Image Cover Sheet

CLASSIFICATION	SYSTEM NUMBER 513419
UNCLASSIFIED	

TITLE

NGSP-Experimental Development Model \(XDM\)

System Number:

Patron Number:

Requester:

Notes:

DSIS Use only:

Deliver to:

This page is left blank

This page is left blank

NGSP – Experimental Development Model (XDM)

Richard E. Vallee[1], D. Bruce Oakley[1], Gavin Hemphill[2], Robert Trider[3], F. Andrew Reid[1]

1. Applied Microelectronics, Inc., 1046 Barrington St., Halifax, Nova Scotia, B3H 2R1 Canada
2. Defence Research Establishment Atlantic, 9 Grove St., Dartmouth, Nova Scotia, B2Y 3Z7 Canada
3. Trider Scientific, 629 Windmill Rd., Unit 9, Dartmouth, Nova Scotia, B3B 1B7 Canada

Abstract

This is the second in a series of three papers which present the Next Generation Signal Processor (NGSP). The first article entitled, NGSP - A Novel Parallel Signal Processing Architecture presents the overall system architecture of the NGSP. This, the second article, presents the experimental development models (XDMs) produced using the NGSP architecture. The third article entitled, NGSP - Integrated Programming Environment (IPE) discusses the software development efforts undertaken to allow execution of signal processing applications using the NGSP XDM systems. The XDM systems were developed in order to demonstrate that the NGSP architecture's performance could scale linearly when used in conjunction with real time signal processing applications. The XDMs developed demonstrated linear scaling to within 3% for sustained performance levels between 400 MFLOPs to 3.2 GFLOPs.

Background

The NGSP is a high performance signal processor which consists of multiple Arithmetic Units (AUs) connected together via a cache-coherent control network, all accessing a common shared memory called the Working Memory[1]. The NGSP architecture was developed for the Defence Research Establishment Atlantic (DREA), for use as a scaleable multi-application signal processor to meet the increasing demands of the Canadian Department of National Defence. The essential attribute of this architecture is its ability to defeat the memory "bottle neck" problem present in most of today's high-performance signal processors.

The NGSP XDMs developed consist of systems with four AUs and eight AUs. The key system components for each system are:

- a cache coherent control network implemented using the Scaleable Coherent Interface (SCI, IEEE Std. 1596-1992) technology,
- a proprietary multi-port collision free Working Memory (WM) system for each XDM unit, and
- multiprocessor AUs to provide the processing elements, using a custom vector engine.

The NGSP consists of the following main elements:

- a Host computer,
- AUs which are multiprocessor signal processing nodes,
- a WM to store the compute data used by the AUs.
- the Integrated Programming Environment (IPE) to provide the software environment[2].

NGSP Host Computer

Each NGSP requires a central supervisory computer, called a "Host Computer", which controls:

- the initialization of each AU,
- the transfer of programs to each AU, and
- the initialization and monitoring of the WM via the VME bus interface.

The NGSP XDM host computer consists of a SPARC CPU-2CE Single Board Computer. The interface between the host and the SCI Control Network is provided by an SBus-to-SCI Adapter Card that provides a non-coherent SCI interface.

NGSP SCI Memory Board (MB)

The SCI MB is a custom general-purpose single board computer equipped with an interface to the SCI ring that was designed to provide a means of supplying memory services to other nodes on an SCI ring. Interpretation of incoming SCI packets, and generation of outgoing SCI packets, is performed by software executing on a high performance RISC processor (MC88110). The software was developed to allow the board to execute the SCI cache coherency protocol, and provide test and debug computing resources for the SCI ring.

NGSP Arithmetic Unit (AU)

Each AU is a multiprocessor signal processing node which is logically partitioned into three sections; the Control Network Interface (CNI), the Vector Engine (VE), and the Data Transfer Unit (DTU). The CNI is responsible for receiving control data from the host computer via the SCI control network. The VE is responsible for performing the vector math processing. The VE program instructions are delivered via the CNI, and the data via the DTU. The DTU is responsible for compute data transfers, as well as providing I/O. The DTU provides a high-bandwidth data link between the VE and WM. An architectural overview of the AU is provided in Figure 1.

a) AU - Control Network Interface (CNI)

The key CNI components are:

- the Control Processor (CP) based on the MC88110,
- the SCI Node Chip (NC),
- the Coherent Memory Controller (CMC),
- the Coherent Memory (CM).

The CP is responsible for executing the NGSP operating environment and preparing tasks for execution by the applications processor (AP) which resides within the VE. The NC supports the SCI standards and protocols that provide the AU access to the control network. The CMC provides CP memory coherence on the AU with all the AUs and the SCI MB on the SCI control network. The CP also acts as the transfer processor (TP) which is responsible for coordinating the DTU data transfers to and from WM. As shown in Figure 1, bus isolators and transceivers on the AU allow the CNI, DTU and VE to operate and perform local data transfers without interfering with each other.

b) AU - Vector Engine (VE)

The key VE components are:

- the Application Processor (AP) based on the MC88110,
- a custom DSP chip set consisting of:
  - the Sequencer (SEQ) ASIC,
  - the Switch (SWX) ASICs,
  - the Vector Math Processor (VMP) ASICs, and
  - the Vector Cache Memory (VCM),
- the DTU data port.

The VE incorporates an MC88110 Motorola RISC microprocessor to act as the AU AP. The computational aspects of the VE are dealt with using the set of three custom ASICs (SEQ, SWX and VMP). The VCM contained within the VE is a high bandwidth data cache. The AP executes algorithms by issuing instructions to the SEQ which is responsible for orchestrating the vector processing.

The SEQ controls the vector pipeline by translating the high level vector instructions issued by the AP into the address and control sequences required by the pipeline at a vector element level. The SEQ has special vector address generating hardware, which generates addresses for three read streams, and two write streams, each with arbitrary strides if desired, and can also produce the striding patterns needed for FFTs and Bitonic sorts.

The VMP performs arithmetic operations on data in the vector pipeline. There are two VMP ASICs, one to handle the real part of the data, and the other the imaginary part. Based on the vector instruction, the SEQ determines the mapping between the read and write operands streams and the VMPs.

c) AU - Data Transfer Unit (DTU)
The key DTU components are:

- the Data Transfer ASIC (DTA),
- HP G-Link Rx/Tx pairs to WM,
- optional AU I/O mezzanine card.

The DTU provides the AU with access to WM and I/O. The DTU consists of a DTA and an optional I/O interface mezzanine board. The DTA acts as the main data transfer controller within the AU and provides direct memory access (DMA) to and from WM and VCM. The DTA provides a list processor interface that allows DMAs to be specified as array memory data transfers with different row and column strides. In addition, the DTA contains a proprietary addressing shuffling scheme known as Perfect Addressing (PA) which allows WM packets to flow collision-free at a sustained data rate of 160 MBytes/sec. WM packet construction and parsing are performed by the DTA. Communication with the WM is via G-Link high speed serial data links. A custom I/O mezzanine provides a buffered, variable length, serial input and output data stream, at rates of up to 32 MBytes/sec using Cypress Hot Link devices. Each AU within the system may support a mezzanine board, thus allowing I/O to be scaled to the needs of almost any application.

NGSP Working Memory (WM)

The WM allows high bandwidth (160 MBytes/sec) collision free shared access to data memory from all AUs simultaneously. The WM consists of two main components, the Switch and Memory Modules (MMs). The interface between the AUs, Switch and MMs, are implemented with high speed serial (1.5 Gbps) interconnects using HP G-Link devices. The WM is a packet based distributed memory system that routes incoming request packets from the AUs to the MMs, and outgoing response packets from the MMs to the outgoing ports connected to AUs. In order to achieve 100% bandwidth, the incoming packets must request the correct MM port in order to be routed through the switch matrix without delay. This is achieved using a proprietary memory addressing scheme known as perfect addressing (PA). The entire WM subsystem is configured and monitored through control and status registers accessed via a VME bus interface. An architectural overview of the NGSP WM is provided in Figure 2.

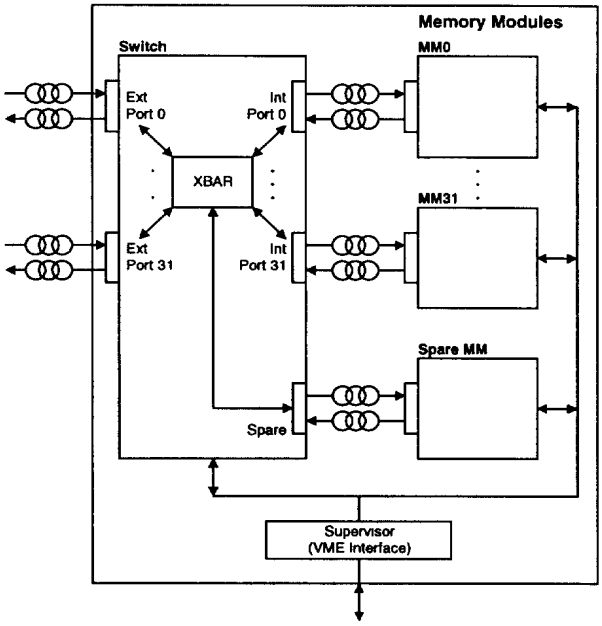

Figure 2: NGSP Working Memory

a) Switch Overview
The Switch is a synchronous 32 channel data switching system. One side of the Switch serves as the 32 port interface to the AUs and the other side serves as the 32 internal port interface to the MMs. The synchronous behaviour of the Switch means that all channels are repositioned in the switch network simultaneously and on fixed time intervals.

The functionality of the Switch is broken up into four modules. Each module contains one 32 by 32 crossbar network, 32 input serial ports, and 32 serial output ports.

The Switch rotates at a rate of one port every 200 nsec, also known as one packet time. The worst case synchronization latency is 4 packet times in the 4AU system (800 nsec), and 8 packet times in the 8AU system (1.6 usec). The WM latency is 9 packets, or 1.8 usec. Two ASICs were developed in support of the WM implementation, a 1 by 33 XBAR switch GaAs ASIC and a 32 deep by 16 bit FIFO GaAs ASIC. The FIFO ASICs are used in conjunction with the G-Link receivers to perform LRC checks and provide sufficient buffering during flow control.

b) Memory Module (MM) Board
The MM provides the memory storage component of the WM system. The memory capacity of a WM system is determined by the number of MMs in the system as well as the memory configuration of each MM. The MM unpacks incoming packets into address (32 bits) and data (64bit words) on write accesses, and packs data (64bit words) into outgoing packets for read accesses.

To achieve both high capacity and high access speed, each MM contains four sets of DRAM arrays that are accessed in parallel. The arrays are configurable for a minimum capacity of 8 MBytes each and a maximum capacity of 64 MBytes. Since all DRAM arrays are accessed in parallel they must be configured identically. Therefore, the MM storage capacity is four times an individual DRAM array capacity (32 MBytes to 256 MBytes). The packet latency through the MM board is 4 packet times or 800 nsecs.

NGSP Chassis

The enclosure for an NGSP 4AU system is a custom 19 inch rack mountable chassis. It has 4 standard VME 6U slots and 16 standard VME slots for 9U cards. The 4AU chassis requires the following backplanes:

- the J1 backplane is a standard 21 slot VME backplane,
- the J2 backplane is a standard 4 slot VME backplane,
- the P2 is a custom backplane for the WM Switch boards to provide inter-Switch communication.
- the P3 is a custom power backplane used to provide power to the NGSP 9U boards.

The chassis provides a total of 2 KW of power. The voltages supplied are ±12, ±5, +3, -2. The chassis also supports fans which provide 400 LFM cooling across the boards housed in the system.

The enclosure for the NGSP 8AU systems consists of two custom 19-inch rack-mountable chassis similar to those described for the NGSP 4AU systems. One chassis is used to house the NGSP 8 port WM, and the second is used to house the 8 AU boards. The AUs and WM are interconnected via coax cables to support the 1.5 Gbps serial data links between AUs and WM.

Discussion

Three NGSP 4AU and two NGSP 8AU systems have been built and integrated by Applied Microelectronics for DREA. The performance figures achieved using the proof of concept systems are shown in Table 1.

References

[1] NGSP - A Novel Parallel Signal Processing Architecture, Bruce Oakley, Richard Vallee, Andrew Reid, Gavin Hemphill, Bob Trider, ICSPAT, September 1998.

[2] NGSP - Integrated Programming Environment (IPE), Chuck Pilkington, Gavin Hemphill, Bob Trider, ICSPAT, September 1998.

Measure	NGSP 4AU XDMs	NGSP 8AU XDMs
Performance Sustained	1.6 GFLOPS	3.2 GFLOPS
Performance Peak	3.2 GFLOPS	6.4 GFLOPS
Word length	64 bits	64 bits
WM Bandwidth (sustained, collision free)	160 MBytes/sec/AU	160 MBytes/sec/AU
WM Size	256 MBytes	1 Gbyte
I/O Bandwidth (raw)	32MBytes/sec/AU	32MBytes/sec/AU
Performance Scaling Linearity	within 3%	within 3%

Table 1: Proof of Concept System Performance Figures

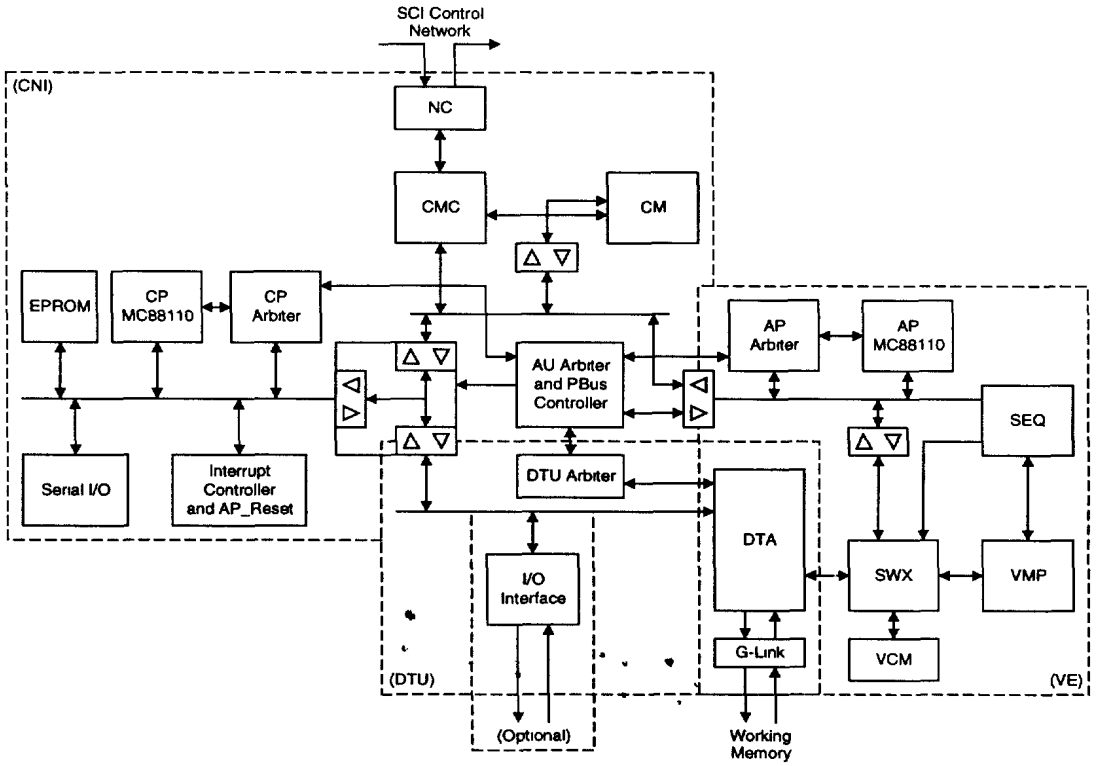

Figure 1: NGSP AU System Architecture

The SEQ issues a pipeline control word for each address set that instructs the SWX data cache interface as to the mapping for a particular group of addresses. The VMP uses a 64 bit data word which can be interpreted as a 64 bit double precision floating point number, or as two 32 bit single precision floating point numbers.

The different number formats supported by the VMP are; IEEE 754 Single Precision Floating Point, 2's Complement, and Signed Magnitude Integer. The VMP resources available for data pipeline processing on a cycle by cycle basis are a Coefficient Generation System (CGS), a Function Evaluation Unit (FEU), and a Multiplier-Accumulator (MA).

In addition, as data passes through the pipeline, the VMP can be made to perform one of several operations such as rounding, sorting, thresholding, conversions, and bitwise logical operations.

The AU SWX is a multi port, data path switch. It is responsible for performing all accesses to the VCM. It allows the DTU the highest priority access, the AP the second highest priority access, and the vector pipeline is given the lowest priority. Since multiple data access paths are concentrated through the SWX, it uses arbitration and pipeline buffering schemes. The SWX is capable of performing single cycle read or write accesses at a rate of 40 MHz. The total SWX bandwidth to the VCM is 1.6 Gbytes/sec.

The VCM is used to store the input and output data processed by the vector pipeline. The VCM is segmented into five cache banks using a hash mod 5 segmentation algorithm which helps reduce cache memory contention from the multiple access ports available through the SWX. The VCM is implemented as 5 banks of 64 bits by 64K, for a total of 2MBytes.

This page is left blank

#513417

This page is left blank