

Scheduling divisible loads on partially reconfigurable hardware

K. N. Vikram and V. Vasudevan

Department of Electrical Engineering, Indian Institute of Technology Madras, Chennai, India - 600036
vikramkn@ieee.org, vinita@iitm.ac.in

Abstract

For a task mapped to the reconfigurable fabric (RF) of a hybrid processor-RF architecture, large speedup can be obtained if multiple processing units (PUs) are used to accelerate the task. The overhead for configuring the PUs can be minimized by using partial run-time reconfiguration (PRTR). We present here the results obtained from the quantitative analysis for a single data-parallel task mapped to the RF of a bus-based hybrid processor architecture.

1. Introduction

Reconfigurable hybrid processor architectures, which constitute of a general purpose processor (GPP) coupled to a reconfigurable fabric (RF), allow flexible implementation of any application. Computationally intensive tasks can be mapped to the RF ; whenever a task is encountered during the run-time of an application, the RF can be configured to execute the required task. Typically applications that are mapped to reconfigurable hybrid processors include signal/image processing, multimedia and computer vision. These applications constitute of *data-parallel* tasks whose input data can be partitioned and processed independently by multiple, functionally equivalent processing units (PUs) configured in the RF. Although larger speedups demand more number of PUs to be used, the reconfiguration overhead can become prohibitive.

Use of partial run-time reconfiguration (PRTR) allows us to configure the PUs sequentially one after the other, instead of using a single bitstream for configuring the entire RF area for the required number of PUs. Since PUs used for data-parallel tasks operate independently, each PU can start functioning as soon as the RF area allocated to it is configured. By overlapping reconfiguration of a PU with load transfer and computation on other PUs, the reconfiguration overhead can be minimized.

In a bus-based hybrid processor architecture [], all PUs use a shared data bus for accessing main memory. Reconfiguration delay and limited data bandwidth are the two ar-

chitectural constraints present in such a system. In order to achieve minimum processing time under these constraints, a systematic technique is required for scheduling and allocating load to the PUs. We have developed a framework for this analysis based on divisible load theory (DLT) [3] in our earlier work [4]. The specific case considered in that work was a situation where the results of the computation could be retained within the local memory of the RF itself. The analysis in [4] is not applicable if the RF needs to be reconfigured to perform another task, in which case the computed results will have to be sent back to main memory. We have addressed this problem here. The problem of scheduling with consideration of result collection has received rigorous treatment in [2], for processors in an arbitrary tree network. The bus network considered in this paper is a specific case of an arbitrary tree network. We therefore use the results in [2] as the foundation for performing the required analysis. PRTR introduces an additional dimension to the problem and gives some interesting results.

2. Analysis and Results

The notation that we use for our computation and communication model is given in Table 1, based on DLT. Using the notation given in Table 1, the time taken to transfer load fraction α_i to PU p_i is $\alpha_i z T_{cm}$, while the time taken by p_i for processing it is $\alpha_i w T_{cp}$. If the result data size is the same as input data size, the time taken to transfer the result from p_i back to memory is also $\alpha_i z T_{cm}$.

Reconfiguration of the PUs is assumed to occur continuously using a separate configuration bus, from p_1 to p_n , where n is the number of PUs used. To obtain the minimum processing time, it is possible to derive the following optimality criteria: (1) Each PU should never be idle between its load transfer and result transfer phases, (2) The data bus and any PU should never simultaneously be idle and (3) Result transfer sequence = Load transfer sequence. These criteria also give us the requirement that the result transfer phase should not have any gaps on the data bus.

Performing a quantitative analysis on our system while enforcing the above criteria gives us the following results.

Table 1. Notation used in the analysis of our system, based on DLT. n is the number of PUs used.

Symbol	Description
α_i	Fraction of total load assigned to PU p_i .
w	Ratio of computation time of a PU for a given load, to the computation time of a standard PU for the same load.
z	Ratio of time taken to transmit a given load on the bus, to the time taken to transmit the same load on a standard bus.
T_{cp}	Time taken to process entire load by the standard PU.
T_{cm}	Time taken to transmit entire load on a standard bus.
T_p	Total processing time, including result collection.
T_r	Time taken to configure / reconfigure a single PU.
σ	$wT_{cp}/(zT_{cm})$
β	$(\sigma + 2)/(\sigma + 1)$

When reconfiguration time $T_r \leq zT_{cm}/n$, load fractions allocated to PUs are equal, and no gaps occur during load transfer. The total processing time is then $T_p = T_r + zT_{cm} + (zT_{cm} + wT_{cp})/n$. When $T_r > zT_{cm}/n$, there are gaps in load transfer, as shown in Fig. 1(a). Then the load fractions and optimum processing time are given by

$$\alpha_i = \beta^{i-1} \alpha_1 - \frac{T_r}{zT_{cm}} (\beta^{i-1} - 1), \quad i = 1, \dots, n$$

$$T_p = T_r \left(1 + \frac{1}{\beta - 1} - \frac{n}{\beta^n - 1} \right) + zT_{cm} \left(1 + \frac{1}{\beta^n - 1} \right)$$

where

$$\alpha_1 = \frac{T_r}{zT_{cm}} - \left(\frac{\beta - 1}{\beta^n - 1} \right) \left(\frac{nT_r}{zT_{cm}} - 1 \right) \quad (1)$$

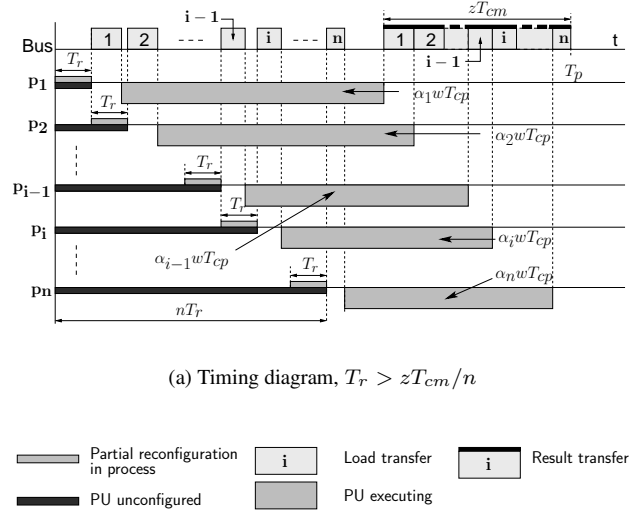
If wT_{cp} is small, some PUs might finish computation even before the end of the load transfer phase. In such cases, an optimal schedule does not exist and we have developed heuristic strategies with the basic idea being that result transfer can occur in the gaps during load transfer.

Fig. 1(c) shows variation of $\phi = T_p/(zT_{cm})$ with the number of PUs, for different values of $\rho = T_r/(zT_{cm})$. The parameter β is based on the computation speed of a 1-D DWT hardware unit. The figure shows that an optimum number of PUs exists for a given ρ , beyond which more PUs do not contribute to speedup.

3. Conclusions


We have presented a theoretical framework for scheduling load for a data-parallel task mapped to the RF of a hybrid processor. The theory gives the maximum speedup that can be obtained, and is also a good approximation when the application load is not arbitrarily divisible.

Acknowledgements: Thanks for financial support.



(a) Timing diagram, $T_r > zT_{cm}/n$

(b) Legend for the timing diagram



(c) Plot of ϕ vs n

Figure 1. Timing diagram for n PUs and plot of ϕ vs n .

References

- [1] G. D. Barlas. Collection-aware optimum sequencing of operations and closed-form solutions for the distribution of a divisible load on arbitrary processor trees. *IEEE Trans. Parallel Distrib. Syst.*, 9(5):429–441, May 1998.
- [2] V. Bharadwaj, D. Ghose, and T. G. Robertazzi. Divisible Load Theory: A new paradigm for load scheduling in distributed systems. *Cluster Computing*, 6(1):7–17, Jan. 2003.
- [3] K. N. Vikram and V. Vasudevan. Mapping data-parallel tasks onto partially reconfigurable hybrid processor architectures. *Submitted to IEEE Trans. VLSI Syst.*, 2005.