Technical Reports in Computer and Information Science Report number 2012:11

# Execution Time Minimization Based on Hardware/Software Partitioning and Speculative Prefetch

by

Adrian Lifa, Petru Eles and Zebo Peng {adrian.alin.lifa, petru.eles, zebo.peng}@liu.se July 9, 2012



Linköpings universitet

Department of Computer and Information Science Linköping University SE-581 83 Linköping, Sweden

Technical reports in Computer and Information Science ISSN: 1654-7233 Year: 2011 Report no. 11

Available online at Linköping University Electronic Press http://www.ep.liu.se/PubList/Default.aspx?SeriesID=2550

© The Author(s)

# Execution Time Minimization Based on Hardware/Software Partitioning and Speculative Prefetch

Adrian Lifa, Petru Eles and Zebo Peng Embedded Systems Laboratory Department of Computer and Information Science Linköping University, S-581 83 Linköping, Sweden {adrian.alin.lifa, petru.eles, zebo.peng}@liu.se

> Technical Report July 2012

#### Abstract

This report addresses the problem of minimizing the average execution time of an application, based on speculative FPGA configuration prefetch. Dynamically reconfigurable systems (like FPGAs) provide both the performance of hardware acceleration and the flexibility and adaptability that modern applications require. Unfortunately, one of their main drawbacks that significantly impacts performance is the high reconfiguration overhead. Configuration prefetching is one method to reduce this penalty by overlapping FPGA reconfigurations with useful computations. In order to make it effective and to avoid very high misprediction penalties, it is important to prefetch the configurations that provide the highest performance improvement, and to do this early enough to hide the reconfiguration overhead. In this report we propose a speculative approach that schedules prefetches at design time and simultaneously performs HW/SW partitioning, in order to minimize the expected execution time of an application. Our method prefetches and executes in hardware those configurations that provide the highest performance improvement. The algorithm takes into consideration profiling information (such as branch probabilities and execution time distributions), correlated with the application characteristics. We demonstrate the effectiveness of our approach compared to the previous state-of-art using extensive experiments, including real-life case studies.

## 1 Introduction

In recent years, FPGA-based reconfigurable computing systems have gained in popularity because they promise to satisfy the simultaneous needs of high performance and flexibility [11]. Modern FPGAs provide support for partial dynamic reconfiguration [18], which means that parts of the device may be reconfigured at run-time, while the other parts remain fully functional. This feature offers high flexibility, but does not come without challenges: one major impediment is the high reconfiguration overhead.

Researchers have proposed several techniques to reduce the reconfiguration overhead. Such approaches are configuration compression [7] (which tries to decrease the amount of configuration data that must be transferred to the FPGA), configuration caching [7] (which addresses the challenge to determine which configurations should be stored in an on-chip memory and which should be replaced when a reconfiguration occurs) and configuration prefetch [4, 8, 9, 10, 13] (which tries to preload future configurations, overlapping as much as possible the reconfiguration overhead with useful computation).

In this report we present a speculative approach to configuration prefetching that implicitly performs HW/SW partitioning and improves the state-of-art<sup>1</sup>. The high reconfiguration overheads make configuration prefetching a challenging task: the configurations to be prefetched should be the ones with the highest potential to provide a performance improvement and they should be predicted early enough to overlap the reconfiguration with useful computation. Therefore, the key for high performance of such systems is efficient HW/SW partitioning and intelligent prefetch of configurations.

## 2 Related Work

The authors of [3] proposed a partitioning algorithm, as well as an ILP formulation and a heuristic approach to scheduling of task graphs. In [2] the authors present an exact and a heuristic algorithm that simultaneously partitions and schedules task graphs on FPGAs. The authors of [5] proposed a CLP formulation and a heuristic to find the minimal hardware overhead and corresponding task mapping for systems with communication security constraints. The main difference compared to our work is that the above papers address the optimization problem at a task level, and for a large class of applications (e.g. those that consist of a single sequential task) using such a task-level coarse granularity is not appropriate. Instead, it is necessary to analyze the internal structure and properties of tasks.

The authors of [12] present a hybrid design-/run-time prefetch scheduling heuristic that prepares a set of schedules at design time and then chooses between them at run-time. This work uses the same task graph model as the ones before. In [7], the author proposes hybrid and dynamic prefetch heuristics that perform part or all of the scheduling computations at run-time and also require additional hardware. One major advantage of static prefetching is that, unlike the dynamic approaches mentioned above, it requires no additional hardware and it generates minimal run-time overhead. Moreover, the solutions generated are good as long as the profile and data access information known at compile time are accurate.

The authors of [16] present an approach for accelerating the error detection mechanisms. A set of path-dependent prefetches are prepared at design time, and at run-time the appropriate action is applied, corresponding to the actual path taken. The main limitation of this work is that it is customized for the prefetch of particular error detection modules.

To our knowledge, the works most closely related to our own are [13], [8] and [9]. Panainte et al. proposed both an intra-procedural [9] and an inter-procedural [10] static prefetch scheduling algorithm that minimizes the number of executed FPGA reconfigurations taking into account FPGA area placement conflicts. In order to compute the locations where hardware reconfigurations can be anticipated, they first determine the regions not shared between any two conflicting hardware modules, and then insert prefetches at the beginning of each such region. This approach is too conservative and a more aggressive speculation could hide more reconfiguration overhead. Also, profiling information (such as branch probabilities and execution time distributions) could be used to prioritize between two non-conflicting modules.

Li et al. continued the pioneering work of Hauck [4] in configuration prefetching. They compute the probabilities to reach any hardware module, based on profiling information [8]. This algorithm can be applied only after all the loops are identified and collapsed into dummy nodes. Then, the hardware modules are ranked at each basic block according to these probabilities and prefetches are issued. The main limitations of this work are that it removes all loops (which leads to loss of path information) and that it uses only probabilities to guide prefetch insertion (without taking into account execution time distributions, for example). Also, this approach was developed for FPGAs with relocation and defragmentation, and it does not account for placement conflicts between modules.

To our knowledge, the state-of-art in static configuration prefetching for partially reconfigurable FPGAs is the work of Sim et al. [13]. The authors present an algorithm that

 $<sup>^{1}</sup>$ Configuration compression and caching are not addressed here, but they are complementary techniques that can be used in conjunction with our proposed approach.



Figure 1: Architecture Model

minimizes the reconfiguration overhead for an application, taking into account FPGA area placement conflicts. Using profiling information, the approach tries to predict the execution of hardware modules by computing 'placement-aware' probabilities (PAPs). They represent the probabilities to reach a hardware module from a certain basic block without encountering any conflicting hardware module on the way. These probabilities are then used in order to generate prefetch queues to be inserted by the compiler in the control flow graph of the application. The main limitation of this work is that it uses only the 'placement-aware' probabilities to guide prefetch insertion. As we will show in this report, it is possible to generate better prefetches (and, thus, further reduce the execution time of the application) if we also take into account the execution time distributions, correlated with the reconfiguration time of each hardware module.

# 3 System Model

#### 3.1 Architecture Model

We consider the architecture model presented in Fig. 1. This is a realistic model that supports commercially available FPGAs (like, e.g., the Xilinx Virtex or Altera Stratix families). Most current reconfigurable systems consist of a host microprocessor connected, either loosely or tightly, to an FPGA (used as a coprocessor for hardware acceleration).

One common scenario for the tightly connected case is that the FPGA is partitioned into a static region, where the microprocessor itself and the reconfiguration controller reside, and a partially dynamically reconfigurable (PDR) region, where the application hardware modules can be loaded at run-time. The host CPU executes the software part of the application and is also responsible for initiating the reconfiguration of the PDR region of the FPGA. The reconfiguration controller will configure this region by loading the bitstreams from the memory, upon CPU requests. While one reconfiguration is going on, the execution of the other (non-overlapping) modules on the FPGA is not affected.

We model the PDR region as a rectangular matrix of configurable logic blocks. Each hardware module occupies a contiguous rectangular area of this matrix. Although it is possible for the hardware modules to be relocated on the PDR region of the FPGA at runtime, this operation is known to be computationally expensive [14]. Thus, similar to the assumptions of Panainte [9] and Sim [13], we also consider that the placement of the hardware modules is decided at design time and any two hardware modules that have overlapping areas are in 'placement conflict'.

## 3.2 Application Model

The main goal of our approach is to minimize the expected execution time of a program executed on the hardware platform described above. We consider a *structured* [1] program<sup>2</sup>,

 $<sup>^2 {\</sup>rm Since}$  any non-structured program is equivalent to some structured one, our assumption loses no generality.

modeled as a control flow graph (CFG)  $\mathcal{G}_{cf}(\mathcal{N}_{cf}, \mathcal{E}_{cf})$ , where each node in  $\mathcal{N}_{cf}$  corresponds to either a basic block (a straight-line sequence of instructions) or a candidate module to be executed on the FPGA, and the set of edges  $\mathcal{E}_{cf}$  corresponds to the possible flow of control within the program.  $\mathcal{G}_{cf}$  captures all potential execution paths and contains two distinguished nodes, root and sink, corresponding to the entry and the exit of the program. The function prob :  $\mathcal{E}_{cf} \to [0, 1]$  represents the probability of each edge in the CFG to be taken, and it is obtained by profiling the application. For each loop header n, iter\_prob<sub>n</sub> :  $\mathbb{N} \to [0, 1]$  represents the probability mass function of the discrete distribution of loop iterations.

We denote the set of hardware candidates with  $\mathcal{H} \subseteq \mathcal{N}_{cf}$  and any two modules that have placement conflicts with  $m_1 \bowtie m_2$ . We assume that all hardware modules in  $\mathcal{H}$  have both a hardware implementation and a corresponding software implementation. Since sometimes it might be impossible to hide enough of the reconfiguration overhead for all candidates in  $\mathcal{H}$ , our technique will try to decide at design time which are the most profitable modules to insert prefetches for (at a certain point in the CFG). Thus, for some candidates, it might be better to execute the module in software, instead of inserting a prefetch too close to its location (because waiting for the reconfiguration to finish and then executing the module on the FPGA is slower than executing it in software). This approach will implicitly generate a HW/SW partitioning of the set  $\mathcal{H}$  of hardware candidates.

The set  $\mathcal{H}$  can be determined automatically (or by the designer) and might contain, for example, the computation intensive parts of the application, identified after profiling. Please note that it is not necessary that all candidate modules from  $\mathcal{H}$  will end up on the FPGA. Our technique will try to use the limited hardware resources as efficiently as possible by choosing to prefetch the modules with the highest potential to reduce the expected execution time. Thus, together with the prefetch scheduling we also perform a HW/SW partitioning of the modules from  $\mathcal{H} \subseteq \mathcal{N}_{cf}$ .

For each node  $n \in \mathcal{N}_{cf}$  we assume that we know its software execution time<sup>3</sup>,  $sw : \mathcal{N}_{cf} \to \mathbb{R}^+$ . For each hardware candidate  $m \in \mathcal{H}$ , we also know its hardware execution time<sup>3</sup>,  $hw : \mathcal{H} \to \mathbb{R}^+$ . The function  $area : \mathcal{H} \to \mathbb{N}$ , specifies the area that hardware modules require,  $size : \mathcal{H} \to \mathbb{N} \times \mathbb{N}$  gives their size, and  $pos : \mathcal{H} \to \mathbb{N} \times \mathbb{N}$  specifies the position where they were placed on the reconfigurable region. Since for all modules in  $\mathcal{H}$  a hardware implementation and placement are known at design time, we also know the reconfiguration duration, which we assume it is given by a function  $rec : \mathcal{H} \to \mathbb{R}^+$ .

### 3.3 Reconfiguration API

We adopt the reconfiguration library described in [14], that defines an interface to the reconfiguration controller, and enables the software control of reconfiguring the FPGA. The architecture described in Sec. 3.1 supports preemption and resumption of hardware reconfigurations. The library defines the following functions to support initialization, preemption and resumption of reconfigurations:

- *load*(*m*): Non-blocking call that requests the reconfiguration controller to start or resume loading the bitstream of module *m*.
- *currently\_reconfig()*: Returns the id of the hardware module being currently reconfigured, or -1 otherwise.
- $is\_loaded(m)$ : Returns true if the hardware module m is already loaded on the FPGA, or false otherwise.
- exec(m): Blocking call that returns only after the execution of hardware module m has finished.

 $<sup>^{3}</sup>$ Instead of a single number, the execution time could be modeled as a discrete probability distribution as well, without affecting our overall approach for configuration prefetching.



Figure 2: Motivational Example

## 3.4 Middleware and Execution Model

Let us assume that at each node  $n \in \mathcal{N}_{cf}$  the hardware modules to be prefetched have been ranked at compile-time (according to some strategy) and placed in a queue (denoted loadQ). The exact hardware module to be prefetched will be determined at run-time (by the middleware, using the reconfiguration API), since it depends on the run-time conditions. If the module with the highest priority (the head of loadQ) is not yet loaded and is not being currently reconfigured, it will be loaded at that particular node. If the head of loadQ is already on FPGA, the module with the next priority that is not yet on the FPGA will be loaded, but only in case the reconfiguration controller is idle. Finally, if a reconfiguration is ongoing, it will be preempted only in case a hardware module with a priority higher than that of the module being reconfigured is found in the current list of candidates (loadQ).

At run-time, once a hardware module  $m \in \mathcal{H}$  is reached, the middleware checks whether m is already fully loaded on the FPGA, and in this case it will be executed there. Thus, previously reconfigured modules are reused<sup>4</sup>. Otherwise, if m is currently reconfiguring, the application will wait for the reconfiguration to finish and then execute the module on FPGA, but only if this generates a shorter execution time than the software execution. If none of the above are true, the software version of m will be executed.

## 4 Problem Formulation

Given an application (as described in Sec. 3.2) intended to run on the reconfigurable architecture described in Sec. 3.1, our goal is to determine, at each node  $n \in \mathcal{N}_{cf}$ , the load Q to be used by the middleware (as described in Sec. 3.4), such that the expected execution time of the application is minimized. This will implicitly also determine the HW/SW partitioning of the candidate modules from  $\mathcal{H}$ .

# 5 Motivational Example

Let us consider the control flow graph (CFG) in Fig. 2a, where candidate hardware modules are represented with squares, and software nodes with circles. The discrete probability

<sup>&</sup>lt;sup>4</sup>Please note that this information is known only at run-time.

distribution for the iterations of the loop a-b, the software and hardware execution times for the nodes, as well as the edge probabilities, are illustrated on the graph. The reconfiguration times are:  $rec(m_1) = 37$ ,  $rec(m_2) = 20$ ,  $rec(m_3) = 46$ . We also consider that hardware modules  $m_1$  and  $m_2$  are conflicting due to their placement  $(m_1 \bowtie m_2)$ .

Let us try to schedule the configuration prefetches for the three hardware modules on the given CFG. If we use the method developed by Panainte et al. [9], the result is shown in Fig. 2b. As we can see, the *load* for  $m_3$  can be propagated upwards in the CFG from node  $m_3$  up to r. For nodes  $m_1$  and  $m_2$  it is not possible (according to this approach) to propagate their *load* calls to their ancestors, because they are in placement conflict. The data-flow analysis performed by the authors is too conservative, and the propagation of prefetches is stopped whenever two *load* calls targeting conflicting modules meet at a common ancestor (e.g. node f for  $m_1$  and  $m_2$ ). As a result, since the method fails to prefetch modules earlier, the reconfiguration overhead for neither  $m_1$ , nor  $m_2$ , can be hidden at all. Only module  $m_3$  will not generate any waiting time, since the minimum time to reach it from r is  $92 > rec(m_3) = 46$ . Using this approach, the application must stall (waiting for the reconfigurations to finish)  $W_1 = 90\% \cdot rec(m_1) + rec(m_2) = 90\% \cdot 37 + 20 = 53.3$  time units on average (because  $m_1$  is executed with a probability of 90%, and  $m_2$  is always executed).

Fig. 2c shows the resulting prefetches after using the method proposed by Li et al. [8]. As we can see, the prefetch queue generated by this approach at node r is  $loadQ : m_2, m_3, m_1$ , because the probabilities to reach the hardware modules from r are 100%, 95% and 90% respectively. Please note that this method is developed for FPGAs with relocation and defragmentation and it ignores placement conflicts. Also, the load queues are generated considering only the probability to reach a module (and ignoring other factors, such as the execution time distribution from the prefetch point up to the prefetched module). Thus, if applied to our example, the method performs poorly: in 90% of the cases, module  $m_1$  will replace module  $m_2$  (initially prefetched at r) on the FPGA. In this cases, none of the reconfiguration overhead for  $m_1$  can be hidden, and in addition, the initial prefetch for  $m_2$  is wasted. The average waiting time for this scenario is  $W_2 = 90\% \cdot rec(m_1) + (100\% - 10\%) \cdot rec(m_2) = 90\% \cdot 37 + 90\% \cdot 20 = 51.3$  time units (the reconfiguration overhead is hidden in 10% of the cases for  $m_2$ , and always for  $m_3$ ).

For this example, although the approach proposed by Sim et al. [13] tries to avoid some of the previous problems, it ends up with similar waiting time. The method uses 'placementaware' probabilities (PAPs). For any node  $n \in \mathcal{N}_{cf}$  and any hardware module  $m \in \mathcal{H}$ , PAP(n, m) represents the probability to reach module m from node n, without encountering any conflicting hardware module on the way. Thus, the prefetch order for  $m_1$  and  $m_2$  is correctly inverted since  $PAP(r, m_1) = 90\%$ , as in the previous case, but  $PAP(r, m_2) = 10\%$ , instead of 100% (because in 90% of the cases,  $m_2$  is reached via the conflicting module  $m_1$ ). Unfortunately, since the method uses only PAPs to generate prefetches, and  $PAP(r, m_3) =$ 95% (since it is not conflicting with neither  $m_1$ , nor  $m_2$ ),  $m_3$  is prefetched before  $m_1$  at node r, although its prefetch could be safely postponed. The result is illustrated in Fig. 2d ( $m_2$  is removed from the load queue of node r because it conflicts with  $m_1$ , which has a higher PAP). These prefetches will determine that no reconfiguration overhead can be hidden for  $m_1$  or  $m_2$  (since the long reconfiguration of  $m_3$  postpones their own one until the last moment). The average waiting time for the application will be  $W_3 = 90\% \cdot rec(m_1) + rec(m_2) =$  $90\% \cdot 37 + 20 = 53.3$  time units.

If we examine the example carefully, we can see that taking into account only the 'placement-aware' probabilities is not enough. The prefetch generation mechanism should also consider the distance from the current decision point to the hardware modules candidate for prefetching, correlated with the reconfiguration time of each module. Our approach, presented in the current report, is to estimate the performance gain associated with starting the reconfiguration of a certain module at a certain node in the CFG. We do this by considering both the execution time gain resulting from the hardware execution of that module (including any stalling cycles spent waiting for the reconfiguration to finish) compared to the software execution, and by investigating how this prefetch influences the execution time of the other

Algorithm 1 Generating the prefetch q	queues
---------------------------------------	--------

1: procedure GENERATEPREFETCHQ 2: for all  $n \in \mathcal{N}_{cf}$  do 3: for all  $\{m \in \mathcal{H} | PAP(n,m) \neq 0\}$  do 4: if  $\overline{G}_{nm} > 0 \lor m$  in loop then compute priority function  $C_{nm}$ 5: 6: end if end for 7: 8:  $loadQ(n) \leftarrow$  modules in decreasing order of  $C_{nm}$ remove all lower priority modules that have area conflicts with higher priority modules in loadQ(n)9: 10: end for 11: eliminate redundant prefetches 12: end procedure

reachable modules. For the example presented here, it is not a good idea to prefetch  $m_3$  at node r, because this results in a long waiting time for  $m_1$  (similar reasoning applies for prefetching  $m_2$  at r). The resulting prefetches are illustrated in Fig. 2e. As we can see, the best choice of prefetch order is  $m_1$ ,  $m_3$  at node r ( $m_2$  is removed from the load queue because it conflicts with  $m_1$ ), and this will hide most of the reconfiguration overhead for  $m_1$ , and all for  $m_3$ . The overall average waiting time is  $W = 90\% \cdot \overline{W}_{rm_1} + rec(m_2) = 90\% \cdot 4.56 + 20 \approx 24.1$ , less than half of the penalties generated by the previous methods (Sec. 6.2 and Fig. 3 explain the computation of the average waiting time generated by  $m_1$ ,  $\overline{W}_{rm_1} = 4.56$  time units).

# 6 Speculative Prefetching

Our overall strategy is shown in Algorithm 1. The main idea is to intelligently assign priorities to the candidate prefetches and determine the load queue (loadQ) at every node in the CFG (line 8). We try to use all the available knowledge from the profiling in order to take the best possible decisions and speculatively prefetch the hardware modules with the highest potential to reduce the expected execution time of the application. The intelligence of our algorithm resides in computing the priority function  $C_{nm}$  (see Sec. 6.1), which tries to estimate at design time what is the impact of reconfiguring a certain module on the average execution time. We consider for prefetch only the modules for which it is profitable to start a prefetch at the current point (line 4): either the average execution time gain  $\overline{G}_{nm}$  (over the software execution of the candidate) obtained if its reconfiguration starts at this node is greater than 0, or the module is inside a loop (in which case, even if the reconfiguration is not finished in the first few loop iterations and we execute the module in software, we will gain from executing the module in hardware in future loop iterations). Then we sort the prefetch candidates in decreasing order of their priority function (line 8), and in case of equality we give higher priority to modules placed in loops. After the loadQ has been generated for a node, we remove all the lower priority modules that have area conflicts with the higher priority modules in the queue (line 9). Once all the queues have been generated, we eliminate redundant prefetches (all consecutive candidates at a child node that are a starting sub-sequence at all its parents in the CFG), as in [8] or [13]. The exact hardware module to be prefetched will be determined by the middleware at run-time, as explained in Sec. 3.4.

## 6.1 Prefetch Priority Function

Our prefetch function represents the priorities assigned to the hardware modules reachable from a certain node in the CFG, and thus determines the loadQ to insert at that location. Considering that the processor must stall if the reconfiguration overhead cannot be completely hidden and that some candidates will provide a higher performance gain than others, our priority function will try to estimate the overall impact on the average execution time that results from different prefetches being issued at a particular node in the CFG. In order to accurately predict the next configuration to prefetch, several factors have to be considered. The first one is represented by the 'placement-aware' probabilities (PAPs), computed with the method from [13]. The second factor that influences the decision of prefetch scheduling is represented by the execution time gain distributions (that will be discussed in detail in Sec. 6.2). The gain distributions reflect the reduction of execution time resulting from prefetching a certain candidate and executing it in hardware, compared to executing it in software. They are directly impacted by the waiting time distributions (which capture the relation between the reconfiguration time for a certain hardware module and the execution time distribution between the prefetch node in the CFG and that module).

We denote the set of hardware modules for which it is profitable to compute the priority function at node n with  $Reach(n) = \{m \in \mathcal{H} \mid PAP(n, m) \neq 0 \land (\overline{G}_{nm} > 0 \lor m \text{ in } loop)\}$ . For our example in Fig. 2a,  $Reach(r) = \{m_1, m_2, m_3\}$ , but  $Reach(m_3) = \emptyset$ , because it does not make sense to reconfigure  $m_3$  anymore (although  $PAP(m_3, m_3) = 100\%$ , we have the average waiting time  $\overline{W}_{m_3m_3} = rec(m_3)$  and  $rec(m_3) + hw(m_3) = 46 + 12 > 50 = sw(m_3)$ ). Thus, we do not gain anything by starting the reconfiguration of  $m_3$  right before it is reached, i.e.  $\overline{G}_{m_3m_3} = 0$ . Considering the above discussion, our priority function expressing the reconfiguration gain generated by prefetching module  $m \in Reach(n)$  at node n is defined as:

$$C_{nm} = PAP(n,m) \cdot \overline{G}_{nm} + \sum_{k \in Mut Ex(m)} PAP(n,k) \cdot \overline{G}_{sk} + \sum_{k \notin Mut Ex(m)} PAP(n,k) \cdot \overline{G}_{nm}^{k}$$

In the above equation,  $\overline{G}_{nm}$  denotes the average execution time gain generated by prefetching module m at node n (see Sec. 6.2), MutEx(m) denotes the set of hardware modules that are executed mutually exclusive with m, the index s in  $\overline{G}_{sk}$  represents the node where the paths leading from n to m and k split, and  $\overline{G}_{nm}^k$  represents the expected gain generated by k, given that its reconfiguration is started immediately after the one for m.

The first term of the priority function represents the contribution (in terms of average execution time gain) of the candidate module m, the second term tries to capture the impact that the reconfiguration of m will produce on other modules that are executed mutually exclusive with it, and the third term captures the impact on the execution time of modules that are not mutually exclusive with m (and might be executed after m). In Fig. 2a, modules  $m_1, m_2$  and  $m_3$  are not mutually exclusive. Let us calculate the priority function for the three hardware modules from Fig. 2a at node r (considering their areas proportional with their reconfiguration time).  $C_{rm_1} = 90\% \cdot 40.4 + 10\% \cdot 36.9 + 95\% \cdot 38 \approx 76.1$ ,  $C_{rm_2} = 10\% \cdot 40 + 90\% \cdot 22.5 + 95\% \cdot 38 \approx 60.3$  and  $C_{rm_3} = 95\% \cdot 38 + 90\% \cdot 1.7 + 10\% \cdot 30.9 \approx 40.7$  (the computation of execution time gains is discussed in Sec. 6.2). As we can see, since  $C_{rm_1} > C_{rm_2} > C_{rm_3}$ , the correct *loadQ* of prefetches is generated at node r. Note that  $m_2$  is removed from the queue because it is in placement conflict with  $m_1$ , which is the head of the queue (see line 9 in Algorithm 1).

## 6.2 Average Execution Time Gain

Let us consider a node  $n \in \mathcal{N}_{cf}$  from the CFG and a hardware module  $m \in \mathcal{H}$ , reachable from n. Given that the reconfiguration of module m starts at node n, we define the average execution time gain  $\overline{G}_{nm}$  as the expected execution time that is saved by executing m in hardware (including any stalling cycles when the application is waiting for the reconfiguration of m to be completed), compared to a software execution of m. In order to compute it, we start with the distance (in time) from n to m. Let  $X_{nm}$  be the random variable associated with this distance. The waiting time is given by the random variable  $W_{nm} = max(0, rec(m) - X_{nm})$ . Note that the waiting time cannot be negative (if a module is already present on FPGA when we reach it, it does not matter how long ago its reconfiguration finished). The execution time gain is given by the distribution of the random variable



Figure 3: Computing the Gain Probability Distribution Step by Step

 $G_{nm} = max(0, sw(m) - (W_{nm} + hw(m)))$ . In case the software execution time of a candidate is shorter than waiting for its reconfiguration to finish and executing it in hardware, then the module will be executed in software by the middleware (as described in Sec. 3.4), and the gain is zero. If we denote the probability mass function (pmf) of  $G_{nm}$  with  $g_{nm}$ , then the average gain  $\overline{G}_{nm}$  will be computed as:

$$\overline{G}_{nm} = \sum_{x=0}^{\infty} \left( x \cdot g_{nm}(x) \right) \tag{1}$$

The discussion is illustrated graphically in Fig. 3, considering the nodes n = r and  $m = m_1$  from Fig. 2a. The probability mass function (pmf) for  $X_{rm_1}$  (distance in time from r to  $m_1$ ) is represented in Fig. 3a and the pmf for the waiting time  $W_{rm_1}$  in Fig. 3b. Note that the negative part of the distribution (depicted with dotted line) generates no waiting time. In Fig. 3c we add the hardware execution time to the potential waiting time incurred. Finally, Fig. 3d represents the discrete probability distribution of the gain  $G_{rm_1}$ . The resulting average gain is  $\overline{G}_{rm_1} = 34 \cdot 18\% + 39 \cdot 42\% + 44 \cdot 6\% + 45 \cdot 34\% = 40.44$  time units.

Before presenting our algorithm for computing the gain distribution and the average gain, let us first introduce a few concepts. Given a control flow graph  $\mathcal{G}_{cf}(\mathcal{N}_{cf}, \mathcal{E}_{cf})$ , we introduce the following definitions [1]:

**Definition 1** A node  $n \in \mathcal{N}_{cf}$  is post-dominated by a node  $m \in \mathcal{N}_{cf}$  in the control flow graph  $\mathcal{G}_{cf}$  if every directed path from n to sink (excluding n) contains m.

**Definition 2** Given a control flow graph  $\mathcal{G}_{cf}$ , a node  $m \in \mathcal{N}_{cf}$  is control dependent upon a node  $n \in \mathcal{N}_{cf}$  via a control flow edge  $e \in \mathcal{E}_{cf}$  if the next conditions hold:

- There exists a directed path P from n to m in  $\mathcal{G}_{cf}$ , starting with e, with all nodes in P (except m and n) post-dominated by m;
- *m* does not post-dominate *n* in  $\mathcal{G}_{cf}$ .

In other words, there is some control edge from n that definitely causes m to execute, and there is some path from n to sink that avoids executing m.

**Definition 3** A control dependence graph (CDG)  $\mathcal{G}_{cd}(\mathcal{N}_{cd}, \mathcal{E}_{cd})$  corresponding to a control flow graph  $\mathcal{G}_{cf}(\mathcal{N}_{cf}, \mathcal{E}_{cf})$  is defined as:  $\mathcal{N}_{cd} = \mathcal{N}_{cf}$  and  $\mathcal{E}_{cd} = \{((n, m), e) \mid m \text{ is control dependent upon n via edge } e\}$ . If we ignore all the backward edges in the CDG we obtain a forward control dependence tree (FCDT) [1].

Fig. 4 shows the FCDT corresponding to the CFG in Fig. 2a (note that the pseudo-edge  $r \rightarrow s$  was introduced in the CFG in order to get all nodes to be directly, or indirectly, control dependent on r).



Figure 4: Forward Control Dependence Tree

Algorithm 2 Computing the average execution time gain				
1: procedure AVGEXECTIMEGAIN $(n, m)$				
2: construct subgraph with nodes between $n$ and $m$				
3: build its FCDT (saved as a global variable)				
4: $X_{nm} \leftarrow ExecTimeDist(n,m)$				
5: $W_{nm} \leftarrow max(0, rec(m) - X_{nm})$				
6: $G_{nm} \leftarrow max(0, sw(m) - (W_{nm} + hw(m)))$				
7: for all $y \in \{t \mid g_{nm}(t) \neq 0\}$ do				
8: $\overline{G}_{nm} \leftarrow \overline{G}_{nm} + g_{nm}(y) \cdot y$				
9: end for				
10: return $\overline{G}_{nm}$				
11: end procedure				

Algorithm 2 presents our method for computing the average gain. Let us consider a node  $n \in \mathcal{N}_{cf}$  and a hardware module  $m \in \mathcal{H}$ . Given that the reconfiguration of module m starts at node n, our algorithm estimates the average execution time gain over a software execution, that results from executing m in hardware (after waiting for its reconfiguration to finish if needed). The first steps are to construct the subgraph with all the nodes between n and m and to build its FCDT, according to Def. 3 (lines 2-3). Then we compute the execution time distribution of the subgraph constructed earlier, representing the distance (in time) between n and m (line 4). Then we compute the waiting time distribution (line 5) and the gain distribution (line 6). Finally we compute the average gain with formula (1) (lines 7-8).

#### 6.3 Execution Time Distribution

Algorithm 3 details our method for computing the execution time distribution between node n and module m. We remind the reader that all the computation is done considering the subgraph containing only the nodes between n and m and its forward control dependence tree. Also, before applying the algorithm we transform all post-test loops into pre-test ones (this transformation is done on the CFG representation, for analysis purposes only). Our approach is to compute the execution time distribution of node n and all its children in the FCDT, using the recursive procedure *ExecTimeDist*. If n has no children in the FCDT (i.e. no nodes control dependent on it), then we simply return its own execution time (line 4). For the root node we convolute<sup>5</sup> its execution time with the execution time distribution of all its children in the FCDT (line 6).

For a control node, the approach is to compute the execution time distribution for all its children in the FCDG that are control dependent on the 'true' branch, convolute this with the execution time of n and scale the distribution with the probability of the 'true' branch, t (line 9). Similarly, we compute the distribution for the 'false' branch as well (line 10) and then we superpose the two distributions to get the final one (line 11). For example, for the branch node c in Fig. 2a, we have  $ex_t(2+3) = 30\%$ ,  $ex_f(2+8) = 70\%$  and, thus, the pmf for the execution time of the entire if-then-else structure is x(5) = 30% and x(10) = 70%.

 $<sup>^{5}</sup>$ This is done because the probability distribution of a sum of two random variables is obtained as the convolution of the individual probability distributions.

1: procedure EXECTIMEDIST( $n, m$ ) 2: $ex_n \leftarrow Exec(n)$ 3: ff $n$ has 0 children in FCDT then 4: $x(ex_n) \leftarrow 100\%$ 5: else if $n$ type = root then 6: $x \leftarrow ex_n + FCDTChildrenDist(n, m, l)$ 7: else if $n$ type = control then 6: $x \leftarrow ex_n + FCDTChildrenDist(n, m, l)$ 10: $ex_l \leftarrow t \times (ex_n + FCDTChildrenDist(n, m, l))$ 11: $x \leftarrow ex_n + ex_T$ 12: else if $n$ type = loop header then 13: $ex_l \leftarrow ex_n + ex_T$ 14: $Truncate(ex_l, rec(m))$ 15: for all i $\in lterations(n)$ do 16: $ex_l \leftarrow titer_prob_n(i) \times [(*^l)ex_{l}]$ 17: $Truncate(ex_l, rec(m))$ 18: $ex_{l} \leftarrow ex_n + ex_T$ 19: $etal \in ex_{l} + ex_{l}$ 10: $break$ $\triangleright$ header executed last time 21: end for 22: end for 23: $x \leftarrow ex_n + ex_T$ 24: end if 25: $Truncate(x, rec(m))$ 26: $return x$ 27: end procedure 28: procedure FCDTCHILDRENDIST( $n, m, label$ ) 29: for all $\in FCDTChildrenDist(n, m)$ 31: $Truncate(dx, rec(m))$ 32: if $nin(y \mid dist(y) \neq 0$ $\geq rec(m)$ then 33: $break$ $\triangleright$ no point to continue 34: end if 35: end procedure 35: procedure TUNCATE(dist, rec) 36: meturn dist 37: end procedure 38: procedure TUNCATE(dist, rec) 39: $y_{min} \leftarrow min(y \mid dist(y) \neq 0$ $\in$ 41: end if 42: end if 43: trunc(y) $\leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \geq rec \end{cases}$ 44: end if 45: return $trunce$ 46: end procedure 47: procedure EXXC(m) 48: if $m.type = hardware$ then 49: return $h(m) + \omega_n \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end procedure	Alge	orithm 3 Computing the execution time distribution	
2. $ex_n \leftarrow Exec(n)$ (b) 3. if $n$ has 0 children in FCDT then 4. $x(ex_n) \leftarrow 100\%$ 5. else if $n.type = root$ then 6. $x \leftarrow ex_n + FCDTChildrenDist(n, m, l)$ 7. else if $n.type = control$ then 8. $(l, f) \leftarrow GetLabels(n)$ $\triangleright$ branch frequencies 9. $ex_t \leftarrow t \times (ex_n + FCDTChildrenDist(n, m, l))$ 10. $ex_t \leftarrow f \times (ex_n + FCDTChildrenDist(n, m, l))$ 11. $x \leftarrow ex_t + ex_f$ 2. else if $n.type = loopheader$ then 12. else if $n.type = loopheader$ then 13. $ex_i(-ex_n + FCDTChildrenDist(n, m, l))$ 14. $Truncat(ex_{it}, rec(m))$ 15. for all i $\in Iterations(n)$ do 16. $ex_i \leftarrow iter_prob_n(i) \times [(*!)ex_{li}]$ 17. $Truncat(ex_{it}, rec(m))$ 18. $ex_{li} \leftarrow ex_{li} + ex_i$ 19. $brack$ $\triangleright$ no point to continue 20. $brack$ $\triangleright$ no point to continue 21. end if 22. end for 23. $x \leftarrow ex_n + ex_{lb}$ $\triangleright$ header executed last time 24. end if 25. $Truncat(x, rec(m)))$ 26. $return x$ 27. end procedure 28. procedure FCDTChildren(n, label) do 30. $dist \leftarrow dist + ExecTimeDist(n, m)$ 31. $Truncat(dist, rec(m))$ 32. if $min(y \mid dist(y) \neq 0 \} \ge rec(m)$ then 33. $brack$ $\triangleright$ no point to continue 34. end if 35. end for 35. end for 36. return dist 37. end procedure 38. procedure MUNCATE(dist, rec) 39. $y_{min} \leftarrow min(y \mid dist(y) \neq 0 \}$ 40. if $y_{min} \ge return$ 41. $trunc(y_{min}) \leftarrow dist(y_{min})$ 42. else 43. $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec$ 44. end if 45. return $func$ 46. end procedure 47. procedure Exec(m) 48. if $m.type = hardmare$ then 47. procedure 47. procedure 47. procedure 47. procedure 47. etturn $w(m)$ 42. else 47. procedure 48. if $m.type = hardmare$ then 44. $return hw(m) + o_m \cdot (sw(m) - hw(m))$ 45. end procedure	1: p	<b>procedure</b> EXECTIMEDIST $(n, m)$	
3: If $r$ has 0 children in FCDT then 4: $x(cx_n) \leftarrow 100^{\infty}$ 5: else if $n.type = root then 6: x \leftarrow ex_n + FCDTChildrenDist(n, m, l)7: else if n.type = control then 8: (l, f) \leftarrow GetLabels(n) ▷ branch frequencies9: ex_l \leftarrow t \times (ex_n + FCDTChildrenDist(n, m, l))10: ex_l \leftarrow f \times (ex_n + FCDTChildrenDist(n, m, l))11: x \leftarrow ex_l + ex_l12: else if n.type = loopheader then13: ex_l \leftarrow ex_l + ex_l14: Truncate(ex_l, rec(m))15: for all \in Ierations(n) do16: ex_i \rightarrow ter_T, root, (l) \times [(*]ex_l]]17: Truncate(ex_l, rec(m))18: ex_{lb} \leftarrow ex_{lb} + ex_{l} ▷ the loop body19: if min(y   ex_l(y) \neq 0} ≥ rec(m) then20: break ▷ header executed last time21: end if22: end for23: x \leftarrow cx_n * ex_{lb} ▷ header executed last time24: end if25: Truncate(x, rec(m))26: return x27: end procedure28: procedure FCDTCHildrenDist(n, m, label)29: for all c \in FCDTCHildrenDist(n, m, label)20: dist(u) \neq 0} ≥ rec(m) then31: Truncate(dist, rec(m))32: if min(y   dist(y) \neq 0} ≥ rec(m) then33: break ▷ no point to continue34: end if35: end for36: return dist37: end procedure38: procedure TRUSCATE(dist, rec)39: ymin \leftarrow mis(y   dist(y) \neq 0]40: if ymin ≥ rec then41: trunc(ymin) \leftarrow dist(ymin)42: else43: trunc(y) \leftarrow \begin{cases} dist(y) : y < rec44: end if45: return trunc46: end procedure47: procedure EXE(m)48: if m.type = hardware then49: return h(m)40: return h(m) + α_m \cdot (sw(m) - hw(m))50: else51: return sw(m)52: end procedure$	2:	$ex_n \leftarrow Exec(n)$	
4: $x(e_n) \leftarrow 100\%$ 5: else if n.type = control then 6: $x \leftarrow e_n * FCDTChildrenDist(n, m, l)$ 7: else if n.type = control then 8: $(t, f) \leftarrow GetLabels(n)$ > branch frequencies 9: $ext \leftarrow t \times (ex_n * FCDTChildrenDist(n, m, t))$ 10: $ex_t \leftarrow f \times (ex_n * FCDTChildrenDist(n, m, f))$ 11: $x \leftarrow ex_t + ex_f$ 2: else if n.type = loopheader then 13: $ex_{1i} \leftarrow ex_n * FCDTChildrenDist(n, m, l)$ 14: $Truncate(e_{x_{1i}}, rec(m))$ 15: for all $\in lterations(n)$ do 16: $ex_i \leftarrow iter.prob_n(i) \times [(*)ex_{1i}]$ 17: $Truncate(e_{x_i, rec(m))}$ 18: $ex_{2ib} \leftarrow ex_{2ib} + ex_i$ > b the loop body 19: fit min{u}   ex_{i}(y) \neq 0] \ge rec(m) then 20: break > branch frequencies 21: end for 22: end for 23: $x \leftarrow ex_n * ex_{1b}$ > b header executed last time 24: end if 25: $Truncate(x, rec(m))$ 26: return x 27: end procedure 28: procedure FCDTChildren, n., habel) 20: for all $c \in FCDTChildren, n., habel$ ) 20: dist $\leftarrow dist \in ExecTimeDist(n, m)$ 31: $Truncate(x, rec(m))$ 32: if min{u}   dist(y) \neq 0] \ge rec(m) then 33: break 34: end if 35: end for 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{nin} \leftarrow min{y}   dist(y) \neq 0]$ 40: $y \le rec$ 41: end if 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec$ 44: end if 45: enturn trunc 46: end procedure 47: procedure EXEC(m) 48: if m.type = hardware then 49: return hw(m) $+ \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return sw(m) 52: end procedure	3:	if n has 0 children in FCDT then	
5: else if $n.ippe = root then 6: x \leftarrow ex, a_r \in PCDTChildrenDist(n, m, l)7: else if n.type = control then  ightarrow if' blocks8: (i, f) \leftarrow GetLabels(n)  ightarrow if' blocks9: ex_i \leftarrow t \times (ex_n * FCDTChildrenDist(n, m, l))10: ex_i \leftarrow ex_i + ex_i12: else if n.type = loop header then13: ex_i \leftarrow ex_i + ex_i14: Truncate(ex_i, rec(m))15: for all \in ltrentons(n) do16: ex_i \leftarrow iter\_prob_n(1 \times [[v^i]ex_{il}]]17: Truncate(ex_i, rec(m))18: ex_{ib} \leftarrow ex_i + ex_i  ightarrow if min [y   ex_i(y) \neq 0] \ge rec(m) then10: break  ightarrow if min [y   ex_i(y) \neq 0] \ge rec(m) then10: break  ightarrow if min [y   ex_i(y) \neq 0] \ge rec(m) then12: end for13: x \leftarrow ex_n * ex_{lb}  ightarrow if min [y   ex_i(x), rec(m)]14: Truncate(x, rec(m))15: for all c \in FCDTChildren(n, label)16: return x17: end procedure28: procedure FCDTCHILDRENDIST(n, m, label)29: for all c \in FCDTChildren(n, label) do30: dist \leftarrow dist + ExecTimeDist(n, m)31: Truncate(dist, rec(m))32: if min [y   dist(y) \neq 0] \ge rec(m) then33: break end if35: end for36: return dist37: end procedure38: procedure TRUNCATE(dist, rec)39: y_{min} \leftarrow min [y   dist(y) \neq 0]41: trunc(y_{min}) \leftarrow dist(y_{min})42: else43: trunc(y \mapsto \left\{ \frac{dist(y) : y < rec}{0 : y \ge rec}44: end if45: return frum runc46: end procedure47: procedure EXEC(m)48: if m.type = hardware then49: return hw(m) + a_m \cdot (sw(m) - hw(m))50: else51: return sw(m)52: end if52: end procedure$	4:	$\mathbf{x}(ex_n) \leftarrow 100\%$	
6: $x \leftarrow ex_n + PCDTChildrenDist(n, m, l)$ 7: else if n:type = control then $\triangleright'if'$ blocks 8: $(t, f) \leftarrow GetLabels(n) \triangleright'if' blocks9: ex_t \leftarrow t \times (ex_n + FCDTChildrenDist(n, m, f))10: ex_t \leftarrow t + ex_t11: x \leftarrow ex_t + ex_t12: else if n:type = loopheader then13: ex_i \leftarrow ex_n + FCDTChildrenDist(n, m, l)14: Truncate(ex_n, rec(m))15: for all i \in iterations(n) do16: ex_i \leftarrow iter_prob_n(i) \times [(e^i)ex_i]17: Truncate(ex_n, rec(m))18: ex_{lb} \leftarrow ex_{lb} + ex_i \triangleright the loop body19: if min{y \mid ex_i(y \neq 0) \ge rec(m) then20: break \triangleright no point to continue21: end if22: end for23: x \leftarrow ex_n + ex_{lb} \triangleright header executed last time24: end if25: Truncate(x, rec(m))26: return \times27: end procedure28: procedure FCDTCHILDRENDIST(n, m, label)20: for all \in FCDTChildrenDist(n, m)31: Truncate(dist, rec(m))32: if min{y \mid dist(y) \neq 0 } ere(m) then33: break \triangleright no point to continue34: end if35: end for36: return dist37: end procedure38: procedure TRUNCATE(dist, rec)39: y_{min} \leftarrow min{y \mid dist(y) \neq 0}40: if y_{min} \ge rec then41: trunc(y_{min}) \leftarrow dist(y_{min})42: else43: trunc(y \leftarrow {dist(y) : y < rec}44: end if45: end procedure45: return trunc46: end procedure47: procedure EXEV(m)48: if rutry = hardware then49: return h(m) + \alpha_m \cdot (sw(m) - hw(m))50: else51: return sw(m)52: end if53: end procedure$	5:	else if $n.type = root$ then	
7: else if $n.type = control then  5: if' blocks 8: (t, f) \leftarrow GetLadels(n)6: ex_t \leftarrow t \times (ex_n * FCDTChildrenDist(n, m, t))10: ex_t \leftarrow t \times (ex_n * FCDTChildrenDist(n, m, t))11: x \leftarrow ex_1 + ex_112: else if n.type = loopheader then13: ex_t \leftarrow ex_n + ex_t13: ex_t \leftarrow ex_n + ex_t14: Truncate(ex_t, rec(m))15: for all i \in lerations(n) do16: ex_t \leftarrow tite_rproh_{0}(1 \times [[e^*]ex_{t}]]17: Truncate(ex_t, rec(m))18: ex_{th} \leftarrow ex_{th} + ex_{t}19: fr min \{y \mid ex_t(y) \neq 0\} \ge rec(m) then20: break21: end if22: end for23: x \leftarrow ex_n * ex_{th}24: end if25: Truncate(x, rec(m))26: return x27: end procedure28: procedure FCDTCHILDRENDIST(n, m, label)29: for all c \in FCDTChildren(n, label) do30: dist \leftarrow dist * ExecTimeDist(n, m)31: Truncate(dist, rec(m))32: if min(y \mid dist(y) \neq 0] \ge rec(m) then33: break34: end if35: end for36: return dist37: end procedure38: procedure TRUNCATE(dist, rec)39: y_{min} \leftarrow then41: trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec43: trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec44: end if45: return frunc46: end procedure47: procedure EXEC(m)48: if m.typ = hadware then49: return hw(m) + \alpha_m \cdot (sw(m) - hw(m))50: else51: return sw(m)52: end if52: end procedure$	6:	$\mathbf{x} \leftarrow ex_n * FCDTChildrenDist(n, m, l)$	
8: $(t, f) \leftarrow cetLabels(n)$ > branch frequencies 9: $ex_t \leftarrow t \times (ex_n + FCDTChildrenDist(n, m, t))$ 10: $ex_t \leftarrow t \times t + ex_t$ 11: $x \leftarrow ex_n + FCDTChildrenDist(n, m, t)$ 12: else if n:type = loopheader then 13: $ex_i \leftarrow ex_n + FCDTChildrenDist(n, m, t)$ 14: $Truncate(ex_i, rec(m))$ 15: for all i $\in Iterations(n)$ do 16: $ex_i \leftarrow iter.prob(i) \times [[v^i]_{otil}]$ 17: $Truncate(ex_i, rec(m))$ 18: $ex_{th} \leftarrow ex_t(y) \neq 0 \ge rec(m)$ then 20: $break$ > the loop body 21: $end$ for 22: $end$ for 23: $x \leftarrow ex_n * ex_{tb}$ > header executed last time 24: $end$ if 25: $Truncate(x, rec(m))$ 26: $return \times$ 27: $end procedure$ 28: $procedure FCDTChildrenDist(n, m, label)$ 20: $for all c \in FCDTChildrenDist(n, m)31: Truncate(dist, rec(m))32: if min\{y \mid dist(y) \neq 0 \} \ge rec(m) then33: break > no point to continue34: end if35: return dist37: end procedure38: procedure TRUNCATE(dist, rec)39: y_{min} \leftarrow min\{y \mid dist(y) \neq 0 \}41: trunc(y_{min}) \leftarrow dist(y_{min})42: else43: trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec44: end if45: return frunc 46: end procedure 47: procedure Exerc(m) 47: procedure Exerc(m) 48: if m.typ = hadware then49: return hw(m) + \alpha_m \cdot (sw(m) - hw(m))50: else51: return sw(m)52: end if53: end procedure$	7:	else if $n.type = control$ then	▷ 'if' blocks
9: $ex_{i} \leftarrow t \times (ex_{n} + FCDTChildrenDist(n, m, t))$ 10: $ex_{f} \leftarrow f \times (ex_{n} * FCDTChildrenDist(n, m, f))$ 11: $x \leftarrow ex_{t} + ex_{t}$ 12: else if $n.type = loop header then 13: ex_{i} \leftarrow ex_{n} + FCDTChildrenDist(n, m, l)14: Truncate(ex_{i}, rec(m))15: for all i \in lterations(n) do16: ex_{i} \leftarrow iter.prob_{n}(i) \times [(e^{i})ex_{l_{i}}]17: Truncate(ex_{i}, rec(m))18: ex_{h} \leftarrow ex_{h} + ex_{i} > the loop body19: if min{y}   ex_{i}(y) \neq 0 \} \geq rec(m)$ then 20: $break$ $>$ header executed last time 21: end if 22: end for 23: $x \leftarrow ex_{n} * ex_{lb}$ $>$ header executed last time 24: end if 25: $Truncate(x, rec(m))$ 26: $return x$ 27: end procedure 28: procedure FCDTCHILDRENDIST(n, m, label) 29: for all $c \in FCDTChildren(n, label)$ do 30: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min{y}   dist(y) \neq 0 \} \geq rec(m) then 33: $break$ $>$ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{mi} \leftarrow \min[y \mid dist(y) \neq 0]$ 40: if $y_{min} \geq rec$ then 41: $trunc(ymin) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec$ 44: end if 45: return trunc 46: end procedure 47: procedure Exec(m) 48: if $m.typ = hardware$ then 49: return $h(m) + \alpha_{m} \cdot (sw(m) - hw(m))$ 50: else 51: return sw(m) 52: end jrocedure	8:	$(t, f) \leftarrow GetLabels(n)$	$\triangleright$ branch frequencies
$\begin{array}{llllllllllllllllllllllllllllllllllll$	9:	$ex_t \leftarrow t \times (ex_n * FCDTChildrenDist(n, m, t))$	-
11: $\mathbf{x} \leftarrow e\mathbf{x}_{1} \leftarrow \mathbf{x}_{1}$ (c) $\mathbf{x}_{1}(\mathbf{x}_{1}, \mathbf{x}_{2}, \mathbf{x}_{3})$ 12: else if $n.type = loop header then 13: e\mathbf{x}_{1} \leftarrow e\mathbf{x}_{n} + FCDTChildrenDist(n, m, l)14: Truncate(e\mathbf{x}_{1}, rec(m))15: for all i \in lterations(n) do16: e\mathbf{x}_{i} \leftarrow iter.prob_{n}(i) \times [(\mathbf{s}^{i})\mathbf{e}\mathbf{x}_{1}]17: Truncate(e\mathbf{x}_{i}, rec(m))18: e\mathbf{x}_{1} \leftarrow e\mathbf{x}_{n} + e\mathbf{x}_{i} \triangleright the loop body19: if \min\{y \mid e\mathbf{x}_{i}(y) \neq 0\} \ge rec(m) then20: break \triangleright no point to continue21: end if22: end for23: \mathbf{x} \leftarrow e\mathbf{x}_{n} \ast e\mathbf{x}_{lb} \triangleright header executed last time24: end if25: Truncate(\mathbf{x}, rec(m))26: return \mathbf{x}27: end procedure28: procedure FCDTCHILDRENDIST(n, m, label)20: for all c \in FCDTCHildren(n, label) do30: dist \leftarrow dist \ast ExecTimeDist(n, m)31: Truncate(dist, rec(m))32: if \min\{y \mid dist(y) \neq 0\} \ge rec(m) then33: break \triangleright no point to continue34: end if35: end for36: return dist37: end procedure38: procedure TRUNCATE(dist, rec)39: y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}40: if y_{min} \ge rec then41: trunc(y_{min}) \leftarrow dist(y_{min})42: else43: trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec44: end if45: return trunc46: end procedure47: procedure EXEC(m)48: if m.type = hardware then49: return hw(m) + \alpha_{m} \cdot (sw(m) - hw(m))50: else51: return sw(m)52: end if52: end jrocedure$	10:	$ex_f \leftarrow f \times (ex_n * FCDTChildrenDist(n, m, f))$	
$\begin{array}{llllllllllllllllllllllllllllllllllll$	11:	$\mathbf{x} \leftarrow ex_t + ex_f$	
13: $ex_{ii} \leftarrow ex_n * FCDTChildrenDist(n, m, l)$ 14: $Truncate(ex_{ii}, rec(m))$ 15: $for all i \in Itrations(n) do$ 16: $ex_i \leftarrow iter.prob_n(i) \times [(*^i)ex_{ii}]$ 17: $Truncate(ex_i, rec(m))$ 18: $ex_{lb} \leftarrow ex_{lb} + ex_i$ ▷ the loop body 19: if min(y   ex_i(y) \neq 0) ≥ rec(m) then 20: break ▷ header executed last time 21: end ff 22: end for 23: $x \leftarrow ex_n * ex_{lb}$ ▷ header executed last time 24: end if 25: $Truncate(x, rec(m))$ 26: return x 27: end procedure 28: procedure FCDTCHildren(n, label) do 30: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(x, rec(m))$ 32: if min{y   dist(y) ≠ 0} ≥ rec(m) then 33: break ▷ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TauxCATE(dist, rec) 39: $y_{min} \leftarrow min{y}   dist(y) \neq 0$ 40: if $y_{min} \ge rethen$ 41: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec$ 41: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if m.type = hardware then 49: return hw(m) + $\alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return sw(m) 52: end for 53: end procedure	12:	else if $n.tupe = loop header$ then	
$\begin{array}{llllllllllllllllllllllllllllllllllll$	13:	$ex_{li} \leftarrow ex_n * FCDTChildrenDist(n, m, l)$	
15: for all i ∈ Iterations(n) do 16: $ex_i + iter prob_n(i) × [[*^i]ex_{li}]$ 17: Truncate(ex_i, rec(m)) 18: $ex_{lb} + ex_i + ex_i$ ▷ the loop body 19: if min{y   $ex_i(y) \neq 0$ } ≥ $rec(m)$ then 20: $break$ ▷ no point to continue 21: end if 22: end for 23: $x \leftarrow ex_n * ex_{lb}$ ▷ header executed last time 24: end if 25: Truncate(x, $rec(m)$ ) 26: return x 27: end procedure 28: procedure FCDTCHILDEENDIST(n, m, label) 29: for all $c \in FCDTCHILDEENDIST(n, m, label)$ 20: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min{y   $dist(y) \neq 0$ } ≥ $rec(m)$ then 33: $break$ ▷ no point to continue 34: end if 35: end for 36: return $dist$ 37: end procedure 38: procedure TRUNCATE( $dist, rec$ ) 39: $y_{min} \leftarrow min{y}   dist(y) \neq 0$ } 40: if $y_{min} \geq rec$ then 41: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec$ 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $mtype = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	14:	$Truncate(ex_{1i}, rec(m))$	
16: $ex_i \leftarrow iter.prob_n(i) \times [(*^i)ex_{li}]$ 17: $Truncate(ax_i, rec(m))$ ▷ the loop body         18: $ex_{lb} \leftarrow ex_{lb} + ex_i$ ▷ the loop body         19:       if min{y   $ex_i(y) \neq 0$ } ≥ $rec(m)$ then       ▷ no point to continue         20: $break$ ▷ no point to continue         21:       end if       ▷ header executed last time         23: $x \leftarrow ex_n * ex_{lb}$ ▷ header executed last time         24:       end if       ▷ header executed last time         25: $Truncate(x, rec(m))$ ▷         26:       return x       ▷         27:       end procedure       ▷         28:       procedure FCDTCHILDRENDIST(n, m, label)       on         0: $dist \leftarrow dist = kzec2^{TimeDist(n, m)}$ ▷         31: $Truncate(dist, rec(m))$ ▷         32:       if min{y   dist(y) ≠ 0} ≥ rec(m) then       ▷         33:       break       ▷ no point to continue         34:       end if       ▷         35:       end for       ▷       is may for a break         36:       if ymin ≥ rec then       ○       is y ≥ rec         41:       trunc(y) ← { dist(y) : y < rec	15:	for all $i \in Iterations(n)$ do	
17:       Truncate[exi, rec(m)]         18:       ext_b ← exi_b + exi,       ▷ the loop body         19:       if min{y   exi(y) ≠ 0} ≥ rec(m) then       ▷ no point to continue         20:       break       ▷ no point to continue         21:       end if       ▷         22:       end of       ▷         23:       x ← exn * ext_b       ▷ header executed last time         24:       end if       ▷         25:       Truncate(x, rec(m))       ▷         26:       return x       ▷         27:       end procedure       ▷         28:       procedure FCDTCHILDRENDIST(n, m, label)       ≥         29:       for all c ∈ FCDTChildren(n, label) do       >         30:       dist ← dist * ExecTimeDist(n, m)       >         31:       Truncate(dist, rec(m))       >         32:       if min{y   dist(y) ≠ 0} ≥ rec(m) then       >         33:       break       >       >         34:       end if       >       >         35:       grocedure       Truncate(dist, rec)       >         36:       return dist       >       >       >         37:       end procedure       0 : y ≥ rec	16:	$ex_i \leftarrow iter\_prob_n(i) \times [(*^i)ex_{li}]$	
$\begin{array}{llllllllllllllllllllllllllllllllllll$	17:	$Truncate(ex_i, rec(m))$	
$\begin{array}{llllllllllllllllllllllllllllllllllll$	18:	$ex_{1b} \leftarrow ex_{1b} + ex_i$	▷ the loop body
$\begin{array}{llllllllllllllllllllllllllllllllllll$	19:	if $\min\{y \mid ex_i(y) \neq 0\} > rec(m)$ then	· ···· ····
21:       end if         22:       end for         23:       x ← ex <sub>n</sub> * ex <sub>lb</sub> 24:       end if         25:       Truncate(x, rec(m))         26:       return x         27:       end procedure         28:       procedure FCDTCHILDRENDIST(n, m, label)         29:       for all c ∈ FCDTCHILDRENDIST(n, m, label) do         30:       dist ← dist * ExecTimeDist(n, m)         31:       Truncate(dist, rec(m))         32:       if min[y   dist(y) ≠ 0] ≥ rec(m) then         33:       break         34:       end if         35:       end procedure         38:       procedure TRUNCATE(dist, rec)         39:       ymin ← min[y   dist(y) ≠ 0]         41:       trunc(ymin) ← dist(ymin)         42:       else         43:       trunc(y) ← { dist(y) : y < rec	20:	break	▷ no point to continue
$ \begin{array}{llllllllllllllllllllllllllllllllllll$	21:	end if	, no point to continue
2: $x \leftarrow ex_n * ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r ex_{lb}$ > header executed last time 2: $dr et a r exp = executed last rectard 2: dr et a r exp = executed last rectard 2: dr et a r exp = executed last rectard 3: dr et a r ertard rectard rectard 3: dr ertard rectard recta$	22:	end for	
$\begin{aligned} find the transformation of transformation of the transformation of transformat$	23:	$\mathbf{x} \leftarrow ex_m * ex_m$	▷ header executed last time
25: Truncate(x, rec(m)) 26: return x 27: end procedure 28: procedure FCDTCHILDRENDIST(n, m, label) 29: for all c ∈ FCDTChildren(n, label) do 30: dist ← dist * ExecTimeDist(n, m) 31: Truncate(dist, rec(m)) 32: if min{y   dist(y) ≠ 0} ≥ rec(m) then 33: break ▷ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: trunc( $y_{min}$ ) ← $dist(y_{min})$ 42: else 43: trunc( $y$ ) ← $\begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	24:	end if	
26: return x 27: end procedure 28: procedure FCDTCHILDRENDIST(n, m, label) 29: for all $c \in FCDTChildren(n, label)$ do 30: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min{ $y   dist(y) \neq 0$ } $\geq rec(m)$ then 33: $break$ $\triangleright$ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow min{\{y   dist(y) \neq 0\}}$ 40: if $y_{min} \geq rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \geq rec$ 44: end if 45: return trunc 46: end procedure 47: procedure Exec(m) 48: if $m.type = hardware$ then 49: $return hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return sw(m) 52: end if 53: end procedure	25:	$Truncate(\mathbf{x}, rec(m))$	
27: end procedure 28: procedure FCDTCHILDRENDIST( $n, m, label$ ) 29: for all $c \in FCDTChildren(n, label)$ do 30: dist $\leftarrow$ dist $\ast ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min{y   dist(y) $\neq$ 0} $\geq$ rec(m) then 33: break ▷ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \geq rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	26:	return x	
28: procedure FCDTCHILDRENDIST $(n, m, label)$ 29: for all $c \in FCDTChildren(n, label)$ do 30: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min $\{y \mid dist(y) \neq 0\} \ge rec(m)$ then 33: $break$ $\triangleright$ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	27: e	and procedure	
28: procedure FCDTCHILDRENDIST $(n, m, label)$ 29: for all $c \in FCDTChildren(n, label)$ do 30: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min{ $y \mid dist(y) \neq 0$ } $\geq rec(m)$ then 33: $break$ $\triangleright$ no point to continue 34: end if 35: end for 36: return $dist$ 37: end procedure 38: procedure TRUNCATE( $dist, rec$ ) 39: $y_{min} \leftarrow min{y \mid dist(y) \neq 0}$ 40: if $y_{min} \geq rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \geq rec \end{cases}$ 44: end if 45: return $trunc$ 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure		F	
29: for all $c \in FCDTChildren(n, label)$ do 30: dist ← dist * ExecTimeDist(n, m) 31: Truncate(dist, rec(m)) 32: if min{y   dist(y) ≠ 0} ≥ rec(m) then 33: break ▷ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if m.type = hardware then 49: return hw(m) + $\alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return sw(m) 52: end if 53: end procedure	28: p	<b>procedure</b> FCDTCHILDRENDIST $(n, m, label)$	
30: $dist \leftarrow dist * ExecTimeDist(n, m)$ 31: $Truncate(dist, rec(m))$ 32: if min{y   $dist(y) \neq 0$ } ≥ $rec(m)$ then 33: $break$ ▷ no point to continue 34: end if 35: end for 36: return $dist$ 37: end procedure 38: procedure TRUNCATE( $dist, rec$ ) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0$ } 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return $trunc$ 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	29:	for all $c \in FCDTChildren(n, label)$ do	
$\begin{array}{llllllllllllllllllllllllllllllllllll$	30:	$dist \leftarrow dist * ExecTimeDist(n, m)$	
32: if $\min\{y \mid dist(y) \neq 0\} \ge rec(m)$ then 33: break $\triangleright$ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if m.type = hardware then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	31:	Truncate(dist, rec(m))	
33: break ▷ no point to continue 34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if m.type = hardware then 49: return hw(m) + $\alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return sw(m) 52: end if 53: end procedure	32:	if $\min\{y \mid dist(y) \neq 0\} \ge rec(m)$ then	
34: end if 35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	33:	break	$\triangleright$ no point to continue
35: end for 36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	34:	end if	
36: return dist 37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	35:	end for	
37: end procedure 38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	36:	return dist	
38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	37: <b>e</b>	and procedure	
38: procedure TRUNCATE(dist, rec) 39: $y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$ 40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure			
$\begin{array}{llllllllllllllllllllllllllllllllllll$	38: p	procedure Truncate(dist, rec)	
40: if $y_{min} \ge rec$ then 41: $trunc(y_{min}) \leftarrow dist(y_{min})$ 42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	39:	$y_{min} \leftarrow \min\{y \mid dist(y) \neq 0\}$	
$ \begin{array}{llllllllllllllllllllllllllllllllllll$	40:	$\mathbf{if}  y_{min} \geq rec  \mathbf{then}$	
42: else 43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	41:	$trunc(y_{min}) \leftarrow dist(y_{min})$	
43: $trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$ 44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	42:	else	
44: end if 45: return trunc 46: end procedure 47: procedure EXEC(m) 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	43:	$trunc(y) \leftarrow \begin{cases} dist(y) : y < rec \\ 0 : y \ge rec \end{cases}$	
45: return trunc 46: end procedure 47: procedure $EXEC(m)$ 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	44:	end if	
46: end procedure 47: procedure $ExEC(m)$ 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	45:	return trunc	
47: procedure $ExEC(m)$ 48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	46: <b>e</b>	nd procedure	
48: if $m.type = hardware$ then 49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	47: p	procedure $Exec(m)$	
49: return $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ 50: else 51: return $sw(m)$ 52: end if 53: end procedure	48:	if $m.type = hardware$ then	
50: else $51:$ return $sw(m)$ $52:$ end if $53:$ end procedure	49:	<b>return</b> $hw(m) + \alpha_m \cdot (sw(m) - hw(m))$	
51:         return sw(m)           52:         end if           53:         end procedure	50:	else	
52: end if 53: end procedure	51:	$\mathbf{return} \ sw(m)$	
53: end procedure	52:	end if	
	53: <b>e</b>	and procedure	

Finally, for a loop header, we first compute the distribution of all its children in the FCDT (which represents the execution time of all the nodes inside the loop body) and then we convolute this with the execution time of the header (line 13). The result will be the execution time distribution of one iteration through the loop  $(ex_{li})$ . Then we use the distribution of loop iterations (available from profiling) to convolute  $ex_{li}$  with itself ((\*<sup>i</sup>) denotes the operation of convolution with itself *i* times). The result is scaled (line 16) with

the probability of *i* iterations to occur  $(iter\_prob_n(i))$  and then superposed (line 18) with the distribution of the loop body computed so far  $(ex_{lb})$ .

Let us illustrate the computation for the loop composed of nodes a - b in our CFG example from Fig. 2a using its corresponding FCDT from Fig. 4. Since in this particular case b is the only node inside the loop,  $ex_{li} = ex_a * ex_b$  gives the probability mass function (pmf)  $ex_{li}(1 + 4) = 100\%$ . Then we convolute  $ex_{li}$  with itself two times, and we scale the result with the probability to iterate twice through the loop,  $iter\_prob_a(2) = 60\%$ . Thus, we obtain  $ex_2(10) = 60\%$ . Similarly,  $ex_4(20) = 20\%$  and  $ex_5(25) = 20\%$ . By superposing  $ex_2$ ,  $ex_4$  and  $ex_5$  we get the pmf for the loop body,  $ex_{lb}$ , which we finally have to convolute with  $ex_a$  to get the pmf of the entire loop: x(10+1) = 60%, x(20+1) = 20% and x(25+1) = 20%. This distribution can be further used in the computation of  $\overline{G}_{rm_1}$ , for example.

The procedure FCDTChildrenDist (line 28) simply convolutes the distributions of all children of n in the FCDT that are control dependent on the parameter edge *label*. In order to speed-up the computation, we exploit the following observation: when computing the pmf for the execution time distributions, we discard any values that are greater or equal than the reconfiguration time, because those components of the distribution will generate no waiting time (one example in lines 31-33).

Procedure Truncate works as follows: if the smallest execution time is already greater than the reconfiguration overhead, we keep only this smallest value in the distribution (line 41). This is done because the distribution in question might be involved in convolutions or superpositions (in procedure ExecTimeDist), and keeping only this minimal value is enough for computing the part of the execution time distribution of interest (that might generate waiting time). Otherwise, we simply truncate the distribution at rec (line 43).

One observation related to the computation of the execution time of a node  $m \in \mathcal{N}_{cf}$ (procedure *Exec* in Algorithm 3, line 47) is that, if m is a hardware candidate  $(m \in \mathcal{H})$  we need to approximate its execution time, since the prefetches for it might be yet undecided and, thus, it is not known if the module will be executed in software or on the FPGA. In order to estimate the execution time, we make use of a coefficient  $\alpha_m \in [0,1]$ . The execution time for a hardware module m will be computed as (line 49):  $exec(m) = hw(m) + \alpha_m \cdot (sw(m) - hw(m))$ . Our experiments have proven that very good results are obtained by setting the value of  $\alpha_m$ , for each hardware module, to the ratio between its own hardware area and the total area needed for all modules in  $\mathcal{H}$ :  $\alpha_m = \frac{area(m)}{\sum_{k \in \mathcal{H}} area(k)}$ .

# 7 Experimental Results

## 7.1 Monte Carlo Simulation

#### 7.1.1 Sampling

In order to evaluate the quality of our prefetch solutions we have used an in-house developed Monte Carlo simulator that produces the execution time distribution of an application considering the architectural assumptions described in Sec. 3.1, 3.3 and 3.4. Each simulation generates a trace through the control flow graph, starting at the *root* node, and ending at the *sink* node (and we record the length of these traces). Whenever a branch node is encountered, we perform a Bernoulli draw (based on the probabilities of the outgoing edges) to decide if the branch is taken or not. At loop header nodes we perform random sampling from the discrete distribution of loop iterations  $(iter_prob_n)$  to decide how many times to loop.

For control nodes, if correlations between two or more branches are known, then they could be captured through joint probability tables. In such a case, whenever we perform a draw from the marginal Bernoulli distribution for a branch, we can compute the conditional probabilities for all the branches correlated with it, based on the joint probability table. Later in the simulation, when the correlated branches are reached, we do not sample their marginal distribution, but instead we sample their conditional distribution based on the outcome of the first branch.

#### 7.1.2 Accuracy Analysis

We stop the Monte Carlo simulation once we reach a satisfactory accuracy for the mean of the execution time distribution. We describe the desired accuracy in the following way: "The mean of the output distribution should be accurate to within  $\pm \epsilon$  with confidence  $\kappa$ ". The accuracy can be arbitrarily precise at the expense of longer simulation times. We will next present an analysis based on confidence intervals [15], to determine the number of samples to run in order to achieve the required accuracy.

Let us assume that  $\mu$  is the actual mean of the true output distribution and  $\hat{\mu}$  is the estimate mean computed by the Monte Carlo simulation. Since each simulation result is an independent sample from the same distribution, using the Central Limit Theorem we have that the distribution of the estimate of the true mean is (asymptotically) given by:

$$\widehat{\mu} = Normal\left(\mu, \frac{\sigma}{\sqrt{N}}\right)$$

where  $\sigma$  is the true standard deviation of the output execution time distribution and N represents the number of samples. The above equation can be rewritten as follows:

$$\mu = Normal\left(\widehat{\mu}, \frac{\sigma}{\sqrt{N}}\right)$$

By considering the required accuracy for our mean estimate and performing a transformation to the standard Normal distribution (i.e. with mean 0 and standard deviation 1), we can obtain the following relationship [15]:

$$\epsilon = \frac{\sigma}{\sqrt{N}} \Phi^{-1} \left( \frac{1+\kappa}{2} \right)$$

where the function  $\Phi^{-1}(\bullet)$  is the inverse of the standard Normal cumulative distribution function. By rearranging the terms and considering that we want to achieve at least this accuracy we obtain the minimum value for the number of samples N:

$$N > \left(\frac{\sigma \Phi^{-1}\left(\frac{1+\kappa}{2}\right)}{\epsilon}\right)^2$$

Please note that we do not know the true standard deviation  $\sigma$ , but for our purpose we can estimate it by taking the standard deviation of the first few samples (for example the first 40).

#### 7.2 Synthetic Examples

In order to evaluate the effectiveness of our algorithm we first performed experiments on synthetic examples. We generated two sets of control flow graphs:  $Set_1$  contains 20 CFGs with ~ 100 nodes on average (between 67 and 126) and  $Set_2$  contains 20 CFGs with ~ 200 nodes on average (between 142 and 268).

The software execution time for each node was randomly generated in the range of 10 to 100 time units. A fraction of all the nodes (between 15% and 25%) were then chosen to become hardware candidates, and their software execution time was generated  $\beta$  times bigger than their hardware execution time. The coefficient  $\beta$  was chosen from the uniform distribution on the interval [3, 7], in order to model the variability of hardware speedups over software. We also generated the size of the hardware modules, which in turn determined their reconfiguration time.

The size of the PDR region available for placement of hardware modules was varied as follows: we summed up all the areas for all hardware modules of a certain application:  $MAX\_HW = \sum_{m \in \mathcal{H}} area(m)$ . Then we generated problem instances by considering the size of the available reconfigurable region corresponding to different fractions of  $MAX\_HW$ : 15%, 25%, 35%, 45%, and 55%. As a result, we obtained a total of  $2 \times 20 \times 5 = 200$ 



Figure 5: Comparison with State-of-Art [13]: Synthetic Benchmarks

experimental settings. All experiments were run on a PC with CPU frequency 2.83 GHz, 8 GB of RAM, and running Windows Vista.

For each experimental setting, we first generated a placement for all the hardware modules, which determined the area conflict relationship between them. Then, for each application we inserted the configuration prefetches in the control flow graph. Finally, we have evaluated the result using the simulator described in Sec. 7.1 that produces the average execution time of the application considering the architectural assumptions described in Sec. 3.1, 3.3 and 3.4. We have determined the result with an accuracy of  $\pm 1\%$  with confidence 99.9%.

As a baseline we have considered the average execution time of the application (denoted as *baseline*) in case all the hardware candidates are placed on FPGA from the beginning and, thus, no prefetch is necessary. Please note that this is an absolute lower bound on the execution time; this ideal value might be unachievable even by the optimal static prefetch, because it might happen that it is impossible to hide all the reconfiguration overhead for a particular application.

First of all we were interested to see how our approach compares to the current stateof-art [13]. Thus, we have simulated each application using the prefetch queues generated by our approach and those generated by [13]. Let us denote the average execution times obtained with  $ex_G$  for our approach, and  $ex_{PAP}$  for [13]. Then we computed the performance loss over the baseline for our approach,  $PL_G = \frac{ex_G-baseline}{baseline}$ ; similarly we calculate  $PL_{PAP}$ . Fig. 5a and 5c show the results obtained (averaged over all CFGs in  $Set_1$  and in  $Set_2$ ). As can be seen, for all FPGA sizes, our approach achieves better results compared to [13]: for  $Set_1$ , the performance loss over ideal is between 10% and 11.5% for our method, while for [13] it is between 15.5% and 20% (Fig. 5a). In other words, we are between 27% and 42.6% closer to the ideal baseline than [13]. For  $Set_2$ , we also manage to get from 28% to 41% closer to the ideal baseline than [13] (Fig. 5c).

One other metric suited to evaluate prefetch policies is the total time spent by the application waiting for FPGA reconfigurations to finish (in case the reconfiguration overhead was not entirely hidden). One major difference between the approach proposed in this report and that in [13] is that we also execute candidates from  $\mathcal{H}$  in software (if this is more profitable than reconfiguring and executing on FPGA), while under the assumptions in [13] all candidates from  $\mathcal{H}$  are executed only on FPGA. Considering this, for each execution in software of a candidate  $m \in \mathcal{H}$ , we have no waiting time, but we do not execute m on FPGA either. In this cases, in order to make the comparison to [13] fair, we penalize our approach with sw(m) - hw(m). Let us define the reconfiguration penalty (RP): for [13]  $RP_{PAP}$  is the sum of all waiting times incurred during simulation, and for our approach  $RP_G$  is the sum of all waiting times plus the sum of penalties sw(m) - hw(m) whenever a module  $m \in \mathcal{H}$  is executed in software during simulation. Fig. 5b and 5d show the reconfiguration penalty reduction  $RPR = \frac{RP_{PAP} - RP_G}{RP_{PAP}}$ , averaged over all CFGs in  $Set_1$  and in  $Set_2$ . As we can see, by intelligently generating the prefetches we manage to significantly reduce the penalty (with up to 40%, for both experimental sets), compared to [13].

Concerning the running times of the heuristics, our approach took longer time than [13] to generate the prefetches: from just  $1.6 \times \text{longer}$  in the best case, up to  $12 \times \text{longer}$  in the worst case, incurring on average  $4.5 \times \text{more}$  optimization time. For example, for the 15% FPGA



Figure 6: Comparison with State-of-Art [13]: Case Study - GSM Encoder

fraction, for the biggest CFG in  $Set_2$  (with 268 nodes), the running time of our approach was 3832 seconds, compared to 813 seconds for [13]; for CFGs with a smaller size and a less complex structure we generated a solution in as low as 6 seconds (vs 2 seconds for [13]).

### 7.3 Case Study - GSM Encoder

We also tested our approach on a GSM encoder, which implements the European GSM 06.10 provisional standard for full-rate speech transcoding. This application can be decomposed into 10 functions executed in a sequential order: Init, GetAudioInput, Preprocess, LPC\_Analysis, ShortTermAnalysisFilter, LongTermPredictor, RPE\_Encoding, Add, Encode, Output. The execution times were derived using the MPARM cycle accurate simulator, considering an ARM processor with an operational frequency of 60 MHz. We have identified through profiling the most computation intensive parts of the application, and then these parts were synthesized as hardware modules for an XC5VLX50 Virtex-5 device, using the Xilinx ISE WebPack. The resulting overall CFG of the application contains 30 nodes. The reconfiguration times were estimated considering a 60 MHz configuration clock frequency and the ICAP 32-bit width configuration interface.

The CFG for the GSM Encoder, as well as the profiling information, was generated using the LLVM suite [6] as follows: *llvm-gcc* was first used to generate LLVM bytecode from the C files. The *opt* tool was then used to instrument the bytecode with edge and basic block profiling instructions. The bytecode was next run using *lli*, and then the execution profile was generated using *llvm-prof*. Finally, *opt-analyze* was used to print the CFGs to .dot files. Profiling was run considering several audio files (.au) as input.

Fig. 8a shows the detailed control flow graph (CFG) for our GSM encoder case study. Nodes are labeled with their ID, their execution time (for example, for node with id = 1, the execution time is 10 time units) and their type: root, sink, control nodes, loop header nodes, basic nodes and hardware candidates (denoted with HW and represented with shaded boxes). We have used two scenarios: one considering 5 nodes as hardware candidates (namely modules with IDs 6, 9, 12, 15 and 22), and another scenario considering 9 nodes as hardware candidates (depicted in Fig. 8a).

We have used the same methodology as for the synthetic examples and compared the results using the same metrics defined above in Sec. 7.2, i.e. performance loss over ideal and reconfiguration penalty reduction (presented in Fig. 6). As can be seen, for the scenario with 5 candidate modules, the performance loss over ideal is between 10.5% and 14.8% for our approach, while for [13] it is between 25.5% and 32.9% (Fig. 6a). Thus, we were from 50% up to 65% closer to the ideal baseline than [13]. The reconfiguration penalty reduction obtained is as high as 58.9% (Fig. 6b). For the setting with 9 hardware candidates, the performance loss over ideal is between 10.5% and 135% (Fig. 6c). Thus we manage to get from 52% up to 59% closer to the ideal baseline than [13]. The reconfiguration penalty reduction of up to 60% (Fig. 6d). The prefetches were generated in 27 seconds by our approach and in 11 seconds by [13].



Figure 7: Comparison with State-of-Art [13]: Case Study - Floating Point Benchmark

## 7.4 Case Study - Floating Point Benchmark

Our second case study was a SPECfp benchmark (FP-5 from [19]), characteristic for scientific and computation-intensive applications. Modern FPGAs, coupled with floating-point tools and IP, provide performance levels much higher than software-only solutions for such applications [17]. In order to obtain the inputs needed for our experiments, we used the framework and traces provided for the first Championship Branch Prediction competition [19]. The given instruction trace consists of 30 million instructions, obtained by profiling the program with representative inputs.

We have used the provided framework to reconstruct the control flow graph (CFG) of the FP-5 application based on the given trace. We have obtained a CFG with 65 nodes, after inlining the functions and pruning all control flow edges with a probability lower than  $10^{-5}$ . Then we used the traces to identify the parts of the CFG that have a high execution time (mainly loops).

Fig. 8b shows the detailed control flow graph (CFG) for our second case study: the floating point benchmark from SPECfp. Nodes are represented as discussed in Sec. 7.3. The software execution times for the basic blocks were obtained by considering the following cycles per instruction (CPI) values for each instruction: for calls, returns and floating point instructions CPI = 3, for load, store and branch instructions CPI = 2, and for other instructions CPI = 1. Similar to the previous experimental sections, we considered the hardware execution time  $\beta$  times smaller than the software one, where  $\beta$  was chosen from the uniform distribution on the interval [3, 7].

We have used two scenarios: one considering as hardware candidates the top 9 nodes with the highest software execution times (namely modules with IDs 5, 11, 17, 21, 22, 32, 42, 45 and 56), and another scenario considering the top 25 nodes (depicted in Fig. 8b). Following the same methodology as described in Sec. 7.2, we compared our approach with [13]. The results are presented in Fig. 7. For example, for the scenario with 9 candidate modules, the performance loss over ideal is between 77% and 90% for our approach, while for [13] it is between 107% and 118% (Fig. 7a). Thus, we are from 20% up to 28% closer to the ideal baseline than [13]. The reconfiguration penalty reduction is as high as 28% (Fig. 7b). For the setting with 25 hardware candidates the reconfiguration penalty reduction increases, up to 31.1% (Fig. 7d). As we can see, for both case studies our approach produces significant improvements compared to the state-of-art. The prefetches for FP-5 were generated in 53 seconds by our approach and in 32 seconds by [13].

# 8 Conclusion

In this report we presented a speculative approach to prefetching for FPGA reconfigurations. Based on profiling information, and taking into account the placement of hardware modules on the FPGA, we compute the probabilities to reach each candidate module from every node in the CFG, as well as the distributions of execution time gain obtained by starting a certain prefetch at a certain node. Using this knowledge, we statically schedule the appropriate prefetches (and implicitly do HW/SW partitioning of the candidate hardware modules) such that the expected execution time of the application is minimized. Experiments have shown significant improvement over the state-of-art.

One direction of future work is to develop dynamic prefetching algorithms, that would also capture correlations, for the case when the profile information is either unavailable, or inaccurate. Another direction is to develop an approach sensitive to the execution context, for applications that exhibit execution phases, in which the branch probabilities differ greatly from the average case.

## References

- A. V. Aho, M. S. Lam, R. Sethi, and J. D. Ullman. Compilers: Principles, Techniques, and Tools. Addison Wesley, 2006.
- [2] S. Banerjee, E. Bozorgzadeh, and N. Dutt. Physically-aware HW-SW partitioning for reconfigurable architectures with partial dynamic reconfiguration. In *Design Automation Conference*, 2005.
- [3] R. Cordone, F. Redaelli, M. A. Redaelli, M. D. Santambrogio, and D. Sciuto. Partitioning and scheduling of task graphs on partially dynamically reconfigurable FPGAs. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, 28(5), 2009.
- [4] S. Hauck. Configuration prefetch for single context reconfigurable coprocessors. In Intl. Symp. on Field Programmable Gate Arrays, 1998.
- [5] K. Jiang, P. Eles, and Z. Peng. Co-design techniques for distributed real-time embedded systems with communication security constraints. In *Design Automation and Test in Europe*, 2012.
- [6] C. Lattner and V. Adve. LLVM: A compilation framework for lifelong program analysis & transformation. In Intl. Symp. on Code Generation and Optimization, 2004.
- [7] Z. Li. Configuration management techniques for reconfigurable computing, 2002. PhD thesis, Northwestern Univ., Evanston, IL.
- [8] Z. Li and S. Hauck. Configuration prefetching techniques for partial reconfigurable coprocessor with relocation and defragmentation. In *Intl. Symp. on Field Programmable Gate Arrays*, 2002.
- [9] E. M. Panainte, K. Bertels, and S. Vassiliadis. Instruction scheduling for dynamic hardware configurations. In *Design Automation and Test in Europe*, 2005.
- [10] —. Interprocedural compiler optimization for partial run-time reconfiguration. The Journal of VLSI Signal Processing, 43(2), 2006.
- [11] M. Platzner, J. Teich, and N. Wehn, editors. Dynamically Reconfigurable Systems. Springer, 2010.
- [12] J. Resano, D. Mozos, and F. Catthoor. A hybrid prefetch scheduling heuristic to minimize at run-time the reconfiguration overhead of dynamically reconfigurable hardware. In *Design Automation and Test in Europe*, 2005.
- [13] J. E. Sim, W.-F. Wong, G. Walla, T. Ziermann, and J. Teich. Interprocedural placement-aware configuration prefetching for FPGA-based systems. In *IEEE Symp.* on Field-Programmable Custom Computing Machines, 2010.
- [14] J. E. Sim. Hardware-software codesign for run-time reconfigurable FPGA-based systems, 2010. PhD thesis, National Univ. of Singapore.
- [15] M. Hollander and D. A. Wolfe. Nonparametric Statistical Methods, 2nd Edition. Wiley-Interscience, 1999.
- [16] A. Lifa, P. Eles, and Z. Peng. Performance Optimization of Error Detection Based on Speculative Reconfiguration. In *Design Automation Conference*, 2011.
- [17] Altera. Taking advantage of advances in FPGA floating-point IP cores white paper WP01116. 2009.
- [18] Xilinx. Partial reconfiguration user guide UG702. 2012.
- [19] The 1st championship branch prediction competition. Journal of Instruction-Level Parallelism, 2004. http://www.jilp.org/cbp.



Figure 8: CFGs for Case Studies