

IAF-98-U.1.04

Custom Computing for Satellite Applications

P. Sutton and N. Bergmann

Cooperative Research Centre for Satellite Systems

Queensland University of Technology

**49th International Astronautical Congress
Sept 28-Oct 2, 1998/Melbourne, Australia**

CUSTOM COMPUTING FOR SATELLITE APPLICATIONS

Peter R. Sutton and Neil W. Bergmann

Cooperative Research Centre for Satellite Systems
Queensland University of Technology
GPO Box 2434, Brisbane 4001, Australia
p.sutton@qut.edu.au, n.bergmann@qut.edu.au

ABSTRACT

Custom computing is an important new technology for satellite applications. Custom computers use a conventional microprocessor in conjunction with reconfigurable hardware based on field programmable gate arrays. Custom computers can be configured on an application-by-application basis and can provide very high performance for many applications, particularly data stream processing. Satellite based custom computers offer increased onboard computing performance with the advantages of dynamic reconfigurability, post-launch hardware configuration upgrades, hardware resource sharing and potential accommodation of failures. A high performance custom computing module is being designed for incorporation into FedSat-1, a small satellite being developed by the Australian Cooperative Research Centre for Satellite Systems.

INTRODUCTION

The push to "Better, Faster, Cheaper" satellites has several implications for the design of onboard computers. Such computers must be designed faster, and they must provide increased performance - as measured by processing capability, mass, volume, power consumption, reliability and flexibility. A potential technology for meeting these requirements is *custom computing*.

Custom computing is a new computing paradigm in which reconfigurable logic in the form of static RAM (SRAM) based Field Programmable Gate Arrays (FPGAs) are used as a coprocessor resource for a conventional microprocessor. Conventional microprocessors process data in an inherently serial fashion. Application algorithms with high degrees of parallelism can be greatly accelerated

by implementing parts of the algorithm in reconfigurable logic. The reconfigurable logic may be programmed with different configurations to suit different applications. Performance exceeding that of a supercomputer at a fraction of the size and cost is possible in some application areas. Speed-up factors of 100-1000 have been reported^{1,2,3} with single chips performing the equivalent of 2 billion operations per second⁴.

This paper examines the applicability and advantages of using this new technology for satellite applications. The paper first presents a review of custom computing, and then describes satellite applications of custom-computing. Advantages and disadvantages of satellite-based custom computers are then described. An outline of an experimental small satellite custom computer payload is given, and some conclusions are drawn.

CUSTOM COMPUTING OVERVIEW

Research into the area of custom computing has been underway for about a decade, enabled by the availability of SRAM-based FPGAs of sufficient size to support useful computing subsystems. This section provides some background to the area of custom computing. FPGA architectures are discussed, followed by a description of the development process for a custom computer. Future directions for custom computing are also discussed.

Field Programmable Gate Arrays

The basis of all custom computers are field programmable gate arrays (FPGAs), specifically static RAM based FPGAs in which configuration information is programmed into the device at system initialisation or at function switching time. Internally, an FPGA typically consists of a two-dimensional array of configurable logic blocks (CLBs); a set of input/output blocks (IOBs) connected to device pins; and a network of interconnecting wires and switches which enable connections to be

Copyright © 1998 by Peter Sutton and Neil Bergmann. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Released to IAF/IAA/AIAA to publish in all forms.

made between CLBs, IOBs and other wires. The CLBs, IOBs and interconnecting wires are programmable - their function or connections are determined by the configuration information.

FPGAs fall into two major architectural groups: coarse-grain and fine-grain. Coarse-grain FPGAs have relatively complex CLBs, typically containing two generic 4-input/1-output logic block, plus two programmable flip-flops. Such structures are well suited to random logic replacement applications. Fine-grain FPGAs have relatively simpler CLBs (typically 2-input/1-output logic function plus one programmable flip-flop) and an interconnection structure which encourages nearest neighbour communications. Both types of FPGA have been successfully used for custom computing platforms.

A recent popular FPGA choice for reconfigurable computing has been the Xilinx XC6200⁶ series of devices*. These devices feature an "open" architecture, that is, the configuration details are available so that anyone may write a tool to configure the device. The chip may also interface directly to a microprocessor bus (even performing its own address decoding) with individual configuration bits directly readable and writable. This feature enables partial reconfiguration, that is, changing the configuration of only part of the device, possibly whilst the remainder of the device is operating. Data registers are also directly readable and writable via the microprocessor bus. This allows data to be read and written whilst the device is operating, for example, changing the filter coefficients in a digital filter. Unlike other Xilinx FPGA product families, it is not possible for a misconfiguration in an XC6200 series device to result in an internal power-ground short circuit.

XC6200 series devices have a register-rich, fine-grain architecture with an emphasis on nearest-neighbour communications. This architecture lends itself to fast systolic array based data-stream processing. Single XC6200 series devices may also be arrayed to form larger "computing surfaces" as in the (ground-based) SPACE-2 architecture⁵.

It is expected that devices with characteristics similar to those of the XC6200 will become more prevalent as custom computers move into the mainstream.

* Now discontinued.

The Development Process

Due to the young age of the custom computing discipline, there is still much research ongoing into how best to program custom computers. Most current approaches design the reconfigurable hardware separately from the software after a manual partitioning of the system. The software is written using a programming language such as C, whilst the hardware is designed using conventional schematic or Hardware Description Language (HDL) approaches. The hardware design must then be mapped to the underlying reconfigurable logic which may involve hardware synthesis followed by place and route.

Future Developments

In the future, it is likely that single chip custom computers will become commercially available. These devices would combine a general purpose microprocessor with reconfigurable logic on a single die. The reconfigurable logic would be used as an "execution unit" in the same way that integer and floating point execution units are used in modern microprocessors. An experimental device of this sort is the hypothetical Garp processor⁷ which combines a standard MIPS processor with a reconfigurable array.

The biggest stumbling block to such devices is the software required to program them. Compiler technology needs to develop so that code can be mapped to variable hardware, not just a fixed instruction set.

Another possible future development is time-multiplexed reconfigurable logic⁸. In this scenario multiple configurations are stored in on-chip memory and can be swapped-in in one memory cycle - perhaps a few tens of nanoseconds.

SATELLITE APPLICATIONS

Custom computers are suited to high bit-rate data processing, particularly for algorithms exhibiting a high degree of parallelism. There are many satellite data processing tasks that fall into this category. Two of these are described below.

Remote-Sensing Data

The data collection ability of many remote-sensing satellites is limited by the downlink capacity. On-board data filtering of data allows the most meaningful data to be downloaded, while on-board data compression allows a greater quantity of data to be downloaded. Custom computing is suited to both of these data processing applications: data filtering and compression both benefit by hardware acceleration. Custom computing allows *applica-*

tion-specific data filtering and compression - the filtering and compression algorithms used will depend on the characteristics of the data. Image processing is one application area where custom computers have shown their worth⁴. Reconfigurable hardware also allow the parameters used in the filtering or compression to be altered to suit current mission requirements.

Adaptive Communications Protocols

Low earth orbit satellites may only be visible to a ground station for periods of ten minutes or so. During this communication period, satellite elevation will range from around 10 degrees above the horizon up to 90 degrees if the satellite passes directly overhead. The maximum channel bit-rate may vary by an order of magnitude for these two elevations. A communications channel which can be adapted to current propagation characteristics may be very useful to maximize data transfer. For example, the number of error detection and correction bits in a data stream may be varied depending upon the current reliability of the channel.

For a high-speed channel, data encoding and decoding must be handled in hardware. Dynamic adaptation is ideally suited to custom computers, where the hardware configuration can be altered to provide the necessary changes.

ADVANTAGES AND DISADVANTAGES

This section describes the advantages and disadvantages of custom computing for use on satellites or other spacecraft.

Advantages

Data Processing

The principal advantage of custom computers is as outlined above - high speed real time data processing is possible with smaller and cheaper hardware resources. Likely application areas include data compression, data selection or filtering, cryptography, pattern matching and signal processing.

Reduced Design Time

Reconfigurable computers enable satellites to be developed and launched in shorter time-frames. Physical hardware design can be frozen early and even reused in many satellites. Hardware configurations can be developed at a later stage - possibly even post-launch.

Hardware Resource Sharing

Because custom computers can be reconfigured, a single onboard reconfigurable hardware resource can provide hardware acceleration to two or more applications. This hardware sharing can take place

on a time or logic basis. Time sharing involves all of the reconfigurable logic being used by a particular application for a given time slice. During a context switch, the hardware can then be reconfigured for use by another application. Logic sharing involves different parts of the reconfigurable logic being used concurrently by two or more applications.

Sharing of such a hardware resource can result in reduced hardware costs (such as weight, volume and power); one reprogrammable custom computer could replace several dedicated hardware devices.

Partial Reconfigurability

Some FPGA devices offer partial reconfigurability, that is, part of the device may be reconfigured whilst the remainder of the device operates normally. This permits hardware sharing as outlined above and also enables immediate changes to existing designs such as changing the constants of a digital filter, or changing the coding parameters in an adaptive communication protocol.

Post-launch Reconfiguration

Another advantage of custom computers for satellite applications is that new hardware configurations can be uploaded to the satellite - in a similar manner to post-launch software updates. Such updates may be used to correct an algorithm or to upload a new algorithm developed post-launch. Reconfiguration will also permit hardware failures to be accommodated, as described next.

Reliability

Custom computers offer the potential for higher reliability computing. Onboard failures may be accommodated by reconfiguring the hardware so as to avoid the failure. If parts of a reconfigurable device or array of devices failed, suitable detection would allow hardware configurations to be changed so as to bypass the failed components or blocks. An interesting topic for future research in this area is that of self-repairing hardware or dynamic self-reconfiguration.

Power Consumption Optimization

The power consumption of FPGAs is very dependent on the configuration. Power can be traded off against application speed and area: lower power consumption can mean larger area and/or slower performance. It is possible to imagine a custom computer using one of two different configurations based on available power or desired speed of processing.

Disadvantages

There are a number of disadvantages to using custom computers on satellites.

The principal disadvantage is the current lack of radiation hardened static RAM based FPGA devices. The need for such devices has been recognised and work is progressing on developing these devices^{9,10}.

Current static RAM based FPGA devices also have some limitations. Most can not be partially reconfigured - the whole device must be reconfigured. This process often takes several milliseconds - too slow for frequent context switches. Many FPGA designs (for example the popular Xilinx 4000 series) are such that a single bit-flip in configuration memory could cause an internal power-ground short circuit within the device. Power monitoring circuitry would be needed to detect such a problem and reset the circuit. New architectures (such as descendants of the Xilinx XC6200 family) are expected to solve these problems.

Another disadvantage is the current lack of suitable design tools for designing application software and hardware configurations in unison. Research is progressing in this field and it is expected that in the next few years such software will be more mature.

EXPERIMENT

We intend to demonstrate the practicality of satellite based custom computing by incorporating a high performance custom computer module into FedSat-1, a satellite being developed by the Australian Cooperative Research Centre for Satellite Systems (CRCSS). The CRCSS consists of several universities, research laboratories and industrial partners.

FedSat-1 will be a small (50kg and 500 x 500 x 500mm) experimental satellite intended to be launched into a low earth orbit in the 2000-2001 timeframe. It will be based on a commercially procured bus with Australian developed research payloads. These payloads will include a GPS receiver, a magnetometer and a communications module, along with the high performance custom computing module (HPC).

Experimental Goals

The aims of the experimental HPC payload are as follows:

- to provide high performance processing onboard FedSat-1;

- to allow experiments to occur on in-orbit reconfiguration of the HPC hardware; and
- to evaluate improvements in satellite performance offered by this approach.

The HPC module will be responsible for several tasks relating to the processing of data from other payloads. Data from the magnetometer payload will require compression and possible selection before transmission to ground. Maximum compression will be achieved by performing this operation in the frequency domain. Depending on the available downlink bandwidth, data dependent selection of the most "interesting" data may be required.

The HPC module will also be responsible for performing position estimation of mobile terrestrial transmitters based on Doppler frequency offset information from the communications module baseband processor and a predicted orbit model updated with GPS data.

In-orbit reconfiguration will be performed from hardware configurations stored onboard the satellite as well as new configurations uploaded from the ground.

Module Architecture

Whilst still in the preliminary design phase, it is anticipated that the HPC module will include a general purpose microprocessor, a digital signal processor and reconfigurable logic in the form of one or more FPGAs. Operation of the module will be controlled by the primary on-board processor. The HPC module will interact with the primary processor to

- obtain experimental data to process;
- to obtain new hardware configurations uploaded via the TT&C link; and
- to return processed data to mass memory storage for later transmission to the ground.

CONCLUSIONS

Custom computing is an up-and-coming technology which promises low cost high performance processing for some applications, particularly data stream processing. Custom computers use reconfigurable logic as a coprocessor resource for a conventional microprocessor. By loading different configurations into the programmable hardware, a custom computer provides dedicated hardware to service the current application. Such hardware can perform the equivalent of billions of operations per second for some applications. There are many satellite applications which could benefit from such processing. These include communications

and remote sensing algorithms, particularly data filtering, compression, encoding and encryption. Custom computers offer advantages such as reduced hardware physical design time, lower mass and volume through hardware resource sharing, optimized power consumption, post-launch hardware configuration upgrades and improved reliability and flexibility through hardware reconfiguration.

The Australian Cooperative Research Centre for Satellite Systems is developing a small satellite which will contain a high performance custom computer. We believe that custom computing using reconfigurable logic is an important technology for future satellite development.

ACKNOWLEDGEMENT

The support of the Australian government through the Cooperative Research Centres Program funding for the CRC for Satellite Systems is gratefully acknowledged.

REFERENCES

1. M. Gokhale, W. Holmes, A. Kopser, S. Lucas, R. Minnich, and D. Sweely, "Building and Using a Highly Parallel Programmable Logic Array," *IEEE Computer*, Jan. 1991, pp. 81-89.
2. J.E. Vuillemin, P. Bertin, D. Roncin, M. Shand, H.H. Touati, and P. Boucard, "Programmable Active Memories: Reconfigurable Systems Come of Age," *IEEE Transactions on VLSI*, Vol. 4, No. 1, March 1996, pp. 56-69.
3. J. Babb, M. Frank, E. Waingold, R. Barua, M. Taylor, J. Kim, S. Devabhaktuni, P. Finch and A. Agarwal, "The RAW Benchmark Suite: Computation Structures for General Purpose Computing," *Proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM'97)*, April, 1997.
4. R. Woods, D. Trainor, and J. Heron, "Applying an XC6200 to Real-Time Image Processing," *IEEE Design and Test of Computers*, vol. 15, no. 1, Jan. 1998.
5. G. Milne, P. Cockshott, G. McGaskill, P. Barrie, "Realising Massively Concurrent Systems on the SPACE machine," *Proceedings of IEEE Workshop on FPGA's for Custom Computing Machines (FCCM'93)*, Napa Valley, April 1993.
6. Xilinx Inc., *XC6200 Field Programmable Gate Arrays Product Description*, version 1.10, San Jose, Calif., April 1997
7. J.R. Hauser and J. Wawrzynek, "Garp: A MIPS Processor with a Reconfigurable Coprocessor," *Proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM'97)*, Napa Valley, April, 1997.
8. S. Trimberger, D. Carberry, A. Johnson, and J. Wong, "A Time-Multiplexed FPGA," *Proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM'97)*, April, 1997.
9. J. McCabe, "Radiation Hard Reconfigurable Field Programmable Array," *Proceedings of 1998 Military and Aerospace Applications of Programmable Devices and Technologies Conference (MAPLD)*, Greenbelt, Maryland, September 1998.
10. P. Alfke and R. Padovani, "Radiation Tolerance of High-Density FPGAs," *Proceedings of 1998 Military and Aerospace Applications of Programmable Devices and Technologies Conference (MAPLD)*, Greenbelt, Maryland, September 1998.