SAND2022-9437C

The Center for Cyber Defenders

Expanding computer security knowledge

FPGA Partial Reconfiguration

Reprogramming FPGA logic at runtime

Michael Carrow

Project Mentor: Rachel Goen and Sean Jensen, Org. 5645

Problem Statement:

- Field Programmable Gate Array (FPGA) devices are very versatile, but they also have many limitations. Each FPGA has a limited number of programmable logic blocks, which can make it difficult to fit all of the desired logic into a design. In addition, the logic design is usually static after initial programming, and cannot be changed without resetting the whole device. This can be especially problematic in the event of a Single Event Upset (SEU) error, in which a radiation induced bit flip occurs in the FPGA's configuration memory, causing unwanted changes to the logic.
- The goal of this project is to use Partial Reconfiguration (PR) to alleviate some of these downsides. PR is an FPGA feature supported by certain advanced devices that allows the user to replace sections of an FPGA design at runtime without interfering with the rest of the running design.

Objectives:

- Determine how Partial Reconfiguration can be used to create more efficient FPGA designs
- Examine the possible applications, capabilities, and limitations of Partial Reconfiguration
- Create usable reference designs and documentation for PR
- Use PR to recover from SEU errors

Approach:

- Used Intel Quartus Prime and a Stratix 10 developer kit to run experiments using PR
- Created multiple test designs that include partially reconfigurable sections
- Implemented multiple versions of the same section, and experimented with reconfiguring those sections at runtime in various situations.
- Programmed an onboard CPU within the FPGA to automatically control PR
- Currently planning an experiment on integrating SEU detection to automatically reset sections that have been affected by SEU

Findings

- PR requires planning ahead of time. Any section that can be partially reconfigured must be defined as a separate design partition.
- FPGA designs can be programmed to automatically use PR to swap out different versions of a logic section at will, allowing a design to reuse the same logic resources to perform multiple functions.
  - Automatic PR can be controlled by an external processor or another section in the same FPGA.
  - PR partition files can be stored on external memory or sent to the device via a variety of communication methods such as ethernet, PCIE, or JTAG.
- With proper planning, PR can be used to fix bugs, update, or change parts of a design in the field without needing to restart the entire system. The desired changes must be inside a predefined reconfigurable partition
- PR could possibly be used to recover from corrupted FPGA logic due to SEU by reloading the original logic on the affected sections

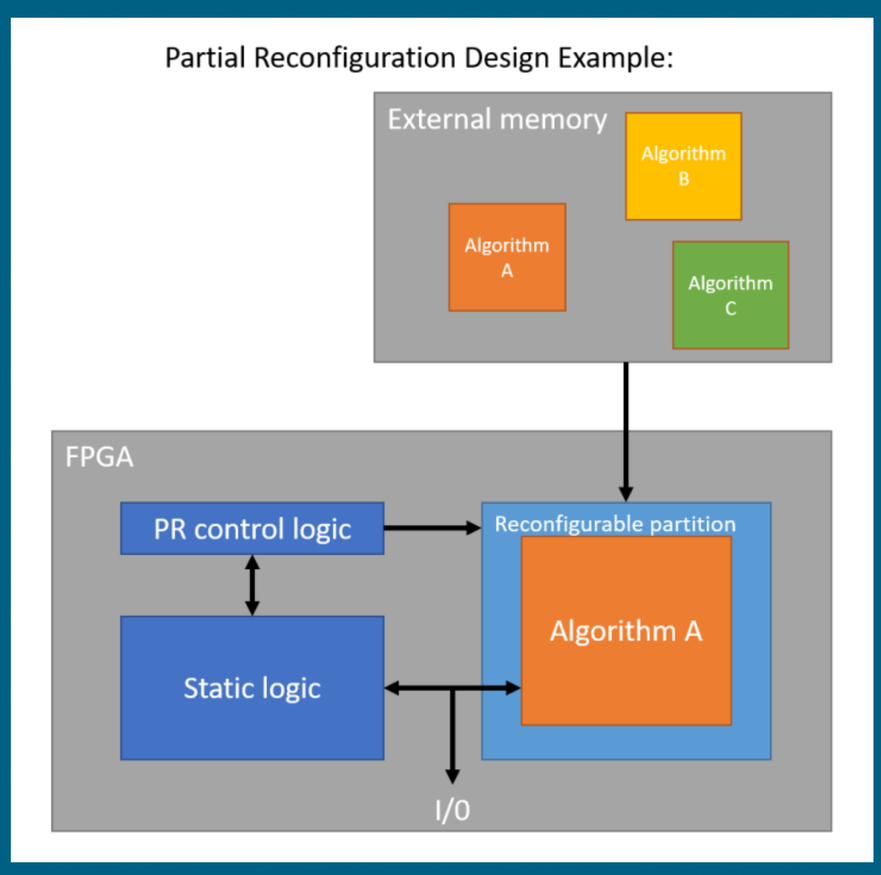

Impact and Benefits:

- Reduced costs by reusing FPGA resources, allowing designs to fit on a smaller FPGA
- Reduced power usage by only running one version of a partition at a time
- In-situ updating or editing of designs in critical systems where a full restart is undesirable

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2017-XXXXC.