Dagstuhl Seminar Proceedings, Abstract, Dynamically Reconfigurable Architectures
J. Becker, J. Teich, G. Brebner, P. Athanas

Dynamic and partial reconfiguration of hardware architectures such as FPGAs and XPPs brings an additional level of flexibility in the design of electronic systems by exploiting the possibility of configuring functions on-demand during run-time. This has led to many new ways of approaching existing research topics in the area of hardware design and optimization techniques. For example, the possibility of performing adaptation during run-time raises questions in the areas of dynamic control, real-time response, on-line power management and design complexity, since the reconfigurability increases the design space towards infinity.

This Dagstuhl Seminar on Reconfigurable Architectures has aimed at raising a few of these topics e.g. on-line placement, pre-routing/on-line routing trade-off, power minimization etc., and also at presenting novel ideas on how to overcome the difficulties introduced in dynamic reconfigurable systems.

The first issue in a reconfigurable system is how to manage the dynamic placement of functional modules on the reconfigurable area. Current reconfigurable architectures are restricted on the granularity of the reconfigurable parts, which results in a limited flexibility in the partitioning of the reconfigurable blocks. Novel methods for overcoming these difficulties had been discussed during the seminar, as well as how future alternative architectures can be adapted for avoiding such problems.

The highest level of flexibility during reconfiguration would be achieved with on-line routing of functional descriptions in order to dynamically place functions on hardware resources that are currently available. On-line routing would make it possible to dynamically adapt the bit width of communication channels. Such a scenario would, however, lead to an increased response time and in some cases even an increased level of power consumption, due to the increased processing required to perform the re-routing. It seems reasonable to introduce dynamic reconfiguration in the system behavior in a way that enables the dynamic adaptation to changing system requirements on e.g. power consumption, performance, accuracy and resource utilization. For example, the placement and routing of a function can be dynamically adapted to provide either high performance or low power, according to the current system workload and power limitations.

These first investigations on dynamically reconfigurable architectures presented and discussed within this Dagstuhl Seminar dealt with clarifying exactly what level of flexibility can be achieved in current architectures, how to further increase this flexibility by exploiting novel methods, and how to integrate and exploit this flexibility in an optimal manner.

The workshop started on Monday morning with the mutual introduction of each participant and by posing exactly one question to the audience on what he/she sees the most important challenge in reconfigurable computing. Many questions dealt with the problem of dynamic reconfiguration. Issues such as need (killer applications), testability, speed and overhead were questioned initially until Andreas Koch found and summarized on Wednesday evening more than ten application areas requiring or benefiting from dynamic hardware reconfiguration. One of these applications was also presented by Oliver Diessel on Thursday morning, namely an industrial project exploiting dynamic hardware reconfiguration, namely positioning of satellite receivers that exploit the diversity in satellite signals to mitigate the effects of interference. Other questions dealt with reconfigurable interconnect, power consumption problems and competitiveness issues of actual reconfigurable devices with respect to ASIC implementations. On the tool side, questions focused on actual problems about stability and support for design flows enabling dynamic hardware reconfiguration. Finally, questions addressed the usefulness of libraries for reconfigurable computing. These questions formed the basis of the Monday evening breakout session where individual groups discussed application domains and module types (such as processing, I/O, memory, communication, operating systems types) that might be reusable in the reconfigurable computing community. A collection of domains and useful module types was summarized and added to the web page http://www.r-space.de under the name ReCoLib. This website is intended to collect library modules as well as benchmarks for reconfigurable computing in the future.

The third day of the seminar (Wednesday) contained just four presentations during the morning (with the presenters invited to be controversial), and a group discussion in the evening. The theme for the day was: “New thought models for reconfigurability and programmability”. The aim was to look beyond immediate aspects of dynamically reconfigurable architectures towards a longer-term research

agenda in the general area of the seminar. In an opening presentation, Gordon Brebner presented arguments in support of entitling the next seminar in the series as “dynamically adaptable behaviors” rather than “dynamically reconfigurable architectures”. The intent was to support a move towards a more applications-centric view of the field. That is, given the relative maturity of research on the architectural side, one should seek natural interpretations of reconfigurability in the needs of applications. The presentation did not advocate ad hoc per-application design approaches, but favored research in domain-specific frameworks that support desirable dynamically adaptable behaviors. In feedback, and on other occasions during the seminar, there seemed to be some consensus on the viewpoint. Rather than perennially chasing one elusive “killer application” for general reconfigurable architectures, we can find a variety of uses by studying a variety of applications. Apt choice of applications is important though, and Peter Athanas made this point strongly, with the presentation titled “The (empty?) Promise of FPGA Supercomputing” conveying his general argument: that there may be serious shortcomings when attempting to use current reconfigurable technologies for high performance computing applications. Hartmut Schmeck contributed a very interesting overview of the new field of “organic computing”, where there is a particular need for properties like robustness, adaptivity, and flexibility. His presentation offered a general systems setting – requiring self-organization with some degree of control - that is perfectly compatible with implementing dynamically adaptable behaviors using dynamically reconfigurable architectures. Finally, Reiner Hartenstein discussed the three paradoxes of Reconfigurable Computing: the low power paradox; the high performance paradox; and the education paradox. The last of these is a critical one for further development of the field. In order for a research breakthrough to have widespread impact, broad education in the new possibilities is essential to encourage thinking away from tradition paradigms. This of course has to be backed up by appropriate software tools, not just physical-design technologies. In the evening, a discussion session was held, titled “Dynamic reconfiguration considered harmful?”. This featured interesting and wide-ranging discussion of the achievements in our area and the practical effect of the research activities. One particular concrete outcome was the collection of a “Top Ten” list of applications of dynamic reconfiguration, collated by Andreas Koch. This served as a very useful assurance that the efforts of researchers in this area have not been in vain. It also confirmed that different behaviors of different applications provide a different scope for the use of reconfigurable architectures.

The gap between tools and available technology became the central topic on Thursday with eight very interesting contributions. Half of these talks dealt with operating systems and task concepts in support of reconfigurable computing devices such as FPGAs. Neil Bergmann, for example, advocated the use of standard Linux processes to enable a codesign for reconfigurable hardware/software designs by introducing ideas how to port typical Linux IPC mechanisms into hardware. Vincent Mooney presented a framework called Delta for generating RTOS for FPGAs including possibly multiple processors, and, Marco Platzner and Florian Dittmann presented algorithms for scheduling reconfigurable hardware modules with real-time constraints on dynamically reconfigurable devices such as the Erlangen Slot Machine (ESM), an FPGA-based platform for supporting dynamic reconfiguration and inter-module communication between modules arranged in slots. The talks by Roger Woods, Andreas Koch addressed questions of compiling data flow application, and C-code specifications respectively onto reconfigurable hardware. Finally, Klaus Waldschmidt introduced interesting perspectives on how reliability affects power consumption for multi-core architectures based on the introduction of novel dynamic power management techniques that avoid thermal cycling, combined with dynamic workload distribution.

On Wednesday afternoon, we had a photo of the participants taken and organized – as is a must when visiting Dagstuhl - a hiking trip – in this case to the Hochwaldalm in Wadrill. The weather luckily was bright and sunny, and, last but not least, a bus returned us home safely in the evening.