

One-step Compilation of Image Processing Applications to FPGAs *

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

Abstract *This paper describes a system for one-step compilation of image processing (IP) codes, written in the machine-independent, algorithmic, high-level single assignment language SA-C, to FPGA-based hardware. The SA-C compiler performs a variety of optimizations, some conventional and some specialized, before generating dataflow graphs and host code. The dataflow graphs are then compiled, via VHDL, to FPGA configuration codes. This paper introduces SA-C and describes the optimization and code generation stages in the compiler. The performance of a target acquisition prescreeener (ARAG-TAP) and the Intel Image Processing Library running on the reconfigurable system are compared to their performance on a PC.*

1 Introduction

Recently, the computer vision and image processing communities have become aware of the potential for masasive parallelism and high computational density in FPGAs. FPGAs have been used for, e.g., real-time point tracking [?], stereo vision [?], color-based detection [?] and image compression [?]. Unfortunately, the biggest obstacle to the more widespread use of reconfigurable computing systems lies in the difficulty of programming them. FPGAs are typically programmed in hardware description languages such as VHDL [25]. These languages require great attention to detail, including issues such as timing and low level synchronization. IP experts are usually not also hardware experts, and prefer a higher level, algorithmic programmng language to express their applications.

The Cameron Project [16, 22] has created a high-level algorithmic language, named SA-C [15], for writing image processing applications and compiling parts




Figure 1. SA-C Compilation system.

of them to FPGAs. The SA-C compiler provides one-step compilation from SA-C code to ready-to-run host executable and FPGA configurations. The compiler uses dataflow graphs to perform a variety of optimizations, some conventional and some unusual.

After parsing and type checking, the SA-C compiler converts the program to a hierarchical graph representation called DDCF (Data Dependence and Control Flow) graphs. DDCF graphs are used in many optimizations, some general and some specific to SA-C and its target platform. 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, which is synthesized and place-and-routed to FPGAs by commercial software tools.

To aid in program development, the SA-C compiler has two other modes. In the first, the entire SA-C code is compiled to a host executable for traditional debugging. In the second, the DFGs are executed by a simulator for validation: the DFG is shown on the screen and the user steps through its execution, watching the values flow through the graph. Figure 1 shows a high-level view of the system. The SA-C compiler can run in stand-alone mode, but it also has been integrated into the Khoros^(TM) [19] graphical software development environment.

The rest of this paper presents an overview of the SA-C language in section 2. Compiler optimizations and pragmas are discussed in section 3. Translations

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

to low-level dataflow graphs and then to VHDL are discussed in sections 4 and 5. The Cameron Project’s test platform is described in section 6, and some applications with performance data are presented in section 7. References to related work are given in section 8, and section 9 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 applications elegantly, and to allow seamless compilation to reconfigurable hardware. Variables in SA-C are associated with wires, not with memory locations. SA-C is a single-assignment side effect free language, each variable’s declaration occurs together with the expression defining its value. This avoids pointer and von Neumann memory model complications and allows for better compiler analysis and translation to DFGs. The IP applications that have been coded in SA-C transform images to images and are easily expressed using single assignment. Data types in SA-C include signed and unsigned integers and fixed point numbers, with user-specified bit widths. SA-C has multidimensional rectangular arrays whose extents 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 by the user.

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. Generators can be combined through **dot** and **cross** products. The **dot** product runs the generators in lock step, whereas **cross** products produce all combinations of components from the generators.

SA-C loops provide a simple and concise way of

processing arrays in regular patterns, often making it unnecessary to create loop nests to handle multi-dimensional arrays or to refer explicitly to the array’s extents or the loop’s index variables. They make compiler analysis of array access patterns significantly easier than in C or Fortran, where the compiler must analyze index expressions in loop nests and infer array access patterns from these expressions. In SA-C, the index generators and the array references have been unified; the compiler can reliably infer the patterns of array access.

A loop can return arrays and reductions built from values that are produced in the loop iterations, including **sum**, **product**, **min**, **max**, **mean**, and **median**. The **histogram** operator returns a histogram of loop body values in a one-dimensional array.

```
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.

Figure 2 shows SA-C code for the Prewitt edge detector, a standard IP operator. The outer **for** loop is driven by the extraction of 3x3 sub-arrays from array `Image`. All possible 3x3 arrays are taken, one per loop iteration. 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.

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.

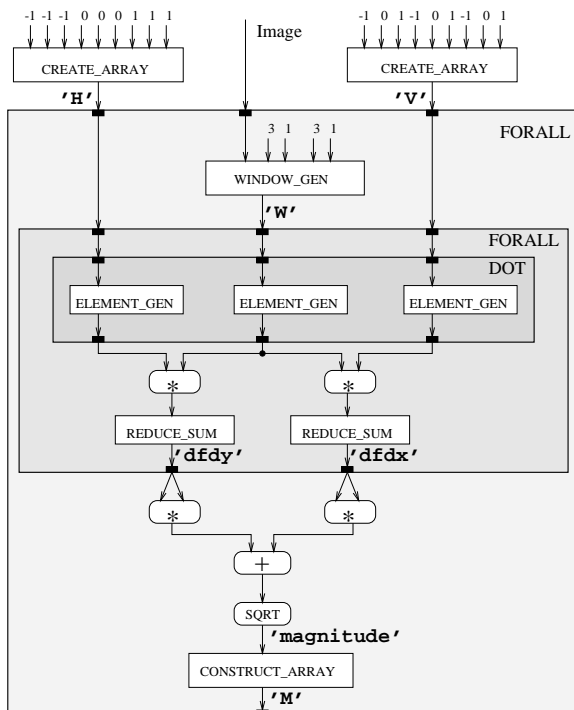


Figure 3. DDCF graph for Prewitt program.

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. The WINDOW_GEN is operating on a two-dimensional image, so it requires window size and step inputs for each of the two 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.

Dataflow analysis is performed on the DDCF graph, and a number of conventional optimizations [3] are performed as DDCF-to-DDCF transformations, including Invariant Code Motion, Function Inlining, Switch Con-

stant Elimination, Constant Propagation and Folding, Algebraic Identities, Dead Code Elimination and Common Subexpression Elimination. Another set of optimizations is more specific to the single assignment nature of SA-C and to the target hardware. Many of these are controlled by user pragmas.

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. This pass is vitally important in this compiler, since it exploits the close association between arrays and loops in the SA-C language. Source array size information can propagate through a loop to its result arrays (downward flow), and result array size information can be used to infer the sizes of source arrays (upward flow). In addition, since the language requires that the generators in a dot product have identical shapes, size information from one generator can be used to infer sizes in the others (sideways flow).

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 implementing 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. As will be shown in the Prewitt example, this optimization may remove entire user defined (or compiler generated) arrays.

N-dimensional Stripmining 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. As an example the following code shows the convolution of an image with a three by three mask:

```
uint8[:,:] Conv (uint8 I[:,:], uint8 M[3,3]) {
    uint16 Res[:,:] =
        for window W[3,3] in I {
            uint16 ip =
                for w in W dot m in M
                    return (sum (w*m));
        } return (array (ip));
    } return (Res);
```

The inner loop computing ip gets unrolled. However, this unrolled loop is still rather small and it is often more efficient to execute a number of these inner loops in parallel. If the outer loop is stripmined with a user specified size of eight by eight, the compiler transforms it to the following:

```

uint16 Res[:,:] =
  for window WT[8,8] in I step (6,6) {
    uint16 ResTile[6,6] =
      for window W[3,3] in WT {
        uint16 ip =
          for w in W dot m in M
            return (sum (w*m));
      } return (array (ip) );
    } return (tile (ResTile));

```

This is important because the two nested inner loops can now be unrolled, generating a larger more parallel circuit. The `tile` loop return operator concatenates equal size N-dimensional sub-arrays into one N-dimensional array. The compiler generates code to compute left over fringes.

Some (combinations of) operators can be inefficient to implement directly in hardware. For example the computation of `magnitude` in Prewitt requires multiplications and square root operators. The evaluation of the whole expression can be replaced by an access to a *Lookup Table*, which the compiler creates by wrapping a loop around the expression, recursively compiling and executing the loop, and reading the results.

The performance of many systems today, both conventional and specialized, is often 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 the `Image` array has extents $d_0 \times d_1$, then `R0` will have extents $(d_0 - 1) \times (d_1 - 1)$ (determined by the number of 2x2 windows that can be referenced in `Image`.) Similarly, the extents of `R1` will be $(d_0 - 2) \times (d_1 - 2)$. This means that these 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`. Sub-arrays of size 2x2 are extracted from the window, and each feeds a copy

of the upper loop body `f`. The values emerging from the copies of `f` feed the lower loop body `g`. This kind of fusion can create large loop bodies since the loop body of the upper loop is replicated in the new loop. These space problems are tackled by Temporal CSE, as described next..

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 [3]. The redundancies are removed by keeping just one of the subexpressions and using its result for all the computations that need it. 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. This is the essence of Temporal CSE performed by the SA-C compiler. This optimization requires some special handling to take care of the problem of getting the registers “primed” with initial values, a level of detail that is beyond the scope of this paper.

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. Figure 4 shows the dataflow graph that results after Temporal CSE and Window Narrowing have been applied to the column-sum example above. The register is produced by the TCSE step. Window narrowing then removes one column of the window.

Another optimization that sets the stage for window narrowing moves the computation of expressions fed by window elements into earlier iterations by mov-




Figure 4. Effect of Temporal CSE and Window Narrowing.

ing the window references rightward, and uses a register delay chain to move the result to the right iteration. This can remove references to left columns of the window, and allows window narrowing. This optimization trades window storage space for register delay chain. Whether this optimization actually provides a gain depends on the individual computation. We call this *window compaction*.

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 makes it relatively painless for a user to experiment with different approaches and evaluate the space-time tradeoffs that inevitably exist.

After multiple applications of fusion and unrolling, the resulting loop often has a long critical path, resulting in a low clock frequency. Adding stages of *pipeline registers* can break up the critical path and thereby boost the frequency. The SA-C compiler uses propagation delay estimates, empirically gathered for each of the DFG node types, to determine the proper placement of pipeline registers. The user specifies the number of pipeline stages.

4 Conversion to DFGs

A dataflow graph (DFG) is a low-level, non-hierarchical and asynchronous program representation.

DFGs are used for mapping SA-C program parts onto reconfigurable hardware. DFGs can be viewed as abstract hardware circuit diagrams without timing considerations taken into account. Nodes are operators and edges are data paths. Dataflow graphs allow token driven simulation, used by the compiler writer and applications programmer for validation and debugging.

After optimizations have been performed on a program's DDCF graph, the compiler searches for loops that meet the criteria for execution on reconfigurable hardware. 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 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.

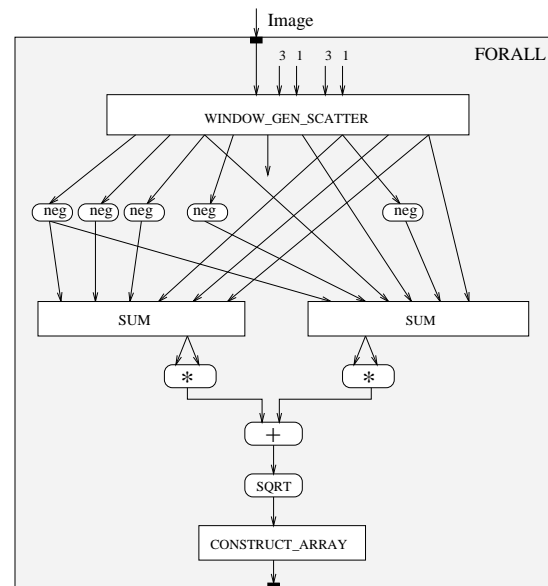


Figure 5. DFG for Prewitt after optimizations.

In the Prewitt program shown earlier, the DDCF graph is transformed to the DFG shown in figure 5. 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 or to the host, as appropriate.

5 Translation of DFGs to VHDL

The DFG to VHDL translator produces synchronous circuits, expressed in VHDL, from asynchronous DFGs [9]. Each DFG *node* corresponds to a VHDL operation; *edges* correspond to the input and output signals that the operation requires. The simple DFG nodes, such as the arithmetic and logical operators, are strictly combinational. Complex nodes like generator nodes and collector nodes are synchronous and contain registers or state machines that use a clock.

The translator handles the two types of nodes differently. The combinational nodes form the inner loop body (ILB) of the circuit, and are translated either into a single VHDL statement or into the instantiation of a pre-built VHDL component. Complex nodes are implemented by selecting the proper VHDL component from a library of pre-built modules; they are parameterized by information gleaned from the DFG. These clock driven nodes are connected together into a “wrapper” around the ILB, resulting in the structure shown in Figure 6. For each iteration of the loop, the generator produces data at the top of the ILB, and the ILB’s output is available a short time later, the exact time being determined by the circuit propagation delay. This delay determines the maximum execution frequency of the resulting circuit.

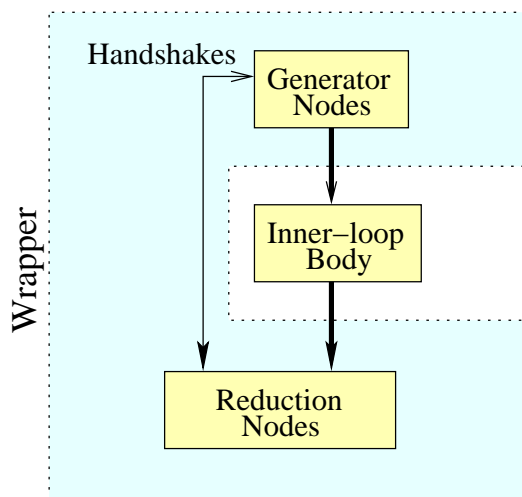


Figure 6. Abstract System Structure

A generator node retrieves the input data in blocks from the external memory in the order required by the ILB, buffers it, and presents it at the proper time to the

inputs of the ILB. A state machine controls the generation of the proper memory access and buffer signals, and provides the timing for the operation of the entire circuit. In the case of a “2D window-generator” node, a 2D buffer array is instantiated, of the same size as the window specified in the generator node. This buffer is formed as a shift register; after the current window of data has been used by the ILB, the oldest column is shifted out, and a new column is shifted in, to form the next window of data. This shift register operation produces the “sliding window” effect required by the SA-C language. Other types of generators provide 1D windows, single array element, and scalars to the ILB in a similar fashion. The generator code also handles the prefetching of data from memory, so the buffers can be kept full and the ILB can be kept as busy as possible.

In a complementary fashion, the reduction nodes take the output from the ILB, buffer this output, and write completed buffers to memory. As with the input, ILB output may be a single element, several single elements, or tiles (sub-arrays.)

The translator selects the proper VHDL components for the generator and reductions and creates VHDL glue logic that “wires” the components together. A parameter file, which specifies the sizes and shapes of the various buffers, as well as any additional timing information required to synchronize the wrapper components, is also created. Finally, the translator generates the script files needed by the commercial VHDL compiler and Place-and-Route tools to synthesize the final design.

6 System Description

As stated earlier, the SA-C system can be used in three modes:

1. Compile the entire program to a host executable. This is useful in the early stages of code development, for quick compiles and functional debugging.
2. Compile the program to dataflow graphs, called by a host executable. The system has a token-driven dataflow simulator that allows validation of the dataflow graphs produced by the compiler. There is an optional “view” mode that displays the DFGs and allows the user to single-step the execution and watch the values flowing through the graph. Figures 7 and 8 show screen captures of the simulator’s main display and its display of the source array with highlighted window.

3. Compile the program to FPGA codes, called by a host executable. The Cameron compiler produces VHDL code [10] that is compiled and place-and-routed by commercial software tools.

The SA-C run-time system has I/O formats that are compatible with standard image processing formats PGM and PPM. This means that after compilation to host and FPGA codes, the program is ready to run immediately on standard image files.

The current test platform in Cameron is the StarFire Board, produced by Annapolis Microsystems [4], which has a single XCV1000-BG560-4 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.

7 Applications

This section reports the performance of SA-C codes running on the StarFire system described earlier.

7.1 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 0.053 seconds as compared to 0.017 seconds when running the equivalent SA-C code on the StarFire. Thus SA-C code for operators

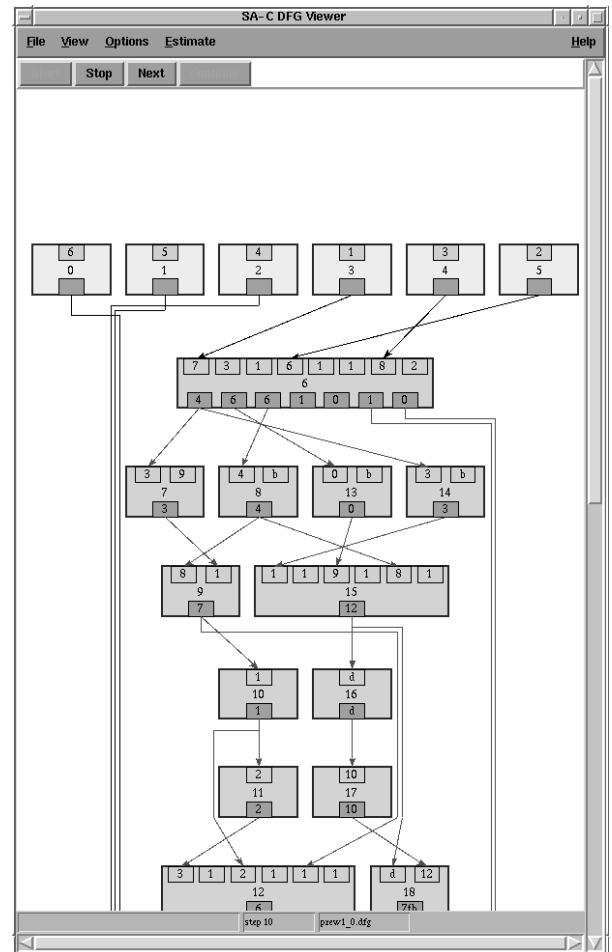


Figure 7. Screen capture of simulator graph.

Generator Node 6							
4	7	4	3	5	4	4	2
1	3	4	6	5	7	6	8
9	0	6	4	5	7	6	3
8	4	6	8	9	7	4	6
7	9	5	3	1	2	2	2
1	2	3	6	9	0	7	5

Figure 8. Screen capture of simulator source array with highlighted 3x1 window.

composed of some low level IPL routines executed on the StarFire board can compete successfully with Pentium executions.

7.2 ARAGTAP

SA-C running on the StarFire compares even better when we look at more complex operations. The ARAGTAP pre-screener 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 comprise down sampling, dilation, erosion, positive differencing, majority thresholding, bitwise And, percentile thresholding labelling, label pruning, and image creation. All these components, apart from 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.

First, the compiler can fuse dilate and erode loops in groups of four. Fusing loop sequences longer than four currently brings the frequency of the resulting FPGA configurations below 25 Mhz, the minimum frequency required by the StarFire. The dilate and erode loops also allow temporal CSE, window compaction. Combined with the fusion of four loops, this results in an FPGA code driven by a 9x1 window generator.

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. Fusing erode and dilate loops into groups of four brings the execution time down to 0.065 seconds: about 50% gain over the unfused FPGA execution time, and a sixteen fold speedup over the Pentium.

8 Related work

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

The Splash-2 [8] 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 [17, 26], pattern matching [28], convolution [27] and image processing [6].

Representing the current state of the art in FPGA-based RCS systems are the AMS WildStar [5] 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 [21] marries an on-board RISC processor with an array of reconfigurable cells (RCs). 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 [13] 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 [20]. Several other hardware projects also involve software development. The RAW project includes a significant compiler effort [2] whose goal is to create a C compiler to target the architecture. For PipeRench, a low-level language called DIL [12] 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 [14] uses SUIF as a front end to compile C to FPGA-based hardware. Handel-C [1] 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 [11] does a similar thing, with particular emphasis on extensions to facilitate the expression of communication between parallel processes. SystemC [31] and Ocap [18] 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 [7, 24] 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 [23] is using Khoros [19] for its input – common glyphs have been written in VHDL, so GUI-based applications can be created in Khoros and mapped to hardware.

9 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. These optimizations are focused on reducing execution time and/or reducing FPGA space. The system has been used in a target acquisition prescreener, as well as in implementing routines from the Intel IPL.

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.

More compiler optimizations are under way, including further manipulations of window generators. Also, stream data structures are being added to the SA-C language, which will allow multiple cooperating processes to be mapped onto FPGAs.

References

- [1] OXFORD hardware compiler group, the Handel Language. Technical report, Oxford University, 1997.
- [2] 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.
- [3] A. Aho, R. Sethi, and J. Ullman. *Compilers Principles, Techniques, and Tools*. Addison-Wesley Publishing Company, 1986.
- [4] Annapolis Micro Systems, Inc., Annapolis, MD. *WILDFORCE Reference Manual*, 1997. www.annapmicro.com.
- [5] Annapolis Micro Systems, Inc., Annapolis, MD. *STARFIRE Reference Manual*, 1999. www.annapmicro.com.
- [6] 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.
- [7] 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.
- [8] D. Buell, J. Arnold, and W. Kleinfelder. *Splash 2: FPGAs in a Custom Computing Machine*. IEEE CS Press, 1996.
- [9] M. Chawathe. Dataflow graph to VHDL translation. Master's thesis, Colorado State University, 2000.
- [10] M. Chawathe, M. Carter, C. Ross, R. Rinker, A. Patel, and W. Najjar. Dataflow graph to VHDL translation. Technical report, Colorado State University, Dept. of Computer Science, 2000.
- [11] M. Gokhale. The Streams-C Language, 1999. www.darpa.mil/ito/psum19999/F282-0.html.
- [12] S. C. Goldstein and M. Budiu. *The DIL Language and Compiler Manual*. Carnegie Mellon University, 1999. www.ece.cmu.edu/research/piperench/dil.ps.
- [13] S. C. Goldstein, H. Schmit, M. Moe, M. Budiu, S. Cadambi, R. R. Taylor, and R. Laufer. Piperench: A coprocessor for streaming multimedia acceleration. In *Proc. Intl. Symp. on Computer Architecture (ISCA '99)*, 1999. www.cs.cmu.edu/~mihaiib/research/isca99.ps.gz.
- [14] 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.
- [15] J. Hammes and W. Böhm. *The SA-C Language - Version 1.0*, 1999. Document available from www.cs.colostate.edu/cameron.
- [16] J. Hammes, R. Rinker, W. Böhm, and W. Najjar. Cameron: High level language compilation for reconfigurable systems. In *PACT'99*, Oct. 1999.
- [17] D. Hoang. Searching genetic databases on Splash 2. In *IEEE Workshop on FPGAs for Custom Computing Machines*, pages 185–192. CS Press, Los Alamitos, CA, 1993.
- [18] IMEC. Ocapi overview. www.imec.be/ocapi/.
- [19] 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.
- [20] 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.
- [21] G. Lu, H. Singh, M. Lee, N. Bagherzadeh, and F. Kurhadi. The Morphosis parallel reconfigurable system. In *Proc. of EuroPar 99*, Sept. 1999.

- [22] W. Najjar. The Cameron Project. Information about the Cameron Project, including several publications, is available at the project's web site, www.cs.colostate.edu/cameron.
- [23] 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. Available from <http://microsys6.engr.utk.edu:80/~bouldin/darpa/mapld2/mapld-paper.pdf>.
- [24] 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.
- [25] D. Perry. *VHDL*. McGraw-Hill, 1993.
- [26] 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, Los Alamitos, CA, 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, Los Alamitos, 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. Available from <http://www.east.isi.edu/SLAAC/>.
- [30] S. K. Skedzielewski and J. R. W. Glauert. *IF1, an Intermediate Form for Applicative Languages*. Lawrence Livermore National Laboratory, reference manual m-170 edition, July 1985.
- [31] SystemC. SystemC homepage. www.systemc.org/.
- [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.