

CS6811: Research Seminar on Reconfigurable Hardware

Lockwood : Spring 2002

Review of: Splash 2 FPGAs in a Custom Computing Machine

- Book by:

- Duncan A. Buell
- Jeffrey M. Arnold
- Walter J. Kleinfelder

–IEEE Computer Society


- Summarized by: Todd Sproull

Overview

- Complete description of Splash 2
 - History / Previous work (Splash I)
 - Hardware Implementation
 - Software Architecture
 - Example Applications
 - Genetic Database search
 - Text searching
 - Fingerprint Matching
 - Image Processing
 - The promises and problems

What is the Splash 2?

- Attached Processor System using Xilinx XC4010 FPGAs as processing elements
- Developed at Center for Computing Sciences
- A follow up to the Splash 1



Splash 1

- History
 - Developed by Dick Kunze and Paul Schneck at Center for Computing Sciences in 1986
- Generalization of P-NAC (Princeton Nucleic Acid Comparator)
 - Used for comparing DNA sequences
- A single Multiwire board plugged into VMEbus of a Sun Workstation
- 32 Xilinx XC3090 FPGA chips
 - Max speed of FPGAs 32 MHz
- Programming Splash 1 was hardware design, completely custom

Splash 2 Goals

- Set of design issues Splash 2 needed to address
 - Programmability
 - Splash 1 was programmed in LDG (Logic Description Generator)
 - Xilinx tools were not very helpful at the time
 - Placement was sometimes done by hand
 - I/O Speed
 - Splash 1 VMEbus can deliver 4Mbytes/sec
 - Memory
 - Data memory in Splash 1 FPGAs not big enough for most applications

Splash 2 Goals


- More Design Issues
 - Multiboard Scalability
 - Moving data from one board to another was a problem with the I/O bottleneck
 - Data Path
 - Splash 1 had only a single data path
 - Linear route through 32 Xilinx chips
 - 32 bit path was sufficient for data, however extra tag bits would have been useful for control information to the FPGAs

Splash 2 Hardware

- **Chip Layout**
 - Consists of 17 Xilinx XC4010 FPGAs
 - Contains 3 lookup tables and 2 flip-flops
 - XC4010 has 20 x 20 array of CLBs
 - Each FPGA connected to 512KB of RAM
 - 16 of the FPGAs are connected in a linear array, the 17th chip provides multicast capabilities
 - FPGAs are interconnected with a 36-bit data path linearly and via a crossbar
- **Computation Models**
 - SIMD or broadcast of data
 - Linear (systolic) model

Splash 2 Hardware

- **System Architecture of Splash 2**
 - Futurebus+ backplane
 - Containing one Interface board
 - 13 Splash 2 Array Boards
 - A SPARCstation 2 host computer and Adapter Board



Software Implementation

- VHDL environment
 - Wanted to move from hardware engineering into more software programming
 - Explored idea of developing a new language based on a subset of C
 - Ultimately wanted a higher-level programming language that can be compiled into a form suitable for CAD tools
 - dbC language tried to accomplish this

Software Implementation

- T2
 - Testing and debugging was performed with T2
 - Runtime Symbolic Debugger
 - Tcl based user interface allowed for simple scripting programs to help debug applications
 - Runtime library
 - Contains built-in procedures allowing for access to Splash 2 hardware resources and runtime software

Data Parallel Programming Model

- Parallel Processing
 - Alternative view for programming on Splash 2
 - Desire to write code in a variant of C
 - Supports bit-oriented data types
 - Massive parallelism of SIMD computation
- Data Parallel C (dbC)
 - Only suitable on Splash 2 for SIMD applications
 - Performance of programs written in dbC was not good enough to be considered a “success” on important applications
 - dbC is a research project in and of itself
 - Described more as a future research project
 - Authors did not adopt dbC as the programming model for Splash 2

dbC Example

```
#include <interproc.h>
typedef poly unsigned Boolean:1;
Boolean a;
#define N 128
#define NPROC 64
unsigned DBC_net = 1; poly unsigned int R:16 = 0;
unsigned DBC_net_shape [1] = (NPROC); int right[1] = {1};
void main(){
  all {
    int b;
    a=0;
    for(b=0;b < N; b++) {
      DBC_write_to_proc(&a, 1, 0);
      R += (a ^ (Boolean) b);
      DBC_net_send(&a, a, right);
      printf("%d \n", DBC_read_from_proc(R, (b%NPROC)));
    } } }
```

-Program computes the cross-correlation of two bit streams

-Compares two bit streams and accumulates a count of the number of bits that have the same value

Applications

- Searching Genetic Databases
- Text Searching
- Fingerprint Matching
- High-Speed Image Processing

Searching Genetic Databases on Splash 2

- A need exists to classify and compare sequences of DNA
 - Original work from Princeton Nucleic Acid Comparator (P-NAC)
 - Lipton and Lopresti developed a special-purpose VLSI chip
- Ported to Splash 1 hardware
- This design was ported to Splash 2 hardware
- A new design was also implemented on Splash 2

Benchmarks

Hardware	Specifics	CUPS – cell updates per second
Splash 2	unidir; 16 boards	43,000M
Splash 2	bidir; 16 boards	34,000M
Splash 2	unidir; 1 board	3,000M
Splash 2	bidir; 1 board	2,100M
Splash 1	bidir; 746 PEs	370M
CM-2	64K nodes	150M
CM-5	32 nodes	33M
SPARC 10	gcc -O2	1.2M
P-NAC		1.1M
SPARC 1	gcc -O2	0.87M
486DX-50	DOS; gcc -O2	0.67M

Text Searching

- Developed a text searching application which tests a stream of words for inclusion and/or exclusion in a predefined list of words, or dictionary
- Pipelined operation with 3 stages
 - Split a word into characters
 - Performed at the interface board
 - Send characters through all the FPGAs
 - Look at each character determine if it is alphabetic character
 - Communicate with other FPGAs to find end of word and determine if word is in the dictionary
 - Count the number of words processed and determine when to write out the value

Fingerprint Matching

- Manual matching of fingerprints is a tedious operation
- Fingerprint is characterized by ridges and valleys, 18 different types of fingerprint patterns have been identified by the FBI
- Problem turns into a parallel pattern matching problem that Splash 2 is able to perform better than a general purpose computer
- SPARC 10 performed 70 matches per second
- Splash 2 performed $2.6 * 10^5$ matches per second

High-Speed Image Processing

- Splash 2 architecture also allows for image processing to be performed at speeds greater than conventional processors
- Several Case Studies Performed
 - Median Filtering
 - Image Pyramid Generation
- Performance
 - Claims made that Splash 2 runs 10 to 100 times faster than the same application written in C on a SPARC 10

Promises and Problems

- Splash 2 hardware Issues
 - High Bandwidth I/O is a Must
 - Memory is a Must
 - Large amounts of memory as close as possible to the FPGAs
 - The Programming Environment is Crucial
 - Users expect an interactive environment
 - A simple debugger
 - Intuitive Simulation tools
- Splash 3?