

Bi-criteria pipeline mappings for parallel image processing

Anne Benoit, Harald Kosch, Veronika Rehn-Sonigo, Yves Robert

► **To cite this version:**

Anne Benoit, Harald Kosch, Veronika Rehn-Sonigo, Yves Robert. Bi-criteria pipeline mappings for parallel image processing. [Research Report] LIP RR-2008-02, Laboratoire de l'informatique du parallélisme. 2008, 2+11p. hal-02102655

HAL Id: hal-02102655

<https://hal-lara.archives-ouvertes.fr/hal-02102655>

Submitted on 17 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Laboratoire de l'Informatique du Parallélisme

École Normale Supérieure de Lyon

Unité Mixte de Recherche CNRS-INRIA-ENS LYON-UCBL n° 5668

*Bi-criteria Pipeline Mappings
for Parallel Image Processing*

Anne Benoit ,
Harald Kosch ,
Veronika Rehn-Sonigo ,
Yves Robert

January 2008

Research Report N° RR2008-02

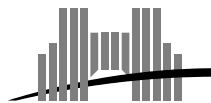
École Normale Supérieure de Lyon

46 Allée d'Italie, 69364 Lyon Cedex 07, France

Téléphone : +33(0)4.72.72.80.37

Télécopieur : +33(0)4.72.72.80.80

Adresse électronique : lip@ens-lyon.fr



Bi-criteria Pipeline Mappings for Parallel Image Processing

Anne Benoit , Harald Kosch , Veronika Rehn-Sonigo , Yves Robert

January 2008

Abstract

Mapping workflow applications onto parallel platforms is a challenging problem, even for simple application patterns such as pipeline graphs. Several antagonistic criteria should be optimized, such as throughput and latency (or a combination). Typical applications include digital image processing, where images are processed in steady-state mode. In this paper, we study the mapping of a particular image processing application, the JPEG encoding. Mapping pipelined JPEG encoding onto parallel platforms is useful for instance for encoding Motion JPEG images. As the bi-criteria mapping problem is NP-complete, we concentrate on the evaluation and performance of polynomial heuristics.

Keywords: pipeline, workflow application, multi-criteria optimization, JPEG encoding

Résumé

L'ordonnement et l'allocation des workflows sur plates-formes parallèles est un problème crucial, même pour des applications simples comme des graphes en pipeline. Plusieurs critères contradictoires doivent être optimisés, tels que le débit et la latence (ou une combinaison des deux). Des applications typiques incluent le traitement d'images numériques, où les images sont traitées en régime permanent.

Dans ce rapport, nous étudions l'ordonnement et l'allocation d'une application de traitement d'image particulière, l'encodage JPEG. L'allocation de l'encodage JPEG pipeliné sur des plates-formes parallèles est par exemple utile pour l'encodage des images Motion JPEG. Comme le problème de l'allocation bi-critère est NP-complet, nous nous concentrons sur l'analyse et évaluation d'heuristiques polynomiales.

Mots-clés: pipeline, application workflow, optimisation multi-critère, encodage JPEG

1 Introduction

This work considers the problem of mapping workflow applications onto parallel platforms. This is a challenging problem, even for simple application patterns. For homogeneous architectures, several scheduling and load-balancing techniques have been developed but the extension to heterogeneous clusters makes the problem more difficult.

Structured programming approaches rule out many of the problems which the low-level parallel application developer is usually confronted to, such as deadlocks or process starvation. We therefore focus on pipeline applications, as they can easily be expressed as algorithmic skeletons. More precisely, in this paper, we study the mapping of a particular pipeline application: we focus on the JPEG encoder (baseline process, basic mode). This image processing application transforms numerical pictures from any format into a standardized format called JPEG. This standard was developed almost 20 years ago to create a portable format for the compression of still images and new versions are created until now (see <http://www.jpeg.org/>). JPEG (and later JPEG 2000) is used for encoding still images in Motion-JPEG (later MJ2). These standards are commonly employed in IP-cams and are part of many video applications in the world of game consoles. Motion-JPEG (M-JPEG) has been adopted and further developed to several other formats, e.g., AMV (alternatively known as MTV) which is a proprietary video file format designed to be consumed on low-resource devices. The manner of encoding in M-JPEG and subsequent formats leads to a flow of still image coding, hence pipeline mapping is appropriate.

We consider the different steps of the encoder as a linear pipeline of stages, where each stage gets some input, has to perform several computations and transfers the output to the next stage. The corresponding mapping problem can be stated informally as follows: which stage to assign to which processor? We require the mapping to be interval-based, i.e., a processor is assigned an interval of consecutive stages. Two key optimization parameters emerge. On the one hand, we target a high throughput, or short period, in order to be able to handle as many images as possible per time unit. On the other hand, we aim at a short response time, or latency, for the processing of each image. These two criteria are antagonistic: intuitively, we obtain a high throughput with many processors to share the work, while we get a small latency by mapping many stages to the same processor in order to avoid the cost of inter-stage communications.

The rest of the paper is organized as follows: Section 2 briefly describes JPEG coding principles. In Section 3 the theoretical and applicative framework is introduced, and Section 4 is dedicated to linear programming formulation of the bi-criteria mapping. In Section 5 we describe some polynomial heuristics, which we use for our experiments of Section 6. We discuss related work in Section 7. Finally, we give some concluding remarks in Section 8.

2 Principles of JPEG encoding

Here we briefly present the mode of operation of a JPEG encoder (see [13] for further details). The encoder consists in seven pipeline stages, as shown in Figure 1. In the first stage, the image is scaled to have a multiple of an 8x8 pixel matrix, and the standard even claims a multiple of 16x16. In the next stage a color space conversion is performed: the colors of the picture are transformed from the RGB to the YUV-color model. The sub-sampling stage is an optional stage, which, depending on the sampling rate, reduces the data volume: as the

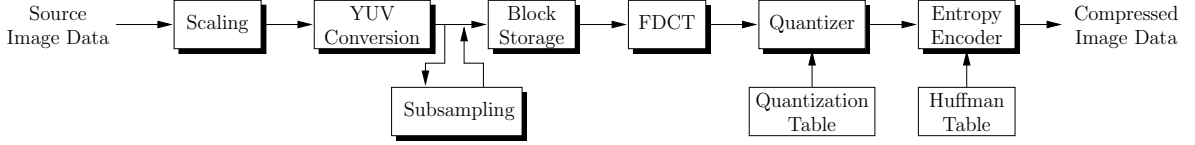


Figure 1: Steps of the JPEG encoding.

human eye can dissolve luminosity more easily than color, the chrominance components are sampled more rarely than the luminance components. Admittedly, this leads to a loss of data. The last preparation step consists in the creation and storage of so-called MCUs (Minimum Coded Units), which correspond to 8x8 pixel blocks in the picture. The next stage is the core of the encoder. It performs a Fast Discrete Cosine Transformation (FDCT) (eg. [14]) on the 8x8 pixel blocks which are interpreted as a discrete signal of 64 values. After the transformation, every point in the matrix is represented as a linear combination of the 64 points. The quantizer reduces the image information to the important parts. Depending on the quantization factor and quantization matrix, irrelevant frequencies are reduced. Thereby quantization errors can occur, that are remarkable as quantization noise or block generation in the encoded image. The last stage is the entropy encoder, which performs a modified Huffman coding: it combines the variable length codes of Huffman coding with the coding of repetitive data in run-length encoding.

3 Framework

3.1 Applicative framework

On the theoretical point of view, we consider a pipeline of n stages \mathcal{S}_k , $1 \leq k \leq n$. Tasks are fed into the pipeline and processed from stage to stage, until they exit the pipeline after the last stage. The k -th stage \mathcal{S}_k first receives an input from the previous stage, of size δ_{k-1} , then performs a number of w_k computations, and finally outputs data of size δ_k to the next stage. These three operations are performed sequentially. The first stage \mathcal{S}_1 receives an input of size δ_0 from the outside world, while the last stage \mathcal{S}_n returns the result, of size δ_n , to the outside world, thus these particular stages behave in the same way as the others.

On the practical point of view, we consider the applicative pipeline of the JPEG encoder as presented in Figure 1 and its seven stages.

3.2 Target platform

We target a platform with p processors P_u , $1 \leq u \leq p$, fully interconnected as a (virtual) clique. There is a bidirectional link $\text{link}_{u,v} : P_u \rightarrow P_v$ between any processor pair P_u and P_v , of bandwidth $b_{u,v}$. The speed of processor P_u is denoted as s_u , and it takes X/s_u time-units for P_u to execute X floating point operations. We also enforce a linear cost model for communications, hence it takes X/b time-units to send (resp. receive) a message of size X to (resp. from) P_v . Communications contention is taken care of by enforcing the *one-port* model [3].

3.3 Bi-criteria interval mapping problem

We seek to map intervals of consecutive stages onto processors [12]. Intuitively, assigning several consecutive tasks to the same processor will increase their computational load, but may well dramatically decrease communication requirements. We search for a partition of $[1..n]$ into $m \leq p$ intervals $I_j = [d_j, e_j]$ such that $d_j \leq e_j$ for $1 \leq j \leq m$, $d_1 = 1$, $d_{j+1} = e_j + 1$ for $1 \leq j \leq m - 1$ and $e_m = n$.

The optimization problem is to determine the best mapping, over all possible partitions into intervals, and over all processor assignments. The objective can be to minimize either the period, or the latency, or a combination: given a threshold period, what is the minimum latency that can be achieved? and the counterpart: given a threshold latency, what is the minimum period that can be achieved?

The decision problem associated to this bi-criteria interval mapping optimization problem is NP-hard, since the period minimization problem is NP-hard for interval-based mappings (see [2]).

4 Linear program formulation

We present here an integer linear program to compute the optimal interval-based bi-criteria mapping on *Fully Heterogeneous* platforms, respecting either a fixed latency or a fixed period. We assume n stages and p processors, plus two fictitious extra stages \mathcal{S}_0 and \mathcal{S}_{n+1} respectively assigned to P_{in} and P_{out} . First we need to define a few variables:

For $k \in [0..n+1]$ and $u \in [1..p] \cup \{\text{in}, \text{out}\}$, $x_{k,u}$ is a boolean variable equal to 1 if stage \mathcal{S}_k is assigned to processor P_u ; we let $x_{0,\text{in}} = x_{n+1,\text{out}} = 1$, and $x_{k,\text{in}} = x_{k,\text{out}} = 0$ for $1 \leq k \leq n$.

For $k \in [0..n]$, $u, v \in [1..p] \cup \{\text{in}, \text{out}\}$ with $u \neq v$, $z_{k,u,v}$ is a boolean variable equal to 1 if stage \mathcal{S}_k is assigned to P_u and stage \mathcal{S}_{k+1} is assigned to P_v : hence $\text{link}_{u,v} : P_u \rightarrow P_v$ is used for the communication between these two stages. If $k \neq 0$ then $z_{k,\text{in},v} = 0$ for all $v \neq \text{in}$ and if $k \neq n$ then $z_{k,u,\text{out}} = 0$ for all $u \neq \text{out}$.

For $k \in [0..n]$ and $u \in [1..p] \cup \{\text{in}, \text{out}\}$, $y_{k,u}$ is a boolean variable equal to 1 if stages \mathcal{S}_k and \mathcal{S}_{k+1} are both assigned to P_u ; we let $y_{k,\text{in}} = y_{k,\text{out}} = 0$ for all k , and $y_{0,u} = y_{n,u} = 0$ for all u . For $u \in [1..p]$, $\text{first}(u)$ is an integer variable which denotes the first stage assigned to P_u ; similarly, $\text{last}(u)$ denotes the last stage assigned to P_u . Thus P_u is assigned the interval $[\text{first}(u), \text{last}(u)]$. Of course $1 \leq \text{first}(u) \leq \text{last}(u) \leq n$.

T_{opt} is the variable to optimize, so depending on the objective function it corresponds either to the period or to the latency.

We list below the constraints that need to be enforced. For simplicity, we write \sum_u instead of $\sum_{u \in [1..p] \cup \{\text{in}, \text{out}\}}$ when summing over all processors. First there are constraints for processor and link usage:

Every stage is assigned a processor: $\forall k \in [0..n+1], \quad \sum_u x_{k,u} = 1$.

Every communication either is assigned a link or collapses because both stages are assigned to the same processor:

$$\forall k \in [0..n], \quad \sum_{u \neq v} z_{k,u,v} + \sum_u y_{k,u} = 1$$

If stage \mathcal{S}_k is assigned to P_u and stage \mathcal{S}_{k+1} to P_v , then $\text{link}_{u,v} : P_u \rightarrow P_v$ is used for this communication:

$$\forall k \in [0..n], \forall u, v \in [1..p] \cup \{\text{in}, \text{out}\}, u \neq v, \quad x_{k,u} + x_{k+1,v} \leq 1 + z_{k,u,v}$$

If both stages \mathcal{S}_k and \mathcal{S}_{k+1} are assigned to P_u , then $y_{k,u} = 1$:

$$\forall k \in [0..n], \forall u \in [1..p] \cup \{\text{in}, \text{out}\}, \quad x_{k,u} + x_{k+1,u} \leq 1 + y_{k,u}$$

If stage \mathcal{S}_k is assigned to P_u , then necessarily $\text{first}_u \leq k \leq \text{last}_u$. We write this constraint as:

$$\forall k \in [1..n], \forall u \in [1..p], \quad \text{first}_u \leq k \cdot x_{k,u} + n \cdot (1 - x_{k,u})$$

$$\forall k \in [1..n], \forall u \in [1..p], \quad \text{last}_u \geq k \cdot x_{k,u}$$

If stage \mathcal{S}_k is assigned to P_u and stage \mathcal{S}_{k+1} is assigned to $P_v \neq P_u$ (i.e., $z_{k,u,v} = 1$) then necessarily $\text{last}_u \leq k$ and $\text{first}_v \geq k + 1$ since we consider intervals. We write this constraint as:

$$\forall k \in [1..n-1], \forall u, v \in [1..p], u \neq v, \quad \text{last}_u \leq k \cdot z_{k,u,v} + n \cdot (1 - z_{k,u,v})$$

$$\forall k \in [1..n-1], \forall u, v \in [1..p], u \neq v, \quad \text{first}_v \geq (k+1) \cdot z_{k,u,v}$$

The latency of schedule is bounded by T_{latency} :

and $t \in [1..p] \cup \{\text{in}, \text{out}\}$.

$$\sum_{u=1}^p \sum_{k=1}^n \left[\left(\sum_{t \neq u} \frac{\delta_{k-1}}{b_{t,u}} z_{k-1,t,u} \right) + \frac{w_k}{s_u} x_{k,u} \right] + \left(\sum_{u \in [1..p] \cup \{\text{in}\}} \frac{\delta_n}{b_{u,\text{out}}} z_{n,u,\text{out}} \right) \leq T_{\text{latency}}$$

and $t \in [1..p] \cup \{\text{in}, \text{out}\}$.

There remains to express the period of each processor and to constrain it by T_{period} :

$$\forall u \in [1..p],$$

$$\sum_{k=1}^n \left\{ \left(\sum_{t \neq u} \frac{\delta_{k-1}}{b_{t,u}} z_{k-1,t,u} \right) + \frac{w_k}{s_u} x_{k,u} + \left(\sum_{v \neq u} \frac{\delta_k}{b_{u,v}} z_{k,u,v} \right) \right\} \leq T_{\text{period}}$$

Finally, the objective function is either to minimize the period T_{period} respecting the fixed latency T_{latency} or to minimize the latency T_{latency} with a fixed period T_{period} . So in the first case we fix T_{latency} and set $T_{\text{opt}} = T_{\text{period}}$. In the second case T_{period} is fixed a priori and $T_{\text{opt}} = T_{\text{latency}}$. With this mechanism the objective function reduces to minimizing T_{opt} in both cases.

5 Overview of the heuristics

The problem of bi-criteria interval mapping of workflow applications is NP-hard [2], so in this section we briefly describe polynomial heuristics to solve it. See [2] for a more complete description or refer to the Web at:

<http://graal.ens-lyon.fr/~vsonigo/code/multicriteria/>

In the following, we denote by n the number of stages, and by p the number of processors. We distinguish two sets of heuristics. The heuristics of the first set aim to minimize the latency respecting an a priori fixed period. The heuristics of the second set minimize the counterpart: the latency is fixed a priori and we try to achieve a minimum period while respecting the latency constraint.

5.1 Minimizing Latency for a Fixed Period

All the following heuristics sort processors by non-increasing speed, and start by assigning all the stages to the first (fastest) processor in the list. This processor becomes *used*.

H1-Sp-mono-P: Splitting mono-criterion – At each step, we select the used processor j with the largest period and we try to split its stage interval, giving some stages to the next fastest processor j' in the list (not yet used). This can be done by splitting the interval at any place, and either placing the first part of the interval on j and the remainder on j' , or the other way round. The solution which minimizes $\max(\text{period}(j), \text{period}(j'))$ is chosen if it is better than the original solution. Splitting is performed as long as we have not reached the fixed period or until we cannot improve the period anymore.

H2-Sp-bi-P: Splitting bi-criteria – This heuristic uses a binary search over the latency. For this purpose at each iteration we fix an authorized increase of the optimal latency (which is obtained by mapping all stages on the fastest processor), and we test if we get a feasible solution via splitting. The splitting mechanism itself is quite similar to **H1-Sp-mono-P** except that we choose the solution that minimizes $\max_{i \in \{j, j'\}} (\frac{\Delta \text{latency}}{\Delta \text{period}(j)})$ within the authorized latency increase to decide where to split. While we get a feasible solution, we reduce the authorized latency increase for the next iteration of the binary search, thereby aiming at minimizing the mapping global latency.

H3-3-Sp-mono-P: 3-splitting mono-criterion – At each step we select the used processor j with the largest period and we split its interval into three parts. For this purpose we try to map two parts of the interval on the next pair of fastest processors in the list, j' and j'' , and to keep the third part on processor j . Testing all possible permutations and all possible positions where to cut, we choose the solution that minimizes $\max(\text{period}(j), \text{period}(j'), \text{period}(j''))$.

H4-3-Sp-bi-P: 3-splitting bi-criteria – In this heuristic the choice of where to split is more elaborated: it depends not only of the period improvement, but also of the latency increase. Using the same splitting mechanism as in **H3-3-Sp-mono-P**, we select the solution that minimizes $\max_{i \in \{j, j', j''\}} (\frac{\Delta \text{latency}}{\Delta \text{period}(i)})$. Here $\Delta \text{latency}$ denotes the difference between the global latency of the solution before the split and after the split. In the same manner $\Delta \text{period}(i)$ defines the difference between the period before the split (achieved by processor j) and the new period of processor i .

5.2 Minimizing Period for a Fixed Latency

As in the heuristics described above, first of all we sort processors according to their speed and map all stages on the fastest processor.

H5-Sp-mono-L: Splitting mono-criterion – This heuristic uses the same method as **H1-Sp-mono-P** with a different break condition. Here splitting is performed as long as we do not exceed the fixed latency, still choosing the solution that minimizes $\max(\text{period}(j), \text{period}(j'))$.

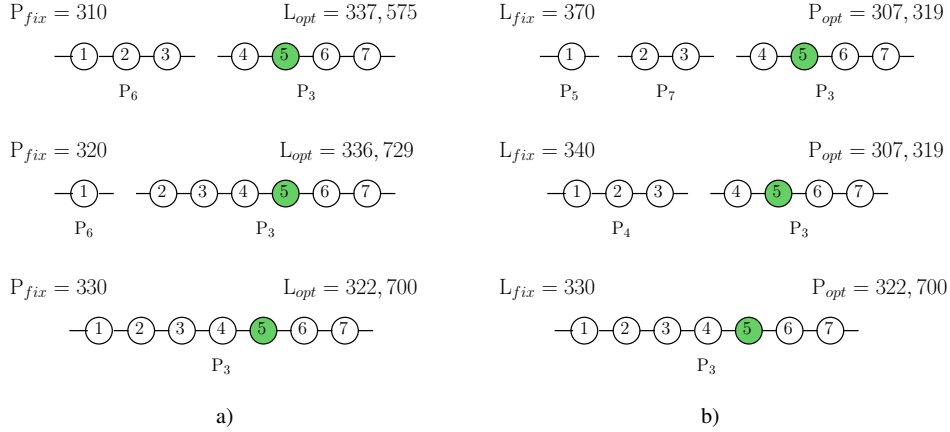


Figure 2: LP solutions strongly depend on fixed initial parameters.

H6-Sp-bi-L: Splitting bi-criteria – This variant of the splitting heuristic works similarly to **H5-Sp-mono-L**, but at each step it chooses the solution which minimizes $\max_{i \in \{j, j'\}} \left(\frac{\Delta latency}{\Delta period(i)} \right)$ while the fixed latency is not exceeded.

Remark In the context of M-JPEG coding, minimizing the latency for a fixed period corresponds to a fixed coding rate, and we want to minimize the response time. The counterpart (minimizing the period respecting a fixed latency L) corresponds to the question: if I accept to wait L time units for a given image, which coding rate can I achieve? We evaluate the behavior of the heuristics with respect to these questions in Section 6.2.

6 Experiments and simulations

In the following experiments, we study the mapping of the JPEG application onto clusters of workstations.

6.1 Influence of fixed parameters

In this first test series, we examine the influence of fixed parameters on the solution of the linear program. As shown in Figure 2, the division into intervals is highly dependant of the chosen fixed value. The optimal solution to minimize the latency (without any supplemental constraints) obviously consists in mapping the whole application pipeline onto the fastest processor. As expected, if the period fixed in the linear program is not smaller than the latter optimal mono-criterion latency, this solution is chosen. Decreasing the value for the fixed period imposes to split the stages among several processors, until no more solution can be found. Figure 2(a) shows the division into intervals for a fixed period. A fixed period of $T_{period} = 330$ is sufficiently high for the whole pipeline to be mapped onto the fastest processor, whereas smaller periods lead to splitting into intervals. We would like to mention, that for a period fixed to 300, there exists no solution anymore. The counterpart - fixed latency - can be found in Figure 2(b). Note that the first two solutions find the same period, but for a different latency. The first solution has a high value for latency, which allows more splits, hence larger communication costs. Comparing the last lines of Figures 2(a) and (b), we

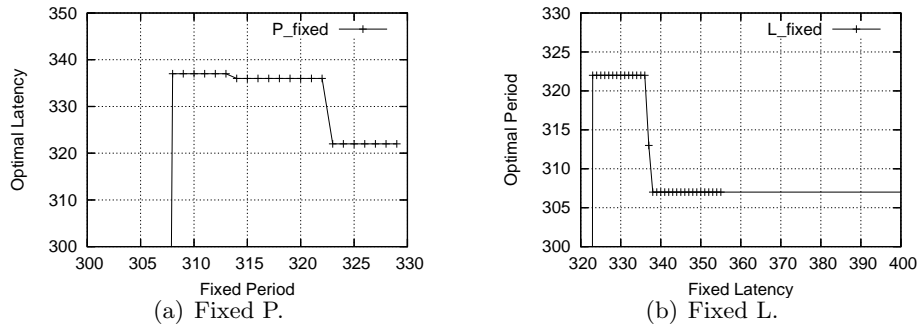


Figure 3: Bucket behavior of LP solutions.

state that both solutions are the same, and we have $T_{\text{period}} = T_{\text{latency}}$. Finally, expanding the range of the fixed values, a sort of bucket behavior becomes apparent: Increasing the fixed parameter has in a first time no influence, the LP still finds the same solution until the increase crosses an unknown bound and the LP can find a better solution. This phenomenon is shown in Figure 3.

6.2 Assessing heuristic performance

The comparison of the solution returned by the LP program, in terms of optimal latency respecting a fixed period (or the converse) with the heuristics is shown in Figure 4. The implementation is fed with the parameters of the JPEG encoding pipeline and computes the mapping on 10 randomly created platforms with 10 processors. On platforms 3 and 5, no valid solution can be found for the fixed period. There are two important points to mention. First, the solutions found by H2 often are not valid, since they do not respect the fixed period, but they have the best ratio latency/period. Figure 5(b) plots some more details: H2 achieves good latency results, but the fixed period of $P=310$ is often violated. This is a consequence of the fact that the fixed period value is very close to the feasible period. When the tolerance for the period is bigger, this heuristic succeeds to find low-latency solutions. Second, all solutions, LP and heuristics, always keep the stages 4 to 7 together (see Figure 2 for an example). As stage 5 (DCT) is the most costly in terms of computation, the interval containing these stages is responsible for the period of the whole application.

Finally, in the comparative study H1 always finds the optimal period for a fixed latency and we therefore recommend this heuristic for period optimization. In the case of latency minimization for a fixed period, then H5 is to use, as it always finds the LP solution in the experiments. This is a striking result, especially given the fact that the LP integer program may require a long time to compute the solution (up to 11389 seconds in our experiments), while the heuristics always complete in less than a second, and find the corresponding optimal solution.

6.3 MPI simulations on a cluster

This last experiment performs a JPEG encoding simulation. All simulations are made on a cluster of homogeneous Optiplex GX 745 machines with an Intel Core 2 Duo 6300 of 1,83Ghz. Heterogeneity is enforced by increasing and decreasing the number of operations a processor

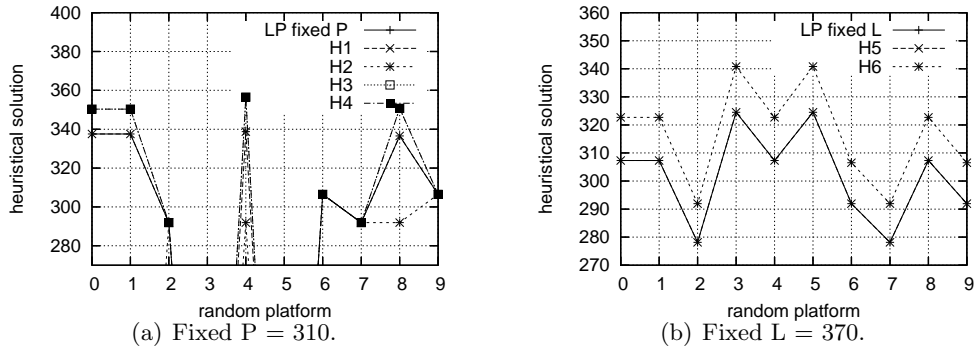


Figure 4: Behavior of the heuristics (comparing to LP solution).

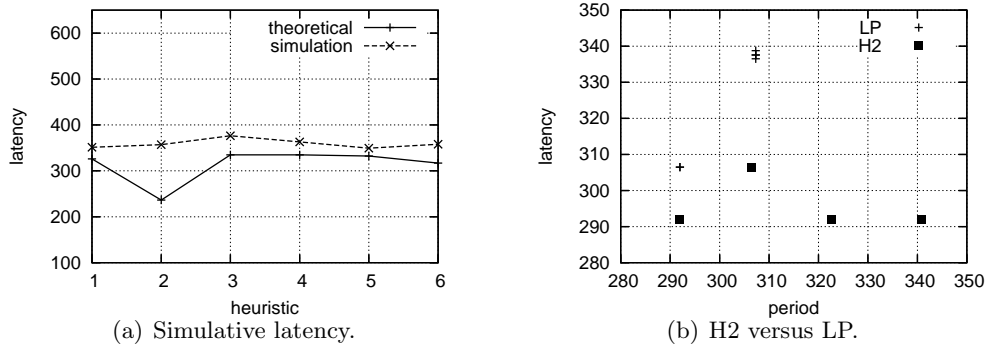


Figure 5: MPI simulation results.

has to execute. The same holds for bandwidth capacities. We call this experiment simulation, as we do not parallelize a real JPEG encoder, but we use a parallel pipeline application which has the same parameters for communication and computation as the JPEG encoder. An mpich implementation of MPI is used for communications.

In this experiment the same random platforms with 10 processors and fixed parameters as in the theoretical experiments are used. We measured the latency of the simulation, even for the heuristics of fixed latency, and computed the average over all random platforms. Figure 5(a) compares the average of the theoretical results of the heuristics to the average simulative performance. The simulative behavior nicely mirrors the theoretical behavior, with the exception of H2 (see Figure 5(b)). Here once again, some solutions of this heuristic are not valid, as they do not respect the fixed period.

7 Related work

The blockwise independent processing of the JPEG encoder allows to apply simple data parallelism for efficient parallelization. Many papers have addressed this fine-grain parallelization opportunity [5, 11]. In addition, parallelization of almost all stages, from color space conversion, over DCT to the Huffman encoding has been addressed [1, 7]. Recently, with respect to the JPEG2000 codec, efficient parallelization of wavelet coding has been introduced [8]. All these works target the best speed-up with respect to different architectures and possible

varying load situations. Optimizing the period and the latency is an important issue when encoding a pipeline of multiple images, as for instance for Motion JPEG (M-JPEG). To meet these issues, one has to solve in addition to the above mentioned work a bi-criteria optimization problem, i.e., optimize the latency, as well as the period. The application of coarse grain parallelism seems to be a promising solution. We propose to use an interval-based mapping strategy allowing multiple stages to be mapped to one processor which allows meeting the most flexible the domain constraints (even for very large pictures). Several pipelined versions of the JPEG encoding have been considered. They rely mainly on pixel or block-wise parallelization [6, 9]. For instance, Ferretti et al. [6] uses three pipelines to carry out concurrently the encoding on independent pixels extracted from the serial stream of incoming data. The pixel and block-based approach is however useful for small pictures only. Recently, Sheel et al. [10] consider a pipeline architecture where each stage presents a step in the JPEG encoding. The targeted architecture consists of Xtensa LX processors which run subprograms of the JPEG encoder program. Each program accepts data via the queues of the processor, performs the necessary computation, and finally pushes it to the output queue into the next stage of the pipeline. The basic assumptions are similar to our work, however no optimization problem is considered and only runtime (latency) measurements are available. The schedule is static and set according to basic assumptions about the image processing, e.g., that the DCT is the most complex operation in runtime.

8 Conclusion

In this paper, we have studied the bi-criteria (minimizing latency and period) mapping of pipeline workflow applications, from both a theoretical and practical point of view. On the theoretical side, we have presented an integer linear programming formulation for this NP-hard problem. On the practical side, we have studied in depth the interval mapping of the JPEG encoding pipeline on a cluster of workstations. Owing to the LP solution, we were able to characterize a bucket behavior in the optimal solution, depending on the initial parameters. Furthermore, we have compared the behavior of some polynomial heuristics to the LP solution and we were able to recommend two heuristics with almost optimal behavior for parallel JPEG encoding. Finally, we evaluated the heuristics running a parallel pipeline application with the same parameters as a JPEG encoder. The heuristics were designed for general pipeline applications, and some of them were aiming at applications with a large number of stages (3-splitting), thus a priori not very efficient on the JPEG encoder. Still, some of these heuristics reach the optimal solution in our experiments, which is a striking result.

A natural extension of this work would be to consider further image processing applications with more pipeline stages or a slightly more complicated pipeline architecture. Naturally, our work extends to JPEG 2000 encoding which offers among others wavelet coding and more complex multiple-component image encoding [4]. Another extension is for the MPEG coding family which uses lagged feedback: the coding of some types of frames depends on other frames. Differentiating the types of coding algorithms, a pipeline architecture seems again to be a promising solution architecture.

References

- [1] L. V. Agostini, I. S. Silva, and S. Bampi. Parallel color space converters for JPEG image compression. *Microelectronics Reliability*, 44(4):697–703, 2004.
- [2] A. Benoit, V. Rehn-Sonigo, and Y. Robert. Multi-criteria Scheduling of Pipeline Workflows. In *HeteroPar'07, Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks (in conjunction with Cluster 2007)*. IEEE Computer Society Press, 2007.
- [3] P. Bhat, C. Raghavendra, and V. Prasanna. Efficient collective communication in distributed heterogeneous systems. *Journal of Parallel and Distributed Computing*, 63:251–263, 2003.
- [4] C. Christopoulos, A. Skodras, and T. Ebrahimi. The JPEG2000 still image coding system: an overview. *IEEE Transactions on Consumer Electronics*, 46(4):1103–1127, 2000.
- [5] J. Falkemeier and G. Joubert. Parallel image compression with jpeg for multimedia applications. In *High Performance Computing: Technologies, Methods and Applications*, Advances in Parallel Computing, pages 379–394. North Holland, 1995.
- [6] M. Ferretti and M. Boffadossi. A Parallel Pipelined Implementation of LOCO-I for JPEG-LS. In *17th International Conference on Pattern Recognition (ICPR'04)*, volume 1, pages 769–772, 2004.
- [7] T. Kumaki, M. Ishizaki, T. Koide, H. J. Mattausch, Y. Kuroda, H. Noda, K. Dosaka, K. Arimoto, and K. Saito. Acceleration of DCT Processing with Massive-Parallel Memory-Embedded SIMD Matrix Processor. *IEICE Transactions on Information and Systems - LETTER- Image Processing and Video Processing*, E90-D(8):1312–1315, 2007.
- [8] P. Meerwald, R. Norcen, and A. Uhl. Parallel JPEG2000 Image Coding on Multiprocessors. In *IPDPS'02, International Parallel and Distributed Processing Symposium*. IEEE Computer Society Press, 2002.
- [9] M. Papadonikolakis, V. Pantazis, and A. P. Kakarountas. Efficient high-performance ASIC implementation of JPEG-LS encoder. In *Proceedings of the Conference on Design, Automation and Test in Europe (DATE2007)*, volume IEEE Communications Society Press, 2007.
- [10] S. L. Shee, A. Erdos, and S. Parameswaran. Architectural Exploration of Heterogeneous Multiprocessor Systems for JPEG. *International Journal of Parallel Programming*, 35, 2007.
- [11] K. Shen, G. Cook, L. Jamieson, and E. Delp. An overview of parallel processing approaches to image and video compression. In *Image and Video Compression*, volume Proc. SPIE 2186, pages 197–208, 1994.
- [12] J. Subhlok and G. Vondran. Optimal latency-throughput tradeoffs for data parallel pipelines. In *ACM Symposium on Parallel Algorithms and Architectures SPAA '96*, pages 62–71. ACM Press, 1996.

- [13] G. K. Wallace. The JPEG still picture compression standard. *Commun. ACM*, 34(4):30–44, 1991.
- [14] C. Wen-Hsiung, C. Smith, and S. Fralick. A Fast Computational Algorithm for the Discrete Cosine Transform. *IEEE Transactions on Communications*, 25(9):1004–1009, 1977.