

ADAPTERS

Sanjaya Kumar
Honeywell Technology Center
(612) 951-7107, skumar@htc.honeywell.com
Open Presentation

Extended Abstract

Introduction

For many types of applications, it is important to develop systems that can adapt to changing computational, environmental, and fault-tolerance requirements. These adaptive computing systems (ACS) hold promise in being able to achieve higher performance and fault-tolerance in a more cost-effective manner than today's systems. However, as architectures for ACS evolve, the complexity of these systems increases. Complex, heterogeneous systems that utilize a diverse collection of reconfigurable computing elements present several challenges to the system designer. Thus, in order to take advantage of these architectures, new tools and technologies are required that enable their development and programming. The ADAPTERS program is addressing the problem of implementing a complex application onto a distributed, heterogeneous collection of resources consisting of general-purpose processors, application-specific processors (for example, digital signal processors), and field-programmable gate arrays (FPGAs), where some FPGAs may support either full or partial reconfiguration. Applications being considered include constant false alarm rate (CFAR), synthetic aperture radar (SAR) automatic target recognition (ATR), missile tracking, and various space applications.

The approach is to develop an integrated co-design environment focusing on three primary technology areas: 1) programming environments, 2) partitioning, mapping, and trade-off analysis techniques, and 3) dynamic reconfiguration of FPGAs. These three technology areas are being developed as the domain-specific application development environment (DADE), partitioning, mapping, and trade-off analysis environment (PMTE), and the dynamic reconfiguration run-time environment (DRRTE), respectively. The DADE supports a unified data-flow programming model, and allows both the software and hardware architecture to be captured. The PMTE consists of two toolsets: 1) a tool that supports design trade-off exploration by considering different partitions and mappings using genetic algorithms, and 2) a system level, token-based hardware/software modeling environment. The DRRTE supports analysis and implementation of real-time systems, expression of mode-based behavior, and the dynamic reconfiguration of FPGAs.

Domain-specific Application Development Environment (DADE)

The DADE supports the representation of software applications and hardware architectures within a system. In addition to general-purpose processors and application-specific processors, the DADE supports the capability of representing application functions to be mapped onto FPGA elements, and the representation of FPGA components within a hardware architecture. Functionality has been provided to support

the expression of VHDL behaviors that can be implemented on FPGAs using the Synopsys Design Compiler. The use of templates, abstractions that allow the representation and customization of various FPGA implementations, is supported as well. These abstractions can be employed to support the representation of applications employing variable precision arithmetic. Also, the Vector Signal Image Processing (VSIP) libraries being developed by Colorado State University have been incorporated into the DADE. A future version of the DADE will demonstrate the ability to express applications employing VSIP primitives and map them onto FPGAs using the compilation tools being developed by Colorado State University. A multi-board CFAR example has been developed to demonstrate the representation of software/hardware architectures (including Virtex FPGAs), and integration of the partitioning and mapping technology with the DADE (explained in more detail below). The representation of a SAR ATR application is underway.

Partitioning, Mapping, and Trade-off Environment (PMTE)

Using the software/hardware architectures captured in the DADE, the PMTE toolset provides an approach for examining several different mappings of software (application) functions onto hardware resources, subject to various constraints, such as size, weight, and power. The environment supports the capability of characterizing FPGA elements, and supporting trade-off analysis between general-purpose processors, application-specific processors, and FPGAs. A CFAR example is used to demonstrate the process of performing trade-offs. The tool supports a variety of analyses: CPU/FPGA loading, bus loading, combined CPU/FPGA/bus loading, and latency.

In addition to the partitioning and mapping technology, a VHDL-based hardware/software system level modeling environment exists that supports the performance analysis of heterogeneous systems which include FPGA elements. The FPGA models can be simulated with general-purpose processor and application-specific processor models. Functionality exists to allow the simulation of time-sharing of tasks on a single FPGA, which supports either full or partial reconfiguration.

Dynamic Reconfiguration Run-Time Environment (DRRTE)

To illustrate the concept of mode-based dynamic reconfiguration, a simple example was developed based on an image compression algorithm containing the following steps: 2D-wavelet, quantization, run-length encoding, and entropy coding. This application was implemented on a PC containing a Pentium and a WILDFORCE board employing Xilinx 4025 and 4013 FPGAs. The 2D-wavelet and quantization functions were downloaded into a single FPGA sequentially, based on the mode of operation, and the remaining functions were implemented on the Pentium.

To further develop these ideas, a missile tracking application from AMCOM is under investigation. It is being implemented on a VME-based platform consisting of Wind River System's VXWorks, a Motorola MVME2604 board (PowerPC 604), and an Annapolis Micro Systems WILDSTAR board with three Virtex FPGAs.