



technology opportunity

FPGAs with Reconfigurable Fault-Tolerant Redundancy

Enables users to choose increased capacity over redundancy



NASA's Johnson Space Center has developed a technology that enables selective reconfiguration of field programmable gate arrays (FPGAs) and similar devices between redundant and non-redundant operation, to meet application needs for the right mix of reliability and high capacity. This innovation allows the flexibility of firmware redundancy while maintaining the efficiency and simplicity of hardware-based redundancy.

Benefits

- **Increased capacity:** User-defined programming provides triple redundancy in only the most critical areas.
- **Improved performance:** Hardware TMR saves at least one "logic level" over firmware redundancy, providing hardware redundancy performance without sacrificing necessary capacity.
- **Lower manufacturing cost for end users:** Higher capacity allows smaller FPGAs for end users, reducing unit and circuit board costs significantly.
- **Lower development cost for end users:** No need to implement TMR into firmware shortens the required design time.
- **High efficiency:** Configuration memory does not have to be separately voted because at least two configuration bit errors are required to circumvent the voters.

Applications

This technology will be beneficial primarily in the high-radiation environments found in the aerospace industry. Key FPGA types include:

- Radiation-tolerant SRAM-based FPGAs
- Radiation-tolerant flash-based FPGAs

For example:

- Signal processing for software radio
- Sensor data analysis
- Automated docking
- Landing sensor data processing (e.g., LiDAR processing)
- Hazard detection and avoidance
- Network packet routing
- Video processing and video display updates
- Communications network infrastructure (e.g., cell phone base stations and high-reliability network data routing)

Technology Details

Typical FPGAs either provide fixed redundant circuitry that enables tolerance of faults such as a single event upset (SEU) or single event transient (SET), or they provide no redundancy at all and the user must program any required redundancy in firmware, causing high overhead costs. Although hardware redundancy provides high performance and assurance, many applications only need to be partially protected from SEU/SET, and other parts of the applications require higher capacity.

How It Works

This innovation uses a hardware implemented voting scheme with two modes—Redundancy Mode, which provides full triple redundancy protection, and Split Mode, which eliminates redundancy protection in selected areas to increase capacity. With Split Mode, end users have the option of choosing the amount of redundancy in their design but can still implement TMR in an efficient manner.

The device has three identical sets of functional units, routing resources, and majority voters that correct errors. It modifies the voter to accept a mode input, which specifies whether ordinary voting is to occur or redundancy is to be split. In Redundancy Mode, the voters work in the usual manner, producing an output corresponding to the two inputs that agree. In Split Mode, each voter selects a different input and conveys this to the output. By changing the operation of the voters, the sections can operate independently.

A single event upset assumes that only one fault will occur in a voting group within one voting cycle, and thus, the fault can be eliminated by majority voting. The only connection between the three sections of the device is through the voters.

The voters also effectively mask errors in the configuration memory because all data passing between sections pass through the voters. At least two configuration bit errors in unrelated (non-adjacent) parts of the device are required to circumvent the voters.

To partition an application, the device is divided into independently configurable blocks. When operating in Split Mode, additional routing resources are used to communicate between the split sections in a block. This additional routing also requires at least two configuration bit errors to generate an error in a redundant block.

Because both the register (storage) elements and combinational logic are redundant, this method protects against SEU in the storage elements and SET arising in the combinational logic.

Why It Is Better

Because up to 95% of the device's capacity can be designed for normal, non-redundant operation in Split Mode, minimal silicon space (or capacity) is used when a device manufacturer includes this capability. Developers have the ease of use of hardware redundancy. The application benefits from the speed and low power of hardware redundancy, from protection against errors in both logic and register elements, and from high capacity when needed for non-redundant implementation.

Patents

Johnson Space Center is seeking patent protection for this technology.

Licensing and Partnering Opportunities

This technology is part of NASA's Innovative Partnerships Program (IPP), which seeks to transfer technology into and out of NASA to benefit the space program and U.S. industry. NASA invites companies to consider licensing this technology for FPGAs with Reconfigurable Fault-Tolerant Redundancy (MSC-24464-1) for commercial applications.

For More Information

If you would like more information or want to pursue transfer of this technology please contact us at:

Innovation Partnerships Office

NASA's Johnson Space Center

Phone: 281-483-3809

Email: jsc-techtran@mail.nasa.gov

Web: <http://technology.jsc.nasa.gov>