X2000: Avionics for A Multi-Mission Spacecraft

Savio Chau Carl Steiner
Jet Propulsion Laboratory
California Institute of Technology

Abstract

The X2000 is a NASA program to develop technologies to enable faster, better and cheaper space missions in the future. The program has planned to have three deliveries. Each delivery is expected to develop many technologies that are an order of magnitude better than the previous delivery. The first delivery of this program will provide an engineering model to multiple missions. Currently, five missions including Pluto/Kuiper Express, Europa Orbiter, Mars Sample Return, Champollion/DS4, and Solar Probe have plans to adopt the X2000. It is expected that more missions including those Earth orbiting missions will adopt the X2000 in the future. The multi-mission requirement has driven the X2000 toward a symmetric, scalable and distributed architecture. In addition, some of the more traditional requirements for spacecraft design such as cost, performance, power, mass, volume, and reliability are exacerbated by the multi-mission requirement. To meet these challenges, the X2000 has chosen a loosely coupled multiprocessor architecture that is based on the IEEE 1394, I2C, and PCI buses. With these standard bus interfaces and advanced packaging concepts, the architecture of the X2000 first delivery will be able to meet all the challenging requirements.

1.0 Introduction

The X2000 is a NASA program to develop technologies to enable faster, better and cheaper space missions in the future. The program has planned to have three deliveries. Each delivery is expected to develop many technologies that are an order of magnitude better than the previous delivery. The technology goals of the first delivery are set relative to the Mars Pathfinder as follows:

Table I: Technology Goals of X2000 1st Delivery

Parameters	Improvements Relative to Mars Pathfinder
Computer Performance	10 to 20 times increase
Avionics Mass	2 times decrease
Avionics Volume	4 times decrease
Digital Electronics Power	10 times decrease
Analog Electronics Power	2 times decrease

Although the X2000 is a technology program the Jet Propulsion Laboratory (JPL) applies techniques of flight project management to develop, integrate, and test these technologies. It is believed that through the rigorous discipline of flight project management, the successful transfer of new technologies to space application can be ensured. Therefore, a set of missions has been chosen as the targets for the first delivery of the X2000. The first delivery will provide an engineering model to all and flight hardware to some of the missions. Currently, five missions

including Pluto/Kuiper Express, Europa Orbiter, Mars Sample Return, DS4 (a comet landing/sample return mission), and Solar Probe have plans to adopt the X2000 engineering model, while flight hardware will be delivered to the DS4 and Europa Orbiter. If the program is successful, it is expected that more missions including many Earth orbiting missions will adopt the X2000 design in the future.

It can easily be seen that the target missions for the X2000 First Delivery have very diverse requirements. For example, the Europa Orbiter has very high radiation environment (~25 Grad), the DS4 lander has to be very low energy/power consumption, the Pluto Express has very long mission life (12 to 15 years), and the Solar Probe has very high thermal requirement. Therefore, the traditional requirement-driven approach for spacecraft design will not be applicable because it is very difficult to establish a set of well-defined requirements. This multi-mission nature has driven the X2000 avionics system toward a symmetric, scalable and distributed architecture. Such architecture will allow addition or removal of modules to/from the system to meet specific mission requirements.

2.0 The X2000 Challenges

In addition to the multimission requirements, many traditional requirements for spacecraft design such as low-cost, high-performance, low-power, low-mass, low-volume, and high-reliability are also exacerbated by ambitious technical goals of the X2000. The challenges that the X2000 faces are explained in details as follows.

Scalability: Since the computation requirements of the missions vary over a wide range, from over 100 MIPS to under 20 MIPS of throughput and over 3 Gbytes to 128 Mbytes of memory, the architecture of the X2000 must be scaleable. In order to achieve this goal, the system must be modularized and the interface of module must be standardized. In addition, for spacecraft design, serial interfaces are preferred over parallel interface since the mass of cable harness can be significantly reduced. This implies that loosely coupled architecture is preferred over the tightly coupled architecture. However, the serial interface must have very high bandwidth in order to accommodate high data rate subsystems and science instruments such as the imagers (e.g., for Europa Orbiter, it is expected to take four frames of image every second, and each frame is about 10 Mbps). Also, the high bandwidth will allow more processors to be added to the system in the future. Finally, in order take full advantage of the loosely coupled architecture, a multi-master bus is preferred over a master-slave or command-response (e.g. the 1553 or 1773 bus) for the serial interface. This is because master-slave architecture is not as scaleable due to the difficulty to balance the workload among the master and slave processors. The multimaster capability also allows some subsystems such as the telecommunication and imagers to initiate data transfer before data such as uplink commands and images are lost.

Low Power and High Performance: Power has always been a premium resource for any deep space missions. This is particularly a critical issue for X2000 since some of the missions such as the Pluto/Kuiper Express travels so far away from the sun, radioactive-thermal power is the only choice of power source. However, due to public concern of safety and mass consideration, it is

desirable to reduce the amount of radioactive material on-board (i.e., the power source of the X2000 is limited to 150 Watts). The dilemma is, while it is desirable to keep the power consumption of each node (i.e., processor or microcontroller) to the minimum, there are subsystems on X2000 need to maintain very high data rates. Thus, if a low power serial interface is chosen, then it will not meet the performance requirement for some of the subsystems. On the contrary, if a high bandwidth serial interface is chosen, then it will not meet the low power requirement. To resolve this dilemma, the X2000 has decided to use a dual bus architecture: a high bandwidth but moderate power bus for high data rate nodes and a low power but relatively low performance bus for engineering and low data rate functions. In addition, it is also important that the high bandwidth bus includes power management features such as sleep mode to conserve power.

Radiation Tolerant: There are drawbacks in applying commercial standards to space missions. In particular, radiation-hardened and flight qualified parts for these standards are not readily available. To solve the first problem, X2000 has decided to use the "system-on-a-chip" approach, in which synthesizable ASIC cores (called intellectual properties or IPs) will be procured and integrated. The integrated design will be fabricated through rad-hard fabrication lines. Hence, it is a critical requirement in the selection of the serial interfaces that synthesizable ASIC cores must be available for the standards.

Low Cost: Low cost perhaps is the foremost important design objective for future NASA missions. Commercial standards and off-the-shelf (COTS) become very attractive due to the availability of their specification, design, software, and test equipment. Widely accepted commercial standards can also reduce mission operation costs since trained personnel are much easier to find. Therefore, the search of the serial interface for X2000 focused mainly on commercial standards. Since a significant portion of the development costs for spacecraft is dedicated to the integration and test. Therefore, the selection of the serial interfaces must also take the cost of integration and test into consideration. Another technique to reduce cost is to minimise the types of components need to be designed. Hence, it is a general guideline of the X2000 that each building block of the system should be as generic as possible.

Low Mass and Volume: Low mass and volume is critical for deep space missions since they can affect the choice of launch vehicles and the flight time of the mission. These will in turn affect the costs of launch and mission operations. The reduction of mass and volume has always been an on-going struggle for spacecraft design and an incentive for technological innovations. The X2000 will further advance the packaging technology of the NMP DSP project at JPL. All the subsystems will be put on one or two stacks of multi-chip-module (MCM) slices. This packaging constraint will also affect the choice of the serial interfaces.

Fault Tolerance: Last but not the least, fault tolerance is extremely important for deep space missions due to the extensive flight time. The problem is that fault tolerance design usually requires additional hardware and software components for error detection and recovery on one hand, and yet it is constrained by the power, mass, volume, and cost on the other. Hence, how to design an architecture that can tolerate as many faults as possible with the minimum redundancy is a challenge. This is particularly true for the serial interface design.

3.0 The X2000 Solution

It is crucial for the X2000 to select bus standards that can meet the multimission requirement. An extensive survey and an industrial workshop had been held at JPL for the bus selections. The high-speed buses that had been surveyed include IEEE 1394 bus, Fiber Channel, Serial Fiber Optic Data Bus (SFODB), Universal Serial Bus (USB), 1773 bus, Fast Ethernet, ATM and other buses. A similar survey also was done for the low power buses. The survey include I2C bus, CAN bus, SPI bus, MicroLan, and the Low Power Serial Bus (a JPL modified low-power 1553 Bus). The IEEE 1394 Bus and the I2C Bus were chosen as the high speed and lower power buses, respectively, and the rationales are given in Tables II and III.

Table II: Selection of High Speed Data Bus for X2000

Bus	Rationale
IEEE 1394✓	Very good data rate, moderate power, strong commercial support, relatively deterministic latency, no rad-hard parts but ASIC core available, although weak in fault tolerance. However, techniques have been developed to compensate for the weakness in fault tolerance
Fiber Channel	Excellent data rate, strong commercial support, flexible protocol & topology, excellent isolation, no rad-hard parts but ASIC core available. However, power is too high and weak in fault tolerance. Optical connector is not compatible with MCM stack configuration, although the flexible protocol and topology are attractive. May be applicable for future deliveries
USB	Very similar to 1394 but much lower data rate (12 Mbps) and yet power is the same as 1394.
Fast Ethernet	Not suitable for real time application due to indeterministic bus latency, also no rad-hard parts or ASIC core, although commercial support is excellent
ATM	Too complicate and high power, no rad-hard parts or ASIC core, although commercial support is excellent
SFODB (IEEE P1393)	Power too high and is not a real standard, although data rate (1 Gbps) and isolation are excellent
Myrinet	Not a widely accepted standard. Power too high, not rad-hard parts or ASIC core, not a real standard, although data rate is extremely high.
SCI	Not sufficient support from the industry, not implemented by any vendor at the time of survey
FDDI	Too complicate and high power, no rad-hard parts or ASIC core, commercial support is giving way to ATM
AS 1773	Protocol well understood by the aerospace industry, data rate is marginally adequate (20 Mbps), not multi-master, limited commercial support, built-in redundancy and excellent isolation, built-in redundancy, relatively high power.
SPI	Relatively low speed (5 Mbps), not enough protocol support (need design in software) for use as the main system bus.

Table III: Selection of Low Power Bus for X2000

Bus	Rationale
I2C	Very low power, multi-master, ASIC core available, adequate data rate (100 to 400 kbps) for low speed data, simple protocol, strong commercial support, rad-hard sensor interface chip has been developing by APL.
CAN	Adequate data rate (up to 1 Mbps). Power is too high for low power bus. ASIC core is available.
J1850	Moderate power, multi-master, data rate may be too low (10 kbps), protocol similar to CAN, fairly strong commercial support, ASIC core or rad-hard parts not found
SPI	Not easily to implement a multimaster bus, no rad-hard parts or ASIC core found.
MicroLan	Rad-hard part and ASIC core are not available, design not transferable to rad-hard foundry

LPSB	Protocol well understood by aerospace industry. Data rate is adequate (1 Mbps), much lower power than standard 1553, not multimaster, ASIC core is difficult to acquire (only 1 vendor found and restricted to specific foundry), built-in redundancy, not compatible with standard 1553 equipment.
Access Bus	Spin off of I2C bus. Power tends to be higher in most implementation.

The resulting architecture of the X2000 is shown in Figure 1 (note: the tree topology of the 1394 Bus is not shown.) A drawback of these commercial buses for space application is that they are designed for high volume commercial market rather than highly reliable applications. For example, the IEEE 1394 Bus employs a tree topology and has no built-in redundancy. The I2C has no built-in redundancy and even less fault detection capability in the protocol. Therefore, the X2000 architecture has to employ redundant buses. Additional layers of protocols are added to the hardware and software to provide the necessary fault protection.

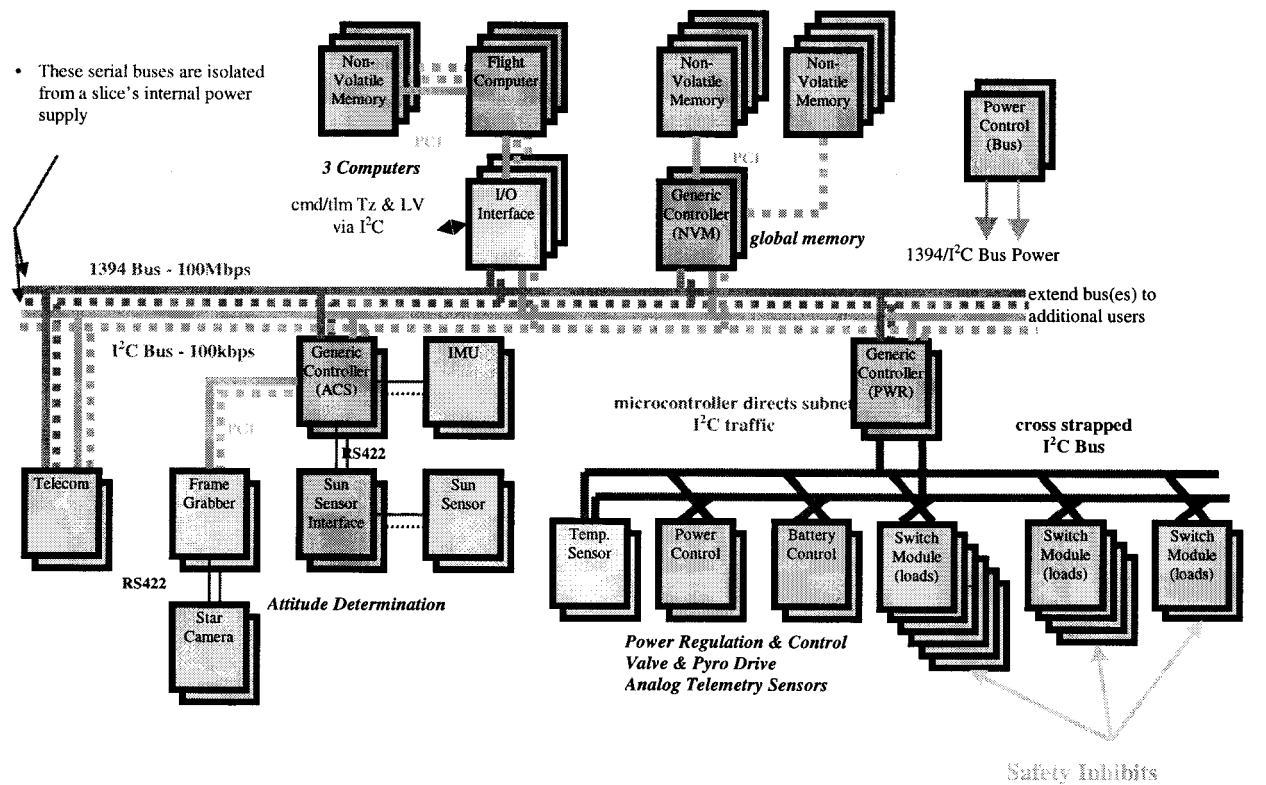

Figure 1: X2000 Avionics System Architecture

In addition to the redundant buses, the different topology of the 1394 and I2C buses also provide some advantage in fault protection. For example, the tree topology of the 1394 Bus is not inherently fault tolerant. If the Physical Layer of one of the branch nodes fails, the 1394 Bus will be partitioned into two subtrees, and it is impossible for a node in a subtree to communicate with a node in another subtree. However, since the I2C Bus has a multi-drop bus topology, it can bypass the failed node in the 1394 Bus to facilitate fault identification and recovery. On the other hand, the multi-drop topology of the I2C Bus is susceptible to shorting to ground or power supply in any node. In this case, the 1394 Bus can be used to facilitate the fault isolation and recovery of the I2C Bus.

Another challenge to this architecture is the subsystem interfaces. Since none of the off-the-shelf instruments has built-in 1394 or I2C bus interfaces, the X2000 architecture will provide a generic microcontroller for that purpose. The generic microcontroller will interface the 1394 and I2C buses with a PCI bus interface, some parallel and serial interfaces, an analog-to-digital converter, and possibly a 1553 interface. With these I/O interfaces, the generic microcontroller will be compatible with most of the off-the-shelf instruments. In addition, the microcontroller will have a separate subsystem I2C Bus to accommodate systems that have large number of low speed devices. The Power Subsystem with large number of power switch modules is a good example of how the subsystem I2C Bus can be used.

As far as the command and data handling concern, the high speed 1394 Bus provides a flexible architecture to accommodate multiple general-purpose flight computers. The software design will ensure that all the general-purpose flight computers are symmetrical. That is, each flight computer will have the capability to take over the tasks of another flight computer. One advantage of this symmetry is fault tolerance. The other advantage is that, during the cruise phase, all the flight computers can be run at lower speed to save power and increase reliability. In the event that a flight computer fails, its tasks can be taken over by other flight computers. It should be noticed that the other flight computers will have to run at higher speed to handle the additional workload, but since the failed computer can be powered off, the overall power consumption is relatively constant. It is fortunate that the 1394 Bus standard supports the shut down of the flight computer while keeping the Physical Layer alive by bus cable power. Hence, the powering off of the flight computer will not cause disturbances to the bus.

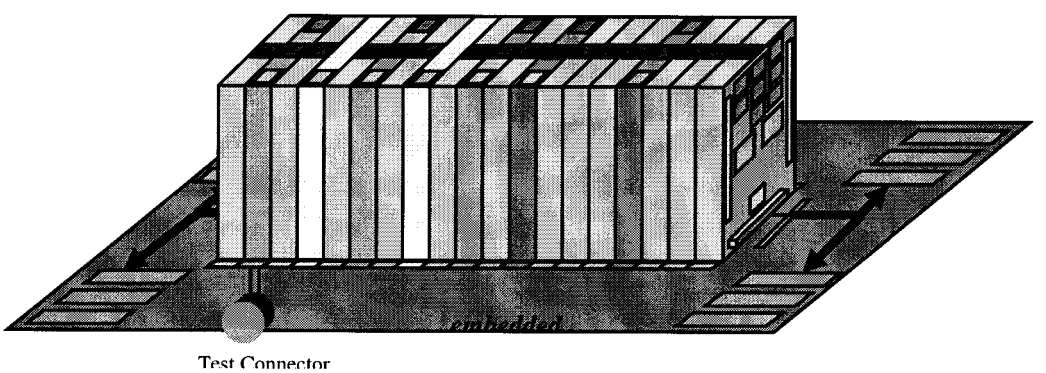

Figure 2: Packaging Concept of X2000 First Delivery

Finally, packaging is a key attribute of this architecture. Due to its modular design with standard interfaces, the X2000 architecture is amenable to a 3-D stacking approach comprising small units called slices (see Figure 2). The size of these slices is about 4 inch square. Bare dies will be used in these slices to increase the circuit density. All slices have the same form factor so that they can be stacked together like a loaf of bread. The stack will be laid horizontally. The bottom surface of this stack will be connected to the structure of the spacecraft. The structure will have traces (called embedded network) similar to those on a circuit board. The traces are used to replace the massive wire harness in traditional spacecraft design. The left and right sides of each slice have vertical connectors, so that multiple slices can be electrically connected together along the sides of the stack. The top surface of the stack will have some connectors for test purpose.

4.0 Conclusion

The X2000 is an ambitious undertaking of technology development. It faces many unprecedented challenges and taking risks in many areas. However, it is Jet Propulsion Laboratory's believe that the technologies developed by the X2000 will enable faster, better, cheaper space missions in the future. Furthermore, these technologies will not only benefit NASA, but it will have significant impact on the commercial aerospace industry as well.

Acknowledgments

The work presented in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the National Aeronautics and Space Administration.

Reference:

[1] IEEE Computer Society, "IEEE Standard for a High Performance Serial Bus,", IEEE Std 1394-1995, Aug 30, 1996

[2] Don Anderson, "FireWire System Architecture," MindShare, Inc., Addison Wesley, 1998.

[3] Philips Semiconductors, "The I 2 C-bus and how to use it (including specifications) 1995 update," Philips Semiconductor web site http://www-us2.semiconductors.philips.com/i2c

[4] Dominique Paret, Carl Fenger, " The I2C Bus, from Theory to Practice,", Addison Wesley, 1994.

[5] PCI Special Interest Group, "PCI Local Bus Specification, Rev 2.1s," 1995.

[6] Tom Shanley, Don Anderson, "PCI System Architecture," MindShare, Inc., Addison Wesley, 1995.

[7] Savio Chau, "X2000 Avionics System Conceptual Design Document," JPL Technical Report, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 1997.