

Advances and Trends in Dynamic Partial Run-time Reconfiguration *

Dirk Koch, Jim Tørresen
Department of Informatics
University of Oslo, Norway

1 Dynamic Partial Runtime Reconfiguration is Becoming Mainstream

The progress in silicon industry has resulted in a tremendous increase in device capacity of FPGAs. As illustrated in Figure 1, the smallest devices of the upcoming Altera Stratix 5 FPGAs as well as the announced Xilinx Virtex-7 FPGAs provide more than double the amount of logic and embedded memory as the flagship devices of the one decade old Stratix or Virtex-II series FPGAs. By passing the one million LUTs border, high density FPGAs are sufficient to host 250 softcore CPUs plus the required peripherals.




Fig. 1. One decade of FPGA evolution. The figure denotes the increase in logic density over time and over the corresponding process technology.

* This work is supported in part by the Norwegian Research Council under grant 191156V30. It summarizes results of the Dagstuhl Seminar 10281 *Dynamically Reconfigurable Architectures*.

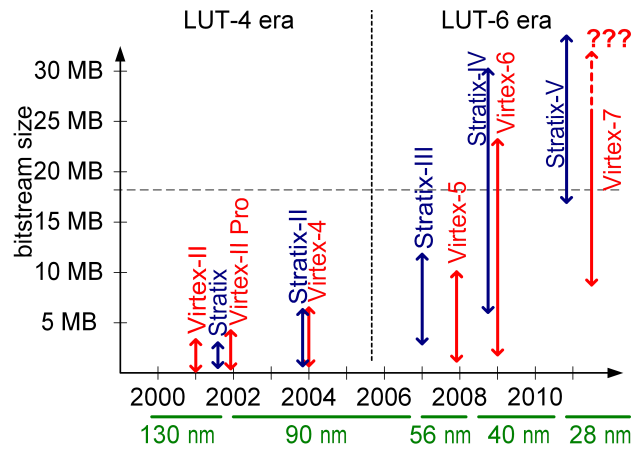


Fig. 2. One decade of FPGA evolution. The figure lists the increase in configuration memory size (total bitstream). By comparison, the largest currently available SRAM device provides 18 MB (Cypress 2009).

As the functionality of the here regarded FPGAs is defined by SRAM cells, the corresponding configuration data has also rose towards tens of megabytes (see Figure 2). Note that the highest capacity FPGAs typically provide more SRAM cells of what can be found in the largest SRAM memories released at the same time.

This progress in device density results in two drivers for introducing dynamic partial runtime reconfiguration in all future high-density devices: 1) increasing vulnerability to SEUs and 2) an increase in the configuration time.

1.1 Increasing Vulnerability to SEUs

By dramatically increasing the total amount of configuration bits as well as by shrinking the configuration memory cells at the same time, FPGAs have become more vulnerable to single event upsets (SEU). As compared to ASICs, a single event upset may not only result in an error in the datapath but much more likely in a malfunction of the circuit. This results from possible SEUs in the configuration SRAM cells that describe the present circuit that has been loaded to the FPGA [1]. Not every SEU results necessary in a malfunction as, for example, not all resources of an FPGA (e.g., logic or routing resources) are used in a particular circuit. However, if SEUs are not corrected, multiple event upsets might occur and therefore dramatically increase the risk of a circuit failure [2].




Fig. 3. One decade of FPGA evolution. The figure lists the configuration time of a full bitstream when considering the fastest possible configuration speed.

For dealing with SEUs, different application and safety scenarios have to be considered [3]. If faults can be temporarily accepted (can be seen as noise), it is sufficient to permanently overwrite the existing configuration (or parts of it) while keeping the device in active operation mode. This process is called *configuration scrubbing*. Note that configuration scrubbing can be combined with other fault tolerant techniques, such as triple modular redundancy (TMR) [4, 5].

In other applications, the present configuration has to be validated (e.g., by reading back the configuration data) before committing the result, like for instance, in banking applications.

We can summarize that *partial runtime reconfiguration is required for implementing safety critical systems*.

1.2 Configuration Bootstrapping

As the configuration bitstream size has enormously increased (see Figure 2), FPGAs vendors have started to introduce faster configuration interfaces in order to avoid an exponential rise in configuration time (see Figure 3). This was mainly achieved by widening the configuration ports (e.g. from 8-bit in Virtex-II to 32-bit in Virtex-4 FPGAs).

However, the initial configuration is typically still provided by a relatively slow flash memory device [6]. This is critical for many applications that demand fast availability after power-up, like for example, a PCIe interfaces implemented on a FPGA [7]. This issue can be solved with

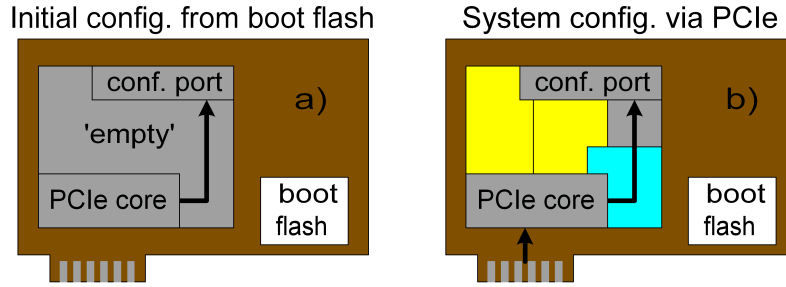


Fig. 4. Configuration bootstrapping example for PCIe solutions. a) After power-up, mainly the PCIe core is configured while leaving the rest of the device empty in order to fulfill fast PCIe core activation. b) Finally, the remaining system is loaded in a second non time-critical phase to the device.

the help of *bootstrapping* where in an initial configuration step only the modules requiring fast availability will be loaded to the device while finishing the configuration in a further step. An example of this procedure is illustrated in Figure 4.

We can summarize that *partial runtime reconfiguration allows to implement faster start-up times*. Together with the increasing vulnerability to SEUs, this forces FPGA vendors to include partial runtime capability in all their future devices. This can be identified by the vendor Altera Inc. that recently announced to provide full support for dynamic partial run-time reconfiguration in their Stratix-V series [8].

2 Context-Switching on FPGAs

With the widely introduction of partial run-time reconfiguration in all future high-density FPGAs, the technical basis for implementing context switching on FPGAs will be provided. In context switching FPGA-based systems, parts of the FPGA fabric will be shared by multiple modules over time. This is similar to multi-context task execution in software systems. However, in the hardware case, we have to distinguish between 1) the context of the FPGA (*device level*) and 2) the context represented inside the memories of the modules (*module level*). As depicted in Figure 5, the device level is given by the present module layout that can vary when (partially) reconfiguring the FPGA. Respectively, the module level is represented by the flip-flop values or the content of RAM blocks within the entire modules located on the device.

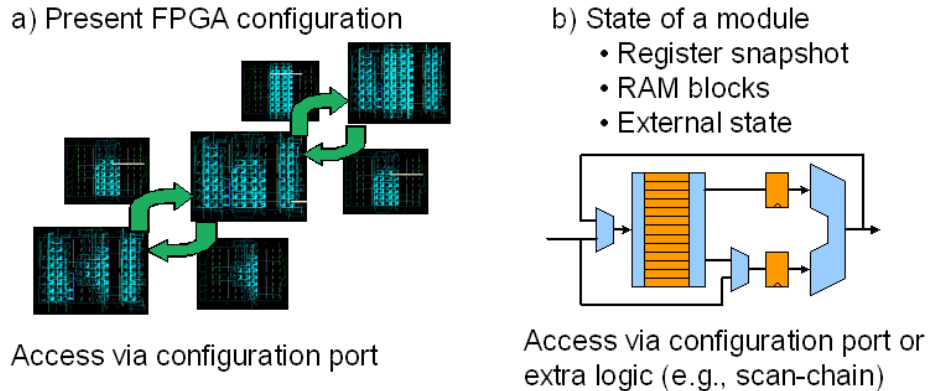


Fig. 5. Context-switching on FPGAs. The context of an FPGA is twofold: on a) the *device level*, it is represented by the present FPGA configuration and on b) the *module level*, by the internal state of the modules.

3 COSRECOS: Making Context-Switching Available

Despite the technical basis and more than two decades of intensive research on exploiting partial run-time reconfiguration, there still exist a wide gap between the possibilities and what is currently available to implement reconfigurable systems. This can be seen by the little acceptance in industry for this topic, and despite the opportunity to save cost, area, and power at the same time, partial runtime reconfiguration is still very exotic.

The aim of the COSRECOS project (Context Switching Reconfigurable Hardware for Communication Systems) [9] is to bridge this gap. By developing novel methodologies and advanced tools, we want to make the implementation of reconfigurable systems and their operation as simple as it is known from the software world. Consequently, we are investigating 1) *design-time aspects*, including models, analysis, debugging, and tools as well as 2) *run-time aspects* where we focus on high-speed reconfiguration, temporal module placement and interprocess communication. Moreover, we will 3) demonstrate our approaches on applications from the networking and general purpose domain.

With this project we want to contribute to the research community with a strong emphasis on active collaborations as well as to enhance acceptance in industry for using dynamic partial run-time reconfiguration.

References

1. Caffrey, M., Graham, P., Johnson, E., Wirthlin, M.: Single-event upsets in SRAM FPGAs. In: Proceedings of the 5th Annual International Conference on Military and Aerospace Programmable Logic Devices (MAPLD). (2002)
2. Gebelein, J., Engel, H., Kebschull, U.: An approach to system-wide fault tolerance for FPGAs. In: Proceedings of the International Conference on Field Programmable Logic and Applications (FPL 2009). (2009) 467–471 ISBN: 978-1-4244-3892-1, DOI: 10.1109/FPL.2009.5272477.
3. Alderighi, M., Casini, F., D'Angelo, S., Pastore, S., Sechi, G.R., Weigand, R.: Evaluation of Single Event Upset Mitigation Schemes for SRAM based FPGAs using the FLIPPER Fault Injection Platform. In: DFT '07: Proceedings of the 22nd IEEE International Symposium on Defect and Fault-Tolerance in VLSI Systems, Washington, DC, USA, IEEE Computer Society (2007) 105–113
4. Heiner, J., Sellers, B., Wirthlin, M.J., Kalb, J.: FPGA Partial Reconfiguration via Configuration Scrubbing. In: Proceedings of the International Conference on Field Programmable Logic and Applications (FPL 2009). (2009) 99–104
5. Lima, F., Carro, L., Reis, R.: Designing fault tolerant systems into SRAM-based FPGAs. In: Proceedings of the 40th annual Design Automation Conference (DAC 03), New York, NY, USA, ACM (2003) 650–655
6. Koch, D., Beckhoff, C., Teich, J.: Hardware Decompression Techniques for FPGA-Based Embedded Systems. ACM Transactions on Reconfigurable Technology and Systems (TRETs) **2** (2009) 1–23
7. Altera Inc.: Introducing Innovations at 28 nm to Move Beyond Moore's Law (2010) available online: www.altera.com/literature/wp/wp-01125-stxv-28nm-innovation.pdf .
8. Altera Inc.: Altera Unveils Innovations for 28-nm FPGAs (2010) available online: http://www.altera.com/corporate/news_room/releases/2010/products-nr-innovating-at-28-nm.html .
9. University of Oslo: The COSRECOS project website (2010) <http://www.matnat.uio.no/forskning/prosjekter/crc/>.