

High Performance Image Processing on FPGAs *

J. Hammes, A.P.W. Böhm, C. Ross, M. Chawathe, B. Draper, W. Najjar
Computer Science Department
Colorado State University
Ft. Collins, CO, U.S.A.

Abstract

This paper presents the high level, single-assignment programming language SA-C and its optimizing compiler targeting reconfigurable systems. SA-C is intended for Image Processing applications. Language features are introduced and discussed. The intermediate forms DDCF, DFG and AHA, used in the optimization and code-generation phases are described. Conventional and reconfigurable system specific optimizations are introduced. The code generation process is described. The performance for these systems is analyzed, using a range of applications.

1 Introduction

The biggest obstacle to the more widespread use of reconfigurable computing systems lies in the difficulty of developing application programs for them. FPGAs are typically programmed using hardware description languages such as VHDL [23]. Application programmers are typically not trained in these hardware description languages and usually prefer a higher level, algorithmic programming language to express their applications.

Turning a computation into a circuit, rather than into a sequence of CPU instructions, may seem to offer obvious performance benefits, but an effective code generation strategy requires an understanding of the fundamental differences between conventional CPUs and FPGAs.

The clock frequency of a Von Neumann machine is typically fixed. The number of instructions in a loop body may vary, and there is usually no ceiling to the number of instructions in a code block. Average memory latency is minimized by caching. Therefore it is important to minimize the number of instructions and cache misses.

On the other hand, an FPGA's frequency is variable, and is determined by the propagation delay of the circuit's critical path. In each clock cycle the entire circuit executes, and an entire code block can be viewed as one large, specialized instruction [7]. However, the available space in the FPGA imposes a ceiling on the size of a code block. Apart from this ceiling, the number of operations executing in one clock cycle does not directly affect the speed of the computation. FPGAs do not have built in cache mechanisms. They are usually connected to a number of local memories so as to maximize throughput. As long as the code block fits on the chip, the governing factors for performance are memory traffic and FPGA clock frequency. Since code frequency is fairly easily increased by the introduction of pipeline registers, an effective optimization strategy is directed almost entirely towards minimization of memory traffic.

The Cameron Project [13] has created a high-level algorithmic language, named SA-C [12], for expressing image processing applications and compiling them to FPGAs. The SA-C compiler provides one-step compilation to host executable and FPGA configurations. After parsing and type checking, the SA-C compiler converts the program to a hierarchical data dependence and control flow (DDCF) graph representation. DDCF graphs are used in many optimizations, some general and

*This work is supported by DARPA under US Air Force Research Laboratory contract F33615-98-C-1319.

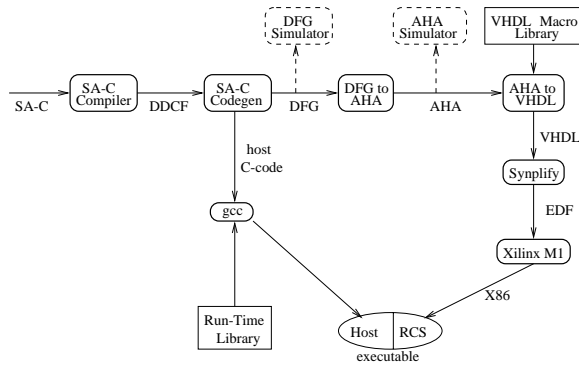


Figure 1. SA-C Compilation system.

some specific to SA-C and its target platform.

Optimizations emphasize a simple goal: once a data item has been moved into the FPGA, as much computation as possible should be done with it. Moreover, if a data item that is produced can be consumed immediately, its memory storage can be avoided completely. Thus producer-consumer loops should be fused wherever possible, and other optimizations should be designed to offset some of the problems that are introduced by fusion.

After optimization, the program is converted to a combination of low-level dataflow graphs (DFGs) and host code. DFGs are then compiled to VHDL code via Abstract Hardware Architecture (AHA) graphs. The VHDL code is synthesized and place-and-routed to FPGAs by commercial software tools. Figure 1 shows a high-level view of the system. For initial debugging the complete SA-C program can be executed on the host. All intermediate graph forms can be viewed, and DFG and AHA can be simulated. The SA-C compiler can run in stand-alone mode, but it also has been integrated into the Khoros^(TM) [16] graphical software development environment.

The rest of this paper is organized as follows. An overview of the SA-C language is presented in section 2. Compiler optimizations and pragmas are discussed in section 3. Translations to low-level dataflow graphs and then to VHDL via AHA are discussed in section 4. Applications and their performance data are presented in section 5. References to related work are given in section 6, and section 7 concludes and describes future work.

2 The SA-C Language

The design goals of SA-C are to have a language that can express image processing (IP) applications elegantly, and to allow seamless compilation to reconfigurable hardware. IP applications are supported by data parallel loops with structured access to rectangular multidimensional arrays. Reconfigurable computing requires fine grain expression level parallelism, which is easily extracted from a SA-C program because of its Single Assignment semantics. Variables in SA-C are associated with wires, not with memory locations.

Data types in SA-C include signed and unsigned integers and fixed point numbers, with user-specified bit widths. The extents of SA-C arrays can be determined either dynamically or statically. The type declaration `int14 M[:,6]` for example, is a declaration of a matrix M of 14-bit signed integers. The left dimension will be determined dynamically; the right dimension has been specified.

The most important aspect of SA-C is its treatment of **for** loops and their close interaction with arrays. SA-C is expression oriented, so every construct (including a loop) returns one or more values. A loop has three parts: one or more generators, a loop body and one or more return values. The generators provide parallel array access operators that are concise and easy for the compiler to analyze. There are four kinds of loop generators: *scalar*, *array-element*, *array-slice* and *window*. The scalar generator produces a linear sequence of scalar values, similar to Fortran's **do** loop. The array-element generator extracts scalar values from a source array, one per iteration. The array-slice generator extracts lower dimensional sub-arrays (e.g. vectors out of a matrix). Finally, window generators allow rectangular sub-arrays to be extracted from a source array. All possible sub-arrays of the specified size are produced, one per iteration. The **dot** product operator combines generators and runs them in lock step. A loop can return arrays and reductions built from values that are produced in the loop iterations, such as **sum**, **product**, **min**, and **max**.

Figure 2 shows SA-C code for the Prewitt edge detector [24], a standard IP operator. The outer for loop is driven by the extraction of 3x3 sub-

```

int16[:,:] main (uint8 Image[:,:]) {
  int16 H[3,3] = {{-1,-1,-1},
                 { 0, 0, 0},
                 { 1, 1, 1}};
  int16 V[3,3] = {{-1, 0, 1},
                 {-1, 0, 1},
                 {-1, 0, 1}};

  int16 M[:,:] =
  for window W[3,3] in Image {
    int16 dfdy, int16 dfdx =
    for h in H dot w in W dot v in V
    return(sum(h*w), sum(v*w));
    int16 magnitude =
    sqrt(dfdy*dfdy+dfdx*dfdx);
  }return(array (magnitude));
}return(M);

```

Figure 2. Prewitt Edge detector code.

arrays from array `Image`. Its loop body applies two masks to the extracted window `W`, producing a magnitude. An array of these magnitudes is collected and returned as the program's result. The shape of the return array is derived from the shape of `Image` and the loop's generator. If `Image` were a 100x200 array, the result array `M` would have a shape of 98x198.

Loop carried values are allowed in SA-C using the keyword `next` instead of a type specifier in a loop body. This indicates that an initial value is available outside the loop, and that each iteration can use the current value to compute a next value.

3 Optimizations and pragmas

The compiler's internal program representation is a hierarchical graph form called the "Data Dependence and Control Flow" (DDCF) graph. DDCF subgraphs correspond to source language constructs. Edges in the DDCF express data dependencies, opening up a wide range of loop- and array-related optimization opportunities.

Figure 3 shows the initial DDCF graph of the Prewitt program of Figure 2. The `forall` and `dot` nodes are compound, containing subgraphs. Black rectangles along the top and bottom of a compound node represent input ports and output ports. The outer `forall` has a single window generator operating on a two-dimensional image, so it requires window size and step inputs for each of the two

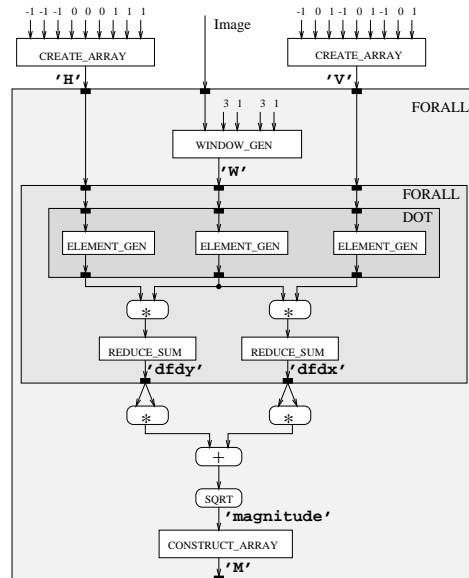


Figure 3. DDCF graph for Prewitt program.

dimensions. In this example, both dimensions are size three, with unit step sizes. The output of the `window_gen` node is a 3x3 array that is passed into the inner `forall` loop. This loop has a `dot` graph that runs three generators in parallel, each producing a stream of nine values from its source array. Each `reduce_sum` node sums a stream of values to a single value. Finally, the `construct_array` node at the bottom of the outer loop takes a stream of values and builds an array with them.

Many IP operators involve fixed size and often constant convolution masks. A *Size Inference* pass propagates information about constant size loops and arrays through the dependence graph. Effective size inference allows other optimizations, such as Full Loop Unrolling and Array Elimination, to take place. *Full Unrolling of loops* with small, compile time assessable numbers of iterations can be important when generating code for FPGAs, because it spreads the iterations in code space rather than in time. Small loops occur frequently as inner loops in IP codes, for example in convolutions with fixed size masks. *Array Value Propagation* searches for array references with constant indices, and replaces such references with the values of the array elements. When the value is a compile time

constant, this enables constant propagation. In the Prewitt example, this optimization removes the arrays *H* and *V* entirely.

Loop Carried Array Elimination represents arrays in loop bodies by a fixed set of registers. The important case is that of a loop carried array that changes values but not size during each iteration. To allocate a fixed number of registers for these arrays their size must not change, and the compiler must be able to infer the size. *N-dimensional Strip-mining* extends stripmining [33] and creates an intermediate loop with fixed bounds. The inner loop can be fully unrolled with respect to the newly created intermediate loop, generating a larger more parallel circuit. The compiler generates code to compute left over fringes.

The performance of many systems is limited by the time required to move data to the processing units. *Fusion* of producer-consumer loops is often helpful, since it reduces data traffic and may eliminate intermediate data structures. In simple cases, where arrays are processed element-by-element, this is straightforward. However, the windowing behavior of many IP operators presents a challenge. Consider the following loop pair:

```
uint8 R0[:,:] =
  for window W[2,2] in Image
    return (array (f(W)));
uint8 R1[:,:] =
  for window W[2,2] in R0
    return (array (g(W)));
```

If *Image* is e.g. a 10x10 array, *R0* is 9x9, so the two loops do not have the same number of iterations. Nevertheless, it is possible to fuse such a loop pair by examining their data dependencies. One element of *R1* depends on a 2x2 sub-array of *R0*, and the four values in that sub-array together depend on a 3x3 sub-array of *Image*. It is possible to replace the loop pair with one new loop that uses a 3x3 window and has a loop body that computes one element of *R1* from nine elements of *Image*.

Common Subexpression Elimination (CSE) is an old and well known compiler optimization that eliminates redundancies by looking for identical subexpressions that compute the same value. This could be called “spatial CSE” since it looks for common subexpressions within a block of code. The SA-C compiler performs conventional spatial

CSE, but it also performs *Temporal CSE*, looking for values computed in one loop iteration that were already computed in previous loop iterations. In such cases, the redundant computation can be eliminated by holding such values in registers so that they are available later and need not be recomputed. Here is a simple example containing a temporal common subexpression:

```
for window W[3,2] in A {
  uint8 s0 = array_sum (W[:,0]);
  uint8 s1 = array_sum (W[:,1]);
} return (array (s0+s1));
```

This code computes a separate sum of each of the two columns of the window, then adds the two. Notice that after the first iteration of the loop, the window slides to the right one step, and the column sum *s1* in the first iteration will be the same as the column sum *s0* in the next iteration. By saving *s1* in a register, the compiler can eliminate one of the two column sums, nearly halving the space required for the loop body. A useful phenomenon often occurs with Temporal CSE: one or more columns in the left part of the window are unreferenced, making it possible to eliminate those columns. *Narrowing* the window lessens the FPGA space required to store the window’s values.

In many cases the performance tradeoffs of various optimizations are not obvious; sometimes they can only be assessed empirically. The SA-C compiler allows many of its optimizations to be controlled by *user pragmas* in the source code. This allows the user to experiment with different approaches and evaluate the space-time tradeoffs.

4 Lower Level Code Generation

A dataflow graph (DFG) is a low-level, non-hierarchical and asynchronous program representation. DFGs can be viewed as abstract hardware circuit diagrams without timing or resource contention taken into account. Nodes are operators and edges are data paths. DFGs have token driven semantics. The SA-C compiler attempts to translate every innermost loop to a DFG. The innermost loops the compiler finds may not be the innermost loops of the original program, as loops may have been fully unrolled or stripmined. In the present

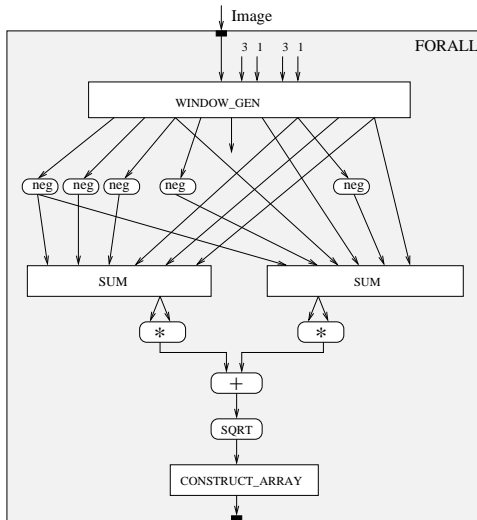


Figure 4. DFG for Prewitt after optimizations.

system, not all loops can be translated to DFGs. The most important limitation is the requirement that the sizes of a loop's window generators be statically known.

In the Prewitt program shown earlier, the DDCF graph is transformed to the DFG shown in Figure 4. The `sum` nodes can be implemented in a variety of ways, including a tree of simple additions. The window generator also allows a variety of implementations, based on the use of shift registers. The `construct_array` node streams its values out to a local memory.

DFGs are translated into a lower level form called Abstract Hardware Architecture (AHA). This is also a graph form, but with nodes that are more fine-grained than DFG nodes and that can be translated to simple VHDL components. AHA graphs have *clocked*, *semi-clocked* and *non-clocked* nodes. The clocked and semi-clocked nodes have internal state but only the clocked nodes participate in the *handshaking* needed to synchronize computations and memory accesses. Some clocked nodes communicate via an arbitrator with a local memory. An AHA graph is organized as a sequence of *sections*, each with a top and a bottom boundary. A section boundary consists of clocked nodes, whereas its internal nodes are non-

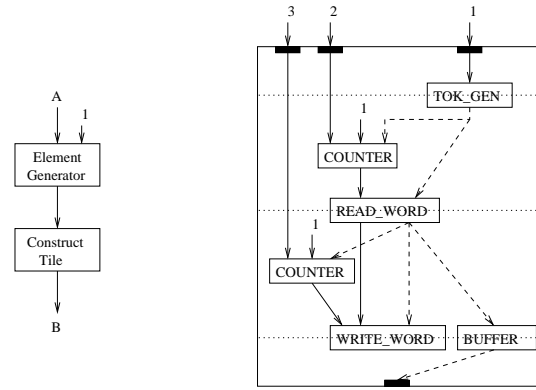


Figure 5. One-dimensional array copy in dataflow and AHA.

clocked or semi-clocked. In the AHA model, a section fires when all clocked nodes at its top boundary can produce new values and all clocked nodes at its bottom boundary can consume new values. This contrasts with DFGs, where each node independently determines when it can fire.

The AHA graph of the Prewitt code is too large and complex to display in this paper (1568 nodes). Figure 5 shows a dataflow graph on the left and an AHA graph on the right of a much simpler code fragment, which copies a one dimensional array `A` to another array `B`:

```
uint32 B[:] = for a in A return(array(a))
```

The DFG consists of an `Element Generator` node extracting elements out of `A`, and a `Construct Tile` node collecting the elements of `B`. The top three nodes in the AHA graph implement the `Element Generator` node, whereas the bottom three AHA nodes implement the `Construct Tile` node. The inputs 1, 2 and 3 in the AHA graph represent the extents of `A`, the start address of `A`, and the start address of `B`, respectively. Dotted horizontal lines represent section boundaries. Upon an input `n`, a `tok_gen` node produces `n+1` control signals (dashed edges): `n` 0-s and a 1. A `counter` is a semi-clocked node, starting at its left input and incrementing with its middle input (1 here). `buffer`, `read_word` and `write_word` speak for themselves.

A *Bitwidth Narrowing* phase is performed just before AHA graphs are generated. *Dead code*

elimination and *graph simplification* sweeps are applied on the AHA graph. A *Pipelining* phase uses node propagation delay estimates to compute the delay for each AHA section, and adds layers of pipeline registers in sections that have large propagation delays. The maximum number of pipeline stages added to the AHA graph is user controlled. Reducing propagation delays is important because it increases clock frequency.

AHA graph simulation allows the user (or, more likely, the compiler or system developer) to verify program behavior. The AHA simulator strictly mimics the hardware behavior with respect to clock cycles and signals traveling over wires. This removes the need for time consuming VHDL simulation and hardware level debugging. AHA-to-VHDL translation is straightforward; AHA nodes translate to VHDL components, which are connected according to the AHA edges.

5 Applications

The current test platform in Cameron is the StarFire Board, produced by Annapolis Micro Systems [2]. The StarFire has a single XCV1000 Virtex FPGA made by Xilinx [34]. The board contains six local memories of one megabyte each. Each of the memories is addressable in 32 bit words; all six can be used simultaneously if needed. The StarFire board is capable of operating at frequencies from 25 MHz to 180 MHz. It communicates over the PCI bus with the host computer at 66 MHz. In our system, the board is housed in a 266-MHz Pentium-based PC. This section reports the performance of SA-C codes running on the StarFire system.

5.1 Intel Image Processing Library

When comparing simple IP operators one might write corresponding SA-C and C codes and compare them on the Starfire and Pentium II. However, neither the Microsoft nor the Gnu C++ compilers exploit the Pentium's MMX technology. Instead, we compare SA-C codes to corresponding operators from the Intel Image Processing Library (IPL). The Intel IPL library consists of a large number of low-level Image Processing operations. Many

of these are simple point- (pixel-) wise operations such as square, add, etc. These operations have been coded by Intel for highly efficient MMX execution. Comparing these on a 450 Mhz Pentium II (with MMX) to SA-C on the StarFire, the StarFire is 1.2 to six times slower than the Pentium. This result is not surprising. Although the FPGA has the ability to exploit fine grain parallelism, it operates at a much slower clock rate than the Pentium. These simple programs are all I/O bound, and the slower clock of the FPGA is a major limitation when it comes to fetching data from memory. However, the Prewitt edge detector written in C using IPL calls and running on the 450 Mhz Pentium II, takes 53 milliseconds as compared to 17 milliseconds when running the equivalent SA-C code on the StarFire. Thus, non I/O bound SA-C programs running on the StarFire board are competitive with their hand optimized IPL counterparts.

5.2 ARAGTAP

SA-C running on the StarFire compares even better when we look at more complex operations. The ARAGTAP pre-screener [26] was developed by the U.S. Air Force at Wright Labs as the initial focus-of-attention mechanism for a SAR automatic target recognition application. Aragtap's components include down sampling, dilation, erosion, positive differencing, majority thresholding, bitwise-and, percentile thresholding, labeling, label pruning, and image creation. All these components, apart from label pruning and image creation, have been written in SA-C. Most of the computation time is spent in a sequence of eight gray-scale morphological dilations, and a later sequence of eight gray-scale erosions. Four of these dilations are with a 3x3 mask of 1's (the shape of a square), the other four are with a 3x3 mask with 1's in the shape of a cross and zeros at the edges.

The dilate and erode loops allow temporal CSE and window narrowing. A straight C implementation of ARAGTAP running on the Pentium takes 1.07 seconds. Running down sampling, eight unfused dilations, and eight unfused erosions, positive differencing, majority thresholding, and bitwise And on the StarFire board and the rest of the code on the PC takes 0.096 seconds, a more than

ten fold speedup over the Pentium. Moreover, the compiler can fuse the dilate and erode loop sequences, each consisting of eight loops, bringing the execution time down to 0.041 seconds: a twenty-six fold speedup over the Pentium.

5.3 Prewitt, Canny, and Wavelet

We have implemented some canonical image processing functions in SA-C. The simplest such function is an edge detection program that calculates the square root of the sum of the squares of responses to horizontal and vertical Prewitt edge masks. Since this same task can be performed using the Intel Image Processing Library (IPL), we are able to compare our results to the results of a hand-optimized Pentium program. The results are discussed in section 5.1.

A more sophisticated edge detector is the Canny operator, which comprises a four step process of image smoothing, computing edge magnitudes and (discretized) orientations, non-maximal suppression in the gradient direction, and hysteresis labeling, which is a connected components algorithm. For a clear and simple explanation of Canny and the reasoning behind it, see [31], Chapter 4. The first three steps of Canny were implemented in SA-C and run on the reconfigurable hardware. Although connected components can be written in SA-C, it can currently not be compiled to FPGAs. Therefore, we assume that the last step will be performed on the host. The compiler performed eight fold vertical stripmining, among other optimizations. The test inputs were eight bit 512x512 images. The SA-C execution time on the FPGA is 26 milliseconds.

Comparing the performance to a Pentium is hard. A version of the same program was written in C, using Intel's IPL whenever possible. The resulting program took 850 milliseconds on an 800MHz Pentium 3. On the other hand, Intel's OpenCV library has a hand-optimized assembly-coded version of Canny that includes the fourth (connected components) step. By setting the high and low thresholds to be the same (so that connected components would not iterate) and timed the OpenCV routine at 135 milliseconds. So SA-C was five times faster than assembly code and thirty times

faster than a C plus IPL implementation.

Wavelets are commonly used for multi-scale analysis in computer vision, as well as for image compression. Honeywell has defined a set of benchmarks for reconfigurable computing systems, including a wavelet-based image compression algorithm [17]. This code takes one image and returns four quarter sized images, which are higher order derivatives of the original. On a 512x512 eight bit input image, the SA-C code takes five milliseconds execution time on the FPGA, whereas on an 800Mhz Pentium the Honeywell C code took 75 milliseconds. So SA-C was fifteen times faster.

6 Related work

Hardware and software research in reconfigurable computing is active and ongoing. Hardware projects fall into two categories – those using commercial off-the-shelf components (e.g. FPGAs), and those using custom designs.

The Splash-2 [6] is an early (circa 1991) implementation of an RCS, built from 17 Xilinx [35] 4010 FPGAs, and connected to a Sun host as a co-processor. Several different types of applications have been implemented on the Splash-2, including searching[14, 25], pattern matching[28], convolution [27] and image processing [3].

Representing the current state of the art in FPGA-based RCS systems are the AMS WildStar[2] and the SLAAC project [29]. Both utilize Xilinx Virtex [34] FPGAs, which offer over an order of magnitude more programmable logic, and provide a several-fold improvement in clock speed, compared to the earlier chips.

Several projects are developing custom hardware. The Morphosys project [19] combines an on-chip RISC processor with an array of reconfigurable cells (RCs). Each RC contains an ALU, shifter, and a small register file.

The RAW Project [32] also uses an array of computing cells, called *tiles*; each tile is itself a complete processor, coupled with an intelligent network controller and a section of FPGA-like configurable logic that is part of the processor data path. The PipeRench [10] architecture consists of a series of *stripes*, each of which is a pipeline stage with an input interconnection network, a lookup-

table based PE, a results register, and an output network. The system allows a virtual pipeline of any size to be mapped onto the finite physical pipeline.

On the software front, a framework called “Nimble” compiles C codes to reconfigurable targets where the reconfigurable logic is closely coupled to an embedded CPU [18]. Several other hardware projects also involve software development. The RAW project includes a significant compiler effort [1] whose goal is to create a C compiler to target the architecture. For PipeRench, a low-level language called DIL [9] has been developed for expressing an application as a series of pipeline stages mapped to stripes.

Several projects (including Cameron) focus on hardware-independent software for reconfigurable computing; the goal – still quite distant – is to make development of RCS applications as easy as for conventional processors, using commonly known languages or application environments. Several projects use C as a starting point. DEFACTO [11] uses SUIF as a front end to compile C to FPGA-based hardware. Handel-C [21] both extends the C language to express important hardware functionality, such as bit-widths, explicit timing parameters, and parallelism, and limits the language to exclude C features that do not lend themselves to hardware translation, such as random pointers. Streams-C [8] does a similar thing, with particular emphasis on extensions to facilitate the expression of communication between parallel processes. SystemC [30] and Ocapi [15] provide C++ class libraries to add the functionality required of RCS programming to an existing language.

Finally, a couple of projects use higher-level application environments as input. The MATCH project [4, 5, 22] uses MATLAB as its input language – applications that have already been written for MATLAB can be compiled and committed to hardware, eliminating the need for re-coding them in another language. Similarly, CHAMPION [20] is using Khoros [16] for its input – common glyphs have been written in VHDL, so GUI-based applications can be created in Khoros and mapped to hardware.

7 Conclusions and Future Work

The Cameron Project has created a language, called SA-C, for one-step compilation of image processing applications that target FPGAs. Various optimizations, both conventional and novel, have been implemented in the SA-C compiler.

The system has been used to implement routines from the Intel IPL, as well as more comprehensive applications, such as the ARAGTAP target acquisition prescreener. Compared to Pentium II/MMX technology built at roughly the same time, the SA-C system running on an Annapolis StarFire board with one Virtex 1000 FPGA has similar performance when it comes to small IPL type operations, but shows speedups up to 25 when it comes to more complex operators such as Canny, Wavelet, and Dilate and Erode sequences. Performance evaluation of the SA-C system has just begun. As performance issues become clearer, the system will be given greater ability to evaluate various metrics including code space, memory use and time performance, and to evaluate the tradeoffs between conventional functional code and lookup tables.

Currently, the VHDL generated from the AHA graphs is behavioral and ignores the structural information available in the AHA graph. We will soon be investigating the use of Relatively Placed Macros (RPM) as a method to make some of the structural information explicit to the synthesis tools. Providing constraints to specify the placement of nodes relative to each other may prove to decrease synthesis and place and route time.

Also, stream data structures are being added to the SA-C language, which will allow multiple cooperating processes to be mapped onto FPGAs. This allows expression of streaming video applications, and partitioning of a program over multiple chips.

References

- [1] A. Agarwal, S. Amarasinghe, R. Barua, M. Frank, W. Lee, V. Sarkar, and M. Srikrishna, D. and Taylor. The RAW compiler project. In *Proc. Second SUIF Compiler Workshop*, August 1997.
- [2] Annapolis Micro Systems, Inc., Annapolis, MD. *STARFIRE Reference Manual*, 1999. www.annapmicro.com.

- [3] P. M. Athanas and A. L. Abbott. Processing images in real time on a custom computing platform. In R. W. Hartenstein and M. Z. Servit, editors, *Field-Programmable Logic Architectures, Synthesis, and Applications*, pages 156–167. Springer-Verlag, Berlin, 1994.
- [4] P. Banerjee et al. A MATLAB compiler for distributed, heterogeneous, reconfigurable computing systems. In *The 8th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, April 2000.
- [5] P. Banerjee, N. Shenoy, A. Choudhary, S. Hauck, C. Bachmann, M. Chang, M. Haldar, P. Joisha, A. Jones, A. Kanhare, A. Nayak, S. Periyacheri, and M. Walkden. MATCH: a MATLAB compiler for configurable computing systems. Technical Report CPDC-TR-9908-013, Center for Parallel and distributed Computing, Northwestern University, August 1999.
- [6] D. Buell, J. Arnold, and W. Kleinfelder. *Splash 2: FPGAs in a Custom Computing Machine*. IEEE CS Press, 1996.
- [7] A. DeHon and J. Wawrzynek. Reconfigurable computing: What, why, and implications for design automation. In *Proceedings of Design Automation Conference (DAC '99)*, pages 610–615, June 1999.
- [8] M. Gokhale et al. Stream oriented FPGA computing in Streams-C. In *The 8th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, April 2000.
- [9] S. C. Goldstein and M. Budiu. *The DIL Language and Compiler Manual*. Carnegie Mellon University, 1999.
- [10] S. C. Goldstein, H. Schmit, M. Moe, M. Budiu, S. Cadambi, R. R. Taylor, and R. Laufer. Pipherench: A coprocessor for streaming multimedia acceleration. In *Proc. Intl. Symp. on Computer Architecture (ISCA '99)*.
- [11] M. Hall, P. Diniz, K. Bondalapati, H. Ziegler, P. Duncan, R. Jain, and J. Granacki. DEFACTO: A design environment for adaptive computing technology. In *Proc. 6th Reconfigurable Architectures Workshop (RAW'99)*. Springer-Verlag, 1999.
- [12] J. Hammes and W. Böhm. *The SAC Language - Version 1.0*, 1999. www.cs.colostate.edu/cameron.
- [13] J. Hammes, R. Rinker, W. Böhm, and W. Najjar. Cameron: High level language compilation for reconfigurable systems. In *PACT'99*, Oct. 1999.
- [14] D. Hoang. Searching genetic databases on Splash 2. In *IEEE Workshop on FPGAs for Custom Computing Machines*, pages 185–192. CS Press, Loa Alamos, CA, 1993.
- [15] IMEC. Ocapi overview. www.imec.be/ocapi/.
- [16] K. Konstantinides and J. Rasure. The Khoros software development environment for image and signal processing. In *IEEE Transactions on Image Processing*, volume 3, pages 243–252, May 1994.
- [17] S. Kumar. A benchmark suite for evaluating configurable computing systems - status, reflections, and future directions. In *FPGA2000 Eighth International Symposium on FPGAs*, 2000. Feb. 10–12, Monterey, CA.
- [18] Y. Li, T. Callahan, E. Darnell, R. Harr, U. Kurkure, and J. Stockwood. Hardware-software co-design of embedded reconfigurable architectures. In *Proc. 37th Design Automation Conference*, 1999.
- [19] G. Lu, H. Singh, M. Lee, N. Bagherzadeh, and F. Kurhadi. The Morphosis parallel reconfigurable system. In *Proc. of EuroPar 99*, Sept. 1999.
- [20] S. Natarajan, B. Levine, C. Tan, D. Newport, and D. Bouldin. Automatic mapping of Khoros-based applications to adaptive computing systems. Technical report, University of Tennessee, 1999.
- [21] The Handel Language. Technical report, Oxford University, 1997.
- [22] S. Periyayacheri, A. Nayak, A. Jones, N. Shenoy, A. Choudhary, and P. Banerjee. Library functions in reconfigurable hardware for matrix and signal processing operations in MATLAB. In *Proc. 11th IASTED Parallel and Distributed Computing and Systems Conf. (PDCS'99)*, November 1999.
- [23] D. Perry. *VHDL*. McGraw-Hill, 1993.
- [24] J. M. S. Prewitt. Object enhancement and extraction. In B. S. Lipkin and A. Rosenfeld, editors, *Picture Processing and Psychopictorics*. Academic Press, New York, 1970.
- [25] D. V. Pryor, M. R. Thistle, and N. Shirazi. Text searching on Splash 2. In *IEEE Workshop on FPGAs for Custom Computing Machines*, pages 172–178. CS Press, Loa Alamos, CA, 1993.
- [26] S. Raney, A. Nowicki, J. Record, and M. Justice. Aragtap atr system overview. In *Theater Missile Defense 1993 National Fire Control Symposium*, Boulder, CO, 1993.
- [27] N. K. Ratha, D. T. Jain, and D. T. Rover. Convolution on Splash 2. In *Proc. IEEE Symp. on FPGAs for Custom Computing Machines*, pages 204–213. CS Press, Loa Alamos, CA, 1995.
- [28] N. K. Ratha, D. T. Jain, and D. T. Rover. Fingerprint matching on Splash 2. In *Splash 2: FPGAs in a Custom Computing Machine*, pages 117–140. IEEE CS Press, 1996.

- [29] B. Schott, S. Crago, C. C., J. Czarnaski, M. French, I. Hom, T. Tho, and T. Valenti. Reconfigurable architectures for systems level applications of adaptive computing. 1997. Available from <http://www.east.isi.edu/SLAAC/>.
- [30] SystemC. SystemC homepage. www.systemc.org/.
- [31] E. Trucco and A. Verri. *Introductory Techniques for 3-D Computer Vision*. Prentice Hall, 1998.
- [32] E. Waingold, M. Taylor, D. Srikrishna, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, P. Finch, R. Barua, J. Babb, S. Amarasinghe, and A. Agarwal. Baring it all to software: RAW machines. *Computer*, September 1997.
- [33] M. Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley Publishing Company, 1996.
- [34] Xilinx, Inc. *Virtex 2.5V Field programmable Gate Arrays: Preliminary Product Description*, Oct. 1999. www.xilinx.com.
- [35] Xilinx, Inc., San Jose, CA. *The Programmable Logic Databook*, 1998. www.xilinx.com.