EI ECTRONIC BRAIN DESIGN FOR ADVANCED MICROSPACECRAFT '

James Wall and Wai - Chi Fang

A dvanced Computing Resources Group and Center for Space Microelectronics Technology
Jet Propulsion I aboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-809, USA

ABSTRACT

A brief history of NASA robotic spacecraft computing architectures is given. Changing needs driving changes in features (10x reduction in cost, various increases in capability, 1(),x to 100x reduction in power, mass and volume) for spacecraft computing are discussed. A scaleable, parallel/neural computing architecture which addresses these changes is proposed I valuation of some of this architecture's components is detailed. I inally, future directions of exploration for this spacecraft computing architecture arc considered,

1. BACKGROUND

1.1 introduction

The current vision for NASA includes a bold increase in mankind's ability to directly gather information about the sun, planets, asteroids, comets and other structures in I arth's solar neighborhood (<60 AU). Also included are deployment of spacecraft for remote sensing of electromagnetic phenomena in and out of our solar system and observation of earth-based phenomena. Most of the NASA community now believes that this vision will he accomplished by a change from the infrequent (1 per S years) rni ssi on based around a large spacecraft (3000 kg) to missions of much greater frequency (3 per year) based around multiple small spacecraft (10 to 600 kg).[2]

The expectation for shrinking NASA budgets imply that this era of higher frequency space missions will be accomplished under a much smaller total budget. Therefore, spacecraft missions must be planned with radically lower design, development, launch, operation and infrastructure costs.[2]

One solution to these problems are inexpensive computing elements which support rapid development of flight hardware and software and allow a shift in operation functionality from the ground to the spacecraft. These elements must also provide more science delivery than earlier systems while commanding much smaller mass, power and volume. resources.

1.2 I listory of NASA Computing Architectures

NASA robotic spacecraft such as Voyager (launch 1977), Galileo (launch 1987) and Cassini (launch 1997) have had linear evolution in their on-board computing capabilities (.S0 KIPS to 0.2 MIPS to 2.() MI 1'S). The computing architectures and mass/power/volume resources have stayed roughly constant, Data bandwidth has been the area of most growth. I light code development costs and operations costs have also largely remained constant. Physically, the computers were built on 15cm x 30cm printed wiring boards placed in a custom backplane.-1css chassis in a radial arrangement which occupied the central core of the spacecraft and allowed for flexible thermal management,

The flight computers for these. three spacecraft perform roughly the. same. common set of tasks[1]: process uplinked data and commands, gather, process and store sci en ce data for do wnl i nk, control tel cinch-y s ystems for downlink sessions, receive engineering and relevant telemetry sensor information for packaging with downlinked science data, receive attitude control sensor information, control attitude, articulation and position control actuators and, in some eases, process algorithms for closed-loop control of these actuator/sensor systems.

The evolution in computing architectures during the twenty-year period from 1972 - 1992 was driven by increased capabili t y and density of VLSI devices, radiation tolerance enhancements and packaging improvements. Other factors were improved knowledge of the space radiation environment, the desire to move from mechanical data

recorders to solid-state data recorders for reliability masons and a desire to provide increased on-board fault management.

1.3 Changes in NASA Spacecraft Computing Needs

The change in the NASA missions from those based on large-class fly-by and orbiter spacecraft to those based on much smaller fly-by, orbiter, lander, rover and acrobot vehicles reflects an renewed interest in in-situ science and sample return efforts[2]. Moving to more, smaller spacecraft allows a mission to be accomplished by physically differentiated spacecraft, 1 for example, a planetary mission may include a lander/rover and multiple acrobots supported by an orbiter.

Reduced operations cost and reduction in overall space/ground telemetry bandwidth also drive a change in spacecraft computing away from the traditional. Solutions include increased on-board autonomy and use of data compression. An interesting challenge is how to use these same techniques to simultaneously increase the amount of gathered science information.

Another area driving the nature of on-board computing is total spacecraft mass and required launch vehicle. Spacecraft computing and telecom have traditionally been the primary users of spacecraft power and spacecraft mass. Therefore, shrinking the power needs of the computing system has a leveraging effect on the power system mass and size[1]. In addition, reducing the mass of the spacecraft computing system has a leveraging effect on the attitude control system power, size and mass. An ideal situation would be for the computing and associated elements to no longer be one of the top two or three most significant contributors to final spacecraft size. Ultimately this would allow these. new spacecraft to be sent on an equivalent orbit by a much smaller (and cheaper) launch vehicle than their heavier predecessors.

2. A PROTOTYPE MICROSPACECRAFT COMPUTING ARCHITECTURE

2.1 Current Work

Several computing systems have been developed for miniature spacecraft[1, 3]. Most use a large amount of lllhc.riled or rXrJnlJmrcird-off- fllc-shelf (COTS) components to achieve low-cost. The primary direction these computing systems have taken is to reduce the packaging mass of traditional spacecraft computing hardware [3]. However, if new capabilities such as synthetic aperture radar (SAR) image processing, on-board autonomy and change-detection-based science data gathering [4, 5] are to be accommodated, space computing engine architectures must at so change. Here we present a neural -processor-bwed computing engine called the. Teraflop Electronics Brain (TEB) which can support the foreseeable processing needs of these new spacecraft applications.

Figure 1. TEB General Purpose Processor MCM

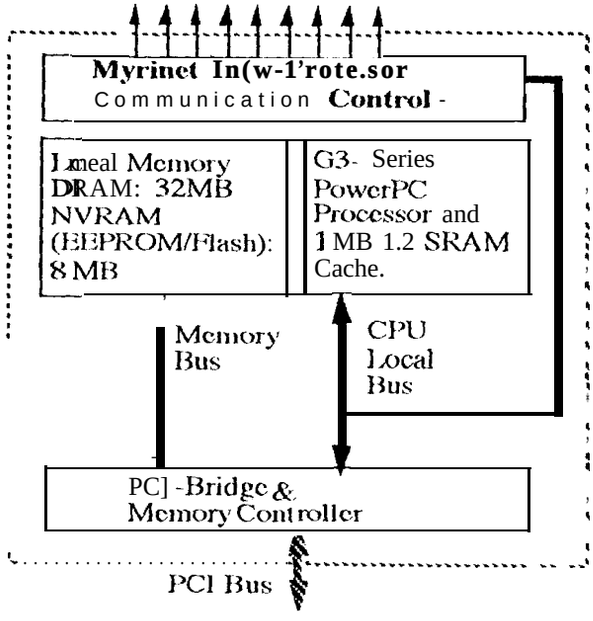

The TEB is based on a scaleable multi-processor architecture, built into a series of multi-chip module (MCM) stacks. Typical stacks have 11 MCMs and provides a peak performance of about 10 GFLOPS with overall dimensions of 10cm x 10cm x 6cm. Four types of MCMs are required: a neural processor, a general purpose. processor, a stack memory and a communication controller. Some intra-stack communication occurs over a PCI bus and the rest of intra- and inter-stack communication occurs over a Myrinet-based communication mesh.

Figure 1 shows the general purpose processor MCM. One or more of these processor MCMs can be included in each stack. This MCM is based around the Motorola G3 series of PowerPC processors (availability 1997, 200 MFLOPS, peak) [6]. Since memory access over the PCI bus requires paged memory or other inefficient access scheme, enough L2 SRAM cache and directly-addressable DRAM has been included on the MCM to allow for sustained local processing of Single Instruction, Single Data (SISD) problems. The powerful

Myracom Corporation's Myrinet-based inter-processor controller allows processor MCMs, both internal and external to each stack, 10 participate in solving complex Multiple-Instruction, Multiple 1 Data (MIMD) problems.

Figure 2 TEB Neural Processor MCM

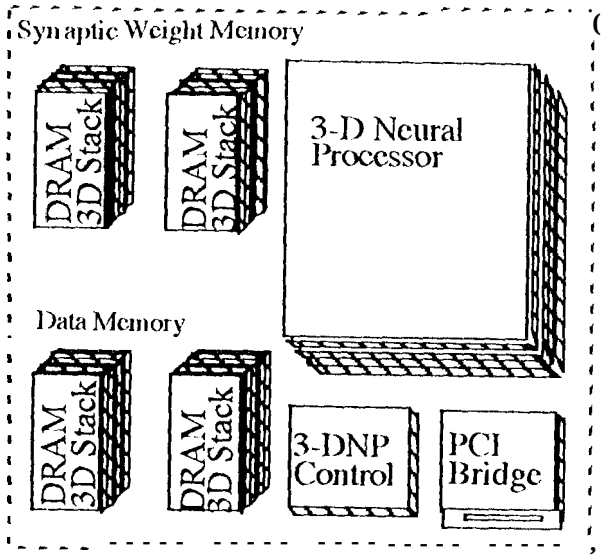

Figure 3 Single MCM Stack of the TEB

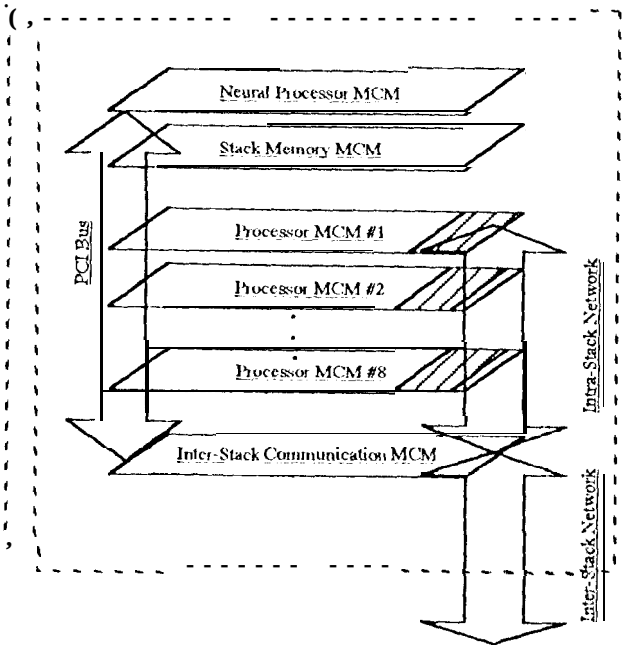

Figure 2 shows the neural processor MCM. It is based on a three-dimensional neural array developed by JPL (1 teraops, peak). It is intended to rapidly process Single Instruction, Multiple Data (SIMD) problems, including those often assigned to a digital signal processor (DSP), although the inclusion of a DSP in the TEB has not been completely ruled out.

The particular neural processor design for the TEB system uses randomly addressable digital synapse weighting. This improves the performance of the neural array when rapid context switching is required. It also supports rapid manufacturing and in-line testing.

Figure 3 shows a 11 MCM 'J' 3D stack. It includes two additional types of MCMs, a mass memory MCM and a communication controller MCM. The memory MCM makes its DRAM available over the PCI bus. The communication controller incorporates a significant amount of cache memory (4MB) and is locally controlled over the PCI bus.

Stack performance (which can be modified bawd on mission needs) is provided by eight general purpose processor MCMs, and one each of the remaining MCM types.

2.2 Evaluation of Components

The neural processor for TEB, the three-dimensional analog neural network (3DANN) has been built under contract to the Ballistic Missile Defense Office by the Neural Processing Group at JPL. It has shown excellent performance in highly paralleled tasks such as image-based pattern recognition. It is currently being used to drive a ground-based template matching system.

The various chip stacks used in TEB have been manufactured by Irvine Sensors Corporation, among others, They have been used in various flight programs, including the Small Satellite Technology Insertion program, the New Millennium program and various commercial ground systems.

The PowerPC family of microprocessors from Motorola and IBM have shown great robustness in responding to challenges in the industry. This validates the TEB present commitment to that processor family. In addition, the PowerPC processors show excellent high-spud arithmetic processing capabilities, initially allowing the TEB architecture to avoid the use of dedicated DSP engines, traditionally required in high-performance guidance, navigation and control systems.

The Myrinet switch is rapidly becoming the de-facto parallel computing inter-procmsor communication controller standard. Its adoption by TEB allows efficient integration with many industry tools and techniques.

3. FUTURE DIRECTIONS

3.1 Architecture Development

The primary new areas of architecture development are three-dimensional topologies, power management, software and simulation/emulation. One intriguing area of research is in the assignment of three-dimensional stack faces to standardized bus functions such as that proposed for Bolotin's SpaceCube architecture[7]. This allows the.rmal and electrical needs to be pre-assigned in the third dimension and supports multiple. buses for fault-tolerance and complex networking.

Most advanced microspacecraft have modest power budgets. Even the classic large spacecraft are constantly struggling with available power versus mission goals. The TEB architecture must be developed with a combination of low-power IC technologies and fine-grained active power management to minimize the overall energy needs of any actual flight computing system, The ultimate goal is to have minimum power primarily be delivered autonomously by the operating system, without application intervention or concern.

Software work includes operating system kernel development to take advantage of the heterogeneous processing environment and the parallel computing engines as well as fault-tolerance capabilities in the real-time, scaleable, heterogeneous parallel processing environment.

A final area of importance, but one needing early emphasis, is the simulation and emulation of a TEB stack rural of a series of stacks. Initial work in parallel processing systems indicates the need for the current typical TEB stack ratio of MCM types, but simulation of this structure against the various applications and rni ssion needs will be needed. 1 'ortunately, the partitioning of this architecture readily supports this type of experimentation.

3,2 Application Development

Science users of these new spacecraft computers must become much more closely coupled to the development environment than they have in the past. As the. science functions migrate on-board[4], they must port their algorithms to the spacecraft and develop new ones to take advantage Of new opportunities. Spacecraft bus engineers for telecom, attitude and command functions must also rework their algorithms to take advantage of the new computing capabilities for such functions as optical navigation, science data edit, telecom lossy and lossless data compression, etc.

3.3 Evaluation Work

Some of the building blocks of this architecture have been developed under the sponsorship of NASA's Advanced 1 light Computing and Remote 1 Exploration and Experimentation programs and the Department of J Defense's Ballistic Missile J Defense Office. The preliminary results show promise in their independent functionality. The primary work ahead will be to combine these elements into the. overall proposed architecture, create. simulations, run benchmarks and optimize for various spacecraft missions and applications,

4. BIBLIOGRAPHY

[1] J arson, W. J. & Wertz, J. R. (eds.) Space Mission Analysis & Design, Klewer Academic Publishers, 1991.

[2] Technology for Small Spacecraft, Aeronautics & Space Engineering Board, National Research Council, published by National Academy Press, 1994.

[3] Alkalaj, L., Jarvis, B. "The Design and Implementation of NASA's Advanced Flight Computing Module," *Proceedings 1995 IEEE Multi-Chip Module Conference MCMC '95,* Santa Cruz, California, pp. 40-44, January 1995.

[4] Stolorz, P. "Remote Exploration and Experimentation: Science-Directed Spaceborne Autonomy Applications", Remote Exploration and Experimentation Study Phase Kickoff Report, Jet Propulsion Laboratory, May 29, 1996.

[5] Elachi, C. Spaceborne Radar Remote Sensing: Applications and Techniques, IEEE Press, 1988.

[6] "Motorola's PowerPC Roadmap", url "http: //www mot .com/SPS/PowerPC/overview/roadmap.html".

[7] Bolotin, G. "Space-Grbc: A Flexible Computer Architecture Based on Stacked Modules," *Proceedings 1995 Multi-Chip Module Conference MCMC '96,* pp. 20-25, February 1996.