Microprocessors and Microsystems 37 (2013) 848-857

Contents lists available at SciVerse ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

Optimizing two-dimensional DMA transfers for scratchpad Based MPSoCs platforms

Selma Saidi^{a,b,*}, Pranav Tendulkar^a, Thierry Lepley^b, Oded Maler^a

^a VERIMAG Lab (CNRS, University of Grenoble), 2 Avenue de Vignate, Gieres, France ^b STMicroelectronics, 12 rue Horowitz, Grenoble, France

ARTICLE INFO

Article history: Available online 17 May 2013

Keywords: Data parallelization Direct memory access (DMA) Double buffering Cell processor MPSoCs

ABSTRACT

Reducing the effects of off-chip memory access latency is a key factor in exploiting efficiently embedded multi-core platforms. We consider architectures that admit a multi-core *computation fabric*, having its own fast and small memory to which the data blocks to be processed are fetched from external memory using a DMA (direct memory access) engine, employing a double- or multiple-buffering scheme to avoid processor idling. In this paper we focus on application programs that process two-dimensional data arrays and we determine automatically the size and shape of the portions of the data array which are subject to a single DMA call, based on hardware and applications parameters. When the computation on different array elements are completely independent, the asymmetry of memory structure leads always to prefer one-dimensional horizontal pieces of memory, while when the computation of a data element shares some data with its neighbors, there is a pressure for more "square" shapes to reduce the amount of redundant data transfers. We provide an analytic model for this optimization problem and validate our results by running a mean filter application on the CELL simulator.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Multiprocessor systems on chip (MPSoCs) such as the CELL processor [1] or the more recent Platform2012 [2] are heterogeneous multi-core architectures, with a powerful host processor and a computation fabric consisting of several smaller cores whose intended role, somewhat similar to graphics processing units (GPUs) [3], is to replace dedicated hardware accelerator. Computationintensive (and parallelizable) parts of the application are offloaded to the multi-cores for execution. Theses parts of the application are often data intensive, operating on large arrays of data initially stored in a remote off-chip memory whose access time is about 100 times slower than that of the cores local memory. A major characteristic of the CELL and P2012 is a software controlled local storage rather than a hidden cache mechanism. Indeed, the local scratchpad memory is usually the only memory directly available to each core, and to access data in the off-chip memory a processor must issue explicit Direct Memory Access (DMA) requests. This simplifies hardware design but at the cost of a greater program complexity for managing explicit memory transfers through the memory hierarchy. To handle these issues new programming models have been introduced such as Sequoia and Cellgen [4-6].

DMA can transfer large amounts of data between memory locations without processor intervention hence it offers another level

E-mail address: Selma.Saidi@imag.fr (S. Saidi).

of parallelism by overlapping computations and data prefetching [1,7]. However, to exploit efficiently these new capabilities, the programmer has to make decisions about the granularity of data transfers and the way they are scheduled to achieve optimal performance. This paper is part of an effort to provide tools that help developers in making such decisions, and ultimately, to automate the whole process of data parallelization for such platforms.

We focus on data parallel applications that exhibit a regular computation and data transfer pattern for which explicit control of data transfers can be more efficient than implicit unpredictable low granularity caching mechanisms [8]. In [9] optimal data transfer granularity for programs that work on one-dimensional arrays of data has been derived. In this paper we tackle the more challenging case of two-dimensional arrays which is common in applications such as video/image processing and certain types of scientific computations and where the transfer using DMA of non-contiguous data blocks is usually required. We extend the DMA performance model defined in [9] to characterize the transfer cost of such blocks. As we explain below, when the computations on array elements are mutually independent, the structure of memories dictates a solution similar to the one-dimensional case. Hence we move onto the situation where the computation for one data block depends also on data from its neighbors. We assume that this shared data is replicated and sent to all processors requiring it for computations. Unlike [9] where the size of replicated data is fixed¹, we capture how the shape of the block







^{*} Corresponding author at: VERIMAG Lab (CNRS, University of Grenoble), 2 Avenue de Vignate, Gieres, France. Tel.: +33 637401143.

^{0141-9331/\$ -} see front matter @ 2013 Elsevier B.V. All rights reserved. http://dx.doi.org/10.1016/j.micpro.2013.04.006

¹ In [9], the size of shared data is always fixed, regardless of the granularity choice, due to the one-dimensional geometry of the array.

influence the size of replicated data and thus the transfer time. We compute optimal size and shape for these data transfers.

While a lot of work has been done in the past to successfully parallelize such applications such as [10–12], contemporary heterogeneous architectures with a limited on-chip memory and a high latency main memory change the formulation and parameters of the problem and call for new solutions in order to achieve high performance on these platforms.

The rest of the paper is organized as follows. In Section 2 we define the applications, the double buffering technique, the DMA cost model and the problem of optimal granularity. In Section 3 we solve the problem for completely independent computations on a single processor. The solution is extended to multi-processors and shared data in Section 4 and then validated experimentally on a simulator of the CELL architecture in Section 5. A discussion of this work and its continuations closes the paper.

2. Preliminaries

2.1. Software pipelining

Consider the sequential algorithm (Program 1) for computing Y = f(X) which *uniformly* applies f to a two-dimensional input array X of $n_1 \times n_2$ elements to produce an output array Y of the same dimension. Array X is stored initially in the off-chip memory and Y is to be written on this memory after computation.

Program 1. Sequential

for $i_1:=1$ to n $_1$ d	0	
for $i_2 := 1$ to n $_2$	do	
$Y(i_1, i_2) := f(X_1)$	$X(i_1, i_2))$	
od		
od		

We assume that an array element represents the minimal gran*ularity* for which the computation of *f* can be carried out. In image processing it can be a pixel, a block or a macroblock. We refer to such granularity as a *basic block*. For the moment we talk only about the logical structure of the array and defer the discussion of the physical memory layout to the sequel. One could handle data transfers at the basic block level by putting explicit DMA calls get and put before and after the computation in the main program loop. However, as is common with other slow memories such as disks, these transfers are applied on a coarser granularity by clustering together several array elements which are brought into a buffer via a single DMA call. We call such clusters super blocks and assume they are rectangular blocks consisting each of $s_1 \times s_2$ basic blocks, see Fig. 1. One can view the super blocks as arranged in an $m_1 \times m_2$ array **X** (and **Y**) with $m_1 = n_1/s_1$ and $m_2 = n_2/s_2$. We use

$$\mathbf{X}(j_1, j_2) = \left\{ X(i_1, i_2) : \frac{(j_1 - 1)s_1 + 1 \leqslant i_1 \leqslant j_1 s_1}{(j_2 - 1)s_2 + 1 \leqslant i_2 \leqslant j_2 s_2} \right\}$$

to denote the set of basic blocks associated with a super block indexed by (j_1, j_2) . It is sometimes more convenient to view twodimensional arrays as one-dimensional and this is done by a flattening function $\phi:[1 \cdots m_1] \times [1 \cdots m_2] \rightarrow [1 \cdots m]$, for $m = m_1 m_2$. We will sometime refer to super block $\mathbf{X}(j_1, j_2)$ as $\mathbf{X}(j)$ for $j = \phi(j_1, j_2)$. We use buffers B_x and B_y for input and output super blocks and rewrite Program 1 as follows.

Program 2. Buffering

for $\mathbf{j} := 1$ to m do	
$get(B_x, \mathbf{X}(j));$	read super block
for $i_1:=1$ to s_1 do	
for $i_2:=1$ to s_2 do	compute for all blocks
$B_y(i_1,i_2) := f(B_x(i_1,i_2))$	in super block j
od	
od	
$put(B_y, \mathbf{Y}(j));$	write super block
od	

In the following, we use comp(j) as a shorthand for the inner double loop.

In program 2, data transfers and computations are performed *sequentially* and the processor is idle during reading and writing. Using *double buffering* for input and output blocks $B_x[0]$, $B_x[1]$, $B_y[0]$ and $B_y[1]$, the processor can work on a super-block *j* residing in one buffer while the DMA brings *in parallel* the super-block *j* + 1 to the other buffer. Program 3 defines a *software pipeline* with 3 stages: input of super block (*j* + 1), computation on super block *j* and output of super block (*j* - 1), see Fig. 2. Reading the first block and writing back the last block are, respectively, the *prologue* and *epilogue* of the pipeline.

Program 3. Double Buffering

c := 0; c' := 1;	
$get(B_x(0), \mathbf{X}(1));$	first read
$get(B_x(1), \mathbf{X}(2)) \parallel comp(1);$	
for $j:=2$ to $m-1$ step 1 do	
$get(B_x(c), \mathbf{X}(j+1)) \ comp(j) \ put(B_y(c'), \mathbf{Y}(j-1));$	
$c:=c\oplus 1;\ c':=c'\oplus 1;$	
od	
$comp(m) \ put(B_y(0), \mathbf{Y}(m-1));$	
$put(B_y(1), \mathbf{Y}(m));$	last write

To reason about the optimal granularity of data transfers we need first to analyze the performance of the pipeline and to this end we need to refine the qualitative description of Fig. 2 which describes the obvious precedences between the computations and data transfers but tells us nothing about their relative durations.

Without loss of generality we assume that a basic (input and output) block consists of a contiguous chunk of *b* bytes,² that the array is organized contiguously in a lexicographic order and that transfer costs in both directions are identical. We denote the transfer time and computation time per super block by *T* and *C*, respectively. Based on the balance between *T* and *C*, the behavior of the software pipeline splits into two cases: the *computation regime*, where C > T and the *transfer regime* when C < T, illustrated in Fig. 3a and b, respectively. It is not hard to see that the total execution time is:

$$m \cdot C + 2T \quad \text{when } C > T \tag{1}$$
$$(m+1) \cdot T \quad \text{when } C < T$$

The relation between the computation time of a super block and its transfer time is not fixed but can be controlled to some extent by varying the *size* and *shape* of the super block. To this end we need to characterize the DMA behavior. In the following we detail the DMA command flow in the CELL architecture on which we base our experiments and derive a model of its performance.

 $^{^2}$ The model can be easily adapted to the case, where it is a $b_1 \times b_2$ rectangle.



Fig. 1. Basic blocks $X(i_1, i_2)$ and super blocks $X(j_1, j_2)$ (logical view).



Fig. 2. A schematic description of a pipeline: Read, compute, and write.



Fig. 3. Pipelined execution using double buffering for one processor in both regimes

2.2. DMA performance model

We start with features that are common to most DMA mechanisms. To copy data from one memory location to another, a processor issues a command to the DMA which takes charge of the transfer. Such a command typically consists of a *source* address, a *destination* address and a block *size*. The command execution decomposes into two major phases:

- 1. *Command initialization* phase: including the command issue time, the time to write the command in the queue and potentially some address translations when virtual memory addressing is used. Note that this phase is *independent* of the amount of data to transfer.
- 2. Data transfer phase: when a command is ready, data transfer begins and the block is divided into smaller packets that travel from source to destination through an interconnect (bus, NoC). The duration of this phase is proportional to the amount of data.

The initialization cost is typically significant and is amortized when the blocks are large. This makes the DMA more attractive for coarse data granularity than for load/store instructions. For one-dimensional data arrays the super blocks are stored as *contiguous* memory segments but this is no more the case for twodimensional data arrays that require usually rectangular data block transfers. Bringing non-contiguous (but regularly structured) pieces of data is possible using *strided* DMA commands that in addition to the source and destination address specify the stride which is an offset to access the next contiguous block in memory. Strided transfers are more expensive than contiguous transfers of the same size.

In the CELL architecture, a strided DMA command is implemented using a *DMA list*, that is, *one* DMA command composed of a list of contiguous transfers. It can be viewed as an array whose entries are pairs consisting of a main memory address and a contiguous transfer size. Each list element can refer to a different location in the external memory. However, all elements of the list move data in the *same* direction (*get* or *put*) and the transferred data should form a contiguous block in local memory.

The cost of transferring a super block of $s_1 \times s_2$ basic blocks (physically a rectangular chunk of memory of s_1 lines and $b \cdot s_2$ columns) can be approximated by the affine function

$$T(s_1, s_2) = I_0 + I_1 \cdot s_1 + \alpha(b \cdot s_1 \cdot s_2)$$

$$\tag{2}$$

This function assumes a fixed initialization $\cot I_0$, independent of the size and form of the super block, an additional initialization $\cot I_1$ associated with each element (contiguous line) in the DMA list, and the transfer itself which is proportional to the size (area) of the super block. Note that in this model we assume a *fixed* transfer latency α . This assumption is imprecise for two major reasons:

- We do not model the characteristics of the external DRAM memory such as the scheduling policy of the memory controller, the effect of page misses and data refreshment latencies.
- The speed of transfer in the interconnect, especially in a multiprocessor setting, depends crucially on the number of simultaneous transfer requests from the processors.

The first issue is too complex to handle precisely as memory controllers vary among vendors. We can assume, however, that page misses are distributed more or less evenly and their effect does not favor or disfavor a specific choice of granularity. Moreover, the preference to contiguous blocks captured by our cost model holds also on the memory controller side. As for the influence of demand patterns on the latency of the interconnect, we will use later a model where α is parameterized by the number p of active processors with $\alpha_p < \alpha_{p'}$ whenever p < p'.

Note that if strided DMA commands are not supported then transferring a rectangular block requires a separate contiguous DMA command for each line of the block. The cost of transferring a super block of $s_1 \times s_2$ basic blocks becomes then: $s_1 \cdot T(1, s_2)$.

We assume the algorithm for computing f to have a fixed (data independent) computation time ω per basic block, once the block is in local memory. This is the time to perform one iteration in Program 1. The cost of computing a super block is therefore

$$C(\mathbf{s}_1, \mathbf{s}_2) = \omega \mathbf{s}_1 \mathbf{s}_2 \tag{3}$$

In the next section, based on (2) and (3), we derive optimal granularity for super blocks in Program 3.

3. Independent computations, single processor

As mentioned previously we can control to some extent the relation between *T* and *C* by controlling the size and shape of super blocks, but which relation is preferred? The answer depends on which resource is more stressed by the application, computation or communication, a fact characterized by

$$\psi = \omega - \alpha b$$

Condition $\psi < 0$ means that regardless of the choice of data granularity, transfer time always dominates computation time. In this case we prefer large data blocks (which corresponds to the maximal buffer size allowed by the local store capacity) to amortize the DMA initialization time and fully utilize the interconnect bandwidth. In the sequel, we focus on the other case where $\psi \ge 0$, that is, more computation than transfer. In this case we opt for a shape that yields a computation regime.

Assuming $\psi > I_1$, the dependence of *T* and *C* on their arguments is illustrated in Fig. 4. The intersection of these two surfaces separates the domain of (s_1, s_2) into two sub-domains, the *computation domain* $T \leq C$ where the computation of a super block dominates the transfer of the next one and the *transfer domain* where T > C, see Fig. 5a.

For the same instance of the problem, computation regime yields a better performance than transfer regime because the processor does not stall between two iterations waiting for data, thereby avoiding idle time. Therefore we orient the super block selection towards shapes (s_1,s_2) such that $C(s_1,s_2) \ge T(s_1,s_2)$ and $m \cdot C(s_1,s_2) + 2T(s_1,s_2)$ is minimal. Since the processor is always



Fig. 4. The dependence of computation *C* and transfer *T* on granularity (s_1, s_2) .

busy, all shapes satisfying C > T admit roughly the same total computation time $m \cdot C(s_1, s_2) \simeq \omega n_1 n_2$ since it is a sequential execution over all the basic blocks. Hence, it remains to optimize the length of the prologue and epilogue $2T(s_1, s_2)$ with computation regime viewed as a constraint. Obvious additional constraints state that a super block is somewhere between a basic block and the full image, provided its size does not exceed the maximum local buffer size *M* imposed by the local store limited capacity. This leads to the following constrained optimization problem.

$$\min T(s_1, s_2) s.t. \ T(s_1, s_2) \leqslant C(s_1, s_2) (s_1, s_2) \in [1..n_1] \times [1..n_2] b \cdot s_1 \cdot s_2 \leqslant M$$
(4)

Note that in addition to these constraints, each specific DMA engine imposes additional constraints on the range possible values of s_1 and s_2 .

Comparing the transfer time of the points (shapes) in the computation domain is not trivial since the computation domain constitutes a partially ordered set where possibly two different shapes *s* and *s'* are such that $s_1 < s'_1$ and $s_2 > s'_2$.³ Observe that the computation domain is convex where for any point *s* inside the domain, we can always find another point *s'* on the boundary such that $s'_2 = s_2$ and $s'_1 < s_1$, see Fig. 5b and hence with a smaller transfer time. Therefore the candidates for optimality are restricted to the *intersection* T = C. These points are of the form $(s_1, H(s_1))$, where

$$H(s_1) = (1/\psi)(I_1 + I_0/s_1)$$

Their transfer time is expressed as a function of the number of clustered horizontal blocks s_1 :

$$T(s_1, H(s_1)) = c(I_0 + I_1 s_1)$$

where *c* is the constant $1 + (\alpha b/\psi)$. This function is linear and monotone in s_1 , means that as we move upwards in the hyperbola *H* the transfer time increases. Hence optimal shape is

$$(\mathbf{s}_1^*, \mathbf{s}_2^*) = (\mathbf{s}_1^*, H(\mathbf{s}_1^*)) = (1, H(1))$$
(5)

which constitutes a contiguous block of one line of the physical data array. This is not surprising as the asymmetry between dimensions in memory access prefers "flat" super blocks with $s_1 = 1$. Without data sharing and memory size constraints the problem becomes similar to the one-dimensional case [9], where it is only the *size* of the super block that needs to be optimized.

For any point $s = (s_1, s_2)$, the area of the rectangle defined by its

³ Note that if these shapes have the same area, then the shape with less lines has a smaller transfer overhead, that is $T(s_1, s_2) < T(s'_1, s'_2)$.



Fig. 5. (a) Computation and transfer domains and (b) optimal granularity candidates and optimal granularity.



Fig. 6. Local memory constraint, (a) excluded solutions and (b) near optimal granularity.

coordinates represents the memory capacity required for this granularity. More generally, the local memory size constraint $s_1 s_2 b \le M$ is represented by the hyperbola shown in Fig. 6, excluding solutions (shapes) above the hyperbola that do not fit in local memory. If optimal granularity (1, H(1)) does not fit in local memory, near optimal solutions that provide good performance results given the local memory budget have to be considered. Such solutions can be other points of the hyperbola H which are, as explained previously, candidates for optimality since they satisfy the equality T = C. However, we know that as we move upwards in the hyperbola *H*, the transfer time increases and consequently computation time which is proportional to the area. Therefore the point (1, H(1))is also the point with minimal transfer time as well as computation time and area (means that it is the closest point of the hyperbola Hto the origin). This proves that if this point does not fit in local memory then no other point (shape) of the hyperbola fits in local memory, which is viewed as the hyperbolas *H* and $s_1s_2b = M$ do not intersect. Optimal shape can therefore be replaced by the point (1,*M*/*b*), see Fig. 6b.

4. Shared Data

So far we considered data independent applications. In the rest of this paper we focus on applications where the computation for each block involves additional *input* data from neighboring blocks. In other words, the computation in the inner loop of Program 1 is replaced by

$$Y(i_1, i_2) = f(V(i_1, i_2))$$

where $V(i_1, i_2)$ denotes a set of basic blocks consisting of $X(i_1, i_2)$ and its neighbors. Without loss of generality, we assume $V(i_1, i_2)$ to be a square around $X(i_1, i_2)$, that is,

$$V(i_1, i_2) = \begin{cases} X(j_1, j_2) : & (i_1 - k/2 \leq j_1 \leq i_1 + k/2) \\ & (i_2 - k/2 \leq j_2 \leq i_2 + k/2) \end{cases}$$

We assume that shared data is *replicated* at each transfer to the local memory. In the sequel we explain how the shape of the block and its replicated area influence the transfer cost and then derive optimal granularity for shared data considering first one processor and then multiple processors.

4.1. Replicated area and transfer cost

To process a super block of shape $s_1 \times s_2$, one needs to load

$$R(s_1, s_2) = (s_1 + k) \times (s_2 + k)$$

basic blocks. In other words, the overhead of replicated external data⁴ is $k(s_1 + s_2) + k^2$ which, among all the super blocks of the same area, is minimal for *square* super blocks as illustrated in Fig. 7. This

⁴ The dark perimeter around the super blocks in Fig. 7.



Fig. 7. Super blocks of different shapes but with the same area $s_1 \times s_2 = 4$. The shaded area represents replicated data overhead for k = 2.

fact is in conflict with the DMA issue overhead, optimized for flat blocks and when both are combined there is a balance to be found between the two.

The DMA transfer cost under sharing becomes

$$T(s_1 + k, s_2 + k) = I_0 + I_1(s_1 + k) + \alpha b \cdot R(s_1, s_2)$$
(6)

Fig. 8a illustrates this function for a fixed value of $\delta = s_1 \times s_2$ along with the DMA issue time overhead optimized for flat block transfer ($s_1 = 1$) and the replicated data transfer overhead optimized for square shapes ($\sqrt{\delta}, \sqrt{\delta}$). Among all combinations (s_1, s_2) satisfying $s_1 \times s_2 = \delta$, the transfer cost is minimal for the point ($s_1^*, \delta/s_1^*$), where $s_1^* = \sqrt{\alpha b k \delta / (l_1 + \alpha b k)}$. This point represents the balance between initialization phase overhead (number of lines) and transfer phase cost (amount of replicated data). In the following section we derive optimal granularity for shared data applications taking this fact into account.

Note that if we look at the computation time of these blocks, then all shapes satisfying $s_1 \times s_2 = \delta$ have approximately the same computation time: $\delta \cdot \omega$, proportional to the area. According to the balance between a block computation time and its transfer time, some shapes will lead to a computation regime and others to a transfer regime as depicted in Fig. 8b, and $(s_1^*, \delta/s_1^*)$ is then the optimal shape for each value of δ (assuming $\delta \cdot \omega \ge T(s_1^* + k, \delta/s_1^* + k)$) since it minimizes the transfer time. In the following, we derive optimal granularity for all shapes yielding a computation regime.

4.2. Optimal granularity for shared data applications

4.2.1. Single processor

With data replication, the constrained optimization Problem 4 becomes

$$\min T(s_1 + k, s_2 + k) s.t. T(s_1 + k, s_2 + k) \leq C(s_1, s_2) (s_1, s_2) \in [1..n_1] \times [1..n_2] b \cdot (s_1 + k) \cdot (s_2 + k) \leq M$$

$$(7)$$

As for independent computations, candidates for optimal granularity are restricted to the points $(s_1, H(s_1))$ satisfying the equality T = C and the problem is reduced to minimizing $T(s_1 + k, H(s_1))$, where

$$H(s_1) = (c_2 s_1 - c_3) / (\psi s_1 - c_1)$$

 c_1, c_2 and c_3 are positive integer constants that depend on I_0, I_1, α, b and k such that

$$\begin{cases} c_1 = \alpha bk \\ c_2 = c_1 + I_1 \\ c_3 = I_0 + I_1 k + \alpha bk^2 \end{cases}$$

 $T(s_1, H(s_1))$ is a second order function with one variable. By computing the derivative, we get one negative point that is not interesting for us and another positive point that is the optimal. To simplify the reading of the formulas, let $\Delta = (c_1/\psi)[1 + D]$, where $D = \sqrt{c_3\alpha/c_1c_2}$, then

$$\begin{cases} s_1^* = \Delta + (c_1/\psi)(1/D) \\ s_2^* = \Delta + (I_1/\psi)(1+D) \end{cases}$$

Fig. 9a illustrates the evolution of the computation domain and the optimal granularity while considering shared data. As discussed in the previous section, we can clearly see that optimal granularity is somewhere between a flat and a square block as s_1^* and s_2^* are both equal to Δ plus a different offset each.

4.2.2. Multiple processors

Given *p* identical processors having the same processing speed and the same local store capacity, the input array is partitioned into *p* chunks of data to be executed in parallel where each processor is responsible of computing its corresponding chunck.

Typically the size of a chunk allocated to a processor is much larger than the local memory capacity, since $p \ll n_1n_2$. Therefore chuncks are further divided into super blocks of (s_1, s_2) basic blocks, see Fig. 10a. Each processor implements double buffering algorithm to fetch the blocks and compute them thereby improving performance by overlapping computations and data transfers. We extend our double-buffer granularity analysis to this case assuming all processors implement the same granularity (s_1, s_2) .

Intuitively, using multiple processors, a conflict arises between computation and data transfers since increasing the number of processors reduces the amount of total work per processor but creates contentions on the shared resources thus increasing the transfer time.

In the analysis, we assume a distributed DMA system where each processor has its own DMA engine. It has the advantage of parallelizing the initialization phase of processors transfer commands which occurs independently on each processor's DMA engine. For this reason we synchronize data transfers of all processors at the beginning of the execution.⁵ Fig. 10b illustrates a pipelined execution using several processors where *p* concurrent transfer requests arrive *simultaneously* to the shared interconnect. Arbitration of these requests is left to the hardware which serves the processors in a low granularity (packet based) round robin fashion. Therefore processors receive their super blocks nearly at the same time and can then perform their computations in parallel. Note that since neighboring data required for the computation of a super block is replicated at each transfer, computations on the different processors are completely

⁵ Note that the gain from overlapping the initialization phase is less significant for large granularities, where the transfer phase time dominates the fixed initialization overhead.



Fig. 8. Given $\delta = s_1 \times s_2$ fixed: (a) DMA transfer cost with replicated area, (b) Balance between computation time and transfer time, as we increase the number of lines in a block.



Fig. 9. Computation domain and optimal granularity considering replication of shared data: (a) Single processor and (b) as we increase the number of processors.



Fig. 10. Double buffering using multiple processors: (a) Data partitioning, shaded area represents replicated data required for the computation of each super block. (b) Pipelined execution.

independent since no synchronization is required to exchange shared.

Increasing the number of processors does not influence the computation time C per super block, however it increases the transfer time because contentions on the shared resources induce a significant overhead that we model by parameterizing the

transfer cost per byte α with the number of active processors such that α_p increases monotonically with p.⁶ We use T_p to denote (6) with α_p replacing α , the total execution time of the pipeline becomes,

⁶ Architectures that have a centralized DMA can avoid this overhead by literally scheduling the transfer at the super block level.

S. Saidi et al. / Microprocessors and Microsystems 37 (2013) 848-857

 Table 1

 Parameters notation.

<i>n</i> ₁ , <i>n</i> ₂	array height and width in number of basic blocks
<i>s</i> ₁ , <i>s</i> ₂	super block height and width in number of basic blocks
b	size in bytes of a basic block
k	shared neighboring data
ω	computation time per basic block
р	number of active processors
Io	DMA initialization cost
I_1	DMA initialization overhead to issue a DMA list element
α	transfer cost in time per byte
α_p	transfer cost in time per byte of <i>p</i> concurrent requests
$T(s_1, s_2)$	transfer time of a super block
$C(s_1, s_2)$	computation time of a super block
M	max local buffer size imposed by the local store capacity

$m/p \cdot C + 2T_p$	when $C \ge T_p$	
$(m/p+1) \cdot T_p$	when $C < T_p$	

Obviously this changes the ratio between the computation time and the transfer time of a super block and consequently the optimal granularity. Fig. 9b shows the evolution of the computation domains and optimal granularity as we increase the number of processors. The reasoning is similar to previously where function *H* becomes H_p thus yielding an optimal data granularity for each value of *p*. Note that the difference between computation and transfer time represented by $\psi = \omega - b\alpha_p$ decreases as we increase *p* reducing the computation domain. Also, beyond some value of *p* we are always in a transfer regime and there is no point in using more processors since our main focus is to optimize processors idle time.

Optimal granularity, which is simply the necessary amount of data needed to hide memory latency, increases as we increase the number of processors since more data needs to be brought to each processor to keep it busy during the time it takes to fetch its next super block *as well as* the next super block of each of the other processors. Note that when optimal granularity does not fit in local memory, decreasing the number of used processors and thus its corresponding optimal granularity can be another way to achieve best performance given the available local memory budget.

Table 1 summarizes the notations for the considered hardware and software parameters.

5. Experiments

We validate our results on the CELL processor whose architecture is shown in Fig. 11. It is by now a decade old architecture, still

Table 2

The transfer time per byte as a function of the number of processes.

	α_p	α_p		
р	min	max	avg	
1	1.13	14.00	2.57	
2	1.78	29.98	4.13	
4	3.97	47.23	11.07	
8	5.43	87.86	18.82	



Fig. 12. Influence of block shape and its replicated data on the transfer time.

favorable for streaming applications. The main features of the architecture include a powerful general purpose processor (PPU – Power Processing Unit) along with eight accelerators (SPU – Syner-gistic Processing Unit). Each SPU has a local scratchpad memory which is the only memory directly accessible using load/store instructions. Data in main memory and other processors memory is accessed using DMA. For more information about the architecture we refer to [13–16].

As an application we use a *mean filter* algorithm that works on a bitmap image of 512×512 pixels. Each pixel is characterized by its intensity ranging over $0 \cdots 255$. The output for a pixel is the average of the value of its neighborhood defined as a square (mask)



Fig. 11. The CELL processor architecture.

centered around it. We have experimented with different mask sizes and focus on the presentation of the results for a 9 × 9 mask, that is, k = 8. In order to use SIMD operations to optimize the implementation of the code, we encode a pixel as an integer (b = 4 bytes). Based on profiling information, we are able to derive the DMA parameter values: fixed initialization cost $I_0 = 108$ and initialization cost per line $I_1 = 50$ cycles. The transfer cost per byte for p processors α_p varies. The minimal, maximal and average values of α_p are shown in Table 2. This variation is mainly due to packet-level arbitration between request of different processors as well as reading and writing of the same processor. We use the average value in our model. The computation workload per basic block is roughly $\omega = 62$ cycles (see remarks at the end of the section).

Note that due to the characteristics of the CELL not all combinations are possible. Indeed a DMA list can hold up to 2 K transfer elements. Each element is a contiguous block transfer with maximum size 16 KBytes (which corresponds in our case to $s_2 = 4096$). Furthermore, the CELL has a strict alignment requirements on 16-byte boundary for both DMA transfers and SPU vector instructions for which the processor is optimized. If this is not taken care of, the DMA engine aligns the data by itself causing erroneous results.

Fig. 12 illustrates the influence of the shape of the block (and its implied replicated area) on the transfer time as explained in Section 4. We consider in this plot different feasible combinations of (s_1, s_2) so that $s_1 \times s_2 = 4096$. A shape (s_1, s_2) yields a block of $s_1 + 8$ lines, each line corresponding to a contiguous transfer of $b \cdot (s_1 + 8)$ bytes. As argued in Section 4, the optimal transfer time is obtained neither for square (64,64) nor the flattest possible (8,512) super blocks and the best trade-off in this case is $(s_1, s_2) = (32, 128)$.

Finally we evaluate the effect of the size and shape of the super blocks and the total execution time of the pipeline for different numbers of processors. Fig. 13a compares the predicted and measured performance for different block shapes where $s_1s_2 = 1024$ while Fig. 13b does the same for $s_1s_2 = 2048$. As one can see, the distance between the predicted and measured values is rather small except for large values of s_1 .

The major reason for the discrepancy between the model and the reality is that $C(s_1, s_2)$ has non negligible component that depends on s_1 for two reasons. The first is due to the overhead at each computation iteration related to the setting required between the outer loop and the inner loop like adjustment of the pointers for every row, pre-calculation of sums of borders, etc. Secondly, the creation of list elements occupies the processor and this overhead is also added to the overall execution time. Fig. 14 combines the measured results for different super block sizes. The measured optimum is obtained for (4,256) while our calculation yield (33,56) whose nearest feasible value is (32,64) whose measured overall performance is less than 10% above the performance for the optimum. The discrepancy can be attributed to the reasons stated above, namely the dependence of *C* on s_1 .

6. Discussion

Adapting array-processing algorithms to multi-core architectures is an activity that will occupy a lot of programmers time in the coming future and it is highly desirable to make it as transparent as possible. Efficient use of the memory hierarchy is crucial for performance on this new class of execution platforms. In this work we have demonstrated how the problem can be approached in a systematic manner for the CELL architecture. Starting from an abstract logical description of the application and the DMA specifications, we could build a model that captures the influence of the size and shape of buffered super blocks on performance. In particular, our model captures the tension between preference of flat super blocks (due to asymmetrical transfer cost) and square super blocks (due to the nature of the application).

We are of course aware of the fact that each real program and each architecture will have its own particularity, more complex and richer in parameters than the model we have built but we



Fig. 14. Observed optimal granularity $s^* = (4, 256)$ and predicted optimal granularity $s^* = (32, 64)$.



Fig. 13. Predicted and measured values for different combinations of $s_1 \times s_2$.

believe that this is a first step toward making such decisions more systematic than by pure trial and error.

Our major observation the experience is that for this type of "collaborative" applications where a large computational task is split, executed and then merged, it is preferable to have a centralized DMA mechanism that can schedule data transfers at the super block rather than the packet level. This way, useless contentions between sub tasks of the *same* application can be avoided. At least in terms of predictability, such policies, used for example in the context of hard real-time systems such as in automotive control [17] will be much simpler. In the future we intend to refine the model and make it more accurate. Then we plan to replicate this work for applications currently being developed for the P2012 architecture where the DMA is more centralized, and the local memory is *shared* between all the cores that reside in the same cluster.

References

- M. Gschwind, The cell broadband engine: exploiting multiple levels of parallelism in a chip multiprocessor, International Journal of Parallel Programming 35 (3) (2007) 233–262.
- [2] D. Melpignano, L. Benini, E. Flamand, B. Jego, T. Lepley, G. Haugou, F. Clermidy, D. Dutoit, Platform 2012, a many-core computing accelerator for embedded SoCs: performance evaluation of visual analytics applications, in: P. Groeneveld, D. Sciuto, S. Hassoun (Eds.), DAC, ACM, 2012, pp. 1137–1142.
- [3] J. Owens, M. Houston, D. Luebke, S. Green, J. Stone, J. Phillips, Gpu computing, Proceedings of the IEEE 96 (5) (2008) 879–899, http://dx.doi.org/10.1109/ JPROC.2008.917757.
- [4] K. Fatahalian, D. Horn, T. Knight, L. Leem, M. Houston, J. Park, M. Erez, M. Ren, A. Aiken, W. Dally, et al., Sequoia: programming the memory hierarchy, in: Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, ACM, 2006 (83–es).
- [5] S. Schneider, J. Yeom, D. Nikolopoulos, Programming multiprocessors with explicitly managed memory hierarchies, Computer 42 (12) (2009) 28–34.
- [6] S. Schneider, J.-S. Yeom, B. Rose, J.C. Linford, A. Sandu, D.S. Nikolopoulos, A comparison of programming models for multiprocessors with explicitly managed memory hierarchies, in: PPOPP, 2009, pp. 131–140.
- [7] J.C. Sancho, K.J. Barker, D.J. Kerbyson, K. Davis, Quantifying the potential benefit of overlapping communication and computation in large-scale scientific applications, in: Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, SC '06, ACM, New York, NY, USA, 2006. http://dx.doi.org/ 10.1145/1188455.1188585.
- [8] J. Sancho, D. Kerbyson, Analysis of double buffering on two different multicore architectures: Quad-core Opteron and the Cell-BE, in: IPDPS 2008. IEEE International Symposium on Parallel and Distributed Processing, IEEE, 2008, pp. 1–12.
- [9] S. Saidi, P. Tendulkar, T. Lepley, O. Maler, Optimizing explicit data transfers for data parallel applications on the cell architecture, ACM Transactions on Architecture Code Optimizer 8 (4) (2012) 37:1–37:20. http://dx.doi.org/ 10.1145/2086696.2086716.
- [10] A. Agarwal, D. Kranz, V. Natarajan, Automatic partitioning of parallel loops and data arrays for distributed shared memory multiprocessors, IEEE Transactions on Parallel Distributed Systems 6 (1995) 943–962.
- [11] C. kin Lee, M. Hamdi, Parallel image processing applications on a network of workstations, Parallel Computing 21 (1) (1995) 137–160. http://dx.doi.org/ 10.1016/0167-8191(94)00068-L, http://www.sciencedirect.com/science/ article/pii/016781919400068L>.
- [12] T. Altilar, Y. Paker, Minimum overhead data partitioning algorithms for parallel video processing, in: Proceedings Domain Decomposition Methods Conference, 2001, pp. 25125–25128.
- [13] IBM, Cell SDK 3.1. <http://www.ibm.com/developerworks/power/cell/>.
- [14] IBM, Cell Simulator, June 2009. http://www.alphaworks.ibm.com/tech/cellsystemsim.
- [15] M. Kistler, M. Perrone, F. Petrini, Cell multiprocessor communication network: built for speed, Micro, IEEE 26 (3) (2006) 10–23, http://dx.doi.org/10.1109/ MM.2006.49.
- [16] J.A. Kahle, M.N. Day, H.P. Hofstee, C.R. Johns, T.R. Maeurer, D. Shippy, Introduction to the cell multiprocessor, IBM Journal of Research and Development 49 (2005) 589–604. http://dl.acm.org/citation.cfm?id=1148882. 1148891>.

[17] H. Kopetz, G. Bauer, The time-triggered architecture, in: Proceedings of the IEEE, 2003, pp. 112–126.



Selma Saidi is a PHD student in Verimag Lab (University of grenoble, France) working in collaboration with STMicroelectronics, she received a computer science engineering degree in 2004 from the university of technology of Algiers, Algeria and then a master degree in 2007. Her research interests include embedded multi-core systems, explicitly managed memory architectures and parallel computing.



Pranav Tendulkar is a Ph.D. student Verimag Laboratory, Gieres, France and University of Joseph Fourier. He received Bachelor of Engineering in electronics and telecommunications engineering from India in 2004. His research interests include algorithm and model design for embedded systems, multicore systems and heterogeneous computing.



Thierry Lepley is an STMicroelectronics Principal Engineer, graduated Electronics and Computer Science from the Grenoble Institute of Technology in 1996. He started research on frequency synthesizers at Carlton University, Ottawa, then moved to the industry with Philips to develop the operating system of smart cards electronic wallet applications. He then entered the compiler field in STMicroelectronics, developing a retargetable optimizing compiler infrastructure for emerging DSP and VLIW processors, with a particular emphasis on instruction scheduling algorithms such as software pipelining. He also developed a compiled simulator

technology for speeding up application simulation and leaded a collaborative research project on real-time adaptive control techniques. Still in STMicroelectronics, Thierry has been responsible from 2009 of the OpenCL activity in the P2012 many-core project, defining the programming methodology and contributing to the development of the compiler, the runtime system, the parallelization of applications. He is also contributing to the OpenCL standardization work-group of the Khronos consortium, representing STMicroelectronics. Author of several scientific publications and industrial patents, his interests are centered on Computing from embedded to HPC, from CPU to GPGPU, covering parallelism at all its granularities, programming models, compilers and computer architecture. His interest also covers the computer vision applicative domain.



Oded Maler is a research director (DR1) at the CNRS (French National Center of Scientific Research), leading the timed and hybrid systems group at VERIMAG, Grenoble. He obtained a B.A. in Computer Science from the Technion, Haifa in 1979, an M.Sc. in Management Science from the University of Tel-Aviv In 1984 and a Ph.D. in Computer Science, from Weizmann Institute, Rehovot. His interests include modeling, simulation and verification for continuous and hybrid systems as well as timing and performance analysis for embedded and other systems.