

# Using Global Code Motions to Improve the Quality of Results for High-Level Synthesis

Sumit Gupta, *Member, IEEE*, Nick Savoiu, *Student Member, IEEE*, Nikil Dutt, *Senior Member, IEEE*,  
Rajesh Gupta, *Senior Member, IEEE*, and Alex Nicolau, *Member, IEEE*

**Abstract**—The quality of synthesis results for most high level synthesis approaches is strongly affected by the choice of control flow (through conditions and loops) in the input description. This leads to a need for high-level and compiler transformations that overcome the effects of programming style on the quality of generated circuits. To address this issue, we have developed a set of speculative code motion transformations that enable movement of operations through, beyond, and into conditionals with the objective of maximizing performance. We have implemented these code transformations, along with supporting code motion techniques and variable renaming techniques, in a high-level synthesis research framework called *Spark*. *Spark* takes a behavioral description in ANSI-C as input and generates synthesizable register-transfer level VHDL. We present results for experiments on designs derived from three real-life multimedia and image processing applications, namely, the MPEG-1, MPEG-2 and GIMP applications. We find that the speculative code motions lead to reductions between 36 % to 59 % in the number of states in the finite state machine (controller complexity) and the cycles on the longest path (performance) compared to the case when only non-speculative code motions are employed. Also, logic synthesis results show fairly constant critical path lengths (clock period) and a marginal increase in area.

**Index Terms**—High-level synthesis, parallelizing compilers, code motions, speculation, embedded systems.

## I. INTRODUCTION

Recent years have seen the widespread acceptance and use of language-level modeling (such as VHDL and Verilog) of digital designs. Increasingly, the typical design process starts with design entry in a hardware description language at the register-transfer level, followed by logic synthesis. Furthermore, with the advent of systems-on-a-chip, system level behavioral modeling in high level languages is being used for initial system specification and analysis. All these factors have led to a renewed interest in high level synthesis from behavioral descriptions [1], [2], [3], [4], [5], [6], [7], [8].

However, current synthesis efforts have several limitations: synthesizability is guaranteed on a small, constrained subset of the input language and the language level optimizations are few and their effects on final circuit area and speed are not well understood. The quality of synthesis results (in terms of circuit delay and area) is adversely affected by the presence

of conditionals and loops. Designers are often given minimal control over the transformations that effect these results. All these factors continue to limit the acceptance of high-level synthesis tools among designers.

To alleviate the problem of poor synthesis results in the presence of complex control flow in designs, there is a need for high-level and compiler transformations that can optimize the synthesis results irrespective of the choice of control flow in the input description. To address this issue, several scheduling algorithms have been proposed that employ beyond-basic-block code motion techniques such as speculation to extract the inherent parallelism in designs and increase resource utilization [4], [5], [6], [7].

Generally, speculation refers to the unconditional execution of operations that were originally supposed to have executed conditionally. However, we found that there are situations when there is a need to move operations *into* conditionals. This may be done by *reverse speculation*, where operations before conditionals are duplicated into *subsequent* conditional blocks and hence, executed conditionally, or this may be done by *conditional speculation*, wherein an operation from after the conditional block is duplicated *up* into *preceding* conditional branches and executed conditionally. Another code motion technique we developed, called *early condition execution*, evaluates conditional checks as soon as their data dependencies are satisfied. In this way, all the operations in the branches of the conditional are ready to be scheduled immediately.

A number of similar code transformations have been proposed for compilers as well. Whereas compilers often pursue maximum parallelization by applying speculative code motions, in high-level synthesis, such code transformations have to be selected and guided based on their effects on the control, interconnect and area costs. An important contribution of our work is in heuristics that select the code transformations so as to improve the overall synthesis results. In some cases, our heuristics actually end up increasing the number of operations by duplicating them into conditional blocks.

We thus present a priority-based global list scheduling heuristic that directs these code motion transformations and obtains significant reductions in schedule lengths and controller complexity. Area overheads are kept in check by using a resource binding technique that minimizes interconnect [9]. We implemented the code motion transformations, the scheduling heuristic, and control and binding passes in a modular and extensible high-level synthesis framework called *Spark* [10]. We also modified several parallelizing compiler transformations [11], [12] for high-level synthesis and imple-

S. Gupta, N. Savoiu, N. Dutt and A. Nicolau are with the Center for Embedded Computer Systems at the University of California, Irvine, Irvine, CA 92697, email: {sumitg, savoiu, dutt, nicolau}@cecs.uci.edu

R. Gupta is with the Department of Computer Science and Engineering at the University of California, San Diego, La Jolla, CA 92093, email: rgupta@ucsd.edu.

mented them within our framework. *Spark* produces synthesizable register-transfer level (RTL) VHDL and thus, enables evaluation of the effects of several coarse and fine-grain optimizations on logic synthesis results. The input language for *Spark* is ANSI-C, currently with the restrictions of no pointers, no unstructured jumps (gotos) and no function recursion.

The rest of this paper is organized as follows: the next section reviews previous related work. In Section III, we present the intermediate representation used by *Spark* and in Section IV, we describe a set of speculative code motions that are useful in high-level synthesis. Next, we present the *Spark* framework and in Section VI, we describe the scheduling heuristic. In Section VII, we present scheduling and logic synthesis results for experiments using the code motions.

## II. RELATED WORK

Early high-level synthesis work focused on data-flow designs and applied optimizations such as algebraic transformations, retiming and code motions across multiplexers for improved synthesis results [13], [14], [15], [16]. Recent work has demonstrated the effectiveness of speculation in improving schedule lengths for designs with control flow. CVLS [4] uses condition vectors to improve resource sharing among mutually exclusive operations. Radivojevic et al. [17] and Haynal [6] present exact symbolic and automata-based formulation for designs with control flow. The “Waveschedule” approach [18] employs speculative execution to minimize the expected number of cycles in the schedule of a design. Santos et al. [19] and Rim et al. [20] support generalized code motions during scheduling of designs with control flow.

A range of code transformation techniques similar to those presented in our work have also been previously developed for high-level language compilers (especially parallelizing compilers) [21], [22], [23]. Although the basic transformations (e.g. dead code elimination, copy propagation) can be used in synthesis as well, other transformations need to be modified for synthesis by incorporating ideas of mutual exclusivity of operations, resource sharing and hardware cost models.

The contributions of this work include: (a) three code motion transformations derived from speculative execution techniques that are specifically targeted for high-level synthesis, (b) a heuristic approach to drive the application of these transformations and (c) a framework that provides a toolbox of code transformations and supporting compiler transformations. This enables the designer to apply heuristics to drive selection and control of individual transformations under realistic cost models for high-level synthesis. The synthesis framework provides a path from an input behavioral description down to RTL code that is synthesizable by logic synthesis tools.

## III. DESIGN DESCRIPTION MODELING

The *Spark* synthesis framework accepts a behavioral description in ANSI-C as input. However, we do not support pointers, unstructured jumps (gotos) and recursive functions for synthesis. The input C code is a sequential description of statements. Statements may be operation expressions, conditional constructs (if-then-else, switch-case), and loop constructs (for, while, do-while loops).

Traditionally, control-data flow graphs (CDFGs) [15], [24] have been the primary model for capturing design descriptions for high-level synthesis. CDFGs consist of operation and control nodes with edges for both data flow and control flow. CDFGs work very well for traditional scheduling and binding techniques. However, we found the abstraction level offered by CDFGs is not enough to enable the range of coarse-grain and fine-grain compiler transformations that we proposed. In particular, loop and conditional structures are not maintained, making it difficult to apply coarse-grain optimizations.

To enable the range of optimizations explored by our work, we use a hierarchical control-flow representation called *Hierarchical Task Graphs* (HTGs) [11], [25] that maintains the control and loop constructs in the design such as if-then-else blocks, for-loops and while-loops.

Of course, several other representation models have been proposed for high-level synthesis [26], [27], [28], [29]. However, HTGs are a convenient representation for designs with considerable control constructs since they maintain a structured view of the design. In fact, we maintain HTGs in conjunction with control flow graphs (CFGs) and data flow graphs (DFGs). Thus, whereas CFGs (and CDFGs) are efficient for traversing the basic blocks in a design, HTGs enable higher order manipulation – for example, they enable coarse-grain code restructuring (such as that done by loop transformations [12]) and also provide an efficient way to move operations across large pieces of code [11] (see Section V-A).

In the next two subsections, we present the data flow graphs, control flow graphs, and the hierarchical task graphs maintained by our framework and how they all tie in together.

### A. Data Flow and Control Flow Graphs

Data dependencies between operations create a partial ordering between the operations. If an operation  $op_j$  reads the result of another operation  $op_i$ , then a *flow dependency* is said to exist between operations  $op_i$  and  $op_j$ . Hence,  $op_j$  can start execution only after  $op_i$  has finished execution. We can define a data flow graph that captures flow dependencies as follows:

**Definition III.1** A *data flow graph* (DFG) is a directed acyclic graph  $G_{DFG}(V_{DFG}, E_{DFG})$ , where the vertices  $V_{DFG} = \{op_i \mid i = 1, \dots, n_{ops}\}$  are the operations in the design and the edges  $E_{DFG}$  represent the flow data dependencies between operations. A directed edge  $e_{ij} = (op_i, op_j)$ , where  $op_i, op_j \in V_{DFG}$ , exists in  $E_{DFG}$  if data produced by operation  $op_i$  is read by operation  $op_j$ .

A sequence of statements or operations from the input description with no conditionals or loops between them is aggregated into a *basic block*. Whereas the input “C” description consists only of operations that execute sequentially, the high-level synthesis scheduler can schedule operations to execute concurrently. We aggregate operations that execute concurrently into *scheduling steps* within basic blocks. These scheduling steps correspond to control steps in high-level synthesis [15] and to VLIW instructions in compilers [25].

This can be formally stated as follows:

**Definition III.2** A *scheduling step*  $step_j = \{op_k | k = 1, 2, \dots, n\}$  is an aggregation of  $n$  operations that execute concurrently in the same cycle or time step.

**Definition III.3** A *basic block*  $bb_i = \{step_j | j = 1, 2, \dots, n_{steps}\}$  is a maximal sequence of  $n_{steps}$  consecutive scheduling steps, where the flow of control enters at the beginning and leaves at the end without halting or possibility of branching except at the end.

The presence of conditional and loop constructs in a design description introduces the notion of *control flow paths*. Control flow can branch into multiple control paths at *fork* (or branch) basic blocks and merge back into a single control flow path at *join* (or merge) basic blocks. We define a control flow graph that captures this control flow information as follows:

**Definition III.4** A *control flow graph* (CFG) is a directed graph  $G_{CFG}(V_{CFG}, E_{CFG})$ , where the vertices  $V_{CFG} = \{bb_i | i = 1, 2, \dots, n_{bbs}\}$  are the basic blocks in the design and the edges  $E_{CFG}$  represent the flow of control between the basic blocks. A directed edge  $e_{ij} = (bb_i, bb_j)$ , where  $bb_i, bb_j \in V_{CFG}$ , exists in  $E_{CFG}$  if it signifies that  $bb_j$  executes after  $bb_i$  has finished execution.  $bb_i$  is denoted as a predecessor of  $bb_j$  and  $bb_j$  as a successor of  $bb_i$ . There exists a unique initial basic block  $bb_0 \in V_{CFG}$  from which all paths in  $G_{CFG}$  originate;  $FirstBB(G_{CFG})$  returns  $bb_0$ .

The mapping of operations in the DFG to basic blocks in the CFG is given by the following definition:

**Definition III.5** There exists a many-to-one mapping of operations in the data flow graph  $G_{DFG}(V_{DFG}, E_{DFG})$  to the basic blocks in the control flow graph  $G_{CFG}(V_{CFG}, E_{CFG})$  given by  $BB_{Ops}: V_{DFG} \mapsto V_{CFG}$ . Thus,  $\forall op_i \in V_{DFG}$ ,  $BB_{Ops}(op_i)$  gives the basic block  $bb_j \in V_{CFG}$  that operation  $op_i$  is mapped to,

### B. HTGs: A Model for Control-Intensive Designs

Whereas CFGs are useful for maintaining the flow of control between basic blocks, we employ hierarchical task graphs to maintain structure of the design description. We define a hierarchical task graph as follows:

**Definition III.6** A *hierarchical task graph* is a hierarchy of directed acyclic graphs  $G_{HTG}(V_{HTG}, E_{HTG})$ , where the vertices  $V_{HTG} = \{htg_i | i = 1, 2, \dots, n_{htgs}\}$  can be one of three types  $Type_{HTG} = \{SN, CN, LN\}$ , corresponding to single, compound, and loop nodes:

- 1) **Single nodes** represent nodes that have no sub-nodes and are used to encapsulate basic blocks.
- 2) **Compound nodes** are recursively defined as HTGs, i.e., they contain other HTG nodes. They are used to represent structures like if-then-else blocks, switch-case blocks or a series of HTGs.
- 3) **Loop nodes** are used to represent the various types of loops (for, while-do, do-while). Loop nodes consist of a loop head and a loop tail that are single nodes and a loop body that is a compound node.

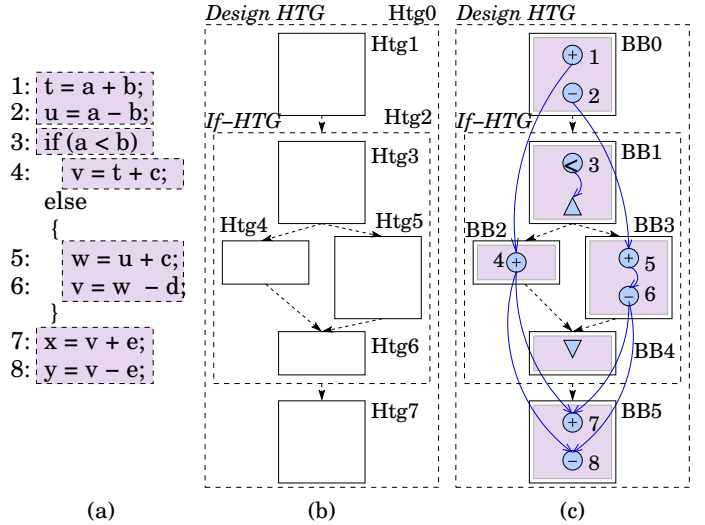


Fig. 1. (a) Example C description. (b) Corresponding Hierarchical Task Graph (HTG) representation. (c) HTG representation with the control and data flow graphs overlaid on top. An empty basic block,  $bb_4$ , is added to the control flow graph in the Join node of the If-HTG node.

The edge set  $E_{HTG}$  in  $G_{HTG}$  represents the flow of control between HTG nodes. An edge  $(htg_i, htg_j)$  in  $E_{HTG}$ , where  $htg_i, htg_j \in V_{HTG}$ , signifies that  $htg_j$  executes after  $htg_i$  has finished execution. Each node  $htg_i$  in  $V_{HTG}$  has two distinguished nodes,  $htg_{start}$  and  $htg_{stop}$ , belonging to  $V_{HTG}$  such that there exists a path from  $htg_{start}$  to every node in  $htg_i$  and a path from every node in  $htg_i$  to  $htg_{stop}$ .

The  $htg_{start}$  and  $htg_{stop}$  nodes for all compound and loop HTG nodes are always single nodes. The  $htg_{start}$  and  $htg_{stop}$  nodes of a loop HTG node are the loop head and loop tail respectively and those of a single node are the node itself. For the rest of this paper, we will denote the top-level HTG corresponding to a design as the **Design HTG**. The design HTG is constructed by creating a compound node corresponding to each control construct in the design.

Consider the sample “C” code in Figure 1(a). The HTG representation for this code is given in Figure 1(b). The HTG representation consists of a compound design HTG ( $htg_0$ ) that encapsulates the HTG nodes,  $htg_1$  to  $htg_7$ , and the control flow edges between the HTG nodes (shown by dashed arrows). The if-then-else control construct from the source code is encapsulated in the compound If-HTG node,  $htg_2$ . The If-HTG node consists of a single node for the conditional check, a compound node for the true/then branch, a compound node for the false/else branch and a single node for the Join node (containing only an empty basic block as explained below). The  $htg_{start}$  node for an If-HTG is the single node with the conditional check and the  $htg_{stop}$  node is the Join node: in Figure 1(b), these correspond to  $htg_3$  and  $htg_6$  respectively. Similarly,  $htg_1$  and  $htg_7$  are the  $htg_{start}$  and  $htg_{stop}$  nodes of the design HTG,  $htg_0$ .

In Figure 1(c), we show how the control flow graph (CFG) and data flow graph (DFG) can be overlaid onto the HTG graph. Basic blocks are denoted by shaded boxes within the HTG nodes ( $bb_0$  to  $bb_5$ ) and operations are denoted by circular nodes with the operator sign within (operations 1 to 8). Dashed lines denote control flow between HTG nodes and solid lines

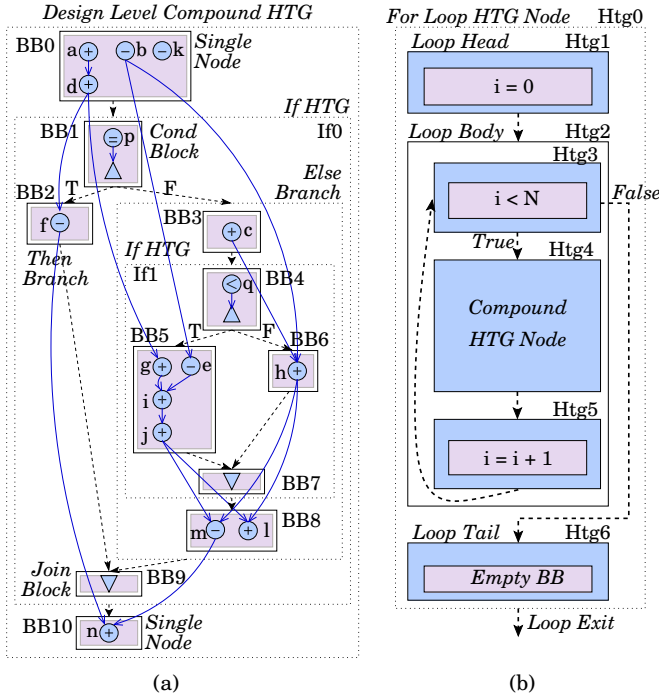


Fig. 2. (a) The hierarchical task graph (HTG) representation of the “waka” benchmark with the control and data flow graphs overlaid on top. (b) The HTG representation of a For-Loop.

denote data flow between operations. A fork in the control flow (i.e. a Boolean condition check) is denoted by a triangle ( $\Delta$ ) and a merge by an inverted triangle ( $\nabla$ ).

Note that, for clarity, the true and false branches in Figure 1(c) are shown as a single node encapsulating a basic block. In practice, the single node is then encapsulated in the compound HTG node that forms the true or false branch.

Also, during the construction of HTGs, we add empty “Join” basic blocks where multiple control flow path merge. In Figure 1(c), basic block  $bb_4$  corresponds to a *Join* basic block that is encapsulated in the single Join node  $htg_6$ . Join HTG nodes serve as the  $htg_{stop}$  nodes of compound HTG nodes and enable an easier and more structured approach to the hierarchical composition of nodes. Detailed notes on HTG construction are presented in [25].

Figure 2(a) illustrates the HTG for the synthetic benchmark “waka” [4] along with the control flow and data flow graphs. This design contains an If-HTG node, whose false/else branch contains another If-HTG node.  $bb_0$  to  $bb_{10}$  denote basic blocks.

We show a For-loop HTG in Figure 2(b). The For-loop HTG,  $htg_0$ , consists of 3 sub-nodes: (a) *Loop head* ( $htg_1$ ): consists of a single node with an optional initialization basic block, (b) *Loop body* ( $htg_2$ ): a compound HTG node containing a single HTG node ( $htg_3$ ) for the conditional check basic block and a compound HTG node ( $htg_4$ ) for the main body of the loop, and an optional single node ( $htg_5$ ) for the loop index increment basic block, and (c) *Loop tail/exit* ( $htg_6$ ): a single node with an empty basic block. There is a backward control flow edge from the end of the loop body to the conditional check single node. Maintaining the loop hierarchy allows us to treat the back edges as implicit self-loops on composite nodes [25]. Therefore, at any hierarchy level, the HTG is a directed acyclic graph.

The relationship between HTGs and CFGs is captured by the following definition:

**Definition III.7** Given  $G_{CFG}(V_{CFG}, E_{CFG})$  and  $G_{HTG}(V_{HTG}, E_{HTG})$ , there exists a one-to-one mapping of the basic blocks in  $G_{CFG}$  to the single nodes in  $G_{HTG}$  given by:  $Htg_{BB} : V_{CFG} \mapsto V_{HTG}$ . Thus,  $\forall bb_j \in V_{CFG}$  there exists a  $htg_i \in V_{HTG}$  such that  $Type_{HTG}(htg_i) = SN$ .

We capture the input description using a design graph that is defined as follows:

**Definition III.8** A *design graph*,  $\mathcal{DG}(V_D, E_D, BB_{Ops}, Htg_{BB})$ , is a layered graph<sup>1</sup>, where the vertex set  $V_D = \{V_{DFG} \cup V_{CFG} \cup V_{HTG}\}$  consists of three layers of nodes corresponding to the nodes of the graphs  $G_{DFG}$ ,  $G_{CFG}$ , and  $G_{HTG}$  and the edge set  $E_D = \{E_{DFG} \cup E_{CFG} \cup E_{HTG}\}$  consists of the edges corresponding to the edges of the graphs  $G_{DFG}$ ,  $G_{CFG}$ , and  $G_{HTG}$ .  $\mathcal{DG}$  also contains the mapping  $BB_{Ops}$  between the operations in  $G_{DFG}$  and the basic blocks in  $G_{CFG}$  and the mapping  $Htg_{BB}$  between the basic blocks in  $G_{CFG}$  and the single nodes in  $G_{HTG}$ .

For clarity, in the rest of this paper, we make several simplifications in the figures used for the examples. We omit the single HTG node that encapsulates basic blocks. Control flow edges in HTG representations are shown to originate from basic blocks and terminate at basic blocks (i.e. these represent the edges from the control flow graph).

#### IV. CODE MOTIONS IN HIGH-LEVEL SYNTHESIS

Computationally expensive portions of several classes of applications – particularly multimedia and image processing applications – are characterized by the presence of a considerable number of unpredictable branches. These control constructs limit the amount of instruction-level parallelism that can be exploited from the input description [30], [31]. There are usually not enough operations available for execution to utilize all the resources in each cycle or scheduling step. Hence, there are a number of idle resources in a basic block.

A resource is said to be *idle* in a scheduling step if there is no operation scheduled to execute on that resource in that scheduling step (the converse of an idle resource is a *busy* resource). Idle resources can be utilized by moving and scheduling operations from subsequent or preceding basic blocks. The candidate operations for these *code motions* are operations whose data dependencies are satisfied, but the conditions under which they execute may not have been evaluated. One of the key enabling transformations for such type of code motions is speculation.

##### A. Using Speculation in High-Level Synthesis

*Speculative execution* or speculation refers to the execution of an operation before the branch condition that controls the operation has been evaluated<sup>2</sup>. In our approach to speculation

<sup>1</sup>A *k-layered graph* is a connected graph in which the vertices are partitioned into  $k$  sets  $L = l_1, \dots, l_k$  and edges run between the vertices of successive layers,  $l_i$  and  $l_{i-1}$ .

<sup>2</sup>*Data speculation* is another type of speculation in which an operation is executed with potentially incorrect operand values.

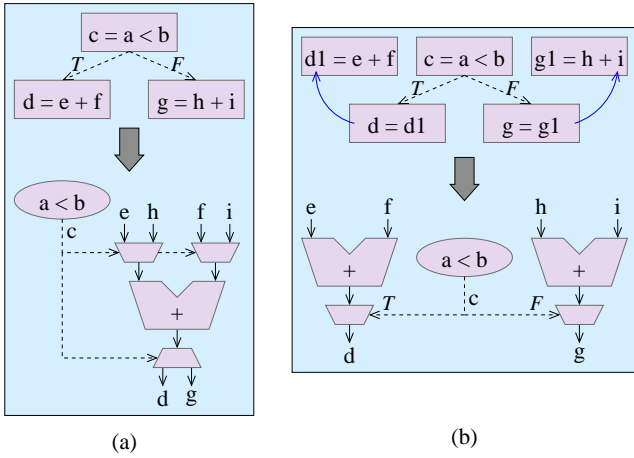


Fig. 3. Extracting the inherent parallelism in a design by speculating the addition operations. This requires an additional resource, but leads to a reduction in the longest path.

for high-level synthesis, we store the result of an speculated path ( $\langle b, g, h \rangle$ ), is placed in the basic block before the If-HTG, whereas operation  $d$ , that is on the longer dependency path ( $\langle d, e, f, h \rangle$ ), is placed in the *true* branch of the If-HTG. If we reverse speculate operation  $b$  into the conditional branches, as shown in Figure 4(b), the adder in basic block  $bb_0$  is left idle (see next paragraph for an explanation of why  $b$  is not duplicated). This enables us to speculatively execute operation  $d$  in  $bb_0$ , as shown in Figure 4(c). The dashed lines in Figure 4 demarcate the state assignments ( $S_0$  through  $S_4$ ) for the three designs. Clearly, the final design in Figure 4(c) after reverse speculation of  $b$  and speculation of  $d$  requires one state less than the original design in Figure 4(a).

Now, consider that an additional adder is available. Then, the operations within the conditional branches can be calculated *speculatively* and concurrently with the calculation of the conditional  $c$ , as shown in Figure 3(b). The corresponding hardware circuit is also shown in this figure. Based on the evaluation of the conditional, one of the results will be discarded and the other committed. It is evident from the corresponding hardware circuits in Figures 3(a) and (b) that as a result of this speculation, the longest path gets shortened from being a sequential chain of a comparison followed by an addition to being a parallel computation of the comparison and the additions.

This example also demonstrates the additional costs of speculation. Speculation requires more functional units and potentially more storage for the intermediate results. Uncontrolled aggressive speculation can also lead to *worse* results due to multiplexing and control overheads. On the other hand, judicious use of speculation can improve resource utilization.

**B. Reverse Speculation**

*Reverse speculation* refers to moving an operation  $op_i$  from its basic block  $bb_j$  into the successors of  $bb_j$ . We employ this code motion to duplicate operation  $op_i$  into the branches of an If-HTG when the If-HTG is the successor of  $bb_j$ . Reverse speculation has been referred to as *lazy execution* [20] and *duplicating down* in past literature [32].

### B. Reverse Speculation

Reverse speculation is useful in instances where an operation inside a branch of an If-HTG is on the longest path

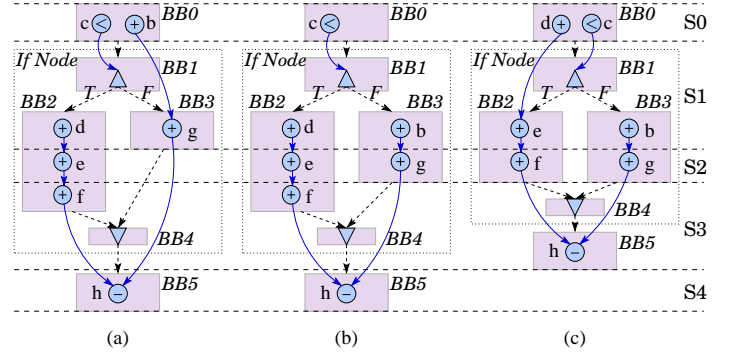


Fig. 4. Reverse Speculation: (a) An example design (b) Operation  $b$  is reverse speculated into the branch of the If-HTG that uses its result, i.e., the *false* branch. (c) Operation  $d$  is now speculated into basic block  $bb_0$ . This reduces the schedule length by one cycle.

through the design, whereas an operation before the If-HTG is not. We demonstrate this by an example in Figure 4(a). In this design, operation  $b$ , that is on the shorter dependency path ( $\langle b, g, h \rangle$ ), is placed in the basic block before the If-HTG, whereas operation  $d$ , that is on the longer dependency path ( $\langle d, e, f, h \rangle$ ), is placed in the *true* branch of the If-HTG. If we reverse speculate operation  $b$  into the conditional branches, as shown in Figure 4(b), the adder in basic block  $bb_0$  is left idle (see next paragraph for an explanation of why  $b$  is not duplicated). This enables us to speculatively execute operation  $d$  in  $bb_0$ , as shown in Figure 4(c). The dashed lines in Figure 4 demarcate the state assignments ( $S_0$  through  $S_4$ ) for the three designs. Clearly, the final design in Figure 4(c) after reverse speculation of  $b$  and speculation of  $d$  requires one state less than the original design in Figure 4(a).

Note that, while applying reverse speculation in the example above, a data dependency analysis determines that the result of operation  $b$  is used only in the *false* branch of the If-HTG. Hence, instead of duplicating  $b$  into both branches, we move  $b$  only into the false branch of the If-HTG, as shown in Figure 4(b). In the general case, reverse speculation leads to duplication of the operation into both the branches of an If-HTG. It is also important to make a distinction between moving operations into a later scheduling step and the downward operation *duplication* done by reverse speculation. When an operation encounters a fork node while being moved down, it has to be duplicated into all the control paths that lead out of the fork node (unless its result is not needed in one of the branches).

### C. Early Condition Execution

We employ reverse speculation by using another novel transformation called *early condition execution*. This transformation attempts to schedule operations such that the conditional check can be evaluated or scheduled as soon as possible. Any operations before the conditional check that are unscheduled are moved into the branches of the If-HTG by reverse speculation. Evaluating a conditional check early using early condition execution resolves the control dependency for operations within branches of the If-HTG. These operations are, thereby, available for scheduling sooner.

Early condition execution is demonstrated by an example in Figure 5(a). In this example, comparison operation  $c$  computes

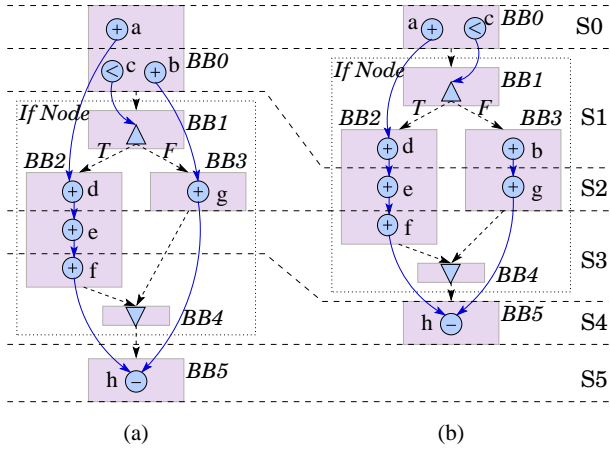


Fig. 5. Early condition execution: (a) Original design (b) Comparison operation  $c$  is scheduled as soon as possible to evaluate the conditional check early. All unscheduled operations before the conditional checks are reverse speculated into the branches of the If-HTG.




Fig. 6. (a) HTG representation of an example. (b) Operations  $x$  and  $y$  are speculated leaving resources idle in the conditional branches. (c) Operation  $z$  is *conditionally speculated* into the conditional branches ( $bb_1$  and  $bb_2$ ).

a conditional that is checked in basic block  $bb_1$  (the Boolean conditional check is denoted by a triangle). We can schedule this comparison operation concurrently with operation  $a$  in state  $S_0$  in basic block  $bb_0$ , as shown in Figure 5(b). Now, the conditional check in basic block  $bb_1$  can be executed “early” in state  $S_1$ . However, operation  $b$  in basic block  $bb_0$  has not been scheduled as yet. Therefore, this operation is reverse speculated into basic block  $bb_3$  (and not into  $bb_2$  since its result is used only in  $bb_3$ ). These code motions lead to an overall shorter schedule length, as shown by the state assignments in Figures 5(a) and 5(b).

#### D. Conditional Speculation

Often design descriptions have instances where there are idle resources in the scheduling steps of the basic blocks that comprise the branches of an If-HTG. Speculating out of If-HTGs also leaves resources idle in the basic blocks of the conditional branches. To utilize these idle resources, we propose *duplicating* operations that lie in basic blocks after the conditional branches up into the basic blocks that comprise the conditional branches. We call this code motion *conditional speculation*. This is similar to the duplication-up code motion used in compilers and the node duplication transformation discussed by Wakabayashi et al. [4].

We demonstrate conditional speculation by the example in Figure 6(a). In this example, operations  $x$  and  $y$  both write to the variable  $a$  in the conditional branches  $bb_1$  and  $bb_2$ . Consider that this design is allocated one adder, one subtractor and

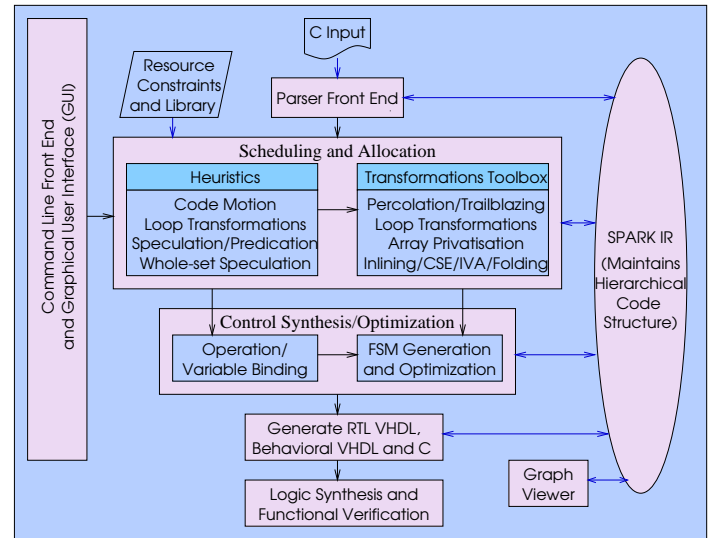


Fig. 7. The *Spark* High Level Synthesis Framework: a complete synthesis system that provides a path from an architectural description in “C” to synthesizable register-transfer level VHDL.

one comparator. Then, operations  $x$  and  $y$  can be speculatively executed as shown in Figure 6(b). The speculation of these operations leaves the resources in basic blocks  $bb_1$  and  $bb_2$  idle. Hence, the operation  $z$  that lies in basic block  $bb_4$  can be duplicated up or *conditionally speculated* into both branches of the If-HTG and scheduled on the idle adder, as illustrated in Figure 6(c). Operation  $z$  is dependent on either the result of operation  $x$  or operation  $y$  depending on how the condition evaluates (since operation  $z$  is dependent on the variable  $a$ ). Hence, the duplicated operations,  $z_1$  and  $z_2$ , directly read the results of operations  $x$  and  $y$  respectively. We have also shown the state assignments ( $S_0$ ,  $S_1$  and  $S_2$ ) for the three designs using dashed lines in Figure 6. Clearly, for this example, this set of code motions leads to a design that requires one less state to execute.

Note that, correctness issues place a number of constraints on the kind of code motions that can be done. We have omitted these for brevity, but they are detailed in [33] and are also dealt with to some extent in [25] and [11].

## V. THE *Spark* HIGH LEVEL SYNTHESIS FRAMEWORK

We have implemented the various speculative code motions in a modular and extensible high-level synthesis framework called *Spark*. *Spark* provides a range of coarse-grain and fine-grain transformations and has been designed to aid in experimenting with new transformations and heuristics that optimize the quality of synthesis results, in terms of circuit delay and area. Figure 7 provides an overview of the *Spark* framework. As shown in this figure, besides the input description, *Spark* also takes additional information as input, such as a hardware resource library, resource and timing constraints, and user directives for the various heuristics and transformations.

The design flow through the *Spark* framework is as follows: *Spark* accepts a behavioral description of a design in ANSIC, creates the intermediate representation that comprises of HTGs, CFGs and DFGs, runs a data dependency analysis pass, schedules the design, binds the resources, performs control

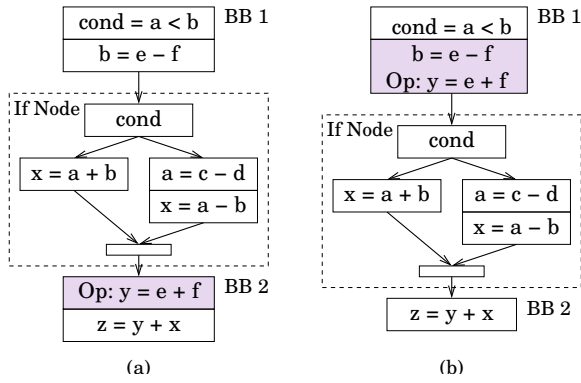


Fig. 8. Trailblazing: Operation  $op1$  is moved from basic block  $bb_2$  to basic block  $bb_1$  across the if-then-else HTG node without visiting each basic block inside the node.

synthesis, and finally generates an output in register-transfer level VHDL.

The core of the synthesis framework has a *transformations toolbox* that consists of a data dependency analysis pass, the *Trailblazing* parallelizing code motion technique [11], dynamic renaming of variables, the basic operations of loop pipelining (or software pipelining), and supporting compiler passes such as constant propagation and dead code elimination. Passes from the toolbox are called by a set of heuristics that guide how the code refinement takes place. Since the heuristics and the underlying transformations that they use are fairly independent, writing new heuristics can be as simple as making calls to the toolbox. The use of transformations can be controlled by the designer using scripts, hence, enabling experimentation with different transformations and heuristics.

A scheduling heuristic employs the techniques from the transformations toolbox to create a scheduled design (described in Section VI). The scheduling phase is followed by a *resource binding and control synthesis* phase. This phase binds operations to functional units, ties the functional units together (interconnect binding), allocates and binds storage (registers), generates the steering logic and generates the control circuits to implement the schedule. The focus of our resource binding approach is to minimize the interconnect (multiplexers and demultiplexers) between functional units and registers [9]. After binding, we generate a finite state machine controller for the scheduled and bound design.

The back-end of the *Spark* framework consists of a register-transfer level (RTL) VHDL generator. This RTL VHDL belongs to the subset of VHDL that is synthesizable by commercial logic synthesis tools and hence, the *Spark* framework integrates into the standard synthesis design flow. This completes the path from architectural design and specification in a high level language such as “C” to synthesizable RTL VHDL code and then down to the synthesized net-list.

#### A. Trailblazing: A Hierarchical Code Motion Technique

The speculative code motions employed by *Spark* are enabled by the *Trailblazing* code motion technique [11]. Trailblazing exploits information about the hierarchical structure of the design that is maintained by hierarchical task graphs (HTGs) [25] to perform efficient code motions across large




Fig. 9. (a) Example C description. (b) Corresponding DFG. (c) DFG after scheduling: ops. 1 and 3 and ops. 2 and 4 are scheduled concurrently. (d) Scheduled output code generated if only flow data dependencies are maintained. (e) Output code generated if non-flow data dependencies are also maintained. Concurrent operations are written in the same line.

pieces of code. The Trailblazing algorithm returns a set of trails or control flow paths that an operation will have to follow in order to move from one basic block to another. To move the operation along these control flow paths, the Trailblazing algorithm employs the use of the speculative code motions presented earlier in this paper. Details about the Trailblazing algorithm are discussed in [11].

To understand the hierarchical moves performed by Trailblazing, consider the example in Figure 8. In this example, we want to move the operation  $Op : y = e + f$  from basic block  $bb_2$  to basic block  $bb_1$ . While moving this operation, Trailblazing encounters the  $htg_{stop}$  node of an If-HTG node. It checks if the moving operation has any dependencies with the If-HTG node. Since, in this example, there are no dependencies, operation  $Op$  is moved across the If-HTG node to  $bb_1$  without visiting each sub-node of the If-HTG, as shown in Figure 8(b).

#### B. Eliminating data dependencies by Dynamic Renaming

There are four types of data dependencies that can exist between operations [32]: *flow* (variable read after write), *anti* (write after read), *output* (write after write) and *input* (read after read). High-level synthesis approaches have traditionally chosen to retain only flow data dependencies. However, this means that the variable names from the original description are discarded, thereby, impairing the ability to correlate the input description with the intermediate representation and the final output code generated after synthesis. This makes it difficult to visualize the effects of applying the various transformations.

To understand the need for maintaining non-flow data dependencies, consider the sample “C” description in Figure 9(a) and its corresponding data flow graph in Figure 9(b). One possible schedule is shown in Figure 9(c), where operations 1 and 3 and operations 2 and 4 are scheduled concurrently. The output code corresponding to this scheduled design, when non-flow data dependencies are not maintained, is given in Figure 9(d). In this output code, we have to create new variables that store the result of each operation in the scheduled design. Clearly, it takes some effort to correlate the operation statements in this output code with the operation statements in the input code (especially operations 2 and 4).

If we maintain non-flow data dependencies as well, we can generate the output code given in Figure 9(e). In this code, the variables that each operation writes to are maintained as per the original code. By inspecting this code, we see that concurrent execution of operations 1 and 3 requires renaming the result variable of operation 3 to  $m1$ . Operation 5 is inserted




Fig. 10. Moving an operation across another operation while eliminating (a) an anti dependency (b) an output dependency and (c) a flow dependency.

in the code to copy the new variable  $m1$  to the variable  $m$  from the original code, as shown in Figure 9(e). Thus, operation 4 can be executed concurrently with operations 2 and 5 using the new variable  $m1$  after employing dynamic variable renaming.

Although non-flow data dependencies can restrict code motions, these can often be resolved by dynamic renaming and combining [34]. Figures 10(a) to (c) demonstrate how one operation can be moved past another one while dynamically eliminating data dependencies. In Figure 10(a) the operation that writes to variable  $y$  is scheduled at an earlier time step by moving only the right hand side of the operation. The result is written to a new destination variable  $y'$  and the original operation is replaced by a copy operation from  $y'$  to  $y$ . Similarly, in Figure 10(b), an output dependency between two operations that write to the same variable  $x$ , can be resolved by creating a new destination variable  $x'$  while moving the operation and replacing the original operation with a copy operation. These copy operations can also be circumvented by a technique known as *combining*. Combining replaces the copy in the operation being moved by the variable being copied. This is demonstrated in Figure 10(c), where the operation  $z = x + 1$  is moved past the copy operation  $x = y$ . The variable  $x$  is replaced with the variable  $y$  on the right hand side of the moving operation.

## VI. PRIORITY-BASED GLOBAL LIST SCHEDULING HEURISTIC

Scheduling is the task of assigning operations to control steps or time intervals while respecting data dependencies and resource constraints (if any) [15]. For the purpose of evaluating the various code motion transformations, we have chosen a *Priority-based* global list scheduling heuristic, although the transformations presented here can be employed by other scheduling heuristics as well. Priority list scheduling works by ordering and scheduling operations based on a *priority* or cost associated with them.

Our objective is to minimize the *longest delay* through the design. Hence, the *priority* of an operation is determined by the length of the data dependency chain from the operation to the outputs of the design. The priority of an operation is calculated as the delay of the resource that the operation can be mapped to summed with the maximum of the priorities of all the operations that use its result. A priority of zero is assigned to operations that produce outputs. A priority of one (or two if the output operations are multiplications) is assigned to operations whose results the output operations depend on and so on. The priority of an operation that creates a conditional

---

```

/* Schedules the Design Graph  $DG$  */
PriorityListScheduling( $DG, \mathcal{R}, CMs$ )

```

```

1:  $Pr \leftarrow$  CalculatePriority( $G_{DFG}$ )
2:  $step_k \leftarrow$  First scheduling step in FirstBB( $G_{CFG}$ )
3: while ( $step_k \neq \phi$ ) {
4:   foreach (resource  $res \in \mathcal{R}$ ) {
5:      $\mathcal{A} \leftarrow$  GetAvailableOps( $DG, step_k, res, CMs$ )
6:     if ( $\mathcal{A} \neq \phi$ ) {
7:       Pick Operation  $op \in \mathcal{A}$  with lowest cost
8:       TrailblazeOp( $op, res, step_k, DG, CMs$ )
9:     }
10:  } /* end foreach */
11:   $step_k \leftarrow$  GetNextSchedulingStep( $G_{HTG}, G_{CFG}, step_k$ )
12: } /* end while */

```

---

Fig. 11. Priority-based List Scheduling Heuristic.

---

```

/* Gets ops for scheduling on  $res$  in  $step_k$  */

```

```

GetAvailableOps( $DG, step_k, res, CMs$ )

```

```

Returns: Available Operations List  $\mathcal{A}$ 

```

```

1:  $\mathcal{A} \leftarrow$  Unscheduled operations in  $G_{CFG}$ 
   that can be scheduled on  $res$ 
2: foreach ( $op_i \in \mathcal{A}$ ) {
3:   if (data dependencies of  $op_i$  cannot be satisfied)
4:      $\mathcal{A} \leftarrow \mathcal{A} - op_i$ 
5:   if (MoveIsNotPossible( $op_i, step_k, CMs$ ))
6:      $\mathcal{A} \leftarrow \mathcal{A} - op_i$ 
7: }
8: Calculate cost of operations in  $\mathcal{A}$ 
9: return  $\mathcal{A}$ 

```

---

Fig. 12. Algorithm to determine the list of *Available* operations.

check is assigned the maximum of the priorities of all the operations in the branches of the If-HTG.

Our scheduling heuristic is presented in Figure 11. The inputs to this heuristic are the unscheduled design graph  $DG$  and the list of resources  $\mathcal{R}$ . Additionally, we can specify a list of allowed code motions,  $CMs$ , (i.e., speculation, reverse speculation, conditional speculation et cetera). This gives us control over the code motions employed while scheduling the design by selecting and de-selecting code motions from  $CMs$ . We can thus analyze the performance-area trade-offs of individual code motions.

The heuristic starts by assigning a priority to each operation in the input description as explained above. Scheduling is done one scheduling step at a time while traversing the basic blocks in the design CFG  $G_{CFG}$ , starting with the first step in the first basic block of  $G_{CFG}$  (line 2 of the algorithm). For each scheduling step ( $step_k$ ) in the basic block, the scheduler iterates over each resource  $res$  in the resource list  $\mathcal{R}$  and calls the *GetAvailableOps* function. This function returns a list of *available* operations  $\mathcal{A}$  that may be scheduled on  $res$  in  $step_k$ . (lines 4 to 5 in the algorithm in Figure 11).

*Available operations* is a list of operations whose data dependencies are satisfied and that can be moved in the design



graph and scheduled on the given resource at the current scheduling step. Pseudo-code for collecting the list of available operations is given in Figure 12. Initially, all unscheduled operations in the design that can be executed on the resource type of  $res$  are added to the available operations list. The unscheduled operations are collected by traversing the basic blocks on the control flow paths from the current basic block being scheduled. However, this basic block traversal algorithm skips over the loop body of any loop nodes it encounters. This is because operations from within loop nodes can only be moved outside the loop body by transformations such as loop-invariant code motion and loop pipelining. Operations whose data dependencies are not satisfied and cannot be satisfied by variable renaming are also removed from this list (lines 3 and 4 in Figure 12).

The available operations algorithm then calls the *MoveIsNotPossible* function (presented in [33]). This function employs *Trailblazing* to determine all the code motions required to move the operation  $op_i$  from its current position to the scheduling step  $step_k$  under consideration. Operations that require code motions that have not been enabled in the list of allowed code motions *CMs* are removed from the available list (lines 5 and 6 of the algorithm).

Finally, the available operations algorithm assigns a cost for each remaining operation in  $\mathcal{A}$ . Currently, this cost is the negative of the operation’s priority, i.e.:

$$cost(op_i) = -Pr(op_i) \quad \forall op_i \in G_{DFG}$$

Since the scheduler picks the operation  $op$  with the *lowest* cost from the available operations list (line 7 of Figure 11), this effectively means that the operation with the highest priority in the available list is chosen. Hence, operations that are on the longest path through the design are favored for scheduling. In this way, the cost function attempts to minimize the longest delay through the design.

It is important to note that minimizing a different cost function such as average delay can be done by incorporating control flow information into the cost function. Also, if we have profiling information about which control paths are more likely to be taken, then we can give operations on those paths a higher priority than operations on less taken paths. Future work entails enhancing the cost function to include hardware (control, interconnect, and area) cost models of the code transformations.

Next, the scheduling heuristic calls *Trailblazing* to move the chosen operation  $op$  to the current scheduling step  $step_k$  (line 8 of algorithm in Figure 11). This scheduling procedure is repeated for all resources in each scheduling step as the basic blocks in the design are traversed from top to bottom by calling function *GetNextSchedulingStep* (presented in [35]). This function traverses the design using topological ordering. Thus, when a fork node of an If-HTG is encountered, the branches of the If-HTG are traversed (scheduled) first and then, the basic blocks past the join node of the If-HTG are traversed [33].

Note that, since operations with higher priorities may be speculated into a basic block, the (lower priority) operations that were originally placed in that basic block by the designer may be left unscheduled. Either new scheduling steps are added to the current basic block to schedule them or if reverse

---

*AllowConditionalSpec*( $op, BBList, bb_{curr}$ )

**Returns:** True if  $op$  should be conditionally speculated

```

1: foreach ( $bb_i$  in  $BBList$ ) {
2:   if ( $bb_i$  is not scheduled)
3:     return false
4:   if ( $\text{FindIdleResInBB}(bb_i, op) = \phi$ ) {
5:     /* Check if we can insert a new step in  $bb_i$  */
6:     if ( $\text{NumSteps}(bb_i) \geq \text{NumSteps}(bb_{curr})$ )
7:       return false /* Cannot insert  $op$  in  $bb_i$  */
8:   }
9: return true

```

---

Fig. 13. Heuristic to determine whether to schedule operation  $op$  into basic block  $bb_{curr}$  by conditionally speculating (duplicating) into the basic blocks in  $BBList$ .

speculation has been enabled, then these operations are reverse speculated into the branches of the subsequent If-HTG.

### A. Scheduling Loops

Scheduling of loops is done by the same procedure outlined above. However, user-specified loop transformations such as loop unrolling are applied first. Also, the scheduler cannot move operations into or out of the loop body. This can only be done by transformations such as loop-invariant code motion or loop pipelining. Hence, the available operations algorithm does not collect unscheduled operations from inside a loop body to schedule them outside the loop body. Also, while scheduling the loop body of a loop node, available operations are collected only from within the loop body.

The *Spark* framework can schedule all types of loops, including those with unknown loop iteration bounds. Our finite state machine (FSM) is generated such that at the end of a loop body, the next state is either the first state in the loop body or the state after the loop body, depending on whether the loop condition is satisfied or not. Hence, loop bounds are not required for generating correct, synthesizable VHDL. However, when the loop bounds are unknown, several loop transformations cannot be applied to the design and we cannot establish the number of cycles that the loop takes to execute.

### B. A Heuristic to Determine Whether to Apply Conditional Speculation

Experimental results have shown us that conditional speculation, when applied unchecked, can lead to increased schedule lengths and interconnect complexity due to operation duplication. Hence, we have developed a heuristic that determines if an operation  $op$  should be conditionally speculated (CS). This heuristic *AllowConditionalSpec*, is outlined in Figure 13. The heuristic starts with the list of basic blocks ( $BBList$ ) into which an operation  $op$  will be duplicated, if it is scheduled in basic block  $bb_{curr}$ . The heuristic returns a true result if it is possible to conditionally speculate  $op$  in each basic block in  $BBList$  and a false result otherwise.

The heuristic iterates over each basic block  $bb_i$  in the list  $BBList$  and returns a false result if any  $bb_i \in BBList$  is unscheduled. This is because without scheduling a basic block  $bb_i$  first, it is not possible to accurately determine if there is

Benchmark	# Ifs	# Lps	# BBs	# Ops	# Resources					
					+−	*	/	<<	==	[ ]
<i>Mpeg-1 pred1</i>	4	2	17	123	2	1	-	2	2	2
<i>Mpeg-1 pred2</i>	11	6	45	287	2	1	-	2	2	2
<i>Mpeg-2 dpframe</i>	18	4	61	260	4	1	-	2	2	2
<i>GIMP tiler</i>	11	2	35	150	3	1	1	2	2	2

TABLE I

CHARACTERISTICS OF THE FOUR DESIGNS USED IN OUR EXPERIMENTS, ALONG WITH THE RESOURCES ALLOCATED FOR SCHEDULING THEM.

an idle resource in  $bb_i$  on which to schedule operation  $op$ . If  $bb_i$  is scheduled, the *AllowConditionalSpec* heuristic calls the *FindIdleResInBB* (presented in [36]) to find an idle resource on which operation  $op$  can be scheduled (line 4 in Figure 13).

If there is no idle resource in  $bb_i$ , then the heuristic checks if it is possible to create a new scheduling step in  $bb_i$  in which  $op$  can be scheduled. We can add a new scheduling step in  $bb_i$  when the number of scheduling steps in  $bb_i$  is less than the number of scheduling steps in  $bb_{curr}$ . This check prevents basic block  $bb_i$  from becoming the critical path in the design with the most number of cycles (or scheduling steps) among its mutually exclusive control paths. Thus, this check ensures that the length of the longest path (path with the maximum number of scheduling steps or cycles) through the If-HTG does not increase by applying conditional speculation. The *AllowConditionalSpec* heuristic returns a true result if it either finds an idle resource or if it is possible to add a new scheduling in each  $bb_i \in BBList$  (line 9 of the algorithm).

## VII. EXPERIMENTAL SETUP AND RESULTS

We used four designs as case studies to study the impact of the proposed code motions: first on the scheduling and controller size results and then on the logic synthesis results. These four designs are derived from three moderately complex real-life applications representative of the multimedia and image processing domains. The designs are: the *pred1* and *pred2* functions from the *Prediction* block of the MPEG-1 algorithm, the *dpframe\_estimate* function from the *Motion Estimation* block of the MPEG-2 encoder algorithm [37] and the *tile* function<sup>3</sup> (with the scale function inlined) from the “tiler” transform of the GIMP image processing tool [38]. Results for more designs are presented in [33].

Table I lists the characteristics of the four designs in terms of the number of if-then-else conditional blocks (If HTGs), loops (Loop HTGs), non-empty basic blocks and the total number of operations in the input description. The number of If HTGs, Loop HTGs, and basic blocks is indicative of the control complexity of the design. All these designs have doubly nested loops. Typical run times of *Spark* to produce the results for these designs are in the range of 5 user seconds on a 1.6 Ghz Linux PC.

Table I also lists the type and quantity of the resources used for scheduling these designs: +− does add and subtract, == is a comparator, \* a multiplier, / a divider, [ ] an array address decoder and << is a shifter. The multiplier (\*) executes in 2

<sup>3</sup>Note that this floating point function has been arbitrarily converted to an integer function for the purpose of our experiments. This does not affect the nature of the data and control flow, but only the data that is processed.

Allowed Code Motions	<i>MPEG-1 pred1</i>		<i>MPEG-1 pred2</i>	
	# States	Long Path	# States	Long Path
Within BBs	71	2009	125	4555
+across HTGs	60(-14.3 %)	1937(-3.6 %)	111(-11 %)	4409(-3.2 %)
+early cond exec	61(+1.7 %)	1937(0 %)	113(-1.8 %)	4409(0 %)
+speculation	56(-8.2 %)	1862(-3.9 %)	104(-8.0 %)	4178(-5.2 %)
+cond spec	40(-28.6 %)	1091(-41 %)	74(-28.8 %)	2575(-38.4%)
Total Reduction	<b>42.9 %</b>	<b>45.7 %</b>	<b>40.8 %</b>	<b>43.5 %</b>

TABLE II

SCHEDULING RESULTS AFTER APPLYING THE SPECULATIVE CODE THE MPEG-1 *pred1* AND *pred2* FUNCTIONS.

Allowed Code Motions	<i>MPEG-2 dpframe</i>		<i>GIMP tiler</i>	
	# States	Long Path	# States	Long Path
Within BBs	77	911	69	6144
+across HTGs	72(-6.5 %)	863(-5.3 %)	66(-4.3 %)	5944(-3.3 %)
+early cond exec	71(-1.4 %)	859(-0.5 %)	63(-4.5 %)	5544(-6.7 %)
+speculation	58(-18.3 %)	607(-29.3 %)	54(-14.3 %)	4734(-14.6 %)
+cond spec	49(-15.5 %)	571(-5.9 %)	31(-42.6 %)	2534(-46.5 %)
Total Reduction	<b>36.4 %</b>	<b>37.3 %</b>	<b>55.1 %</b>	<b>58.8 %</b>

TABLE III

SCHEDULING RESULTS AFTER APPLYING THE SPECULATIVE CODE THE *dpframe\_estimate* AND *tiler* DESIGNS.

cycles and the divider (/) in 5 cycles. All other resources are single cycle.

### A. Effects on Performance and Controller Size

Tables II and III list the scheduling results for the four designs as each code motion is enabled. These scheduling results are in terms of the number of states in the finite state machine controller (denotes controller complexity) and the cycles on the longest path in the design (i.e. execution cycles). The longest path through an if-then-else construct is the number of scheduling steps through the longest branch. For loops, the longest path length of the loop body is multiplied by the number of loop iterations (loop bounds are known for all four designs).

The rows in Tables II and III present results with each code motion enabled incrementally, i.e., these signify the “allowed code motions” while determining the available operations (see Section VI) and do *not* represent an ordering of code motions. We first allow code motions only within basic blocks (first row) and then, then across hierarchical blocks as well, i.e., across entire if-then-else conditionals and loops (second row), then with reverse speculation and early condition execution also enabled (third row), then with speculation (fourth row), and finally, with conditional speculation enabled as well (fifth row). The percentage reductions of each row over the previous row are given in parentheses.

The results in Tables II and III demonstrate that enabling just the non-speculative code motions across hierarchical blocks of code (second row) leads to modest improvements: ranging from 4 % to 14 % in the number of FSM states and 3 to 5 % in the cycles on the longest path. Early condition execution (that employs reverse speculation) also leads to little or no improvements. (see third row in Tables II and III).

The largest improvements in performance (cycles) are obtained by employing speculation and conditional speculation. Speculation reduces the number of states and the number of cycles by 8 to 18 % and 3 to 29 % respectively. Conditional speculation leads to even more impressive improvements of 15

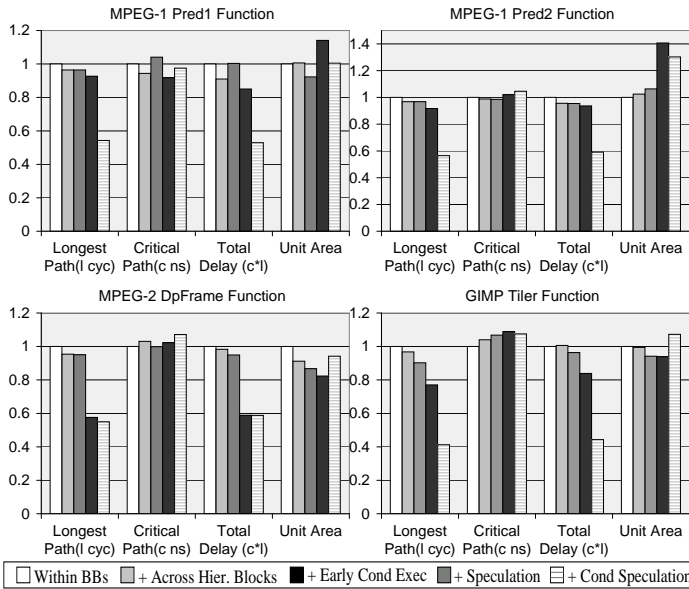


Fig. 14. Effects of Speculative Code Motions on logic synthesis results for *pred1*, *pred2*, *dpframe\_estimate*, and *tiler* designs.

to 42 % in the number of states and 5 to 46 % in the number of cycles, as shown in the fifth row of Tables II and III.

The total reductions for the two metrics with all the code motions enabled over code motions only within basic blocks are given in the last row of Tables II and III. These improvements are in the range of 36 to 55 % and 37 to 58 % in the number of states and cycles respectively. These substantial gains demonstrate that speculative code motions are essential in designs with moderate control-flow. We also find that to be truly effective all the code motions must be applied together, thus, giving the scheduler the maximum freedom to choose from among the candidate operations. Note that, when code motions only within basic blocks are enabled, our global list scheduling heuristic reduces to the classical list scheduling approaches presented for data flow only designs [15], [16].

### B. Effects on Area and Clock Period

In order to gauge the control and interconnect costs of the code motions, we synthesized the register-transfer level (RTL) VHDL generated by *Spark* after scheduling and resource binding using the Synopsys *Design Compiler* logic synthesis tool. We used the LSI-10K synthesis library for technology mapping.

The logic synthesis results are summarized in the graphs in Figure 14 for the four designs. We map three metrics in these graphs: the critical path length (in nanoseconds), the unit area (in terms of synthesis library used) and the maximum delay through the design. The *critical path length* is the length of the longest combinational path in the netlist as reported by the static timing analysis tool. The critical path length dictates the clock period of the final design. The maximum delay is the product of the longest path length (in cycles) and the critical path length (in ns) and signifies the maximum input to output latency of the design.

The bars in the graphs in Figure 14 present results for code motions within basic blocks only (first bar), for code motions across hierarchical blocks enabled as well (second bar), early condition execution enabled also (third bar),

Benchmark	# BBs	Resources	Schedule Length				
			CVLS [4]	HRA [39]	Radiv [17]	Santos [5]	<i>Spark</i> Ours
kim [39]	7	2+,1-,2==	6	7	6	6	6
parker[40]	20	2+,3-,5==	4	NA	4	4	4
waka [4]	9	1+,1-,2==	7	7	7	7	7
rotor [17]	11	2+,-2*,1[ ]	NA	NA	8	8	8

TABLE IV

COMPARISONS WITH PREVIOUS WORK USING CLASSICAL HIGH-LEVEL SYNTHESIS BENCHMARKS. NA = RESULTS ARE NOT AVAILABLE.

speculation enabled as well (fourth bar), and finally with conditional speculation enabled as well (fifth bar). Thus, the fifth bar represents the results with all the speculative and non-speculative code motions enabled and the second bar represents the results when only the non-speculative code motions are enabled.

The graphs in Figure 14 show that the total delay is almost halved when all the code motions are enabled over when code motions only within basic blocks are allowed. The results in these graphs also show that the critical path lengths remain fairly constant as the code motions are enabled. This means that the clock period does not increase by applying these code motions. The constant critical path length, coupled with large decreases in cycles on the longest path, leads to large decreases in the total delay through the circuit.

However, code motions such as speculation and conditional speculation can lead to an increase in area, as we can see from graphs in Figure 14. This area increase is due to the increasing complexity of the interconnect (multiplexers and associated control logic) that is a product of the shorter schedule lengths produced by the speculative code motions. Shorter schedule lengths mean that resource utilization and resource sharing increases and this leads to an increase in the complexity of the multiplexers and associated control logic. This complexity increase is particularly large due to conditional speculation because it duplicates operations and thus, more operations are mapped to the same number of resources as before.

Note that critical path lengths remain fairly constant despite the increases in interconnect complexity, since these increases are counter-balanced by decreases in the controller size. We found that critical paths in our designs typically start in the control logic that generates the select signals for the multiplexers, continue through the multiplexers, the functional unit, a de-multiplexer, and finally write to an output register. Even though the code motions increase the size and complexity of the multiplexers and de-multiplexers, they also lead to fewer states in the FSM controller. Thus, the size and complexity of the controller decreases.

We also keep area and critical path overheads in check by employing an interconnect minimizing resource binding methodology [9] that aids the logic synthesis tool to optimize away a lot of the complexity of the interconnect. The area may still increase sometimes since we direct the logic synthesis tool to sacrifice area in a bid to achieve shorter critical path lengths.

### C. Comparison with Previous Work

In Table IV, we compare our scheduling results with the CVLS approach [4], the HRA approach [39], the *exact* approach presented by Radivojevic [17] and the approach pre-

sented by Santos et al. [5]. The classical high-level synthesis benchmarks used for comparisons in these works are used here as well. These are: *kim* [39], *parker* [40], *waka* [4] and *rotor* [17]. The columns present the number of basic blocks, the resources used for scheduling and the longest path length (cycles) of the schedule produced by each approach. For these small benchmarks, nearly all the approaches (including *Spark*) are able to achieve the shortest scheduling length. We are unable to compare area and control costs, since these have not been published in previous work.

### VIII. CONCLUSIONS AND FUTURE WORK

We presented a set of speculative code motions that reorder, speculate, and sometimes even increase the number of operations in a behavioral description so as to achieve higher quality of synthesis results. These code motions are essential for minimizing the effects of the choice of control flow or programming style in high level languages. We presented a scheduling heuristic that guides these code motions and improves scheduling results (schedule length and FSM states) and logic synthesis results (circuit area and delay) for moderately complex real-life designs by up to 50 % in performance and controller size, when compared to list scheduling techniques that allow code motions only within basic blocks. We implemented the code motions and heuristics in the *Spark* high-level synthesis framework that provides a platform for applying a range of coarse-grain and fine-grain code optimizations aimed at improving synthesis results. Future work entails developing more comprehensive cost models that incorporate the control and interconnect costs of the code motions.

### ACKNOWLEDGMENTS

This work is supported by the Semiconductor Research Corporation under Task I.D. 781.001. We would like to thank all the reviewers for their comments and for helping us significantly improve the quality of this publication.

### REFERENCES

- [1] K. Wakabayashi. C-based synthesis experiences with a behavior synthesizer, "Cyber". In *Design, Automation and Test in Europe*, 1999.
- [2] Celoxica. DK design suite. <http://www.celoxica.com>.
- [3] Forte Design Systems. Behavioral design suite. <http://www.forteds.com>.
- [4] K. Wakabayashi and H. Tanaka. Global scheduling independent of control dependencies based on condition vectors. In *Design Automation Conference*, 1992.
- [5] L.C.V. dos Santos. *Exploiting instruction-level parallelism: a constructive approach*. PhD thesis, Eindhoven University of Technology, 1998.
- [6] S. Haynal. *Automata-Based Symbolic Scheduling*. PhD thesis, University of California, Santa Barbara, 2000.
- [7] G. Lakshminarayana, A. Raghunathan, and N.K. Jha. Wavesched: A novel scheduling technique for control-flow intensive designs. *IEEE Transactions on CAD*, May 1999.
- [8] S. Gupta, N.D. Dutt, R.K. Gupta, and A. Nicolau. SPARK: A high-level synthesis framework for applying parallelizing compiler transformations. In *International Conference on VLSI Design*, 2003.
- [9] S. Gupta, N. Savoie, N.D. Dutt, R.K. Gupta, and A. Nicolau. Conditional speculation and its effects on performance and area for high-level synthesis. In *International Symposium on System Synthesis*, 2001.
- [10] SPARK Parallelizing High-Level Synthesis Framework website. <http://www.cecs.uci.edu/~spark>.
- [11] A. Nicolau and S. Novack. Trailblazing: A hierarchical approach to Percolation Scheduling. In *International Conference on Parallel Processing*, 1993.
- [12] S. Novack and A. Nicolau. An efficient, global resource-directed approach to exploiting instruction-level parallelism. In *Conference on Parallel Architectures and Compilation Techniques*, 1996.
- [13] M. Potkonjak and J. Rabaey. Optimizing resource utilization using transformations. *IEEE Transactions on CAD*, March 1994.
- [14] R. Walker and D. Thomas. Behavioral transformation for algorithmic level IC design. *IEEE Transactions on CAD*, October 1989.
- [15] D. D. Gajski, N. D. Dutt, A. C-H. Wu, and S. Y-L. Lin. *High-Level Synthesis: Introduction to Chip and System Design*. Kluwer Academic, 1992.
- [16] G. De Micheli. *Synthesis and Optimization of Digital Circuits*. McGraw-Hill, 1994.
- [17] I. Radivojevic and F. Brewer. A new symbolic technique for control-dependent scheduling. *IEEE Transactions on CAD*, January 1996.
- [18] G. Lakshminarayana, A. Raghunathan, and N.K. Jha. Incorporating speculative execution into scheduling of control-flow intensive behavioral descriptions. In *Design Automation Conference*, 1998.
- [19] L.C.V. dos Santos and J.A.G. Jess. A reordering technique for efficient code motion. In *Design Automation Conference*, 1999.
- [20] M. Rim, Y. Fann, and R. Jain. Global scheduling with code-motions for high-level synthesis applications. *IEEE Transactions on VLSI Systems*, September 1995.
- [21] J. Fisher. Trace scheduling: A technique for global microcode compaction. *IEEE Transactions on Computers*, July 1981.
- [22] A. Nicolau. Uniform parallelism exploitation in ordinary programs. In *International Conf. on Parallel Processing*, 1985.
- [23] K. Ebcioğlu and A. Nicolau. A global resource-constrained parallelization technique. In *3rd International Conference on Supercomputing*, 1989.
- [24] A. Orailoglu and D.D. Gajski. Flow graph representation. In *Design Automation Conference*, 1986.
- [25] M. Girkar and C.D. Polychronopoulos. Automatic extraction of functional parallelism from ordinary programs. *IEEE Trans. on Parallel & Distributed Systems*, Mar. 1992.
- [26] M. C. McFarland. The Value Trace: A database for automated digital design. Technical Report DRC-01-4-80, Carnegie-Mellon University, Design Research Center, 1978.
- [27] V. Chaiyakul, D.D. Gajski, and L. Ramachandran. Minimizing syntactic variance with assignment decision diagrams. Technical Report ICS-TR-92-34, UC Irvine, 1992.
- [28] D. Ku and G. De Micheli. Relative scheduling under timing constraints. In *Design Automation Conference*, 1990.
- [29] A.A. Kountouris and C. Wolinski. Hierarchical conditional dependency graphs as a unifying design representation in the CODESIS high-level synthesis system. In *Intl. Symposium on System Synthesis*, 2000.
- [30] D.W. Wall. Limits of instruction-level parallelism. In *International Conference on Architectural Support for Programming Languages and Operating System (ASPLOS)*, 1991.
- [31] M. S. Lam and R. P. Wilson. Limits of control flow on parallelism. In *International Symposium on Computer Architecture*, 1992.
- [32] S. S. Muchnick. *Advanced Compiler Design and Implementation*. Morgan Kaufmann, 1997.
- [33] S. Gupta. *Coordinated Coarse-Grain and Fine-Grain Optimizations for High-Level Synthesis*. PhD thesis, University of California, Irvine, 2003.
- [34] S.-M. Moon and K. Ebcioğlu. An efficient resource-constrained global scheduling technique for superscalar and VLIW processors. In *International Symposium on Microarchitecture*, 1992.
- [35] S. Gupta, N.D. Dutt, R.K. Gupta, and A. Nicolau. Coordinated parallelizing compiler optimizations and high-level synthesis. Technical Report CECS-TR-02-35, Center for Embedded Computer Systems, Univ. of California, Irvine, 2002.
- [36] S. Gupta, N.D. Dutt, R.K. Gupta, and A. Nicolau. Dynamic conditional branch balancing during the high-level synthesis of control-intensive designs. In *Design, Automation and Test Conference*, 2003.
- [37] C. Lee, M. Potkonjak, and W. H. M.-Smith. UCLA Mediabench benchmark suite. <http://cares.icsl.ucla.edu/MediaBench/>.
- [38] GNU Image Manipulation Program. <http://www.gimp.org>.
- [39] T. Kim, N. Yonezawa, J.W.S. Liu, and C.L. Liu. A scheduling algorithm for conditional resource sharing - a hierarchical reduction approach. *IEEE Transactions on CAD*, April 1994.
- [40] A.C. Parker, J. Pizarro, and M. Mlinar. MAHA: A program for datapath synthesis. In *Design Automation Conference*, 1986.