ADSP-21000 core (system on-chip), adding a dual ported on-chip SRAM and integrated I/O peripherals supported by a dedicated I/O bus. ADSP-210xx have a 25ns instruction cycle time operating at 40 MIPS. The SHARC core is compatible with the ADSP-21020.

2. BUS AND NON-BUS STRUCTURES

There are bus and nonbus approaches for data acquisition and control architectures. The bus one is based on parallel µP-system interconnections for SMP; the non-bus one can be based on switches and SCI standards for Massively Parallel Processing (MPP), Supercomputers and high-speed DACs.

CROSSBAR SWITCHES FOR DATA ACQUISITION can provide [1] very high-speed data processing on distributed systems and networks. There are packet switching, message passing and shared memory µP-systems for pipelined data processing and MPP. A crossbar interconnection is one of the ways for supporting multiple µPs sharing common memory having access time comparable with µP cycle times. As seen in the results of AT&T's Crossbars, these systems allow one to utilize maximum configurations and provide up to 30% improved performance over normal switching systems. The probability of successful connection of a µP to any memory unit is a function of the ratio of the number of µPs (N) and memory units (M), which function decreases to some asymptotic value as N/M. Crossbar-based systems provide, for instance, 15% better CPI value (the number of cycles per instruction) for four individual µPs having four memory units, as compared to similar bus-configured systems. Switch based systems can be constructed as either small or large networks with packet switching [4].

SMALL-SCALE NETWORKS are useful in a wider range of systems than MPP, particularly applications in which I/O is dominant. Fast buses are not very cost effective because every bus slot must have all the high-speed electronics, though they may use it only a fraction of the time. In small systems many parts can be used simultaneously by different processors. Each part can be made to a narrow width and the requirements for a low-cost high-speed interface can be met by full duplex serial connections. One such approach is DS-link (from SGS-Tompson Microelectronics) for Transputers, an asynchronous serial protocol. A serial bit-stream together with its clock is coded onto 2 wires. Each DS-link requires 2 signals in each direction.

A PACKET SWITCH based on STC104 is a 'wormhole' router with a 32-bit duplex DS-link interface, supporting locally adaptive routing. It consists of 32 link slices, each of which provides an input and output port. The IMS 9000 transputer integrates 4 serial link interfaces together with a superscalar µP, FPU, cache and memory interface [5]. The virtual channel processor accepts high-level communication commands from the µP and translates them into sequences of packets. A packet starts with a header and contains at most 32 bytes of data. The STS101 implements a DS-link and provides a parallel interface to other 16/32-bit µPs (synchronously or asynchronously).

A LARGE NETWORK with packet switching is under research and development at CERN for future generation experiments. The IEEE P13335 standard covers the connections, cables, and electrical and logical protocols for serial interconnections, operating at speeds of 100 Mbit/s and 1 Gbit/s over copper/optical fibers. This standard has been developed in the Open Microprocessor Systems Initiative/Heterogeneous Interconnect Project (OMI/HIC). It has a goal of constructing modular, scalable, parallel low-cost systems, which support high-performance communications and transparent implementations of high level protocols. The standard is based on the STC101 and STC104 packet switches, connected to 32 DC bi-directional link ports via a 32*32 non-blocking crossbar switch [6]. The ROUTER processes up to 200 Mpackets/s with a maximum bandwidth of 300 MB/s and latency of less than 1 µs. As result they construct very large networks with different topologies (mesh, grid, tree, Clos) and investigate traffic, deliverery of packets and latency. Packets are dispatched according to some predetermined schedule. A cost-effective solution for loading, starting and monitoring the links is the T225 transputer, handling up to 4 links at a time at full bandwidth. The OS-links from that controller are used to connect nodes to the host. One implementation is being made in a study for a second level trigger. The samples of data from a detector are buffered in memory units and passed to a µP farm. A local system uses a TMS320C40 DSP for combined buffer management and local processing. This system is built around DSP-links with switching by DS-links. The TMS320C44 DSP can receive data from up to 4 DSP-links. This special processing machine is under construction and has 3 stages of routing chips. The first and the third stages of switching have 112 DS-links connected to a T9000. The second stage supports 64 concurrent circuits.

SCI-based modular systems can join many PCs or workstations and clusters of µPs in Multi-Processor (SCI-LAMP) Architectures. Each node in this network has a local µP, cache, memory and I/O. The local memory is distributed into private and exported portions. Private memory is used only locally, and does not participate in the SCI protocol. Conversely exported memory can be seen by other nodes as part of a single global physical address space supported by SCI.

SCHEDULING of a distributed LAN is traditionally based on a centralized controller (such as a Condor OS) and message passing. Alternatively SCI can provide physically shared memory and cache-coherence among workstations across a network. Scheduling parallel jobs on a network of workstations introduces an additional level of complexity . The fast process migration and choosing of idle µPs are important to parallel performance. The number of µPs in multiprocessor systems assigned to an application might have to change when another application arrives or departs or when the degree of parallelism changes within an application. Better utilization than for fixed static partitioning is achieved with dynamic partitioning of µPs to parallel jobs. If the application cannot adjust its number of tasks during execution, several tasks from the same application may have to share a µP, introducing context switching and other related overhead. An approach to a distributed dynamic µP allocation scheme with even-loading considerations for such jobs in an SCI-LAMP system has been proposed by Li [8].

3. SCI-BASED DISTRIBUTED SYSTEMS

Development of the first modular systems, beginning from simple register-oriented interfaces, was based on micro-computers and input-output devices for measurement, acquisition and control tasks. Bus-oriented architecture was the rule for both instrumentation and data processing. A single controller system was the rule for parallel bus (CAMAC-DW/BHW, HP-IB) and serial loop (CAMAC-SHL, HP-IL) structures. Many existing register-oriented protocols were used for measurement (HP-IB, HP-IL), data acquisition (CAMAC), control (FIELDBUS), robotics (BITBUS) and Instrumentation. Input-output subsystems [9] must be developed according to measurement, acquisition and control requirements (short address fields, direct control, data block acquisition and a simple data transfer protocol for connecting different existing and new devices). Next generation systems need a new approach to system construction based on distributed memory and modern microprocessors [10,11].

The new generation distributed microprocessor systems can be divided as follows into main subsystems: a) distributed data-flow processing interconnects; b) distributed memory links; c) distributed input-output/instrumentation links. The traditional means of interconnecting is parallel bus or serial interface. But many new possibilities can be used once one has ring interconnections. The SCI-type systems are one of the best ways to construct a shared memory system. The distributed memory subsystem can be based on a RAMLINK, which is now under development. It is a single controller subsystem with a simpler protocol than SCI, which can be used as an interface for control of object-oriented single-controller systems (fieldbus/fieldring).

The resulting document on SCI was approved by the IEEE on March 18, 1992. The study group concluded that any practical solution would involve the use of packet-based signaling over many independent point-to-point links, would eliminate the bus bottleneck problem and would address the problem of how to maintain cache coherence. The old concept of a single- or few-processor supercomputer has become uneconomical. The cost per computing operation is less with µP technology, so the goal of system research for some years has been to divide problems and spread them across large numbers of inexpensive distributed µPs. The simplest way is loosely coupled µPs, using many small ones each with its own memory, that pass data via "message passing".

The general model will be as in a shared memory system, which may be divided into pieces and distributed as are the µPs. To reduce the effective access time to this memory cache memories keep copies of the most needed data near each µP (copies of parts of shared memory). Thus it is easy to pass data in shared memory. This model needs "cache coherent" protocols to support distributed µPs with shared memory. Among different ways of interconnecting are ringlets and switch connections. To support cost-effectively efficient coherence protocols, µPs are expected to manage cache fault handling; infrequent and complex updates would be done in software, but frequent cache updates would be done in hardware.

To reduce the cost of low-end systems, SCI's nodes support a ringlet connection. This requires some extra circuitry for address recognition and buffering in each node, but makes it possible to build inexpensive SCI systems. A wide range of applications can cover the whole range from high-end µPs and their I/O systems to workstations and LANs. Very efficient LANs have been built using this model, and the usual layers of network protocols can be added on top for compatibility with existing software. A distributed SCI system shares a 64-bit address space, where the high-order 16 bits are used to route packets to the appropriate nodes. The interconnections can be powerful switches or simple bridges that route packets between ringlets. The interconnects merely forward packets, and need know nothing about cache coherence. The initial links were 1 GB/s (16-bit 2 ns rate, differential ECL) and 1 Gbit/s (fiber optics up to a few km, coaxial cable up to about 20 m).

SCI packet protocols maintain cache coherence, providing for flow control, mutual exclusion, and forward-progress guarantees needed to support multiprocessor system software and applications. Support for message-passing is an important subset of these goals. The CSR work (IEEE Std 1212-1991) was started as part of SCI. A joint approach for Futurebus+ and SCI was developed, a modular metric mechanical packaging standard IEEE Std 1301-1991, which provides physical interchangeability for SCI modules from multiple vendors. A bridge between SCI and VME (IEEE1014) will be one of the resulting commercial products.

SCI is a replacement for computer buses, intended for the next generation of systems ranging from multiprocessors to workstations and PCs. But present data transfer devices are poorly matched to high performance memory systems. The low bandwidth at the I/O of the RAM packages forces one to use a wide array of RAM switch interleaving. The most effective possibility uses LVDS to connect RAMLINK chips in small rings with one controller per ring, which can interface a ring to SCI and schedules all ring activity. The controller uses a split-response packet protocol. The response time may be given either as a precise time or as an upper bound. The RAMLINK protocols simplified the SCI protocol, which can run at 500 MB/s.

Interfaces between the SCI-based tightly-coupled portion of a system and other modular systems (VME, FB+, PCI), serial buses or SCI-type I/O buses are based on bridges. Local I/O devices should share µP memory address space, but peripherals may be located on separate buses. Closely coupled bridges are intended to interconnect multiple workstations and servers within a LAN. A module can consist of several linc (interface chip) interfaces (one for each bus) and could be a block in a µP. One could also be a switching component for a wireless bus.

To develop a version of the RT-SCI standard optimized for real-time applications is one of the new problems of distributed systems. For RT systems latency is more important than throughput. The SCI mechanism can be replaced by a strict priority one which ensures meeting application deadlines under specified conditions. Some applications can rely on Rate Monotony Scheduling, which requires that latency be a function of priority. Thus the next levels of IEEE standards are being developed [12-14].

- **P1596.1**: The specification defines SCI/VME bridge architectture for interfacing VME buses to an SCI. This will provide I/O support for SCI systems via VME.
- **P1596.2**: Cache Optimizations for Large Numbers of Processors (kiloprocessors), using the SCI-developed tree-structured coherence directories and fast data distribution mechanisms for SCI systems.
- **P1596.3**: Low-Voltage Differential Interface for SCI, specifies low-voltage differential signals for high speed communication between CMOS, GaAs and BiCMOS logic arrays, used to implement SCI. The project defined new signals that are appropriate to CMOS and other technologies to be differential with 0.25 V swing, centered on +1 V, at 2 ns signaling rate. SCI adopted the serial encoding that Hewlett-Packard devised originally for Serial HIPPI links.
- **P1596.4**: RAMLINK (High-Bandwidth Memory Interface), based on SCI Signaling Technology, permits access to the large memory chips. This work is converging toward point-to-point links at 500 MB/s using a simplified packet protocol.
- **P1596.5**: Data Transfer Formats Optimized for SCI, defined a set of data types and formats that will work efficiently on SCI for transferring data among heterogeneous µPs in a multiprocessor SCI system.
- **P1596.6**: SCI/RT Study Group, considers ways of getting deterministic behavior in SCI systems for real-time applications.

By this time SCI has been adopted by many computer companies for internal use. Dolphin Server Technology has been formed to market SCI chips, development tools and high level models. Dolphin Interconnect Solutions' SCI-SBus adapter uses programmed I/O for messages smaller than 64 bytes. For larger messages a DMA engine conducts the data transfer . As the message size increases, a peak is reached at 85 MB/s for 64 kB messages. In its clustering of QUAD-P6 nodes, Siemens Nixdorf has announced its use of SCI as well. Among other companies which are ready to use SCI-type systems are Cray, Intel, AT&T and GIS. Cray Research's I/O channel and SCX system interconnect are moving to SCI. A new supercomputer-class system interconnect provides high performance, scalable, reliable inter-system and system to peripheral communication (as an open standard). The SCX is ring-based and has been enhanced to address the reliability, flexibility and performance needs of distributed supercomputer environments.

Much modern computing consists of a number of modular systems (PC/workstation, clusters of workstations), connected by networks. One of the best possible architectures for the nodes in SCI is that of SMP-based workstations.

SUMMARY

Using the SCI standard for control systems is still expensive, and the RAMLINK protocol can be used for subsystems with distributed memory on the first level of hierarchical µP/DSP-based control systems. The complex integration of data acquisition and control can be based on the same architecture of distributed memory and direct interconnections between µP- and DSP-based nodes on a network, simplifying programming. But SCI is possible within a collaboration, using for the R&D program experimental objects of a small accelerator. The different networks for experiment data acquisition and control can be based on SCI-networks, wich connect ring- and bus-based µP subsystems.

REFERENCES

1. Motorola M56000, M96000 Digital Signal Processing, Motorola.

2.- TMS320C50, TMS320C40 User's Guide, Texas Instruments;
- MVP: The dawn of a new-era in DSP - introducing TMC320C8X,Texas Instruments.
- Digital Control Applications with the TMS320; Application Notes. 1991; Texas Instruments.

3. ADSP2100, ADSP21000 notes, Analog Devices Inc.

4. O. Panfilov. Comparative performance analysis of crossbar switches and multiple buses. AT&T/GIS. ICSNET95 - the 11-th International Symposium. St. Petersburg. 1995.

5. Peter Tompson. Small-scale Networks for General Systems Interconnect. SGS-Tompson Electronics Limited.,UK., The 11-th International Symposium ICSNET'95. St. Petersburg. 1995.

6. INMOS Limited (1993). The 9000 transputer Hardware Reference Manual. SGS-Tompson Microelectronics Limited . Bristol. UK.

7. R.W. Dobson, A. Martin, S. Haab, R. Heeley, M. Zhu, J. Renner Hansen. Realization of a 1000 Node High-speed Packet Switching Network. The 11-th International Symposium ICSNET'95. St. Petersburg. 1995.

8. Saravanan Agasaveeran, Qiang Li. Distributed Job for Scheduling in SCI Local-Area Multiprocessor. The Fourth International Workshop on SCI-based High-performance Low-cost Computing. 1995.

9. V.I. Vinogradov. Parallel-pipeline Architecture for Real-time Data-flow Acquisition in Modular Systems. IEEE Nuclear and Plasma Sciences Society. 7-th conference on computer applications in Nuclear, Particle and Plasma Physics RT'91. KFA, Germany, 1991.

10.D.B. Gustavson, V.I. Vinogradov. Advanced Modular System Architecture with Distributed Memory Based on SCI. The 10-th International Symposium on problems of modular systems and networks (jubilee). ICSNET'93 Proceedings of Plenary reports, Puskin, Russia, 1993.

11.D.B. Gustavson, V.I. Vinogradov. Status and Development Problems of Distributed Modular Systems Based on SCI. The 11-th International Symposium. ICSNET'95 Abstracts. St. Petersburg, 1995.

12. IEEE 1596 - SCI Standard.

13. IEEE P1596.4 - RAMLINK Proposal.

14. The Fourth International Workshop on SCI-based High- Performance Low-cost Computing. Proceedings. 1995.

Table 1. Comparison of fixed point DSP basic parameters

Parameters	ADSP2100	TMS320C50	DSP56000
-Instruction execution	50ns	35/50ns	50ns(40MHz)
- MIPS	20	30	30
- RAM, kB	2/P,1/D	10	50
- Com. Port	2	1	1

Table 2. Comparision of floating point DSPs.

Parameters	ADSP-21060	TMS32C40/C32	C80	DSP96002	i860
-Number of µP	1	1	4DSP+ +1RISC	1	1
-Instr.exec. time,ns	20	60/50	-	50	20
- MIPS	40	25/20	2BOPS	20	30
- MFLOPS	120	50/60	-	60	100
- RAM (32 bit), K	128	2	4*2K-DSP 4K-RISC	2	8
- Number of reg.	32(16)	8	-	8	30
- Com. Port	2	4/2	-	0	0
- DMA bandwidth MB/s	240	100	400	80	0
- FFT complex,ns	0.46	1.54	-	1.05	0.75
- Divide 32 bit,ns	150	360	-	350	440
float.point,cycles	6	9	-	7	22