

RECENT PROGRESS IN SPACE-BASED MICRO-CONTROLLER DESIGN

Keith Avery, Nichols Research Corporation
David Alexander, Mission Research Corporation

ABSTRACT

This paper examines several recent space micro-controller designs employing printed wiring boards (PWBs) and multi-chip modules (MCMs) undertaken at Phillips Laboratory. In each case, a micro-controller implementation resulted in a design to minimize the size, weight, and power of mission-essential processing functions. MCMs were indicated in some cases to afford the most efficient and compact design.

INTRODUCTION

Though it is obvious that space systems benefit from low-power electronics for processing, integrated "off-the-shelf" designs are difficult to find that can meet aggressive volume, mass, and power goals, even for modest processing requirements. This observation, combined with a seemingly ubiquitous need for simple, efficient processors, led the Phillips Laboratory (PL) to develop a family of low-power micro-controller solutions for space missions. This paper reviews the requirements that led to these approaches, the designs that have resulted, and the future directions of this research.

Basic System Representation

In most cases of interest, a self-contained micro-controller-based data command and acquisition system has been a recurring theme. A canonical representation of a typical micro-controller-based design is shown in Figure 1. The most generic controller, in addition to analog, digital, and serial communication interface (SCI) input/output (I/O), has power input and a control interface. As a "black box" module, the controller does not provide direct access

to memory, address, and data busses. Instead, the basic controller provides a rich set of interfaces to accommodate the most commonly needed functions, such as data acquisition, analog voltage generation, binary level setting and control, and telemetry. The control interface permits special features inherent to a particular controller design, such as power savings, diagnostics, and code download / verification.




Figure 1. Typical micro-controller system.

One debatable aspect of the Figure 1 scheme is that it implies a closed system form, as no provisions for memory expansion and additional peripheral interconnections are permitted. For a large class of applications, such as simple instrument control, motor control, and distributed health and status monitoring, a closed architecture offers a simpler interface, leading to a potentially simpler integration, particularly when common protocols are available. Applications requiring more intimate access to internal data busses or more expansive memory are better served by a general purpose processing system with a bus-oriented interface, such as VME and Futurebus. The advantage of the simple, closed architecture shown in Figure 1 is that it is conducive to miniaturization and power optimization.

Performance Levels. The highest utility instances of the Figure 1 architecture are not always the highest performance cases. The center of focus in much of PL's current space microelectronics research is geared towards the development of high performance processing solutions. The current state-of-the-art in radiation hardened processing is represented by the Lockheed Martin RAD6000 and the Honeywell RH32 microprocessors, both 32 bit architectures with 20-30 million instruction per second (MIPS) performance range. These processors are powerful candidates for space systems, but do not address simpler requirements. In fact, simple, low-power architectures of modest performance and memory capacity have wide ranging appeal for low-end data acquisition, especially in a distributed processing paradigm. It does not make sense in modern system design, for example, to directly interface a thermometer to a RAD-6000 (32-bit processor), nor is it attractive to create an elaborate data acquisition nexus containing hundreds of monitor and control points, multiplexed to a single processor. On the other hand, a hierarchical processing architecture, shown in Figure 2, permits improved system performance, in a way analogous to the "span of control" in a personnel organization. In the hierarchical processing model, a powerful central system processor administers complex data acquisition and control through a network of low-end micro-controller nodes, similar to that shown in Figure 1. Each controller in this network performs autonomous processing and /or local data acquisition / control, periodically coordinating with processors at levels immediately above and below its position in the hierarchy. Such a network permits improved system fault tolerance and real time performance through single-point failure reduction, lower process and interface contention, and reduced cabling.

The notion of a processing hierarchy suggests an accelerated benefit in the

reduction of size, weight, and power of controllers represented in Figure 1. Even in single node applications, reduction of size, weight, and power is paramount in many space applications. Our first experience with the need for a simple controller required a multi-experiment payload that could operate within a total power ceiling of five watts orbit average, and nearly every experiment required analog sampling. When it was found that no known "off the shelf" solution could meet this requirement in an acceptable size, weight, power, and cost envelope, it became increasingly clear that low-end controllers were needed, and these controllers were not available to the best of our knowledge. Thematically, the basic need for such low-end controllers recurred in several other applications.

Case for distributed processing. Multiple processors exist in all but the most trivial spacecraft, and often even in a single subsystem or payload, several processors are often required. The analogy of the personal computer is not altogether inappropriate, where besides a single, high-performance




Figure 2. Hierarchical processing model.

central processor, a number of processors are present in the disk control, video display, mouse, and keyboard. This obvious example is an analogy of modern systems, and space is no exception, yet a disproportional amount of small investment occurs in the low end

processors, even though they outnumber the high performance processors in a system application. If efficient, low end processors could be developed and applied within a system, the benefit of size, weight, and power reduction is multiplied by the number of instances in that system.

The COTS Dilemma with Simple Control Processors. In a recent space experiment development program, local PL researchers attempted to secure a compact controller which required radiation hardness in a compact (6 x 6 inch), low power (< 1W) form. The closest available solution found in a market survey at that time in ready-to-use form was a large (SEM-E or VME), power hungry (> 3 W) and expensive (> \$200K). While a commercial designer in the same solution would face a "decision paralysis" level of options, it appeared that, in fact, nothing really existed for space designers in terms of compact, low end processors. So we started designing our own.

The remainder of this paper is dedicated to discussing these developments and possible future directions in micro-controller research.

RECENT DESIGN ACTIVITIES

MAPLE 1

A concept for a series of space experiments involving the demonstration of new technologies in microelectronics, advanced packaging, and micro-electromechanical systems (MEMS) technologies was established under the name "Microsystems And Packaging of Low-power Electronics" (MAPLE). The concept for MAPLE involved creating a rapid prototyping carrier system into which emerging technologies could be incorporated as sub-experiments. MAPLE-1 was the first installment in the MAPLE series, which became a payload in Phillips Laboratories MightySat I. The concept for MAPLE-1 included five sub-experiments: (1) packaging reliability, (2) field programmable gate arrays, (3) MEMS accelerometers; (4) dosimetry/temperature,

and (5) solid state recorder. Each sub-experiment was designed as a stack-able segment; the collection of MAPLE-1 experiments are shown in Figure 3, and all segments were stacked together with a controller as a sixth segment, as shown in Figure 4.

Main controller. The low-power micro-controller architecture for MAPLE 1 was relatively simple, based on a radiation-hardened version of the Intel 8051 manufactured by Sandia National Laboratory (SNL), designated SA3865, and several Texas Instruments SIMOX SRAM's (two 8Kx8 components). All sub-experiment to controller communications were handled with simple 2400 baud RS-232 serial communications interfaces. The micro-controller, which operated from a 5.0688 MHz crystal, gathered data from each sub-experiment for subsequent transmission to the spacecraft. A cumbersome "legacy" real-time clock (Holt Industries HI-8000) became




Figure 3. Set of MAPLE-1 experiments on stack-able frames.

a requirement to facilitate time-stamping.

Sub-experiment controllers. As an ambitious approach to build extremely low-power intelligence into MAPLE-1 sub-experiments, a decision was made to use a PIC 16C74 (Microchip Technologies, Chandler AZ) micro-controller, which embodies a complete system of the form of Figure 1, on a single monolithic IC. The rationale for the use of the PIC in the five experiments was simple: (1) reduced parts count; (2) low power, (3) high level of functionality (built in digital I/O, analog-to-digital conversion, and SCI), (4) built-in

watch dog timer (WDT), and (5) a suggestion that the same part had been investigated for space use by another organization. Since the latter information was only substantiated by telecon with little




Figure 4. MAPLE-1 micro-controller, 8051-based, in 6x6 in printed wiring board. specific data, PL conducted an independent low dose rate radiation test on a PIC16C74 circuit in a Cesium source, duplicating the total ionizing radiation dose expected to be inflicted upon MAPLE during its orbit. The PIC performed flawlessly in these tests. Though not a validation of the PIC's radiation-hardness (the tests did not exceed 20 kilorads total ionizing dose), the test indicated that for simple, non-critical experiments in shuttle orbits, a PIC might be applied with care. Another often controversial aspect of PIC controllers is that they employ plastic encapsulation, which is often summarily ruled out for space applications. Further investigations revealed that the Sumitomo molding material used in over 80% of all plastic parts was found acceptable in at least limited application in shuttle missions (the parent satellite is designed to be deployed from the shuttle bay).

MAPLE 2

MAPLE-2 was designed to be a single-experiment "MAPLE", in this case an

experiment to examine radiation degradation of two different types of MEMS accelerometers, including a 3-axis configuration of Analog Device ADXL-50 devices, along with a special single-axis tunneling transducer from JPL. The MAPLE-2 parent spacecraft was the Space Test Research Vehicle (STRV) 2 satellite, which is designed for launch in a highly elliptic orbit. The harsher radiation environment eliminated the PIC micro-controllers from contention. Though the MAPLE-1 central controller was adequate in terms of radiation-hardness, it also could not be used, since: (1) it did not possess analog conversion facilities; (2) it did not have the correct SCI; and (3) the mass and volume constraints were too severe. As such, a new micro-controller design was required, which is shown in simplified form in Figure 5. This design employed the same basic components of the MAPLE-1 master controller, but also included: (1) a hardened analog-to-digital converter (Harris HS-9008); (2) a "legacy" serial communications interface; and (3) a switching dc-dc regulator to convert from isolated spacecraft +28 VDC. Due to the




Figure 5. MAPLE-2 architecture.

arcane "legacy" SCI, which required an inordinately high data transfer rate, the 8051 was operated at 16 MHz. The analog-to-digital converter (ADC) featured a flash architecture at eight bits resolution. Though rated at 400 mW, the MAPLE-2 design

operated at a much lower frequency than normally used for such a high performance converter, resulting in a much lower average power dissipation.

DEEP SPACE II

Extraordinary design demands for a low-power micro-miniature data handling system to be used in the NASA New Millennium Program's Deep Space 2 micro-probe resulted in the design of a new breed of highly integrated stand-alone controller module, which like the PIC, embodied a Figure 1 architecture within a single package. As such, the design embodied many of the features of the MAPLE-1 and MAPLE-2 designs, which provided inspiration for a new compact design. Referred to as the Advanced Instrument Controller (AIC), the design is based around a very small (25mm x 30mm) patterned

overlay multi-chip module (MCM), which contains a customized 8051 8-bit micro-controller, 128 kilobytes of static RAM, 128 kilobytes of non-volatile storage, on-board startup and clock circuitry, and a compact arrangement of large amounts of built in peripheral support: 32 digital discrete input/output (I/O), 32 analog inputs (aggregate 25 kilo-samples at 12-bit resolution), eight analog outputs, and six serial ports. Total module power consumption is designed to be less than 50mW at five megahertz. The module itself weighs less than five grams and is designed to function after 30,000-G shock loads and will operate at temperatures as low as -120 °C. Part of the power reduction strategy involves optimizing the design of two integrated circuits for the reduced capacitance represented by millimeter scale

interconnections within the interior of the MCM. The approach chosen will result in a package-less, few-chip MCM with relatively low cost potential as compared to the more cumbersome conventional and MCM configurations typically found in contemporary space-grade computers.

AIC Design. Many obvious constraints are imposed in designing a system controller for the DS-2 interplanetary microprobe. Despite the severe mechanical requirements, energy availability for the DS2 mission is such a difficult constraint that power management became a singularly dominant philosophy in the design. The overall design chosen for the AIC, shown in Figure 6, is based on the desire to provide high utility at modest performance levels with minimum size, weight, and power. By achieving this in a small profile, it is possible to use many AICs in larger systems to dramatically simplify complex control, particularly related to health and status monitoring. Following is a discussion




Figure 6. Block diagram of the Advanced Instrument Controller (AIC).

of some key features of this novel space processor architecture.

CPU- The central processing unit, designed by Boeing (Kent, WA), is binary compatible with the Intel 8031 / 8051 8-bit micro-controller, but with enhanced features and peripheral capability. A special provision exists to power-down the EEPROM after boot-up to conserve power, which requires establishing a second isolated address and data bus. The additional analog and digital functionality is implemented in special function registers and memory mapping. To further enhance the performance of the 8051, portions of the AIC8051 were expanded to 12-bit data widths, improving throughput for higher data sampling rates. The need to support dual voltage levels —5V and 3.3V— placed further complexity requirements on the design of the AIC8051. Another important feature of the CPU is the provision for non-volatile data write-back to the EEPROM, implemented as a normally unused 8051 op code.

Analog - An analog support integrated circuit (design by Technology Associates Group, CA) operating on a 5V supply provides a customized low-power extension for analog input and output from the AIC. A total of up to 32 analog input voltages with 0-4.096V span can be directly interfaced to the AIC at 12-bit resolution. Eight analog outputs are available at 10-bit resolution, and each output can also be sampled through the internal ADC under CPU control.

Other components and features- Existing memory components were used in the AIC, based on the ability to function at -130°C. The AIC contains high frequency (10 MHz) and low frequency (200 Hz) internal oscillators, and has provisions for external crystal drive. The AIC is programmable *in situ* through external control, by asserting a PROG signal upon power-up.




Figure 7. Advanced Instrument Controller (AIC) (a) floor plan before and after 1st metal layer. (b) Mock-up used in high-g cannon tests.

Packaging Approach. The AIC construction approach employs a tightly-coupled MCM design based on a version of the high density interconnect (HDI) process¹ that employs plastic substrates. The floor plan of the AIC is shown in Figure 7.

ELECTRONICS TESTBED (ETB)

The previous MAPLE experiments operated under the concept of a master controller, arbitrating and gathering data from a group of sub-experiments. Another series of such controllers, referred to Electronic Test Beds (ETBs), similarly served the function of a sort of satellite “service provider”, providing power, communications, scheduling, and data gathering support. PL had been commissioned to establish an ETB for STRV1-d to support a total of 10 sub-experiments. Like MAPLE-1, the STRV1-d ETB had to communicate with multiple experiments. Like MAPLE-2, the ETB also had to provide analog conversion support and direct power conversion support between spacecraft voltage level (+28V) to sub-experiment required voltage levels (+5V, +12V, -12V). Unique in the case of this ETB, a larger amount of data storage would be required, as much as one megabyte, which violated a previous tenet regarding memory expansion in simple controller architectures. The constraint arose in this case due to the very real limits of STRV1-d spacecraft data storage. Furthermore, the radiation


environment of STRV1-d was specified as a harsh (high radiation), geo-transfer orbit.

The ETB architecture for a data handling system (DHS) could not employ the AIC directly, since the AIC was not hardened for this environment (Mars transit was projected to be relatively radiation-benign). Instead, PL leveraged previous work in MAPLE-1 and MAPLE-2 for building blocks, from Deep Space II for gate array concepts, and from previous work in "hardening by design" (HBD), a technique which permitted the use of certain commercial semiconductor fabrication lines to achieve 100 kilo-rad performance. The resulting ETB micro-controller architecture, shown in Figure 8, extends the use of a SNL 8051 through a gate array designed with HBD methods to expand I/O, particularly for providing an SCI to each of ten sub-experiments and a spacecraft. RS-422 protocol is used throughout at 9600 baud for sub-experiments and at 19,200 baud for the spacecraft. A converter-per-sub-experiment approach is used, with an isolated converter for the controller itself. In this manner, failed sub-

experiments can be shut down without affecting other experiments. Current monitoring for the ETB and accessory analog I/O is provided with dual HS-9008 converters. The memory store is expanded to 1Mx8 (eight Lockheed Martin 128Kx8 bulk-hardened SRAMs) through bank-switching in the gate array. As an enhanced data reliability feature, a shadow memory bank is provided on-board using commercial EEPROM. Due to possible degradation of the EEPROM, it is interfaced using a sub-experiment interface to provide isolation in case of a radiation-induced failure.

A COMPARISON AND FUTURE DIRECTIONS

A simple comparison of the controllers discussed is presented in Table 1. The MAPLE-1, MAPLE-2, and ETB designs represent a trend towards increased functionality in hardened designs, while the Deep Space II represents a trend in ultimate miniaturization and power reduction. A highly desirable feature of these systems is robust analog functionality and a lot of SCI



2.1-13

Table 1. Comparison of recent PL micro-controllers

Feature	MAPLE-1 Master	MAPLE-1 Sub-experiment	MAPLE-2	ETB	Deep Space II
CPU	8051 rad-hard	PIC 16C74	8051 rad-hard	8051 rad-hard	8051 enhanced
Power	0.3 W	0.2W	0.6W	1.5-3W	0.05W
SCI ports	five	one	one	eleven	six
RAM	16K	192 bytes	8K	1024K	128K
EEPROM	32K PROM	4K PROM	8K PROM	32K PROM	128K EEPROM
Weight	300 grams	200 grams	300 grams	1200 grams	3 grams
Analog channels	none	eight	four	sixteen	32
Radiation	~100Krads	<23Krads	~100Krads	~100Krads	<20Krads

ports. Only MAPLE-2 and the MAPLE-1 sub-experiments had fewer than five SCI ports, primarily because those designs were lowest tier sub-ordinates in a three-level hierarchy of spacecraft processors.

Future directions for micro-controller work at PL are expected to bring down the large gap in mass and power consumption between rad-hard and non-rad-hard controllers. One fundamental weakness has been the lack of robust analog input and output circuitry. Better low-power rad-hard converters, and converters with programmable front end functionality (gain, offset, filtering), is desirable if it can be achieved efficiently in a controller architecture. An increased number of SCIs has been clearly beneficial and will likely continue. Improved provisions for networking multiple controllers is clearly an important area for future research, since lower power controllers will find their way into more second and third tier subsystems. As such, improved network support will promote more effective overall autonomy and control approaches. If power consumption can be made dramatically low, i.e., less than a few milliwatts, then it is possible to consider large networks of controllers in satellites for fine grain health and status monitoring. It is possible to even consider wireless controllers in single and multiple configurations, where they may receive communications, even power, through RF links. Smart power is yet another consideration for robust controllers. Other intriguing possibilities for future

controllers includes improved support for "plug and play", embedded programmable logic arrays (i.e., field programmable gate arrays), and even a programmable internal switch matrix for signal re-routing.

The most important observation in our research is that controllers have been neglected in research programs, in part possibly due to their more modest processing requirements, which are not stressing cases for high-performance benchmarks. The oversight seems unwarranted, as an active and diverse application base for modest controllers exist. In an attempt to address this gap, at least as far as PL programs are concerned, we have designed and constructed a number of efficient space-based controllers. Further research will be devoted to further improvements in size, weight, and power of low-end controllers to complement the work in high-end processors, as both are required in complex systems.

ACKNOWLEDGEMENTS

Authors would like to thank Todd Goforth, Maxwell Technologies, Inc., for his contributions with MAPLE-1.

¹ James C. Lyke, "Two- and Three-dimensional High Performance, Patterned Overlay Multi-chip Module Technology", published in *Proceeding of NASA Technology 2002 Conference*, presented at Technology 2002, Dec 1992.