EXECUTIVE SUMMARY

PROJECT NAME: F6-10, Adaptively Fault-Tolerant Reconfigurable Architectures

INVESTIGATOR(S): Dr. Alan George, Dr. Ann Gordon-Ross, and Dr. Herman Lam, UF site of CHREC

PROJECT DESCRIPTION

The requirements of advanced computing systems for space as well as other high-performance embedded computing (HPEC) systems differ from traditional HPC systems. While high performance is obviously desired, other factors such as reliability and power consumption can be even more important. Meeting all of these factors while minimizing cost is a significant challenge, especially as processing requirements of future systems continue to scale higher. Commercial-off-the-shelf (COTS) microprocessors, FPGAs, and other fixed-logic or reconfigurable-logic multicore devices can provide required performance at a significantly reduced cost, but at the expense of less radiation-tolerant and reliable components.

This project will investigate and enhance COTS-based, fault-tolerant system architectures based upon a variety of software and hardware techniques, with an emphasis upon system-level fault tolerance and testing. Our concept of reconfigurable fault tolerance (RFT) exploits the ability of each FPGA to partially reconfigure itself while in operation, changing modes of fault tolerance concomitant with changes in radiation hazards, allowing real-time adaptation to environmental factors that may affect system reliability while avoiding the low performance of worst-case-scenario designs. By modifying the amount of redundant logic within one FPGA or a set of FPGAs, systems can dynamically tradeoff reliability for increased performance. Our concept of hybrid fault tolerance (HFT) is similar to RFT except that it features spatial instead of temporal mixes of varying modes of fault tolerance to meet the heterogeneous needs of missions and applications.

To validate results of a variety of fault detection and mitigation methods, sophisticated testing tools are required, which will be able to reach across hybrid systems. Our concept of comprehensive fault injection (FI) for HFT exploits a variety of debugging and programming techniques available for FPGA and microprocessor-based systems to form a portable FI framework which can be applicable to a large range of devices, such as various FPGAs, soft-core or hard-core processors, and many-core processors including the Tilera TILE64 and the OPERA Maestro.

Additionally, this project will explore amenability of adapting a NASA Dependable Multiprocessor SIFT methodologies and framework to many-core devices, such as TILE64 or Maestro, as well as investigate reconfigurable fault-tolerant networking for space-based applications by performing analytical and experimental studies on Time-Triggered Ethernet.

EXPERIMENTAL PLAN

Comprehensive fault injection for HFT task will explore methodologies and procedures for combining FI methods to evaluate susceptibility of hybrid systems comprised of FPGAs, traditional CPUs, and emerging many-core devices. Our Simple Portable Fault Injector (SPFI) was introduced in CHREC project F6-09 as a proof-of-concept set of disjoint tools. F6-10 will extend and unify SPFI to support multiple devices including but not limited to FPGAs, embedded and standalone CPUs, as well as many-core devices. In addition, statistical and experimental methods for result aggregation and subcomponent testing will be explored in order to provide better means for improved system-level reliability prediction.

RFT was originally introduced in CHREC project F4-08 as a proof-of-concept for partial reconfiguration, and extended to provide several FT modes (unprotected, self-checking pair (SCP), triple-modular redundancy (TMR), or algorithm-based fault tolerance (ABFT)) in project F6-09. F6-10 will augment the RFT framework to provide FPGA services to allow for state saving and recovery using hardware checkpointing and rollback techniques on FPGA systems. User-directed and automated methods of saving FPGA state in arbitrary designs will be explored and analyzed.

Techniques from the NASA Dependable Multiprocessor project and experiences at Florida will be adapted to many-core architectures, using the TILE64 as a prototype platform. A variety of architecture choices will be explored, including task and hardware granularity and network communication protocols.

HOW THIS PROJECT IS DIFFERENT

The HFT task leverages previous work in fault-tolerant systems, such as the NASA Dependable Multiprocessor developed by Honeywell and the University of Florida, which pioneered and featured a fault-tolerant software framework to mitigate COTS radiation susceptibility for a suite of FT modes and space-based applications for a space supercomputer. While many other research groups have studied fault-tolerant approaches for FPGAs (e.g. scrubbing and TMR), the RFT task focuses on providing low-overhead, system-level coverage using adaptive fault tolerance with partial reconfiguration. Additionally, software-based fault injection approaches for alternative computing devices has not been previously studied.

POTENTIAL MEMBER COMPANY BENEFITS

- Influence over project direction, ensuring relevance to interesting systems, applications, and problems
- Access to research exploring novel ways to provide fault tolerance in FPGAs and RC systems

EXPECTED DELIVERABLES

- Prototype of comprehensive fault injection framework for systems with hybrid architectures
- Check pointing framework allowing for automated fault recovery in RFT system
- Several scholarly conference and/or journal publications

PROJECT BUDGET: 3.5 memberships