

A Compiler Intermediate Representation for Reconfigurable Fabrics

Zhi Guo · Betül Buyukkurt · John Cortes ·
Abhishek Mitra · Walid Najjar

Received: 2 April 2008 / Accepted: 23 July 2008 / Published online: 19 September 2008
© Springer Science+Business Media, LLC 2008

Abstract Configurable computing relies on the expression of a computation as a circuit. Its main purpose is the hardware based acceleration of programs. Configurable computing has received renewed interest with the recent rapid increase in both size and speed of FPGAs. One of the major obstacles in the way of wider adoption of (re)configurable computing is the lack of high-level tools that support the efficient mapping of programs expressed in high-level languages (HLL) to reconfigurable fabrics. The major difficulty in such a mapping is the translation from a temporal execution model to a spatial execution model. An intermediate representation (IR) is the central structure around which tools such as compilers and synthesis tools are built. In this paper we propose an IR specifically designed for reconfigurable fabrics: CIRRF (Compiler Intermediate Representation for Reconfigurable Fabrics). We describe the design of CIRRF and its initial implementation as part of the ROCCC compiler for translating C code to VHDL. CIRRF is designed to support the creation of a datapath and the scheduling of operations on it. It provides support for buffers, look-up tables,

Z. Guo (✉)
Brocade Communications Systems, San Jose, CA, USA
e-mail: zguo@brocade.com

B. Buyukkurt · J. Cortes · A. Mitra · W. Najjar
Department of Computer Science and Engineering, University of California,
Riverside, Riverside, CA, USA
e-mail: abuyukku@cs.ucr.edu

J. Cortes
e-mail: jcortes@cs.ucr.edu

A. Mitra
e-mail: amitra@cs.ucr.edu

W. Najjar
e-mail: najjar@cs.ucr.edu

predication and pipelining in the datapath. One of the important features of CIRRF, and ROCCC, is its support for the import of pre-designed IP cores into the original C source code allowing the user to leverage the huge wealth of existing IP cores while programming the configurable platform using a HLL. Using experiments and examples we show that CIRRF is a solid foundation to generate high-performance hardware.

Keywords Configurable computing · Intermediate representation · FPGA · VHDL

1 Introduction

The main problem standing in the way of wider acceptance of reconfigurable computing platforms is their programmability. Currently, application developers must have extensive hardware expertise, in addition to their application area expertise, if they are to develop efficient designs that can fully exploit the potential of FPGA-based configurable platforms. Designing and mapping large applications onto FPGAs is a long and tedious task that involves a large amount of low-level design in a Hardware Description Language (HDL). This poses two problems: traditional application developers are typically not HDL designers, and HDLs are not well suited to algorithm implementation. Several projects have looked at the translation of traditional programming languages, such as C/C++ or Java, to HDLs for mapping onto FPGAs or other similar fabrics. This is a challenging task. The FPGA is an amorphous mass of logic onto which the compiler must create a datapath and schedule the computation. Such a task requires the harnessing of technologies developed for parallelizing compilers as well as those developed for high-level synthesis. The fundamental differences between the spatial computing model and the temporal, or von Neumann, model are:

- Spatial computing is inherently parallel while temporal computing is sequential.
- Temporal computing relies on two centralized storage locations that are both explicitly addressed by the code: the register file and the memory. In spatial computing, storage is distributed throughout the circuit and is accessed implicitly rather than explicitly. Furthermore, it is the task of the compiler to explicitly create the storage on the FPGA and schedule its accesses.
- Scheduling in temporal computing is driven by control flow, while in spatial computing it is driven by data flow.

The main challenge in translating from a HLL to an HDL is in overcoming these fundamental differences. Optimizing compilers for traditional processors have benefited from several decades of extensive research that has led to efficient tools. Similarly, electronic design automation (EDA) tools have also benefited from several decades of research and development leading to powerful tools that can translate VHDL and Verilog code, and recently SystemC code, into efficient circuits. However, little work has been done to combine these two approaches into one integrated compilation tool where HLL are translated into a high-performance circuit.

At the heart of each compiler or synthesis tool is an intermediate representation (IR) around which the tool is built. In this paper we propose CIRRF (Compiler Intermediate Representation for Reconfigurable Fabrics), an IR designed for the compilation of

traditional imperative, high-level languages, targeting reconfigurable devices. CIRRF is intended to be an open halfway-point representation between a high-level language and a specific reconfigurable platform. A front-end tool would translate C/C++, FORTRAN, Java or SystemC to CIRRF. Back tools would map CIRRF to a specific target. Loop and array transformations are dealt with in the front-end tools; target-specific optimizations are implemented in the back tools. CIRRF is designed to be both language and target independent. It is the intermediate representation of the ROCCC compiler (Riverside Optimizing Compiler for Configurable Computing). CIRRF differs from traditional compiler IRs in that it supports concurrency, both explicitly and implicitly, as well as the instantiation of and accesses to on-chip storage structures. It records information about loop types, memory interfacing, instruction predication and pipelining. Special instructions for efficient datapath generation are introduced. ROCCC does not support pointers and memory allocation.

The ROCCC compiler is designed to generate VHDL from C. However, not all application algorithms can be efficiently described by C. Furthermore, industry has invested tremendous financial and technical efforts on pre-designed intellectual property (IP) cores for FPGA-based platforms that are not only very efficient but have been thoroughly tested and verified. These IP cores come in the form of synthesizable HDL code or even lower level descriptions. They vary drastically with respect to their control and timing protocol specifications which intended to be interfaced to HDL-based designs. Compilers for FPGA-based reconfigurable systems must therefore leverage that huge wealth of IP designs by allowing the user to import these into high-level language (HLL) source codes. To do so would require a wrapper structure that would hide the timing and stateful nature of the IP cores and make each look, to the HLL compiler, as an un-timed side-effect free function call.

We propose a mechanism for the automatic generation of such a wrapper using ROCCC. Users provide the high-level description of a wrapper, which is based on C with timing information. CIRRF records the timing information so that from the compiler's point of view, the wrapper described in C is essentially a timed control flow graph. The compiler's back-end converts this timed CFG into predicated DFG and eventually generates an IP wrapper in VHDL. Notice that a normal C code input to ROCCC does not have any timing implication, while a wrapper in C is a special case, which does have timing information.

The rest of this paper is organized as follows: Sect. 2 presents CIRRF's architecture and the method we build CIRRF; Sect. 4 extends CIRRF to support the compiler for IP wrapper generation. Results are given in the subsections of Sect. 2 and Sect. 4. Section 5 reviews related work; and Sect. 6 concludes the paper.

2 CIRRF Intermediate Representation

The major objective of CIRRF is to separate the language specific compiler concerns from those of the target platform. The architecture of FPGA-based platforms can vary widely in the number and types of FPGAs, the number and size of on-board memory modules, the bus width connecting the FPGAs to the memories, etc. The compiler must therefore generate code that can take advantage of the unique features of each

platform. This approach allows for an easy targeting or re-targeting of new or modified platforms.

In this section we describe the overall structure of CIRRF and the two levels currently implemented as part of the ROCCC compiler tool. We present ROCCC’s workflow to build the CIRRF IR. The front-end optimizes the user-input code and generates Hi-CIRRF by adding macros into the source (Sect. 2.3). Starting from a conventional CFG, the back-end first constructs data flow for do-all loops (Sect. 2.4.1), then converts non-do-all nodes into data flow nodes using predication (Sect. 2.4.3).

2.1 CIRRF Architecture and Execution Model

ROCCC’s intermediate representation, CIRRF, is built using SUIF2 [1] and Machine-SUIF [2]. SUIF2 consists of high-level statements, such as loop statements and if-else statements, while Machine-SUIF consists of virtual machine instructions. CIRRF consists of two distinct but equivalent representations, as shown in Fig. 1. The Hi-CIRRF, built using SUIF2, is essentially C code augmented with macros. The Lo-CIRRF format, extended from Machine-SUIF, is semantically similar to assembly code. The advantage of this approach, which is commonly used in various compiler IRs, is that it allows the user to have multiple levels of entry into the IR. Functional debugging, for example, would be a lot easier at the Hi-CIRRF level. A functional simulation of the generated code would be feasible at the Lo-CIRRF level.

2.1.1 Hi- and Lo-CIRRF

The Hi-CIRRF representation is generated after the high-level compiler transformations have been applied. It has a C syntax that has been augmented with macros. The Lo-CIRRF serves as a platform for pipelining, interfacing to memory and generating VHDL. Its semantic and syntax are similar to assembly language and rely on statically single-assigned virtual registers.

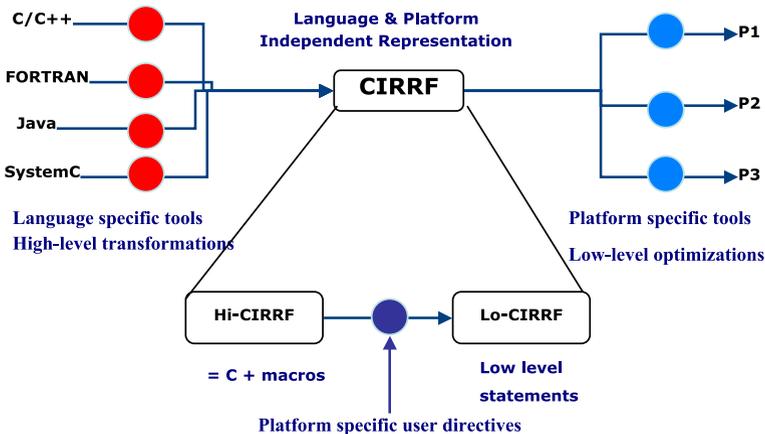


Fig. 1 Overview of CIRRF

The macros in Hi-CIRRF are used to:

- Instantiate and access on-chip buffers.
- Implement special operations that eliminate recurrence or enforce a pipeline delay for IP wrapper generation.
- Invoke bitwise and arithmetic operations, such as bit-insert, bit-extract, and minimum of two values.
- Invoke look-up tables and IP cores.

In Lo-CIRRF the code is similar to an assembly code. It is implemented as a control and data flow graph (CDFG) with the following characteristics:

- Virtual statically single-assigned registers.
- Register name indicates type (signed, unsigned) and bit size.
- Predicated instructions.
- Pipelining information for each instruction.
- Loop nodes that specify complete loop information such as loop type (parallel or sequential), nesting level, index and stride.

2.1.2 Execution Model

CIRRF is essentially intended to represent loop nests that would be mapped to hardware. Its execution model, shown in Fig. 2, is very simple: it consists of an input data memory which can be on or off-chip, a data fetch engine that collects the data into an input buffer. The data is then pushed, every cycle, into the pipelined datapath. The same structure is replicated at the output side. Note that, unlike in a von Neumann architecture, the datapath does not fetch the data, rather at each iteration, the correct set of data is selected by a controller from the input buffer and pushed into the datapath. The CIRRF execution model can therefore be considered a *decoupled* model of execution.

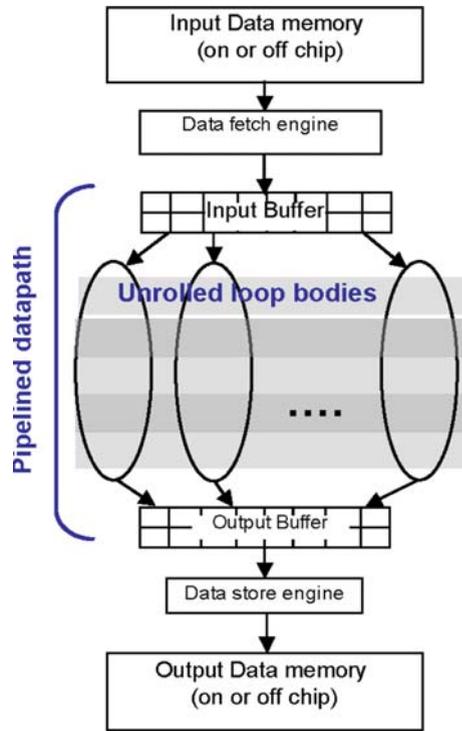
2.2 On-Chip Buffers

One of the distinguishing features of spatial computing representation is that storage needs to be explicitly architected by the code while in temporal computing it is implicitly available in the form of registers and memory (cache, main memory etc). Furthermore, the accesses to memory for both data reads and writes have to be generated as part of the computation.

Buffers are therefore first class citizens in CIRRF. They play several roles including:

- Interface to memory for both reads and writes.
- Interface between two segments of the datapath that operate in a producer consumer relationship.
- Cache data fetched from memory for future reuse to reduce the number of memory accesses.
- Hold run-time constants.

Fig. 2 Execution model of CIRRF



The following is a partial list of buffers that are part of the CIRRF design to date:

- *Memory interface.* The *mem_read_fifo_buffer* and *mem_write_fifo_buffer* serve as interfaces between the memory, off or on chip, and the datapath. They are parameterized in both width and size and are accessed in one cycle.
- *Data reuse.* The *smartbuffer* [3] is designed to facilitate the reuse of data fetched from off-chip memory. It is particularly well suited for window-based operation as is very common in signal and image processing algorithm where every data sample, or pixels, participates in the computations of several windows. The deallocation and replacement of a data value in the *smartbuffer* is scheduled by the compiler.
- *Run-time constants.* The ROCCC compiler folds compile-time constants into the computation. However, run-time constants need to be made available to the computation. Two such structures are built into CIRRF.
 - The *constant_buffer* consists of a small set of scalar values that are directly written to registers in the datapath.
 - The *one_time_buffer* holds an array of constant values. It operates just like the *mem_write_fifo_buffer* but is read only once. It is used in the cases where the number of constants is too large to be enumerated by the compiler.

Figure 3 shows a matrix multiplication example with the *mem_read_fifo_buffer* and the *one_time_write_buffer*.

```

for(i=0;i<N;i=i+1){
  sum = 0;
  for(j=0;j<N;j=j+1)
    sum=sum+A[i][j]*B[j];
  C[i] = sum;
}

```

(a) The original C code of a matrix multiplication.

```

for(i=0;i<N;i=i+4){
  constant_read_buffer(i);
  sum1=0; sum2=0; sum3=0; sum4=0;
  for(j=0;j<N;j=j+2){
    mem_read_fifo_buffer(A, i, j,
      a1,0,0,a2,0,1,a3,1,0,a4,1,1,
      a5,2,0,a6,2,1,a7,3,0,a8,3,1);
    mem_read_fifo_buffer(B, j,b1,0,b2,1);
    sum1=sum1+a1*b1;    sum1=sum1+a2*b2;
    sum2=sum2+a3*b1;    sum2=sum2+a4*b2;
    sum3=sum3+a5*b1;    sum3=sum3+a6*b2;
    sum4=sum4+a7*b1;    sum4=sum4+a8*b2;
  }
  one_time_write_buffer(C, i,
    sum1, 0, sum2, 1, sum3, 2, sum4, 3);
}

```

(b) The C code with the buffer macros. The *i* loop is unrolled four times and the *j* loop is unrolled twice.

Fig. 3 An example of the *mem_read_fifo_buffer* and the *one_time_write_buffer*

- *Producer/Consumer*. Buffers are also generated in CIRRF to implement on-chip producer consumer relationships between segments of the datapath: *pc_fifo_buffer*.

2.3 Building Hi-CIRRF

The ROCCC system performs the following loop transformations: invariant code motion; partial and full loop unrolling; loop peeling; loop un-switching; loop tiling; strip-mining; loop fusion; constant propagation of scalars and constant array masks; constant folding; elimination of algebraic identities; copy propagation; dead and unreachable code elimination; code hoisting; code sinking; scalar renaming; and division by constant approximation using shifts and additions. ROCCC generates reduction on scalars that accumulate values through associative and commutative operations on themselves. It also carries out the following hardware-specific analysis and transformations:

- *Scalar replacement*. The front-end decouples a do-all loops' array accesses from computation. Figure 4a shows the original C code of a gray scale transformation example. After undergoing scalar replacement, the computation is isolated from memory accesses (Fig. 4b) by a smart buffer. The smart buffer will be synthesized on configurable fabrics as the interface with memory. One important characteristic of smart buffers is that they reuse input data between iterations and push one iteration's input data initiatively to the datapath, rather than being accessed by the datapath [3]. The syntax of a two-dimensional smart buffer macro is:

```

smartbuffer2(input_array_name, address_index_1, address_index_2,
  scalar_1, offset_1_1, offset_1_2, scalar_2, offset_2_1, offset_2_2);

```

```

for (i=0; i<62; ++i) {
  for(j=0; j<62; ++j) {
    sum = (a[i][j] + a[i][j+1]) + (a[i+1][j] + a[i+1][j+1]);
    if(sum > 170)
      b[i][j] = 255;
    else {
      if(sum < 85)
        b[i][j] = 0;
      else
        b[i][j] = 127;
    }
  }
}

```

(a) The original C

```

for (i=0; i<62; ++i) {
  for(j=0; j<62; ++j) {
    smartbuffer2(a, i, j, x1, 0, 0, x2, 0, 1, x3, 1, 0, x4, 1, 1);
    sum = (x1 + x2) + (x3 + x4);
    if(sum > 170)
      tmp = 255;
    else {
      if(sum < 85)
        tmp = 0;
      else
        tmp = 127;
    }
    memory_write_fifo_buffer2(b, i, j, tmp, 0, 0);
  }
}

```

(b) The intermediate C code.

Fig. 4 A gray scale transformation example in C. (a) The C code sums the gray scale values in a 2×2 window in the input image (array a) and assigns one of three possible values to the pixel in the output image (array b) depending on the value of *sum*. (b) The intermediate C code emitted by front-end. The highlighted segments are created by scalar replacement

For example, in the smart buffer macro in Fig. 4 b, the last three parameters (*x4*, 1, 1) stands for: $x4 = a[i + 1][j + 1]$; Similarly, the syntax of a two-dimensional memory write FIFO buffer macro is:

```

mem_write_fifo_buffer2(output_array_name, address_index_1,
address_index_2, scalar_1, offset_1_1, offset_1_2, scalar_2,
offset_2_1, offset_2_2);

```

The memory write FIFO buffer macro in Fig. 4b stands for $b[i + 0][j + 0] = tmp$; Currently we have the following constraints on buffer macros. An array can only appear in at most one buffer macro. The address indexes of buffers are also the loop counters. The operator between an address index variable and the offset can only be either addition or subtraction.

- Feedback variable detection. The compiler detects scalar recurrence between adjacent iterations. For example, for a loop having a statement $sum = sum + a[i]$, to eliminate the loop-carried dependency, the compiler replaces the sum on the left and the sum on the right with *store2next()* macro and *load_previous()* macro, respectively. These macros guide the backend to instantiate a feedback register to store the current sum for the next iteration.

The output from the front-end is in the forms of both an IR file and an intermediate C with macros. Users could do further optimizations and add pragmas onto the intermediate C.

2.4 Building Lo-CIRRF

The backend first constructs a conventional CFG (Fig. 5a). The compiler finds loops and loop-depth. Loop types are recovered from pragmas provided by the user: Currently, these include one-dimensional do-all loop, two-dimensional perfect nested do-all loop, and non-do-all loop. The pre-process phase of the back-end converts macros in Hi-CIRRF into corresponding instructions. Particularly, buffer macros are converted into buffer instructions and put into separated nodes, as shown in Fig. 5b for example. The smart buffer instruction is shown below.

$$smb2 \text{ main.x1, main.x2, main.x3, main.x4} \leftarrow \\ \text{main.a, } \$vr35.s32, \$vr75.s32, 0, 0, 0, 1, 1, 0, 1, 1;$$

where *main.x1* through *main.x4* are the buffer's output ports, *main.a* is the input memory name, *vr35* and *vr75* are the address index variable (also loop counters). The integers are four pairs of offsets. We categorize basic nodes into two types: parallel *do-all* nodes and sequential *non-do-all* nodes. The compiler's back-end generates datapath using different pipelining and scheduling mechanisms accordingly.

2.4.1 Building Lo-CIRRF for Parallel Loops

For a parallel loop, or do-all node, ROCCC exploits both instruction-level and loop-level parallelisms. The compiler first performs if-conversion to eliminate control flow within the loop body. It then walks through the loop body in a depth-first order and adds extra nodes for joint nodes that have more than two predecessors. Figure 5c shows that a new node, node 16, is added as node 10's predecessor.

In order to allow the datapath to implement multiple concurrent loop iterations, the compiler groups the instructions inside a node into execution levels so that each level is an instantiation of one iteration. Statically single-assigned variables are added to

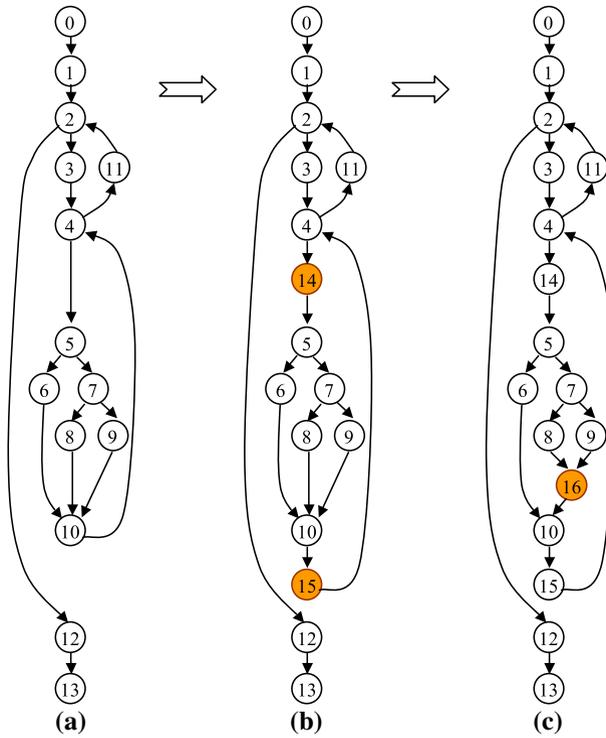


Fig. 5 CFG of the gray scale transformation example. (a) The original CFG. (b) The CFG after pre-processing: highlighted nodes are added. Node 14 is the smart buffer node and node 15 is the memory write FIFO buffer node. (c) The CFG right before performing if-conversion. Node 16 is added to make each joint node have only two predecessor nodes

duplicate a variable if that variable's definition reaches more than one level of execution lower. Therefore, the definition of each value is at a level strictly higher than that of its use. Multiplexers are added, and predictors are duplicated and propagated along with execution levels, as shown in Fig. 6. Notice that each execution level, in the data flow in Fig. 6, represents a single iteration at a given execution phase. Then the compiler pipelines the loop body by automatically inserting latches in some execution levels of the data-flow. Instructions belonging to the same execution level are either all latched or all non-latched. Multiple consecutive execution levels may be assigned to the same pipeline stage.

The performance of the generated do-all loop datapaths is described in [4]. The parameters of the datapath, such as the number of pipeline stages, are known at compile time and determined by the user. A parameterized do-all loop controller schedules the fully pipelined datapath's operation.

2.4.2 Building Lo-CIRRF for Multi-Cycle Operation Support

The ROCCC compiler supports multi-cycle pipelined instructions. This requires synchronization of the dataflow at all levels and prior knowledge of latency. To accomplish

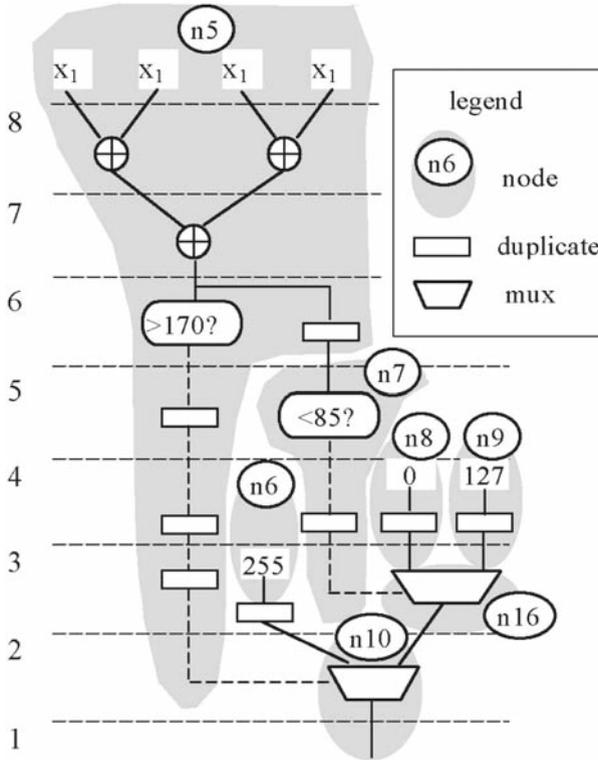


Fig. 6 DFG of the gray scale transformation example. The numbers on the left side are the execution levels. The dashed lines carry speculators, while the solid lines carry data values. The nested if-else statements are converted into a data-flow free of control. Speculators are duplicated along the execution levels. Each execution level is an instantiation of one iteration. Notice that latches can be added between any execution levels

this, we introduce an extra step in the compilation process after pipelining all the one-cycle operations. In this extra step, we process the DFG for the bottom up all stages that contain generic functions. If the generic function requires more than one cycle the compiler pushes all lower stages by equivalent amount of remaining stages. If data will not pass through the generic function then pipelined copies (*mov* instructions) are placed to propagate data from before the pipeline stage of the generic function until the end.

With the generic multi-cycle function added to ROCCC, we have added support for IP core math functions, most notably, the support for floating point functions (ADD, SUB, MUL, DIV). Floating point instructions are preprocessed and converted to generic multi-cycle IP cores.

2.4.3 Building Lo-CIRRF for Sequential Loops

In a sequential loop, or non-do-all node, only one iteration is executed at a time. The compiler schedules instructions into different execution levels in a manner similar to

a do-all node. But the definition of a variable does not have to be constrained to one level above. Multiple instructions might belong to the same execution level and can be executed simultaneously to exploit instruction level parallelism.

The compiler utilizes predication to schedule the execution of non-do-all nodes' instructions. Each pipeline stage is guarded by a predicator. Lo-CIRRF records predicated instructions in the following format:

$$ADD \ \$vr4.s16, \$vr3.s16, \$vr2.s16, \$vr1.u1$$

$vr4$ is the destination operand and $vr3$ and $vr2$ are the source operands. $vr1$ is the predicator, which is also a source operand. Predicators are passed inside basic nodes for scheduling purpose. A special instruction, PFW (predicator forward), is created to pass a predicator from the current stage to the next stage, which may be or may not be in the same node:

$$PFW \ \$vr2.u1, \$vr1.u1$$

$vr1$ and $vr2$ are two predicators. The instructions guarded by $vr2$ are one pipeline stage later than the ones guarded by $vr1$. Their types are $u1$, which stands for unsigned one-bit. To convert CFG to DFG, the branch instructions of basic nodes are replaced by Boolean instructions, whose destination operands are evaluated by this basic nodes' successor nodes.

The back-end IR construction phases described so far translate both do-all nodes and non-do-all loop nodes into a DFG. Essentially, the original CFG now is a DFG, in which do-all loop nodes, if any, are connected together by non-do-all nodes. Then the compiler's VHDL generator emits VHDL code for the entire DFG, including buffers.

2.5 Case Study

In this case study, besides reporting the synthesis results of the gray scale transformation example discussed in previous sections, we examine CIRRF on another application, an alternative finite impulse response filter (FIR).

Figure 7 shows the original C code. In even outer iterations flag is one, while in odd outer iterations flag is zero. Therefore the two do-all inner loops (the two highlighted regions) are executed alternately. Each of these two inner loops is a 5-tap FIR. The upper FIR reads array a and writes array b , while the lower FIR reads array c and writes array d . ROCCC's front-end performs scalar replacement, and instantiates a one-dimensional smart buffer macro and a one-dimensional memory write FIFO buffer macro into each inner do-all loop, as described in Sect. 4.

The back-end first builds a CFG from the input. For this example, there are two do-all loops nested inside the outer non-do-all loop. The back-end scans the whole CFG and transforms the two inner 5-tap FIRs into data flow. Each FIR's loop body is aggressively pipelined, and the resulting datapath has a throughput of one iteration per clock cycle. These two FIRs are controlled by two do-all loop controllers, as shown in Fig. 8a. The back-end then converts the outer loop and the rest of the CFG

```

void alternative_fir5() {
    int i,j, m;
    int a[256], b[256];
    int c[256], d[256];
    int flag;

    flag = 1;
    for (m = 0; m < 10; m = m + 1) {
        if(flag == 1) {
            for(i = 1; i < 251; i = i + 1)
                b[i] = (3 * a[i-1] + 5 * a[i]) +
                    (7 * a[i+1] + 9 * a[i+2]) + 11 * a[i+3];
            }
            else {
                for(j = 1; j < 251; j = j + 1)
                    d[j] = (3 * c[j-1] + 5 * c[j]) +
                        (7 * c[j+1] + 9 * c[j+2]) + 11 * c[j+3];
            }
        flag = flag ^ 1;
    }
    return;
}

```

Fig. 7 An alternative FIR example in C. The first highlighted segment is a 5-tap FIR reading array *a* and writing array *b*, while the second highlighted segment is a 5-tap FIR reading array *c* and writing array *d*. *flag* alternates the execution of these two segments

into a DFG by predicating all the non-do-all nodes. We list the instructions of these nodes in Fig. 8b. Node 2 and node 10 are the head and tail nodes of the outer loop, respectively. The first instruction of node 2, the *ior* instructions, produces the predicator (*vr1321*) for the two instructions below it (*pfw* and *sle*) by examining a valid output predicator from either node 1 (not shown), the first active node; or node 11, the loop tail. Figure 8c depicts node 2's circuit in detail. Guarded by *vr1321*, the *sle* instruction asserts its destination operand when the outer loop is done, or de-asserts its destination operand when needing to execute a new outer loop iteration. Node 2's *pfw* instruction forwards a valid *vr1321* to the two successor nodes, node 3 and node 11, for their predicator evaluation. Node 3 enables one of the two FIRs by either asserting or de-asserting *vr1325*, depending the value of *flag* (*vr78*). Node 10 is activated by the done signal from one of the FIRs' loop controller and updates the value of *flag* (*vr230*) and the loop counter *m* (*vr233*). Node 11 indicates the completion of the whole procedure.

Table 1 shows the synthesis results of the gray scale transformation example discussed in previous sections and the alternative FIR. The target FPGA is the Xilinx xc2v8000-5 with 46592 slices. The generated VHDL is synthesized and placed-and-routed using the Xilinx tool chain. The second and the third rows are the datapath's bit-size and BRAM bus's bit-size. *slices* and *clock* are collected from place-and-route reports. The last row is the number of do-all loop iterations executed per clock cycle. For gray scale transformation, the do-all loop's loop body has nested if-else statements, as shown in Fig. 4. After undergoing if-conversion, in the IR right before VHDL

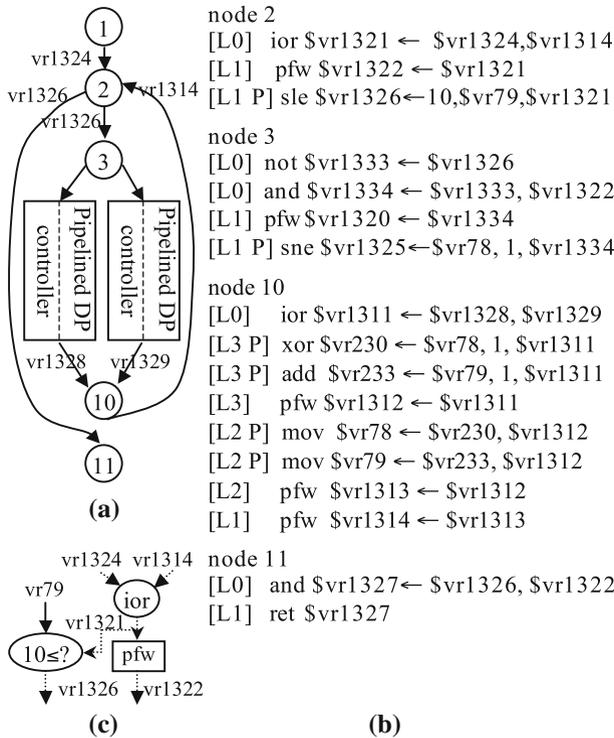


Fig. 8 The DFG and IR of alternative FIR. (a) is the DFG. Node 3’s successors are two do-all loops controlled by their loop controllers. For simplicity, we are not presenting the details of these two loops. Operands’ data types are not shown. Node 1 through 3 and node 10 through 12 are non-do-all nodes, scheduled by predicators that are carried by the edges. (b) The DFG IR of node 2, 3, 10 and 11. The *L* field is the pipeline stage (latch level) of the instruction within its node. An instruction with a *P* field is predicated by its last source operand, which is the predicator. (c) Shows node 2’s circuit, in which a solid lines carry data and a dashed lines carry predicators

Table 1 Synthesis results of case study examples

	Gray scale	FIR
Datapath bitsize	16	8
Memory bus bitsize	16	8
Slices	318	531
Clock (MHz)	59.7	100
Iterations per cycle	0.5	1

emission (Fig. 6), the control flow is eliminated and the resulting data flow is capable of executing one iteration each clock cycle. Notice we configure the BRAM’s data bus (the third row) to have the same bit-size as that of the data elements (pixels), and each iteration needs four (2×2) pixels. Though the smart buffer reuses one column of the pixels loaded in previous iteration, it still needs two cycles to load the remaining two new pixels. This explains why, for the gray scale transformation example, the number of iterations per cycle is 0.5. For the alternative FIR, the VHDL generator

generates one smart buffer and one memory write FIFO buffer for each of the two do-all loops according to the buffer representations in the IR. When either one of the two do-all loops is active, the corresponding smart buffer exports one window of data (five elements) to the datapath every clock cycle, and therefore the circuit executes one iteration per cycle. *slices* consists of the hardware of two do-all loops (the datapath, the buffers and the controller for each FIR) and the hardware of the non-do-all nodes, as shown in Fig. 8a.

3 On-Chip Buffers—Implementation and Evaluation

The on-chip buffer macro is one of the features of CIRRF. In [3] and [5] we introduced the smart buffer approach to input data reuse. The smart buffer is ideally suited for operations that involve a heavy reuse of fetched data, windowing operations on images being an obvious example. In order to support more memory access patterns this mechanism is extended by adding more types of FIFO buffers as described in Sect. 2.2. In this section we describe the implementation of various on-chip FIFO buffer structures and report on the evaluation of their area and speed.

3.1 Implementation of FIFO Buffers

ROCCC supports both pre-designed VHDL library FIFO buffers and compiler generated buffers. The smart buffer falls in the later category. This section describes the three VHDL library FIFO buffers and an improvement to the smart buffer design.

The three FIFO buffers in the ROCCC VHDL library are push stack FIFO buffer, circular FIFO buffer and hybrid FIFO buffer. The first two are built with pure logic, while the hybrid FIFO buffer uses BRAM for storage. They are instantiated by the compiler when generating VHDL code. Each of these FIFOs have *registered outputs*—the data output and data output assertion are synchronized together instead of having to wait one cycle after assertion for the data output to be valid.

A template that all the library fifos follow is show in Fig. 9a: each fifo has *datain*, *dataout*, and handshaking signal ports. Figure 9b–d are different implementations based on that template.

- *Push stack FIFO buffer*. Shown in Fig. 9b. Items are pushed onto the top of the stack buffer (like a real FIFO system) and popped out from the bottom.
- *Circular FIFO buffer*. Shown in Fig. 9c. It uses a single buffer and two array pointers, *write_ptr* and *read_ptr*, for keeping track of where the data has been written and where it can be read from. Due to the implementation, if a size of N elements is needed to be stored in the FIFO, $N+1$ elements must be allocated for the buffer.
- *Hybrid FIFO buffer*. Shown in Fig. 9d, its component's layout is actually similar to that of the logic FIFO—circular. Read and Write pointers are used, but the buffer is implemented with BRAM instead of a set of register buffers. To improve the latency, a mini buffer (a logic FIFO) is used to cache a few elements from the top of the FIFO. BRAM FIFOs do not spend FPGA fabrics on storage and are area efficient.

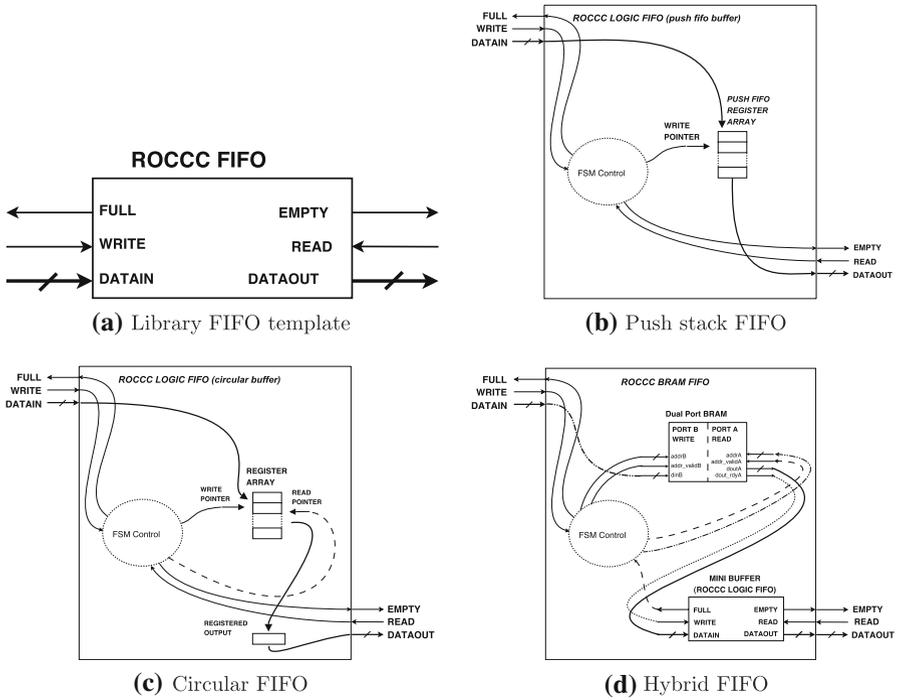


Fig. 9 Library FIFO buffers

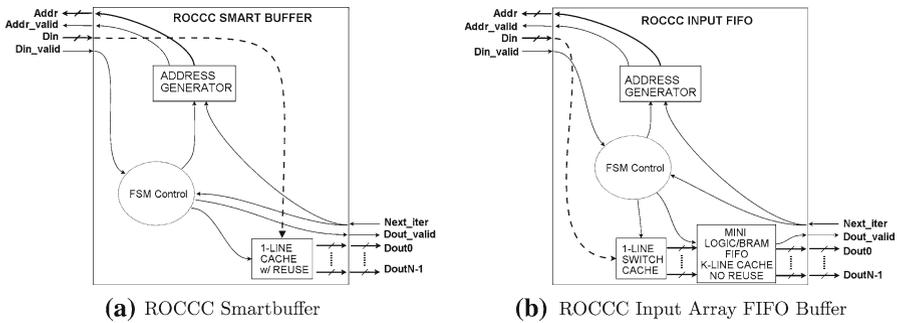


Fig. 10 Smartbuffer and input array FIFO buffer

The implementation of the smart buffer (Fig. 10a) is designed so that if no reuse is needed, the buffer generated acts like a FIFO. In such a case, the compiler builds a small cache within the smart buffer, now used as a FIFO, so it can work in the background during pipeline stalls, memory fetch stalls, or buffer synchronization stalls. This structure is called the ROCCC Input Array FIFO Buffer (Fig. 10b).

A couple of assumptions that are required before designing this input array FIFO buffer:

- The input array FIFO buffer must *not* be used in cases where re-use is needed. If this is done then extra logic will be used for storing the same element and thus wasting resources.
- The desired latency and acceptable area cost of the input FIFO. If an area efficient system is desirable then the use of the smartbuffer with no-reuse (which downgrades to a FIFO) would be best to use.

If the system can afford extra area cost to allow decrease of execution time, and no-reuse is needed, then the input array FIFO should be generated.

3.2 Evaluation

In this section we evaluate the performance of the various FIFO implementations and their tradeoffs.

3.2.1 ROCCC Lib FIFO Buffer Experimental Results

The library FIFOs area and timing are as followings in Fig. 11a and b, respectively. The figures suggest the BRAM FIFO is one of the most efficient in area and with a decent clock speed. The shortcoming with this type of FIFO is the latency is about four cycles (from the time when data is inputted to when it is outputted). For the logic FIFOs the latency is about 1 cycle (which is a good as it can get), but as can be seen in the figures the area usage is more but the clock speed suffers. If we compare the two logic FIFOs we get that the push stack implementation is more area efficient than the circular implementation. This is most likely due to the routing of the outputs of each of the registers in the circular buffer to the output as opposed to routing only the last register to the output in the push stack implementation.

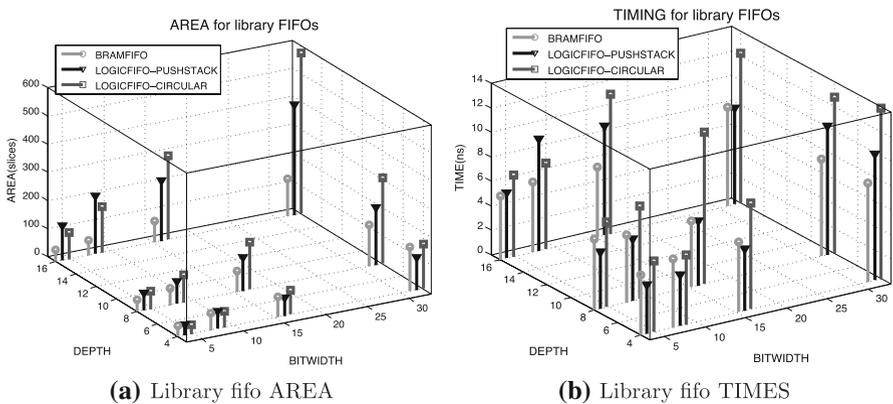


Fig. 11 3D plots of library FIFOs

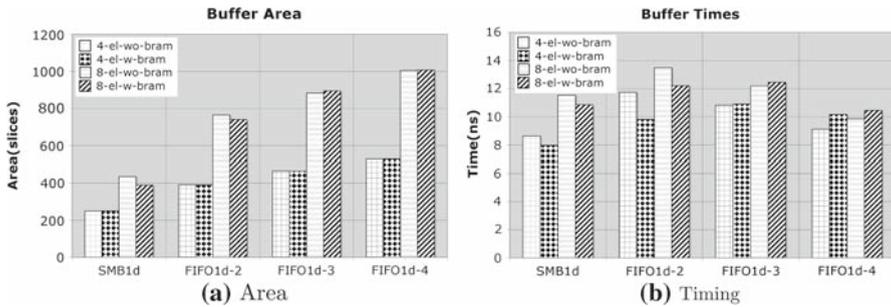


Fig. 12 Input array buffers—area and timing plots

3.2.2 Input Array FIFO Buffer Experimental Results

Figure 12a and b show how the area or timing is affected by the type and depth of the buffer. Each buffer holds 4 or 8 elements (32 bits each) per line of cache. There is data for the SMB1D (smartbuffer 1-dimension), which acts as a fifo for 1 cache line. The rest are the input array FIFO buffer implementations of different depths (depths 2, 3, and 4), which are listed horizontally on the table. So the table lists buffers for sizes 1 to 4, but the size 1 implementation is the smartbuffer, and the rest are the FIFOs so that one can compare the relative growth in area and time. On top of all the implementation we can include the buffers with some BRAM or not, so that the synthesis tools knows to use the internal board components or port map all I/O pins of the buffer to the I/O pines of the board. As can be seen, offers some area discrepancy.

The plots show in Fig. 12 exhibit a linear growth in area for different cache sizes of 2–4 elements. If the compilation process requires optimal area then a cache size of 1 is considered and the smartbuffer is generated. The timing on the other hand exhibits a decrease in clock rate, faster frequency, from cache level 2–4 (all the input array fifos), due to the way the cache is implemented (in this case generated with a cache built from the library ROCCC logic fifo—stack). All things considered, if area can be sacrificed the compilation process may consider to reduce stalls by generating the input array FIFO buffer.

4 Interface Synthesis

Pre-designed IP core represent a huge intellectual and financial wealth that high-level compilation tools targeting FPGAs should not ignore. The ROCCC compiler does support the import of pre-designed IP cores into C source codes. Most often, the interface to these cores is timed and requires several cycles of synchronization and handshaking. These characteristics do not fit well with the C semantics. In this section we describe our approach which is to generate a wrapper that would make the IP core look and behave like a C function. The workflow is shown in Fig. 13. Taking the high-level wrapper abstractions as input, ROCCC generates synthesizable wrappers in VHDL separately and these wrappers are instantiated as components in the outer circuit. Notice that an IP core is not necessarily a mandatory element of the main

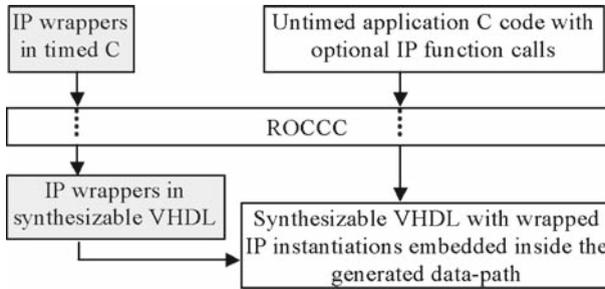


Fig. 13 ROCCC’s workflow when IP function call present. Notice that an IP core is not necessarily a mandatory element of the main untimed application C code. The grayed out part on the left only exists when there is an IP instantiation in the source code

untimed application C code. The grayed out part on the left only exists when there is an IP instantiation in the source code.

We start with an example of a 16 samples complex FFT, in Sect. 4.1, taken from the Xilinx website that we use to demonstrate our approach in the remainder of this section.

4.1 An IP Core Example

The grayed out part of Fig. 14 is a 16-point complex Fast Fourier Transform core (FFT16). Pins *di_r* and *di_i* are respectively the real and imaginary serial data input, *xk_r* and *xk_i* are the output. *ce*, clock enable, must be asserted only when the core is active. *start* must be asserted two clock cycles ahead of the first pair of input data. *done* is asserted when the first pair of output data is ready. *fwd_inv* selects between forward or inverse FFT. *scale_mode* selects from two scale-coefficients: 1/16 or 1/32.

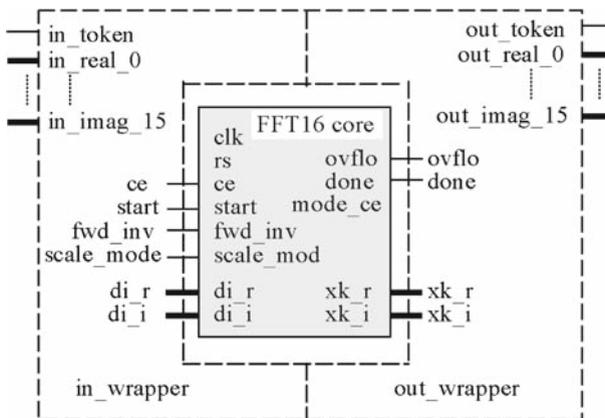


Fig. 14 The grayed out square is the FFT16 IP core. A wrappers interface consists of one or more data ports and one token (either input or output token)

The *ovflo* pin indicates the core has generated an arithmetic overflow. *mode_ce* input indicates when to sample *fwd_inv* and *scale_mode*.

4.2 High-Level Wrapper Abstraction

An IP core requires a wrapper for both its input and output interfaces. In some cores these two interfaces have common signals that handle synchronization and handshaking. In our implementation this role is covered by the outer circuit within which the core is embedded.

Figure 15 lists the code for input wrapper of FFT16's in C. We use pointer type to distinguish output signals from input signals in the function declaration. The input set, which communicates with the outside, is composed of one token and several data variables. The output wrapper, not shown here, has the same structure. Thus both the input and output interfaces have the same structure as shown in Fig. 16.

By its very nature, an interface to an embedded core must support timed activity. In Fig. 15, the function call *wait_cycles_for(n)* indicates the statements following it must be executed *n* cycles later. Any statements between two adjacent *wait_cycles_for* calls

Fig. 15 Timed high-level abstraction of FFT16's input wrapper in C

```

void in_fft16 (int in_token, /*the core's input predictor*/
int real_0, ..., int real_15, /*16 real-component inputs*/
int imag_0, ..., int imag_15, /*16 imaginary-component inputs*/
int* CE, int* SCALE_MODE, /*pointers are output*/
int* START, int* FWD_INV, int* DI_R, int* DI_I)
{
    int real_reg_0, ..., real_reg_15; /*internal registers to*/
    int imag_reg_0, ..., imag_reg_15; /*store the input data*/

    *SCALE_MODE = 1;
    *FWD_INV = 1;

    if(in_token == 1) {
        wait_cycles_for(1);
        real_reg_0 = real_0;
        .....
        real_reg_15 = real_15;
        imag_reg_0 = imag_0;
        .....
        imag_reg_15 = imag_15;
    }

    *START = 1; /*assert start signal in this cycles*/
    *CE = 1; /*assert ce signal in this cycles*/

    wait_cycles_for(1);
    *START = 0; /*de-assert start signal in this cycles*/

    wait_cycles_for(1);
    *DI_R = real_reg_0;
    *DI_I = imag_reg_0;
    .....
    wait_cycles_for(1);
    *DI_R = real_reg_15;
    *DI_I = imag_reg_15;

    wait_cycles_for(69);
    *CE = 0; /*de-assert ce signal 69 cycles later*/
}

```

store the 16 pairs of input data into internal registers in this cycle

export the 16 pairs of data into the core serially in 16 consecutive cycles

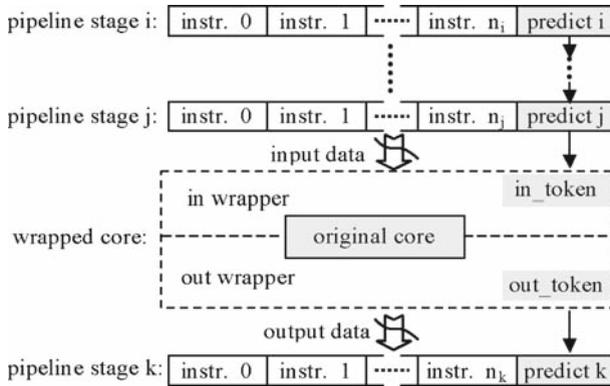


Fig. 16 The execution model of a wrapped IP core inside the predicated datapath. From the outside, a wrapped IP core has an identical predication mechanism as other predicated instructions

must be executed in one clock cycle. For example, FFT16’s timing protocol requires that the *start* signal be high for one clock cycle, two clock cycles ahead of the first pair of input data. In order to describe this timing requirement, the user assigns *start = 1*, calls *wait_cycles_for(1)*, de-asserts *start*, calls *wait_cycles_for(1)* again, and begins assigning input data into the core, as shown in Fig. 15. The timing of signal *ce*, which needs to be asserted for a 87-cycle period, is expressed in the same way. Parallel to serial conversion is also describes naturally in the timed C code. At the beginning of the function body, *scale_mode* and *fwd_inv* are statically assigned to *high*. If desired, they can also be assigned by the wrapper’s input signals at run-time in the same way as assigning *start* or *ce*. That way, the FFT16 core can be easily switched between a forward FFT and an inverse FFT.

The wrapped FFT16 core is instantiated in high-level C in the form of:

```
FFT16(di_r[0], di_i[0], ..., di_r[15], di_i[15], *xk_r[0], *xk_i[0], ...,
      *xk_r[15], *xk_i[15]);
```

This wrapping approach keeps the original IP core’s functionality to a great extent and still stays at high-level. The wrapper plays a role of a bridge between the timing diagram in an IP core’s data-sheet and the automatically generated synthesizable wrapper in VHDL.

4.3 Wrapper Synthesis

The timed high-level wrapper, the code in Fig. 15 for example, is passed through the ROCCC compiler in Fig. 1 as user input. Currently, the front-end does not do any optimizations on IP wrappers. The back-end first gets the control flow graph (CFG) of a wrapper, and converts the CFG into static single assignment (SSA) CFG. Starting from this SSA-CFG, the back-end constructs the DFG (Fig. 17)]. First, the *pre-process* pass converts *wait_cycles_for(n)* function calls into instruction *WCF n*, where *n* is an immediate operand. When building the data-flow, the compiler replaces each *WCF n*

Fig. 17 Wrapper pipelining and scheduling heuristic

```

for( each basic node b in cfg)
{
  for( each “WCF n” instruction instr in b)
    if( n > 1) {replace instr with n “WCF 1” instructions}
  assign the instructions between two adjacent WCF
    instructions into the same pipeline stage
  replace all “WCF 1” instructions in b into “PFW $vr1, $vr2”
  for( each PFW instruction instr in b) {
    guard all instructions at the same pipeline stage as instr
    using instr’s source operand as the predicator
  }
  if( b ends in a conditional branch instruction instr)
    convert instr to a Boolean instruction
}
{ add more combinational Boolean instructions if necessary to
  pass predicators appropriately from predecessor nodes to
  successor nodes.
}

```

into n consecutive *WCF 1* instructions. Thus a *WCF* instruction has a clear hardware timing meaning, passing the predicator to the next pipeline stage. Essentially, at this point of the compilation, the wrapper is a timed CFG in CIRRF. The compiler forces all instructions between two adjacent *WCFs* to be in the same pipeline stage, as in the pseudo-code in Fig. 17.

This constraint ensures that the back-end’s pipelining consists with the high-level C’s timing semantics, and thereby satisfies the IP core’ timing requirement.

In the scheduling process, *WCF* instructions are replaced by *PFW* (predicator forward) instructions. *PFW* instructions pass predicators through the data-flow, while other pipelined instructions are guarded by predicators.

The IR, right before VHDL emission of the FFT16 input wrapper, is shown in Fig. 18. The IR records the predicated hardware actions with cycle-level timing constraints. In front of an instruction, the *L* field is the latch-level, namely, at which pipeline stage the instruction is executed. Instructions with a zero latch-level are combinational logic or even just wires if the opcode is *mov*. An instruction with a Boolean *P* field is guarded by its predicator, which is the last source operand. For IP wrappers, a *str* (store) instruction with zero address offset is treated as a *mov* instruction, whose destination is the operand that the pointer is pointing to. From line 16 through line 33, the anticipated hardware does the following: monitoring the assertion of the predicator from outside and passing it (line 16 and 17), storing all the 16 pairs of input data into internal registers and asserting *start* and *ce* (line 18 through line 21), one cycle later (line 22) de-asserting *start*, one more cycle later (line 24) starting feeding input data into the core pair by pair serially, after 69 more cycles waiting (line 30 through line 31), de-asserting *ce*.

The compiler’s very last pass emits the VHDL code for the wrappers. The combinational instructions become combinational logic in hardware and pipelined instructions become sequential logic. From the point of view of outside, the generated wrappers (the wrappers of FFT16 in Fig. 14, for example) have unified interface: one input

Fig. 18 Back-end IR of FFT16's input wrapper. The L field is the latch-level, a P field marks a predicated instruction. Line 12 and 13 configure the IP core. Line 18 through line 19 latch the 32 input data elements into internal registers. Line 20 and 21 assert *start* and *ce*, while line 23 de-asserts *start* (one cycle later). *ce* is de-asserted in line 33, 85 cycles later. Line 24 through line 29 export the 16 pairs of input data elements into the IP serially

```
(1) .in_fft16
(2) Node 1, {0} {2, 3}
(3) [L0] mov $vr85.u1 <- in_fft16.in_token
(4) [L0] mov $vr84.u16 <- in_fft16.real_0
.....
(5) [L0] mov $vr53.u16 <- in_fft16.imag_15
(6) [L0] mov $vr52.p1 <- in_fft16.CE
(7) [L0] mov $vr51.p1 <- in_fft16.SCALE_MODE
(8) [L0] mov $vr50.p1 <- in_fft16.START
(9) [L0] mov $vr49.p1 <- in_fft16.FWD_INV
(10) [L0] mov $vr48.p16 <- in_fft16.DI_R
(11) [L0] mov $vr47.p16 <- in_fft16.DI_I
(12) [L0] str 0($vr51.p16) <- 1 /*configure SCALE_MODE*/
(13) [L0] str 0($vr49.p16) <- 1 /* configure FWD_INV*/
(14) [L0] sne $vr559.u1 <- $vr85.u1, 1 /*set if not equal*/

(15) Node 2, {1} {3}
(16) [L0] not $vr560.u1 <- $vr559.u1
(17) [L87] pfw $vr471.u1 <- $vr560.u1

(18) [L87, P] mov $vr167.u16 <- $vr84.u16, $vr560.u1
..... /* latch the input data to the internal registers. */
(19) [L87, P] mov $vr198.u16 <- $vr53.u16, $vr560.u1

(20) [L87, P] str 0($vr50.p1) <- 1, $vr560.u1 /*assert START*/
(21) [L87, P] str 0($vr52.p1) <- 1, $vr560.u1 /*assert CE*/
(22) [L86] pfw $vr472.u1 <- $vr471.u1
(23) [L86, P] str 0($vr50.p1) <- 0, $vr471.u1 /*de-assert START*/

(24) [L85] pfw $vr473.u1 <- $vr472.u1
(25) [L85, P] str 0($vr48.p16) <- $vr167.u16, $vr472.u1
(26) [L85, P] str 0($vr47.p16) <- $vr183.u16, $vr472.u1
.../*export the 16 pairs of data elements serially to the IP core*/
(27) [L70] pfw $vr488.u1 <- $vr487.u1
(28) [L70, P] str 0($vr48.p16) <- $vr182.u16, $vr487.u1
(29) [L70, P] str 0($vr47.p16) <- $vr198.u16, $vr487.u1

(30) [L69, E69] pfw $vr489.u1 <- $vr488.u1
..... /* wait for 69 clock cycles */
(31) [L2] pfw $vr556.u1 <- $vr555.u1
(32) [L1] pfw $vr557.u1 <- $vr556.u1
(33) [L1, P] str 0($vr52.p1) <- 0, $vr556.u1 /*de-assert CE*/

(34) Node 3, [-1], {2, 1} {4}
(35) in_fft16_no_while_iTmp0:
(36) [L0] ior $vr558.u1 <- $vr557.u1, $vr559.u1
(37) [L0] ret $vr558.u1
```

predicator and input data ports at input side, and one output predicator and output data ports at output side. Fig. 16 shows a wrapped IP core embedded in a compiler-generated outer circuit. The wrapped IP core has an identical interface as that of other regular predicated instructions.

4.4 Experimental Results

We have used five Xilinx IP cores, shown in Table 2, in our experimental evaluation. *Cordic* performs a rectangular-to-polar vector translation. The input is a vector

Table 2 Synthesis results of the wrappers for five Xilinx IPs

		Cordic	10b/8b	DCT8	FFT16	RS encode
Input wrapper	Area (slice)	2	2	55	532	53
	Area (%)	0.3	5.9	6.7	24	64
	Additional cycles	1	1	1	1	1
Output wrapper	Area (slice)	2	2	426	290	9
	Area (%)	0.3	5.9	52	13	11
	Additional cycles	1	1	1	1	1
Total	Total area (slice)	663	34	817	2183	83
	Clock (MHz)	123	223	68.7	45.0	96.4
	Total cycles	23	3	23	200	20

in Cartesian coordinate and the outputs are the magnitude and the angle in a polar coordinate. *10b/8b* decodes 10-bit symbols into 8-bit bytes and an accompanying *K* bit. *DCT8* performs a one-dimensional 8-point discrete cosine transform. *FFT16* is the IP core shown in Fig. 14. *RS encode* is a (15, 13) Reed-Solomom encoder. It has a 13-symbol code block and outputs 15 symbols, including 13 original data symbols followed by two check symbols. In Table 2, *Total Area* is the total circuit including the input and output interfaces and the IP core itself. *Area (slice)* and *Area (%)* are the area utilization of the wrappers in number of slices and in percentage with respect to the entire circuit, respectively. *Additional cycles* is the number of extra clock cycles after the addition of the wrappers. *Total Cycles* is the total number of clock cycles to compute one set of input data. The input data size in *DCT8* is 8-bit while its output data size is 19-bit. *RS encode*'s input and output data sizes are 4-bit. Both *Cordic* and *FFT16*'s input and output sizes are 16-bit. The target architecture is the Xilinx Virtex-II XC2V8000-5 FPGA having 46,592 slices.

Cordic has only two inputs and two outputs and a simple handshaking protocol. *10b/8b* has an 8-bit input and a 1-bit special character indicator (the *K* bit). *DCT8*'s input wrapper latches all eight 16-bit input data. These are fed serially into the IP core. The wrapper asserts the *new_data* signal to be high during the data transmission and de-asserts it right after the transmission, following the timing requirement of the *DCT8* IP core. The output wrapper monitors the *output_ready* signal from the core and starts receiving the eight serial output data elements once it is asserted high. On the next clock cycle after all the eight output elements have been collected, the wrapper exports them all in parallel. *FFT16* requires similar serial to parallel and parallel to serial conversions, except that the IP imports and exports data in pairs, one real component and one imaginary component. *FFT16*'s input timing is different in the way that *start* and *ce* (clock enable) have certain cycle-level specifications described in the previous section. The generated interface meets all those timing requirements. The *FFT16* core's overflow output pin, *OVFLO*, is duplicated and exported by the wrapper to the outside datapath for further use.

In *RS_encode*'s output, the first 12 data elements are the data symbols that were fed to the IP. From the point of view of the outside datapath, these data are known and do not necessarily need to be recovered from the IP core again, and only the two check symbols, which follow the first 12 data elements, are needed. The *RS_encode* IP core

uses output signal *info* to indicate the present of the check symbols. The generated wrapper monitors *info*'s de-assertion and latches the check symbols in an appropriate timing.

These five examples illustrate CIRRF's capability to describe various timing protocols of IP cores. ROCCC wraps these IPs so that they have unified outside interface. The execution time overhead at both the input side and output side for these five examples is one clock cycle. The area of wrappers accounts for 2–64% of the corresponding wrapped cores. Most of the wrappers area cost comes from the registers used to do serial to parallel and parallel to serial conversion. Compared to the capacity of modern FPGAs, this overhead is quite small.

5 Related Work

We have grouped the related work discussions in two sections addressing the intermediate representation and the IP core wrapping respectively.

5.1 Intermediate Representations

Several projects have worked on reconfigurable compilers. These projects either target regular configurable devices and generate HDLs, or target special, coarse-grained, configurable architectures.

The Streams-C compiler [6] relies on the CSP model for communication between processes, both hardware and software, and can meet relatively high-density control requirements. Streams-C has three distinct objects—processes, streams and signals—in the user-input abstraction. Abstract Syntax Tree (AST) is used to partition a process into the datapath, encompassing basic blocks and pipeline blocks, and control flow. A state machine is generated for the control flow in the AST. User-defined input or output streams form the interfaces with memories.

Trident [7] uses LLVM (Low Level Virtual Machine [8]) as a C/C++ front-end to produce low-level object code. The low-level object code is transformed into a predicated IR.

SA-C's [9] input is a single-assignment high-level synthesizable language. The SA-C compiler translates loops into a data-dependence and control-flow (DDCF) graph. A DDCF graph is flattened into a token-driven data-flow graph. The DFG is eventually translated into an abstract hardware architecture graph (AHA), which includes timing information.

The customizable hardware compiler in [10] takes the Cobble language as source and produces the target language, Pebble. Cobble is based on a subset of C with extensions for synchronous parallel threads and channels for synchronous communication between them. Cobble is a timed language since it has timing semantics. Pebble is a variant of structural VHDL. Similar to VHDL, Pebble has GENERATE statements, which provide conditional compilation and recursive definition. Using the Visitor design pattern, the compiler's AST can be extended in terms of input languages, custom compile schemes and transformations.

DEFACTO [11, 12] system takes C as input and generates behavioral VHDL code. The behavioral VHDL code is then synthesized by the Synopsys Behavioral Compiler or the Monet behavioral compiler from Mentor Graphics. DEFACTO is built on SUIF.

SPARK [13] is another C to VHDL compiler. Its optimizations include code motion, variable renaming, etc. The transformations implemented in SPARK reduce the number of states in the controller FSM, and the number of cycles in the longest path. SPARK encapsulates basic blocks into Hierarchical Task Graphs (HTGs).

The GARP [14] compiler is designed for the GARP reconfigurable architecture, and generates GARP configuration file. The GARP compiler forms a *hyperblock* in a DFG by joining all frequent-executed basic blocks of the loop body.

Pegasus [15] is the IR of the CASH compiler. CASH generates data-flow machines implemented as asynchronous circuits. Pegasus decomposes a Control Flow Graph (CFG) into hyperblocks, and hyperblocks are connected by merge and other specialized nodes.

5.2 IP Wrapping

Substantial amounts of effort have been devoted on standardizing or interfacing pre-designed IP cores. Companies and organizations tried to define IP bus standards. For example, VSIA [16] specifies interface standards that allow IP cores to fit into virtual sockets. Cores are designed using a standard internal interface and wrappers have to be provided to retarget cores into other buses.

Several projects focus on bus wrapping that connects IP cores with microprocessors. Glue logic is generated in [17] to connect processors to peripheral devices and hardware co-processors. A prefetching technique is introduced in [18] to improve bus wrapper's performance.

A Meta-RTL is proposed in [19] that raises the abstraction level and reuse IPs by extending traditional HDLs. Meta-programming [20] is a proposed customization model for IP wrapping using UML class diagrams.

Trident [7] is a compiler framework for floating point algorithms. The floating point units are pre-designed IP units with known pipeline delay. Users can select floating point units from a VHDL library.

In [21] the authors describe a system level approach for interfacing IP blocks generated by the behavioral synthesis tool itself. The I/O pins and timing information is fixed and known by the tool. This information, however, is not visible at the C level and the user cannot modify it.

6 Conclusion

We have presented CIRRF, an intermediate representation for compiling high-level languages to reconfigurable fabrics. CIRRF is designed to support spatial computations as opposed to temporal computations as in traditional compiler IRs. Features that distinguish CIRRF from traditional compiler IRs include (1) Support for declaring and accessing on-chip storage, (2) Preserving loop semantics and parameters in the IR. In addition, we have designed CIRRF to support the seamless import of IP cores into

original C source codes thereby allowing the user to re-use a very large amount of pre-designed IP cores.

We have described two levels of CIRRF: Hi-CIRRF and Lo-CIRRF. Hi-CIRRF consists of C code that has been augmented with macros that support the spatial computing model. The macros are used to record information associated with buffers, pipelining, look-up tables and special operations etc. Lo-CIRRF decomposes conventional CFGs into parallel, do-all loop nodes, and sequential, non-do-all nodes. The loop body instructions of a do-all loop are placed into execution levels. Each execution level is an instantiation of one iteration at different execution phases. One or multiple consecutive execution levels are assigned into one pipeline stage. Lo-CIRRF provides a platform for the compiler to aggressively pipeline do-all loops. Non-do-all nodes are predicated in Lo-CIRRF and predicates are passed within and between nodes. The CFG is therefore transformed into a DFG. We have shown, through case studies, how CIRRF models the application examples and provides a solid foundation for the compiler to generate efficient hardware.

We extended CIRRF to support the automated generation of IP core wrapper. As the input to the ROCCC system, IP wrappers are written in high-level timed C by the user. Clock cycle delays are described as function calls and users do not have to implement any cycle-level details in the input abstraction. CIRRF records the IP wrapper as a timed CFG in the IR. Constrained by the delay macros, ROCCC converts the wrapper from control flow graph to data flow graph. The compiler schedules pipelined instructions using predication. Wrapped IP cores have unified interface compared with the outer predicated circuit that also generated by ROCCC.

The wrappers of the IP core examples meet the various timing protocol requirements, and unify the IP core interface with the outer compiler-generated circuit. The results show that the execution time and area overhead are reasonable low.

References

1. Aigner, G., Diwan, A., Heine, D.L., Lam, M.S., Moore, D.L., Murphy, B.R., Sapuntzakis, C.: An Overview of the SUIF2 Compiler Infrastructure. Computer Systems Laboratory, Stanford University
2. Smith, M.D., Holloway, G.: An Introduction to Machine SUIF and its Portable Libraries for Analysis and Optimization. Division of Engineering and Applied Sciences, Harvard University
3. Guo, Z., Buyukkurt, B., Najjar, W.: Input data reuse in compiling window operations onto reconfigurable hardware. In: ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems, pp. 249–256. ACM Press, New York, NY, June 2004.
4. Guo, Z., Buyukkurt, B., Najjar, W., Vissers, K.: Optimized generation of data-path from C codes for FPGAs. In: ACM/IEEE Conference on Design, Automation and Test in Europe, pp. 112–117. IEEE Computer Society, Washington, DC, February 2005
5. Guo, Z., Najjar, W., Buyukkurt, B.: Efficient hardware code generation for fpgas. *ACM Trans. Arch. Code Optim.* **5**(1), 1–26 (2008)
6. Gokhale, M.B., Stone, J.M., Arnold, J., Lalinowski, M.: Stream-oriented FPGA computing in the streams-c high level language. In: IEEE Symposium on FPGAs for Custom Computing Machines (FCCM 2000) (2000)
7. Tripp, J., Peterson, K., Ahrens, C., Poznanovic, J., Gokhale, M.: Trident: an FPGA compiler framework for floating-point algorithms. In: FPL '05: Proceedings of the 2005 Conference on Field Programmable Logic and Applications (2005)
8. LLVM: The LLVM Compiler Infrastructure. Information about the LLVM Project is available at the project's web site, www.llvm.org

9. Najjar, W.A., Böhm, A.P.W., Draper, B.A., Hammes, J., Rinker, R., Beveridge, R., Chawathe, M., Ross, C.: From algorithms to hardware—a high-level language abstraction for reconfigurable computing. *IEEE Comput.* **36**(8), 63–69 (2003)
10. Todman, T., de Coutinho, J.G.F., Luk, W.: Customisable hardware compilation. *Supercomputing* **32**(2), 119–137 (2005)
11. Diniz, P.C., Hall, M.W., Park, J., So, B., Ziegler, H.E.: Bridging the gap between compilation and synthesis in the DEFACTO system. In: Dietz, H.G. (ed.) *LCPC*, Volume 2624 of *Lecture Notes in Computer Science*, pp. 52–70. Springer (2001)
12. So, B., Hall, M.W., Diniz, P.C.: A compiler approach to fast hardware design space exploration in FPGA-based systems. In: *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation, PLDI*, pp. 165–176 (2002)
13. Gupta, S., Dutt, N.D., Gupta, R.K., Nicolau, A.: Spark: a high-level synthesis framework for applying parallelizing compiler transformations. In: *VLSI '03: Proceeding of the 2003 International Conference on VLSI Design* (2003)
14. Callahan, T.J., Hauser, J.R., Wawrzynek, J.: The Garp architecture and C compiler. *IEEE Comput.* **33**, 62–69 (2002)
15. Budiu, M., Goldstein, S.C.: Pegasus: an efficient intermediate representation. Technical report, Carnegie Mellon University (2002)
16. VSIA.: Information about Virtual Socket Interface Association (VSIA) is available at <http://www.vsi.org/>
17. Chou, P., Ortega, G., Borriello, G.: Interface co-synthesis techniques for embedded systems. In: *ICCAD '95: Proceedings of the 1995 International Conference on Computer Aided Design* (1995)
18. Lysecky, R., Vahid, F.: Prefetching for improved bus wrapper performance in cores. *ACM Trans. Des. Autom. Electron. Syst.* **7**, 58–90 (2002)
19. Zhu, J.: Raising the abstraction level of RTL design. In: *DATE '01: Proceedings of the Conference on Design, Automation and Test in Europe*. IEEE Computer Society (2001)
20. Stuijkys, V., Damasevicius, R.: Soft IP customisation model based on metaprogramming techniques. *Informatika* **15**, 111–126 (2004)
21. Mukherjee, R., Jones, A., Banerjee, P.: System level synthesis of multiple ip blocks in the behavioral synthesis tool. In: *PDCS '03: Proceeding of the 2003 International Conference on Parallel and Distributed Computing and Systems*, Marina Del Rey, CA (2003)