ANNE BENOIT, École Normale Supérieure de Lyon ÜMIT V. ÇATALYÜREK, The Ohio State University YVES ROBERT, École Normale Supérieure de Lyon and University of Tennessee Knoxville ERIK SAULE, The Ohio State University

A large class of applications need to execute the same workflow on different datasets of identical size. Efficient execution of such applications necessitates intelligent distribution of the application components and tasks on a parallel machine, and the execution can be orchestrated by utilizing task, data, pipelined, and/or replicated parallelism. The scheduling problem that encompasses all of these techniques is called *pipelined workflow scheduling*, and it has been widely studied in the last decade. Multiple models and algorithms have flourished to tackle various programming paradigms, constraints, machine behaviors, or optimization goals. This article surveys the field by summing up and structuring known results and approaches.

Categories and Subject Descriptors: F.2.2 [Nonnumerical Algorithms and Problems]: Sequencing and scheduling; C.1.4 [Parallel Architectures]: Distributed architectures

General Terms: Algorithms, Performance, Theory

Additional Key Words and Phrases: Workflow programming, filter-stream programming, scheduling, pipeline, throughput, latency, models, algorithms, distributed systems, parallel systems

ACM Reference Format:

Benoit, A., Catalyurek, U. V., Robert, Y., and Saule, E. 2013. A survey of pipelined workflow scheduling: Models and algorithms. ACM Comput. Surv. 45, 4, Article 50 (August 2013), 36 pages. DOI: http://dx.doi.org/10.1145/2501654.2501664

1. INTRODUCTION

For large-scale applications targeted to parallel and distributed computers, finding an efficient task and communication mapping and schedule is critical to reach the best possible application performance. At the heart of the scheduling process is a graph, the *workflow*, of an application: an abstract representation that expresses the atomic computation units and their data dependencies. Hence, the application is partitioned into tasks that are linked by precedence constraints, and it is described by, usually,

A. Benoit and Y. Robert are currently affiliated with the Institut Universitaire de France.

© 2013 ACM 0360-0300/2013/08-ART50 \$15.00 DOI: http://dx.doi.org/10.1145/2501654.2501664

This work was partially supported by the DOE grant DE-FC02-06ER2775, by AFRL/DAGSI Ohio Student-Faculty Research Fellowship RY6-OSU-08-3, by the NSF grants CNS-0643969, OCI-0904809, and OCI-0904802, and by the French ANR StochaGrid and RESCUE projects.

Authors' addresses: A. Benoit, Laboratoire de l'Informatique du Parallélisme, École Normale Supérieure de Lyon, 46 allée d'Italie 69364 Lyon, Cedex 07, France; emial: Anne.Benoit@ens-lyon.fr; Ü. V. Çatalyürek, Department of Biomedical Informatics and Department of Electrical and Computer Engineering, The Ohio State University, 3190 Graves Hall, 333 W. Tenth Ave., Colombus, OH 43210; email: unit@bmi.osu.edu; Y. Robert, Laboratoire de l'Informatique du Parallélisme, École Normale Supérieure de Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France; email: Yves.Robert@ens-lyon.fr; E. Saule (corresponding author), Department of Biomedical Informatics, The Ohio State University, 3190 Graves Hall, 333 W. Tenth Ave., Colombus, OH 43210; email: esaule@bmi.osu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

a Directed Acyclic Graph (also called DAG), where the vertices are the tasks and the edges represent the precedence constraints. In classical workflow scheduling techniques, there is a single dataset to be executed, and the goal is to minimize the *latency* or *makespan*, which corresponds to the total execution time of the workflow, where each task is executed once [Kwok and Ahmad 1999b].

The graphical representations are not only used for parallelizing computations. In the mid-70s and early 80s, a graphical representation called dataflow [Dennis 1974, 1980; Davis 1978] emerged as a powerful programming and architectural paradigm. Lee and Parks [1995] present a rigorous formal foundation of dataflow languages, for which they coined the term *dataflow process networks* and presented it as a special case of Kahn Process Networks (KPN) [Kahn 1974]. In KPN, a group of deterministic sequential tasks communicate through unbounded first-in, first-out channels. As a powerful paradigm that implicitely supports parallelism, dataflow networks (hence KPNs) have been used to exploit parallelism at compile time [Ha and Lee 1997] and runtime [Nikolov et al. 2008].

With the turn of the new millennium, Grid computing [Foster et al. 2001] emerged as a global cyber-infrastructure for large-scale, integrative e-Science applications. At the core of Grid computing sit Grid workflow managers that schedule coarse-grain computations onto dynamic Grid resources. Yu and Buyya [2005] present an excellent survey on workflow scheduling for Grid computing. Grid workflow managers, such as DAGMan [Tannenbaum et al. 2001] (of the Condor project [Litzkow et al. 1988; Thain et al. 2002]), Pegasus [Deelman et al. 2005], GrADS [Berman et al. 2001], Taverna [Oinn et al. 2006], and ASKALON [Fahringer et al. 2005], utilize DAGs and *abstract workflow languages* for scheduling workflows onto dynamic Grid resources using performance modeling and prediction systems like Prophesy [Taylor et al. 2003], NWS [Wolski et al. 1999], and Teuta [Fahringer et al. 2004]. The main focus of such Grid workflow systems is the discovery and utilization of dynamic resources that span over multiple administrative domains. It involves handling of authentication and authorization, efficient data transfers, and fault tolerance due to the dynamic nature of the systems.

The main focus of this article is a special class of workflow scheduling that we call *pipelined workflow scheduling* (or in short *pipelined scheduling*). Indeed, we focus on the scheduling of applications that continuously operate on a stream of *datasets*, which are processed by a given wokflow, and hence the term *pipelined*. In steady state, similar to dataflow and Kahn networks, datasets are pumped from one task to its successor. These datasets all have the same size, and they might be obtained by partitioning the input into several chunks. For instance, in image analysis [Sertel et al. 2009], a medical image is partitioned in tiles, and tiles are processed one after the other. Other examples of such applications include video processing [Guirado et al. 2005], motion detection [Knobe et al. 1999], signal processing [Choudhary et al. 2000; Hartley et al. 2009], databases [Chekuri et al. 1995], molecular biology [Rowe et al. 2003], medical imaging [Guirado et al. 2006], and various scientific data analyses, including particle physics [Deelman et al. 2003], earthquake [Kim et al. 2004], weather, and environmental data analyses [Rowe et al. 2003].

The pipelined execution model is the core of many software and programming middlewares. It is used on different types of parallel machines such as SMP (Intel TBB [Reinders 2007]), clusters (DataCutter [Beynon et al. 2001], Anthill [Teodoro et al. 2008], Dryad [Isard et al. 2007]), Grid computing environments (Microsoft AXUM [Microsoft 2009], LONI [MacKenzie-Graham et al. 2008], Kepler [Bowers et al. 2006]), and more recently on clusters with accelerators (see, for instance, DataCutter [Hartley et al. 2008] and DataCutter-Lite [Hartley and Çatalyürek 2009]). Multiple models and algorithms have emerged to deal with various programming paradigms, hardware constraints, and scheduling objectives.

It is possible to reuse classical workflow scheduling techniques for pipelined applications, by first finding an efficient parallel execution schedule for one single dataset (makespan minimization), and then executing all the datasets using the same schedule. one after the other. Although some good algorithms are known for such problems [Kwok and Ahmad 1999a; 1999b], the resulting performance of the system for a pipelined application may be far from the peak performance of the target parallel platform. The workflow may have a limited degree of parallelism for efficient processing of a single dataset, and hence the parallel machine may not be fully utilized. Rather, for pipelined applications, we need to decide how to process multiple datasets in parallel. In other words, pipelined scheduling is dealing with both intra-dataset and inter-dataset parallelism (the different types of parallelism are described shortly in more details). Applications that do not allow the latter kind of parallelism are outside the scope of this survey. Such applications include those with a feedback loop such as iterative solvers. When feedback loops are present, applications are typically scheduled by software pipelining, or by cyclic scheduling techniques (also called cyclic PERT-shop scheduling, where PERT refers to Project Evaluation and Review Technique). A survey on software pipelining can be found in Allan et al. [1995], and on cyclic scheduling in Levner et al. [2010].

To evaluate the performance of a schedule for a pipelined workflow, various optimization criteria are used in the literature. The most common ones are: (i) the *latency* (denoted by \mathcal{L}), or *makespan*, which is the maximum time a dataset spends in the system, and (ii) the *throughput* (denoted by \mathcal{T}), which is the number of datasets processed per time unit. The period of the schedule (denoted by \mathcal{P}) is the time elapsed between two consecutive datasets entering the system. Note that the period is the inverse of the achieved throughput, hence we will use them interchangeably. Depending on the application, a combination of multiple performance objectives may be desired. For instance, an interactive video processing application (such as SmartKiosk [Knobe et al. 1999], a computerized system that interacts with multiple people using cameras) needs to be reactive while ensuring a good frame rate; these constraints call for an efficient latency/throughput trade-off. Other criteria may include reliability, resource cost, and energy consumption.

Several types of parallelism can be used to achieve good performance. If one task of the workflow produces directly or transitively the input of another task, the two tasks are said to be *dependent*; otherwise they are *independent*. *Task parallelism* is the most well-known form of parallelism and consists in concurrently executing independent tasks for the same dataset; it can help minimize the workflow latency.

Pipelined parallelism is used when two dependent tasks in the workflow are being executed simultaneously on different datasets. The goal is to improve the throughput of the application, possibly at the price of more communications, hence potentially a larger latency. Pipelined parallelism was made famous by assembly lines and later reused in processors in the form of the instruction pipeline in CPUs and the graphic rendering pipeline in GPUs.

Replicated parallelism can improve the throughput of the application, because several copies of a single task operate on different datasets concurrently. This is especially useful in situations where more computational resources than workflow tasks are available. Replicated parallelism is possible when reordering the processing of the datasets by one task does not break the application semantics, for instance, when the tasks perform a stateless transformation. A simple example of a task allowing replicated parallelism would be computing the square root of the dataset (a number), while computing the sum of the numbers processed so far would be stateful and would not allow replicated parallelism.

Finally, *data parallelism* may be used when some tasks contain inherent parallelism. It corresponds to using several processors to execute a single task for a single dataset. It is commonly used when a task is implemented by a software library that supports parallelism on its own, or when a strongly coupled parallel execution can be performed.

Note that task parallelism and data parallelism are inherited from classical workflow scheduling, while pipelined parallelism and replicated parallelism are only found in pipelined workflow scheduling.

In a nutshell, the main contributions of this survey are the following: (i) proposing a three-tiered model of pipelined workflow scheduling problems; (ii) structuring existing work; and (iii) providing detailed explanations on schedule reconstruction techniques, which are often implicit in the literature.

The rest of this article is organized as follows. Before going into technical details, Section 2 presents a motivating example to illustrate the various parallelism techniques, task properties, and their impact on objective functions.

The first issue when dealing with a pipelined application is to select the right model among the tremendous number of variants that exist. To solve this issue, Section 3 organizes the different characteristics that the target application can exhibit into three components: the workflow model, the system model, and the performance model. This organization helps position a given problem with respect to related work.

The second issue is to build the relevant scheduling problem from the model of the target application. There is no direct formulation going from the model to the scheduling problem, so we cannot provide a general method to derive the scheduling problem. However, in Section 4, we illustrate the main techniques on basic problems, and we show how the application model impacts the scheduling problem. The scheduling problems become more or less complicated depending upon the application requirements. As usual in optimization theory, the most basic (and sometimes unrealistic) problems can usually be solved in polynomial time, whereas the most refined and accurate models usually lead to NP-hard problems. Although the complexity of some problems is still open, Section 4 concludes by highlighting the known frontier between polynomial and NP-complete problems.

Finally, in Section 5, we survey various techniques that can be used to solve the scheduling problem, that is, to find the best parallel execution of the application according to the performance criteria. We provide optimal algorithms to solve the simplest problem instances in polynomial time. For the most difficult instances, we present some general heuristic methods which aim at giving good approximate solutions.

2. MOTIVATING EXAMPLE

In this section, we focus on a simple pipelined application and emphasize the need for scheduling algorithms.

Consider an application composed of four tasks, whose dependencies form a linear chain: a dataset must first be processed by task t_1 before it can be processed by t_2 , then t_3 , and finally t_4 . The computation weights of tasks t_1 , t_2 , t_3 , and t_4 (or task weights) are set respectively to 5, 2, 3, and 20, as illustrated in Figure 1(a). If two consecutive tasks are executed on two distinct processors, then some time is required for communication in order to transfer the intermediate result. The communication weights are set respectively to 20, 15, and 1 for communications $t_1 \rightarrow t_2$, $t_2 \rightarrow t_3$, and $t_3 \rightarrow t_4$ (see Figure 1(a)). The communication weight along an edge corresponds to the size of the intermediate result that has to be sent from the processor in charge of executing the source task of the edge to the processor in charge of executing the sink task of the edge, whenever these two processors are different. Note that since all



Fig. 1. Motivating example.

input datasets have the same size, the intermediate results when processing different datasets also are assumed to have identical size, even though this assumption may not be true for some applications.

The target platform consists of three processors with various speeds and interconnection bandwidths, as illustrated in Figure 1(b). If task t_1 is scheduled to be executed on processor P_2 , a dataset is processed within $\frac{5}{1} = 5$ time units, while the execution on the faster processor P_1 requires only $\frac{5}{10} = 0.5$ time units (task weight divided by processor speed). Similarly, the communication of a data of weight *c* from processor P_1 to processor P_2 takes $\frac{c}{1}$ time units, while it is ten times faster to communicate from P_1 to P_3 .

ACM Computing Surveys, Vol. 45, No. 4, Article 50, Publication date: August 2013.

First examine the execution of the application when mapped sequentially on the fastest processor, P_3 (see Figure 1(c)). For such an execution, there is no communication. The communication weights and processors that are not used are shaded in grey in the figure. On the right, the processing of the first dataset (and the beginning of the second one) is illustrated. Note that because of the dependencies between tasks, this is actually the fastest way to process a single dataset. The latency is computed as $\mathcal{L} = \frac{5+2+3+20}{20} = 1.5$. A new dataset can be processed once the previous one is finished, hence the period $\mathcal{P} = \mathcal{L} = 1.5$.

Of course, this sequential execution does not exploit any parallelism. Since there are no independent tasks in this application, we cannot use task parallelism here. However, we now illustrate pipelined parallelism: different tasks are scheduled on distinct processors, and thus they can be executed simultaneously on different datasets. In the execution of Figure 1(d), all processors are used, and we greedily balance the computation requirement of tasks according to processor speeds. The performance of such a parallel execution turns out to be quite bad, because several large communications occur. The latency is now obtained by summing up all computation and communication times: $\mathcal{L} = \frac{5}{10} + 20 + 2 + 15 + \frac{3}{10} + \frac{1}{10} + \frac{20}{20} = 38.9$, as illustrated on the right of the figure for the first dataset. Moreover, the period is not better than the one obtained with the sequential execution presented previously, because communications become the bottleneck of the execution. Indeed, the transfer from t_1 to t_2 takes 20 time units, and therefore the period cannot be better than $20: \mathcal{P} \geq 20$. This example of execution illustrates that parallelism should be used with caution.

However, one can obtain a period better than that of the sequential execution as shown in Figure 1(e). In this case, we enforce some resource selection: the slowest processor P_2 is discarded (in grey) since it only slows down the whole execution. We process different datasets in parallel (see the execution on the right): within one unit of time, we can concurrently process one dataset by executing t_4 on P_3 , and another dataset by executing t_1 , t_2 , t_3 (sequentially) on P_1 . This partially sequential execution avoids all large communication weights (in grey). The communication time corresponds only to the communication between t_3 and t_4 , from P_1 to P_3 , and it takes a time $\frac{1}{10}$. We assume that communication and computation can overlap when processing distinct datasets, and therefore, once the first dataset has been processed (at time 1), P_1 can simultaneously communicate the data to P_3 and start computing the second dataset. Finally, the period is $\mathcal{P} = 1$. Note that this improved period is obtained at the price of a higher latency: the latency has increased from 1.5 in the fully sequential execution to $\mathcal{L} = 1 + \frac{1}{10} + 1 = 2.1$ here.

to $\mathcal{L} = 1 + \frac{1}{10} + 1 = 2.1$ here. This example illustrates the necessity of finding efficient trade-offs between antagonistic criteria.

3. MODELING TOOLS

This section gives general information on the various scheduling problems. It should help the reader understand the key properties of pipelined applications.

All applications of pipelined scheduling are characterized by properties from three components that we call the *workflow model*, the *system model*, and the *performance model*. These components correspond to "which kind of program we are scheduling", "which parallel machine will host the program", and "what are we trying to optimize". This three-component view is similar to the three-field notation used to define classical scheduling problems [Brucker 2007].

In the example of Section 2, the workflow model is an application with four tasks arranged as a linear chain, with computation and communication weights; the system model is a three-processor platform with speeds and bandwidths; and the performance



Fig. 2. The components of the workflow model.

model corresponds to the two optimization criteria, latency and period. We present in Sections 3.1, 3.2, and 3.3 the three models; then Section 3.4 classifies work in the taxonomy that has been detailed.

3.1. Workflow Model

The workflow model defines the program that is going to be executed; its components are presented in Figure 2.

As stated in the Introduction, programs are usually represented as Directed Acyclic Graphs (DAGs) in which nodes represent computation tasks and edges represent dependencies and/or communications between tasks. The shape of the graph is a parameter. Most program DAGs are not arbitrary but instead have some predefined form. For instance, it is common to find DAGs that are a single linear chain, as in the example of Section 2. Some other frequently encountered structures are fork graphs (for reduce operations), trees (in arithmetic expression evaluation; for instance in database [Hasan and Motwani 1994]), fork-join, and series-parallel graphs (commonly found when using nested parallelism [Blelloch et al. 1994]). The DAG is sometimes extended with two zero-weight nodes, a source node, which is made a predecessor of all entry nodes of the DAG, and a sink node, which is made a successor of all exit nodes of the DAG. This construction is purely technical and allows for faster computation of dependence paths in the graph.

The weight of the tasks are important because they represent computation requirements. For some applications, all the tasks have the same computation requirement (they are said to be unit tasks). The weight of communications is defined similarly; it usually corresponds to the size of the data to be communicated from one task to another when mapped on different processors. Note that a zero weight may be used to express a precedence between tasks, when the time to communicate can be ignored.

The tasks of the program may themselves contain parallelism. This adds a level of parallelism to the execution of the application, that is called data parallelism. Although the standard model only uses sequential tasks, some applications feature parallel tasks. Three models of parallel tasks are commonly used (this naming was proposed by Feitelson et al. [1997] and is now commonly used in job scheduling for production systems): a *rigid* task requires a given number of processors to execute; a *moldable* task can run on any number of processors, and its computation time is given by a speedup function (that can either be arbitrary or match a classical model such as the Amdahl's law [Amdahl 1967]); and a *malleable* task can change the number of processors it is executing on during its execution.

The task execution model indicates whether it is possible to execute concurrent replicas of a task at the same time or not. Replicating a task may not be possible due to an internal state of the task; the processing of the next dataset depends upon the result of the computation of the current one. Such tasks are said to be *monolithic*;



Fig. 3. The components of the system model.

otherwise they are *replicable*. When a task is replicated, it is common to impose some constraints on the allocation of the datasets to the replicas. For instance, the dealable stage rule [Cole 2004] forces datasets to be allocated in a round-robin fashion among the replicas. This constraint is enforced to avoid out-of-order completion and is quite useful when, say, a replicated task is followed by a monolithic one.

3.2. System Model

The system model describes the parallel machine used to run the program; its components are presented in Figure 3 and are now described in more detail.

First, processors may be identical (homogeneous), or instead they can have different processing capabilities (heterogeneous). There are two common models of heterogeneous processors. Either their processing capabilities are linked by a constant factor, that is, the processors have different speeds (known as the *related* model in scheduling theory and sometimes called heterogeneous uniform), or they are not speed related, which means that a processor may be fast on a task but slow on another one (known as the *unrelated* model in scheduling theory and sometimes called completely heterogeneous). Homogeneous and related processors are common in clusters. Unrelated processors arise when dealing with dedicated hardware or from preventing certain tasks to execute on some machines (to handle licensing issues or applications that do not fit in some machine memory). This decomposition in three models is classical in the scheduling literature [Brucker 2007].

The network defines how the processors are interconnected. The topology of the network describes the presence and capacity of the interconnection links. It is common to find *fully connected networks* in the literature, which can model buses as well as Internet connectivity. Arbitrary networks whose topologies are specified explicitly through an interconnection graph are also common. In between, some systems may exhibit structured networks such as chains, 2D meshes, 3D torus, etc. Regardless of the connectivity of the network, links may be of different types. They can be homogeneous—transport the information in the same way—or they can have different speeds. The most common heterogeneous link model is the *bandwidth* model, in which a link is characterized by its sole bandwidth. There exist other communication models such as the *delay* model [Rayward-Smith 1987], which assumes that all the communications are completely independent. Therefore, the delay model does not require communications to be scheduled on the network but only requires the processors to wait for a given amount of time when a communication is required. Frequently, the delay between two tasks scheduled on two different processors is computed based on the size of the message and the characteristics (latency and bandwidth) of the link between the processors. The LogP (Latency, overhead, gap, and Processor) model [Culler et al. 1993] is a realistic



Fig. 4. The components of the performance model.

communication model for fixed-size messages. It takes into account the transfer time on the network, the latency of the network, and the time required by a processor to prepare the communication. The LogGP model [Alexandrov et al. 1995] extends the LogP model by taking the size of the message into account using a linear model for the bandwidth. The latter two models are seldom used in pipelined scheduling.

Some assumptions must be made in order to define how communications take place. The *one-port* model [Bhat et al. 2003] forbids a processor to be involved in more than one communication at a time. This simple, but somewhat pessimistic, model is useful for representing single-threaded systems; it has been reported to accurately model certain MPI implementations that serialize communications when the messages are larger than a few megabytes [Saif and Parashar 2004]. The opposite model is the *multiport* model that allows a processor to be involved in an arbitrary number of communications simultaneously. This model is often considered unrealistic since some algorithms will use a large number of simultaneous communications, which induce large overheads in practice. An in-between model is the *k-port* model where the number of simultaneous communications must be bounded by a parameter of the problem [Hong and Prasanna 2003]. In any case, the model can also limit the total bandwidth that a node can use at a given time (that corresponds to the capacity of its network card).

Finally, some machines have hardware dedicated to communication or use multithreading to handle communication; thus they can compute while using the network. This leads to an *overlap* of communication and computation, as was assumed in the example of Section 2. However, some machines or software libraries are still monothreaded, and then such an overlapping is not possible.

3.3. Performance Model

The performance model describes the goal of the scheduler and tells from two valid schedules which one is better. Its components are presented in Figures 4.

The most common objective in pipelined scheduling is to maximize the throughput of the system, which is the number of datasets processed per time unit. In permanent applications such as interactive real-time systems, it indicates the load that the system can handle. Recall that this is equivalent to minimizing the period, which is the inverse of the throughput.

Another common objective is to minimize the latency of the application, which is basically defined as the time taken by a single dataset to be entirely processed. It measures the *response time* of the system to handle each dataset. The objective chosen to measure response time is most of the time the maximum latency, since the latency of different datasets may be different. Latency is mainly relevant in interactive systems. Note that latency minimization corresponds to *makespan* minimization in DAG scheduling, when there is a single dataset to process. Other objectives have also been studied. When the size of the computing system increases, hardware and software become more likely to be affected by malfunctions. There are many formulations of this problem (see Besseron et al. [2009] for details), but most of the time it reduces to optimizing the probability of correct execution of the application, which is called the *reliability* of the system [Girault et al. 2009]. Another objective function that is extensively studied is the *energy consumption*, which has recently become a critical problem, both for economic and environmental reasons [Mills 1999]. It is often assumed that the speed of processors can be dynamically adjusted [Jejurikar et al. 2004; Wang et al. 2010], and the slower a processor is, the less energy it consumes. Different models exist, but the main parameters are how the energy cost is computed from the speed (the energy cost is usually quadratic or cubic in the speed) and whether possible speeds are given by a continuous interval [Yao et al. 1995; Bansal et al. 2007] or by a discrete set of values [Okuma et al. 2001; Prathipati 2004].

The advent of more complex systems and modern user requirements increased interest in the optimization of several objectives at the same time. There are various ways to optimize multiple objectives [Dutot et al. 2009], but the most classical one is to optimize one of the objectives while ensuring a given threshold value on the other ones. Deciding which objectives are constrained and which one remains to optimize makes no theoretical difference [T'kindt and Billaut 2007]. However, there is often an objective that is a more natural candidate for optimization when designing heuristics.

3.4. Placing Related Work in the Taxonomy

The problem of scheduling pipelined linear chains, with both monolithic and replicable tasks, on homogeneous or heterogeneous platforms has extensively been addressed in the scheduling literature [Lee et al. 1998; Subhlok and Vondran 1996; Benoit and Robert 2008, 2010]. Lee et al. [1998] propose a three-step mapping methodology for maximizing the throughput of applications comprising a sequence of computation stages, each one consisting of a set of identical sequential tasks. Subhlok and Vondran [1996] propose a dynamic programming solution for optimizing latency under throughput constraints for applications composed of a linear chain of data-parallel tasks. Benoit and Robert [2008] address the problem of mapping pipelined linear chains on heterogeneous systems. Benoit and Robert [2010] explore the theoretical complexity of the bi-criteria optimization of latency and throughput for chains and fork graphs of replicable and data-parallel tasks under the assumptions of linear clustering and round-robin processing of input datasets.

Other works that address specific task graph topologies include Choudhary et al. [1994], which proposes a scheme for the optimal processor assignment for pipelined computations of monolithic parallel tasks with series-parallel dependencies, and focuses on minimizing latency under throughput constraints. Also, Hasan and Motwani [1994] (extended in Chekuri et al. [1995]) discuss throughput optimization for pipelined operator trees of query graphs that comprise sequential tasks.

Pipelined scheduling of arbitrary precedence task graphs of sequential monolithic tasks has been explored by a few researchers. In particular, Jonsson and Vasell [1996] and Hary and Özgüner [1999] discuss heuristics for maximizing the throughput of directed acyclic task graphs on multiprocessor systems using point-to-point networks. Yang et al. [2003] present an approach for resource optimization under throughput constraints. Suhendra et al. [2006] propose an integrated approach to optimize throughput for task scheduling and scratch-pad memory allocation based on integer linear programming for multiprocessor system-on-chip architectures. Guirado et al. [2005] propose a task mapping heuristic called EXPERT (EXploiting Pipeline Execution under Time constraints) that minimizes latency of streaming applications while

satisfying a given throughput constraint. EXPERT identifies maximal clusters of tasks that can form synchronous stages that meet the throughput constraint, and maps tasks in each cluster to the same processor so as to reduce communication overhead and minimize latency.

Pipelined scheduling algorithms for arbitrary DAGs that target heterogeneous systems include the work of Beynon [2001], which presents the Filter Copy Pipeline (FCP) scheduling algorithm for optimizing latency and throughput of arbitrary application DAGs on heterogeneous resources. FCP computes the number of copies of each task that is necessary to meet the aggregate production rate of its predecessors and maps these copies to processors that yield their least completion time. Later on, Spencer et al. [2002] proposed Balanced Filter Copies, which refines Filter Copy Pipeline. Banerjee et al. [1995] and Ranaweera and Agrawal [2001] address the problem of pipelined scheduling on heterogeneous systems. Ranaweera and Agrawal [2001] use clustering and task duplication to reduce the latency of the pipeline while ensuring a good throughput. However, these works target monolithic tasks, while Spencer et al. [2002] target replicable tasks. Finally, Vydyanathan et al. [2007] (extended in Vydyanathan et al. [2010]) address the latency optimization problem under throughput constraints for arbitrary precedence task graphs of replicable tasks on homogeneous platforms.

An extensive set of papers dealing with pipelined scheduling is summed up in Table I. Each paper is listed with its characteristics. Since there are too many characteristics to present, we focus on the main ones: structure of the precedence constraints, type of computation, replication, performance metric, and communication model. The table is sorted according to the characteristics so that searching for papers close to a given problem is made easier. Different papers with the same characteristics are merged into a single line.

The structure of the precedence constraints (the *Structure* column) can be a single chain (C), a structured graph such as a tree or series-parallel graph (S), or an arbitrary DAG (D). Processing units have computation capabilities (the *Comp.* column) that can be homogeneous (H), heterogeneous related (R), or heterogeneous unrelated (U). Replication of tasks (the *Rep.* column) can be authorized (Y) or not (N). The performance metric to compare the schedules (the *Metric* column) can be the throughput (T), the latency (L), the reliability (R), the energy consumption (E), or the number of processors used (N). The multi-objective problems are denoted with an & so that T&L denotes the bi-objective problem of optimizing both throughput and latency. Finally, the communication model (the *Comm*. column) can follow the model with only precedence constraints and zero communication weights (P), the one-port model (1), the multiport model (M), the k-port model (k), the delay model (D), or can be abstracted in the scheduling problem (abstr). When a paper deals with several scheduling models, the variations are denoted with a slash (/). For instance, paper Benoit et al. [2008] deals with scheduling a chain (C) on either homogeneous or heterogeneous related processors (H/R) without using replication (N) to optimize latency, reliability, or both of them (L/R/L&R) under the one-port model (1).

4. FORMULATING THE SCHEDULING PROBLEM

The goal of this section is to build a mathematical formulation of the scheduling problem from a given application. As explained shortly, it is a common practice to consider a more restrictive formulation than strictly necessary, in order to focus on more structured schedules that are likely to perform well.

We outline some principles in Section 4.1, and then we detail a few examples to illustrate the main techniques in Section 4.2. Finally we conclude in Section 4.3 by highlighting the known frontier between polynomial and NP-complete problems.

Reference	Structure	Comp.	Rep.	Metric	Comm.
[Bokhari 1988][Han et al. 1992] [Iqbal 1992][Manne and Olstad 1995]	С	н	Ν	Т	Р
[Nicol 1994][Pınar and Aykanat 2004] [Lee et al. 1998]	С	н	Ν	Т	Р
[Devi 2009]	С	н	Ν	T&L	Р
[Moreno et al. 2008]	С	Н	Y	Т	Р
[Benoit and Robert 2008]	С	H/R	Ν	Т	1
[Agrawal et al. 2008]	С	H/R	Ν	T/L	М
[Benoit et al. 2010]	С	H/R	Ν	T/L/E	1
[Benoit et al. 2008]	С	H/R	Ν	L/R/L&R	1
[Benoit and Robert 2009]	С	H/R	Y/N	T/L/T&L	1
[Benoit et al. 2009b]	С	R	Ν	T&L	1
[Benoit et al. 2009c]	С	R	Ν	\mathbf{L}	1
[Benoit et al. 2007]	С	R	Ν	T/L/T&L	1
[Agrawal et al. 2010]	С	R	Ν	T/L/T&L	1/M
[do Nascimento et al. 2005]	С	R	Y	T&N	Μ
[Benoit et al. 2009a]	С	R	Y	Т	1/M
[Kijsipongse and Ngamsuriyaroj 2010]	С	R	Y/N	Т	М
[Benoit and Robert 2010]	C/S	H/R	Y/N	T/L&T	Р
[Hasan and Motwani 1994] [Chekuri et al. 1995]	s	Н	Ν	Т	м
[Choudhary et al. 1994]	s	Н	Ν	T&L	Р
[Jonsson and Vasell 1996]	D	Н	Ν	Т	М
[Hary and Özgüner 1999]	D	Н	Ν	T&L	М
[Guirado et al. 2005]	D	Н	Ν	T&L	D
[Knobe et al. 1999]	D	Н	Ν	T&L	Р
[Vydyanathan et al. 2007]	D	Н	Y	T&L	М
[Vydyanathan et al. 2010]	D	Н	Y	T&L	k
[Subhlok and Vondran 1995]	D	Н	Y/N	Т	abstr
[Subhlok and Vondran 1996]	D	Н	Y/N	T&L	abstr
[Ranaweera and Agrawal 2001]	D	H/U	Ν	T&L	М
[Taura and Chien 1999]	D	R	Ν	Т	М
[Yang et al. 2003]	D	R	Ν	T&N	D
[Beynon 2001][Spencer et al. 2002]	D	R	Y	Т	М
[Banerjee et al. 1995]	D	U	Ν	Т	D
[Suhendra et al. 2006]	D	U	Ν	Т	М

Table I. Papers on pipelined scheduling, with characteristics of the scheduling problems

4.1. Compacting the Problem

One way to schedule a pipelined application is to explicitly schedule all the tasks of all the datasets, which amounts to completely unrolling the execution graph and assigning a start-up time and a processor to each task. In order to ensure that all dependencies and resource constraints are fulfilled, one must check that all predecessor relations are satisfied by the schedule, and that every processor does not execute more than one task at a given time. To do so, it may be necessary to associate a start-up time to each communication and a fraction of the bandwidth used (multiport model). However, the number of tasks to schedule could be extremely large, making this approach highly impractical.

To avoid this problem, a solution is to construct a more compact schedule which hopefully has some useful properties. The overall schedule should be easily deduced from the compact schedule in an incremental way. Checking whether the overall schedule is valid or not, and computing the performance index (e.g., throughput, latency) should be easy operations. To make an analogy with compilation, this amounts to transitioning from DAG scheduling to loop nest scheduling. In the latter framework, one considers a loop nest, that is, a collection of several nested loops that enclose a sequence of scalar statements. Each statement is executed many times, for each value of the surrounding loop counters. Compiler techniques such as Lamport hyperplane vectors, or space-time unimodular transformations [Wolfe 1989; Darte et al. 2000; Kennedy and Allen 2002] can efficiently expose the parallelism within the loop nest, by providing a linear or affine closed-form expression of scheduling dates for each statement instance within each loop iteration. On the contrary, a DAG schedule would completely unroll all loops and provide an exhaustive list of scheduling dates for each statement instance.

The most common types of schedules that can be compacted are *cyclic* schedules. If a schedule has a period \mathcal{P} , then all computations and communications are repeated every \mathcal{P} time units: two consecutive datasets are processed in exactly the same way, with a shift of \mathcal{P} time units. The cyclic schedule is constructed from the *elementary* schedule, which is the detailed schedule for one single dataset. If task t_i is executed on processor j at time s_i in the elementary schedule, then the execution of this task t_i for dataset x will be executed at time $s_i + (x - 1)\mathcal{P}$ on the same processor j in the cyclic schedule. The elementary schedule is a compact representation of the global cyclic schedule, while it is straightforward to derive the actual start-up time of each task instance, for each dataset, at runtime. The relation between cyclic and elementary schedule will be exemplified in Sections 4.2.1 and 4.2.2.

With cyclic schedules, one dataset starts its execution every \mathcal{P} time units. Thus, the system has a throughput $\mathcal{T} = 1/\mathcal{P}$. However, the latency \mathcal{L} of the application is harder to compute; in the general case, one must follow the entire processing of a given dataset (but all datasets have the same latency, which helps simplify the computation). The latency \mathcal{L} is the length of the elementary schedule.

Checking the validity of a cyclic schedule is easier than that of an arbitrary schedule. Intuitively, it is sufficient to check the datasets released in the last \mathcal{L} units of time, in order to make sure that a processor does not execute two tasks at the same time and that a communication link is not used twice. Technically, we can build an operation list [Agrawal et al. 2010] whose size is proportional to the original application precedence task graph, and does not depend upon the number of datasets that are processed.

A natural extension of cyclic schedules are *periodic* schedules, which repeat their operation every K datasets [Levner et al. 2010]. When K = 1, we retrieve cyclic schedules, but larger values of K are useful to gain performance, in particular through the use of replicated parallelism. We give an example in which the throughput increases when periodic schedules are allowed. Suppose that we want to execute a single task of weight 1, and that the platform consists of three different-speed processors P_1 , P_2 , and P_3 with speeds 1/3, 1/5, and 1/8, respectively. For a cyclic schedule, we need to specify on which processor the task is executed, and the optimal solution is to use the fastest processor, hence leading to a throughput $\mathcal{T} = 1/3$. However, with the use of replication, within 120 time units, P_1 can process 40 datasets, P_2 can process 24 datasets, and P_3 can process 15 datasets, resulting in a periodic schedule with K = 40 + 24 + 15 = 79 and a throughput $\mathcal{T} = 79/120$, about twice that of the cyclic schedule. Of course, it is easy to generalize the example to derive an arbitrarily bad throughput ratio between cyclic and periodic schedules. Note, however, that the gain in throughput comes with a price: because of the use of replication, it may become very difficult to compute the

ACM Computing Surveys, Vol. 45, No. 4, Article 50, Publication date: August 2013.



Fig. 5. An instance of the chain scheduling problem.

throughput. This is because the pace of operation for the entire system is no longer dictated by a single critical resource, but instead by a cycle of dependent operations that involves several computation units and communication links (refer to Benoit et al. [2009a] for details). Periodic schedules are represented in a compact way by a schedule that specifies the execution of K datasets similarly to the elementary schedule of a cyclic schedule.

Other common compact schedules consist in giving only the fraction of the time each processor spends executing each task [Beaumont et al. 2004; Vydyanathan et al. 2010]. Such representations are more convenient when using linear programming tools. However, reconstructing the actual schedule involves advanced concepts from graph theory and may be difficult to use in practice (although it can be done in polynomial time) [Beaumont et al. 2004].

4.2. Examples

The goal of this section is to provide examples to help the reader understand how to build a scheduling problem from the workflow model, system model and performance model. We also discuss how the problem varies when basic assumptions are modified.

4.2.1. Chain on Identical Processors with Interval Mapping. We consider the problem of scheduling a linear chain of n monolithic tasks onto m identical processors (with unit speed), linked by an infinitely fast network. For $1 \le i \le n$, task t_i has a weight p_i , and hence a processing time p_i on any processor. Figure 5 presents an instance of this scheduling problem with four tasks of weights $p_1 = 1$, $p_2 = 2$, $p_3 = 4$, and $p_4 = 3$. When scheduling chains of tasks, several mapping rules can be enforced.

The one to one manning who ensures that each task is monred to a different pro-

- —The *one-to-one mapping* rule ensures that each task is mapped to a different processor. This rule may be useful to deal with tasks having a high memory requirement, but all inter-task communications must then be paid.
- —Another classical rule is the *interval mapping* rule, which ensures that each processor executes a set of consecutive tasks. Formally, if a processor executes tasks $t_{i_{begin}}$ and $t_{i_{end}}$, then all tasks t_i , with $i_{begin} \leq i \leq i_{end}$, are executed on the same processor. This rule, which provides an extension of one-to-one mappings, is often used to reduce the communication overhead of the schedule.
- -Finally, the *general mapping* rule does not enforce any constraint, and thus any schedule is allowed. Note that for a homogeneous platform with communication costs, Agrawal et al. [2008] showed for the throughput objective that the optimal interval mapping is a 2-approximation of the optimal general mapping.

In this section, we consider interval mappings. Therefore, a solution to the scheduling problem is a partition of the task set $\{t_1, \ldots, t_n\}$ into m sets or intervals $\{I_1, \ldots, I_m\}$, where I_j $(1 \le j \le m)$ is a set of consecutive tasks. Note that one could want to have fewer intervals than processors, leaving some processor(s) completely idle, but here we assume that all the processors are used to make the notation simpler. The length of an interval is defined as the sum of the processing time of its tasks: $L_j = \sum_{i \in I_j} p_i$, for $1 \le j \le m$. Processors are identical (with unit speed), so that all mappings of intervals onto processors are identical too.



Fig. 6. The solution of optimal throughput to the instance of Figure 5 using an interval mapping on two processors.

In this case, the intervals form a compact representation of the schedule. The *elementary* schedule represents the execution of a single dataset: task t_i starts its execution at time $s_i = \sum_{i' < i} p_{i'}$ on the processor in charge of its interval. An overall schedule of period $\mathcal{P} = \max_{1 \le j \le m} L_j$ can now be constructed: task t_i is executed at time $s_i + (x - 1)\mathcal{P}$ on the *x*-th dataset. A solution of the instance of Figure 5 on two processors that use the intervals $\{t_1, t_2, t_3\}$ and $\{t_4\}$ is depicted in Figure 6, where the boxes represent tasks and datasets are identified by colors. The schedule is focused on the cyan dataset (the labeled tasks), which follows the green one (partially depicted) and precedes the red one (partially depicted). Each task is periodically scheduled every 7 time units (a period is depicted with dotted lines). Processor 2 is idle during 4 time units within each period.

One can check that such a schedule is valid: the precedence constraints are respected, two tasks are never scheduled on the same processor at the same time (the processor in charge of interval I_j executes tasks for one single dataset during L_j time units, and the next dataset arrives after $\max_{j'} L_{j'}$ time units), and the monolithic constraint is also fulfilled, since all the instances of a task are scheduled on a unique processor.

To conclude, the throughput of the schedule is $\mathcal{T} = \frac{1}{\mathcal{P}} = \frac{1}{\max_{1 \le j \le m} L_j}$, and its latency is $\mathcal{L} = \sum_{1 \le i \le n} p_i$. Note that, given an interval mapping, \mathcal{T} is the optimal throughput since the processor for which $L_j = \max_{j'} L_{j'}$ will never be idle, and it is the one that defines the period. Note also that the latency is optimal over all schedules, since $\sum_{1 \le i \le n} p_i$ is a lower bound on the latency.

For such a problem (no communication, identical processors, linear dependency graph, no replication, interval mapping), the problem of optimizing the throughput is reduced to the classical chains-on-chains partitioning problem [Pinar and Aykanat 2004], and it can be solved in polynomial time using, for instance, a dynamic programming algorithm.

4.2.2. Chain on Identical Processors with General Mapping. This problem is a slight variation of the previous one: solutions are no longer restricted to interval mapping schedules, but any mapping may be used. By suppressing the interval mapping constraint, we can usually obtain a better throughput but the scheduling problem and schedule reconstruction become harder, as we illustrate in the following example.

The solution of a general mapping can be expressed as a partition of the task set $\{t_1, \ldots, t_n\}$ into *m* sets $\{A_1, \ldots, A_m\}$, but these sets are not enforced to be intervals anymore. The optimal period is then $\mathcal{P} = \max_{1 \le j \le m} \sum_{i \in A_j} p_i$.

We present a generic way to reconstruct from the mapping a cyclic schedule that preserves the throughput. A core schedule is constructed by scheduling all the tasks according to the allocation without leaving any idle time and, therefore, reaching the optimal period. Task t_i in set A_j is scheduled in the core schedule at time $s_i = \sum_{i' < i, i' \in A_j} p_{i'}$. A solution of the instance presented in Figure 5 is depicted in Figure 7 between the



Fig. 7. The solution of optimal throughput to the instance of Figure 5 using a general mapping on two processors.



Fig. 8. A solution of the same throughput with Figure 7, but with a better latency.

dotted lines (time units 0 to 5); it schedules tasks t_1 and t_3 on processor 1, and tasks t_2 and t_4 on processor 2.

The notion of core schedule is different than the notion of elementary schedule. Informally, the elementary schedule describes the execution of a single dataset while the tasks in the core schedule may process different datasets.

The cyclic schedule is built so that each task takes its predecessor from the previous period: inside a period, each task is processing a different dataset. We can now follow the execution of the *x*-th dataset: it starts being executed for task t_i at time $s_i + (i+x-1)\mathcal{P}$, as illustrated for the white dataset (x = 0) in Figure 7. This technique produces schedules with a large latency, between $(n-1)\mathcal{P}$ and $n\mathcal{P}$. In the example, the latency is 20, exactly 4 times the period. In Figure 7, the core schedule is given between the dotted lines (from time step 0 to 5). The elementary schedule is given by restricting the figure to the white dataset (i.e., removing all other datasets).

The strict rule of splitting the execution in n periods ensures that no precedence constraint is violated. However, if the precedence constraint between task t_i and task t_{i+1} is respected in the core schedule, then it is possible to schedule both of them in a single time period. Consider the schedule depicted in Figure 8. It uses the same allocation as the one in Figure 7, but tasks t_2 and t_4 have been swapped in the core schedule. Thus, tasks t_1 and t_2 can be scheduled in the same period, leading to a latency of 13 instead of 20.

Note that the problem of finding the best general mapping for the throughput maximization problem is NP-complete: it is equivalent to the 2-PARTITION problem [Garey and Johnson 1979] (consider an instance with two processors).

4.2.3. Chain with a Fixed Processor Allocation. In the previous examples, we have given hints of techniques to build the best core schedule, given a mapping and a processor allocation, in simple cases with no communication costs. In those examples, we were able to schedule tasks in order to reach the optimal throughput and/or latency.

Given a mapping and a processor allocation, obtaining a schedule that reaches the optimal latency can be done by greedily scheduling the tasks in the order of the chain. However, this may come at the price of a degradation of the throughput, since idle times may appear in the schedule. We can ensure that there will be no conflicts if the period equals the latency (only one dataset in the pipeline at any time step).



Fig. 9. A series-parallel graph and its binary decomposition tree.

If we are interested in minimizing the period, the presence of communications makes the problem much more difficult. In the model without computation and communication overlap, it is actually NP-hard to decide the order of communications (i.e., deciding the start time of each communication in the core schedule) in order to obtain the minimum period (see Agrawal et al. [2010] for details). If computation and communication can be overlapped, the processor works simultaneously on various datasets, and we are able to build a conflict-free schedule. When a bi-criteria objective function is considered, more difficulties arise, as the ordering of communications also becomes vital to obtain a good trade-off between latency and period minimization.

4.2.4. Scheduling Moldable Tasks with Series-Parallel Precedence. Chains are not the only kind of precedence constraints that are structured enough to help derive interesting results. For instance, series-parallel graphs [Valdes et al. 1982] are defined by composition. Given two series-parallel graphs, a series composition merges the sink of one graph with the root (or source) of the other one; a parallel composition merges the sinks of both graphs and the roots of both graphs. No other edges are added or removed. The basic series-parallel graph is composed of two vertices and one edge. Figure 9 gives an example of a series-parallel graph. The chain of length three (given as t_2 , t_5 and t_7 in Figure 9) is obtained by composing in series the chain of length two with itself. The diamond graph (given as t_1 , t_3 , t_4 , and t_6 in Figure 9) is obtained by composing in parallel the chain of length three with itself. In parallel computing, series-parallel workflow graphs appear when using nested parallelism [Blelloch et al. 1994; Blikberg and Sørevik 2005].

Choudhary et al. [1994] consider the scheduling of series-parallel pipelined precedence task graphs, composed of moldable tasks. A given processor executes a single task and communications inside the moldable task are assumed included in the parallel processing times. There is no communication required between the tasks, just a precedence constraint. Provided a processor allocation, one can build an elementary schedule by scheduling the tasks as soon as possible. Since a processor is only involved in the computation of a single task, this elementary schedule reaches the optimal latency (for the processor allocation). Moreover, the elementary schedule can be executed with a period equal to the length of the longest task, leading to a cyclic schedule of optimal throughput (for the processor allocation).

Since the application task graph is a series-parallel graph, the latency and throughput of a solution can be expressed according to its Binary Decomposition Tree (BDT) [Valdes et al. 1982]. Each leaf of the BDT is a vertex of the graph and each



Fig. 10. An arbitrary DAG (the task weight is the label next to the task).

internal node is either a series node S(l, r) or a parallel node P(l, r). A series node S(l, r) indicates that the subtree l is a predecessor of the subtree r. A parallel node P(l, r) indicates that both subtrees l and r are independent. A Binary Decomposition Tree is depicted in Figure 9.

In BDT form, the throughput of a node is the minimum of the throughputs of the children of the node: $\mathcal{T}(S(l,r)) = \mathcal{T}(P(l,r)) = \min(\mathcal{T}(l),\mathcal{T}(r))$. The expression of the latency depends on the type of considered node. If the node is a parallel node, then the latency is the maximum of the latencies of its children: $\mathcal{L}(P(l,r)) = \max(\mathcal{L}(l),\mathcal{L}(r))$. If it is a series node, the latency is the sum of the latencies of its children: $\mathcal{L}(S(l,r)) = \mathcal{L}(l) + \mathcal{L}(r)$.

4.2.5. Arbitrary DAGs on Homogeneous Processors. Many applications cannot be represented by a structured graph such as a chain or a series-parallel graph. Arbitrary DAGs are more general but at the same time they are more difficult to schedule efficiently. Figure 10 presents a sample arbitrary DAG.

Scheduling arbitrary DAGs poses problems that are similar to those encountered when scheduling chains. Consider first the case of one-to-one mappings, in which each task is allocated to a different processor. A cyclic schedule is easily built by scheduling all tasks as soon as possible. Task *i* is scheduled in the cyclic schedule on processor *i* at time $s_i = \max_{i' \in pred(i)} s_{i'} + p_{i'}$. This schedule can be executed periodically every $\mathcal{P} = \max_i p_i$ with throughput $\mathcal{T} = \frac{1}{\max_i p_i}$. The latency is the longest path in the graph, that is, $\mathcal{L} = \max s_i + p_i$. A schedule built in such a way does not schedule two tasks on the same processor at the same time; indeed, each task is executed during each period on its own processor, and its processing time is smaller or equal to the period. Under the one-to-one mapping constraint, this schedule is optimal for both objective functions. The solution for the graph of Figure 10 is presented in Figure 11, with a latency $\mathcal{L} = 11$ and a throughput $\mathcal{T} = \frac{1}{8}$.

When there is no constraint enforced on the mapping rule, problems similar to those of general mappings for linear chains appear (see Section 4.2.2): we cannot easily derive an efficient cyclic schedule from the processor allocation. Establishing a cyclic schedule that reaches the optimal throughput given a processor allocation is easy without communication cost, but it can lead to a large latency. Similarly to the case of chains, a core schedule is obtained by scheduling all the tasks consecutively without taking care of the dependencies. This way, we obtain the optimal period (for this allocation) equal to the load of the most loaded processor. The cyclic schedule is built so that each task takes its data from the execution of its predecessors in the



Fig. 11. One-to-one mapping of the instance of Figure 10 with $\mathcal{L} = 11$ and $\mathcal{T} = \frac{1}{8}$. Tasks t_1 and t_6 have computation time 0, therefore they are omitted.



Fig. 12. A general mapping solution of the instance of Figure 10 with $\mathcal{L} = 15$ and $\mathcal{T} = \frac{1}{8}$. Tasks t_1 and t_6 have computation time 0, therefore they are omitted.



Fig. 13. A general mapping solution of the instance of Figure 10 with $\mathcal{L} = 11$ and $\mathcal{T} = \frac{1}{8}$. Tasks t_1 and t_6 have computation time 0, therefore they are omitted.

last period. Therefore, executing a dataset takes as many periods as the depth of the precedence task graph. On the instance of Figure 10, the optimal throughput on two processors is obtained by scheduling t_2 alone on a processor. Figure 12 presents a cyclic schedule for this processor allocation according to this generic technique, leading to a latency $\mathcal{L} = 15$. Note that t_5 could be scheduled in the same period as t_3 and in general this optimization can be done by a greedy algorithm. However, it does not guarantee to obtain the schedule with the optimal latency, which is presented in Figure 13 and has a latency $\mathcal{L} = 11$. Indeed, contrarily to linear pipelines, given a processor allocation, obtaining the cyclic schedule that minimizes the latency is NP-hard [Rayward-Smith et al. 1995].

The notion of interval mapping cannot be directly applied on a complex DAG. However, we believe that the interest of interval mapping schedules of chains can be transposed to *convex clustered schedules* on DAGs. In a convex clustered schedule, if two tasks are executed on one processor, then all the tasks on all the paths between these two tasks are scheduled on the same processor [Lepere and Trystram 2002]. Convex clustered schedules are also called *processor-ordered schedules*, because the graph of the inter-processor communications induced by such schedules is acyclic [Guinand et al. 2004]. The execution of the tasks of a given processor can be serialized and executed without any idle time (provided their execution starts after all the data have arrived). This leads to a reconstruction technique similar to the one applied on chains of tasks



Fig. 14. A general mapping solution of the instance of Figure 10 with $\mathcal{L} = 14$ and $\mathcal{T} = \frac{1}{7}$ when task t_2 is replicable. Tasks t_1 and t_6 have computation time 0, therefore they are omitted.

following the interval mapping rule. Two processors that are independent in the interprocessor communication graph can execute their tasks on a given dataset during the same period in any order, without violation of the precedence constraints. Such a construction leads to the optimal throughput for a given convex clustered mapping of tasks to processors, and to a latency $\mathcal{L} \leq x\mathcal{P} \leq m\mathcal{P}$, where *x* is the length of the longest chain in the graph of communications between processors.

Algorithms to generate convex clustered schedules based on recursive decomposition have been proposed for classical DAG scheduling problems [Pecero-Sanchez and Trystram 2005]. In pipelined scheduling, heuristic algorithms based on stages often generate convex clustered schedules such as Banerjee et al. [1995] and Guirado et al. [2005]. However, the theoretical properties of such schedules have never been studied for pipelined workflows.

4.2.6. Scheduling Arbitrary DAGs on Homogeneous Processors with Replication. A task is replicable if it does not contain an internal state. It means that the same task can be executed at the same time on different datasets. Replication allows one to increase the throughput of the application. (We point out that this is different from duplication, which consists in executing the same task on the same dataset on multiple different processors. Redundantly executing some operations aims at either reducing communication bottlenecks or increasing reliability.) On the instance presented in Figure 10, only two processors can be useful: the dependencies prevent any three tasks from being executed simultaneously, so a third processor would improve neither the throughput nor the latency for monolithic tasks. However, if task t_2 is replicable, the third processor could be used to replicate the computation of this task, therefore leading to the schedule depicted in Figure 14.

Replicating t_2 leads to a periodic schedule that executes two datasets every 14 time units (K = 2). Its throughput is therefore $\mathcal{T} = \frac{2}{14} = \frac{1}{7}$, which is better than without replication. The latency is the maximum time a dataset spends in the system. Without replication, all the datasets spend the same time in the system. With replication, this statement no longer holds. In the example, the cyan dataset spends 11 time units in the system whereas the green one spends 14 time units. The latency of the schedule is therefore $\mathcal{L} = 14$. If task t_4 was replicable as well, two copies could be executed in parallel, improving the throughput to $\mathcal{T} = \frac{2}{11}$ and the latency to $\mathcal{L} = 11$. A fourth processor could be used to pipeline the execution of t_4 and reach a period of $\mathcal{P} = 8$ and, hence, a throughput of $\mathcal{T} = \frac{1}{4}$.

A schedule with replication is no longer cyclic but instead is periodic, with the definitions of Section 4.1. Such a schedule can be seen as a pipelined execution of an unrolled version of the graph. The overall schedule should be specified by giving a periodic schedule of length ℓ (the time between the start of the first task of the first dataset of the period and the completion of the last task of the last dataset of the period), detailing how to execute K consecutive datasets and providing its period \mathcal{P} . Verifying that the schedule is valid is done in the same way as for classical elementary schedules: one needs to expand all the periods that have a task running during the schedule, that is to say the tasks that start during the elementary schedule and within the ℓ time units before. Such a schedule has a throughput of $\mathcal{T} = \frac{K}{\mathcal{P}}$, and its latency should be computed as the maximum latency of the datasets in the elementary schedule.

Note that if all tasks are replicable, the whole task graph can be replicated on all the *m* processors. Each processor executes sequentially exactly one copy of the application. This leads to a schedule of latency and period $\mathcal{P} = \mathcal{L} = \sum_{i} p_i$, and a throughput of $\mathcal{T} = \frac{m}{\sum_{i} p_i}$.

When using replication, it is possible that dataset *i* is processed before its predecessor i-1. This behavior mainly appears when processors are heterogeneous. The semantic of the application might not allow datasets to be processed in such an *out-of-order* fashion. For instance, if a task is responsible for compressing its input, providing the datasets out-of-order will change the output of the program. One can then either impose a delay or use some other constraints, such as for instance the dealable stage constraint [Cole 2004].

4.2.7. Model Variations. In most cases, heterogeneity does not drastically change the scheduling model. However, the compact schedule description must then contain the processor allocation, that is, it must specify which task is executed on which processor. Otherwise the formulations stay similar.

A technique to reduce latency is to consider *duplication* [Ahmad and Kwok 1998; Vydyanathan et al. 2010]. Duplicating a task consists in executing the same task more than once on different processors for every dataset. Each task receives its data from one of the duplicates of each of its predecessors. Hence, this allows more flexibility for dealing with data dependency. The idea is to reduce the communication overheads at the expense of increasing the computation load. Another goal is to increase the reliability of the system: whenever the execution of one duplicate would fail, that of another duplicate could still be successful. The major difference of duplication as compared to replication is the following: with duplication, a single dataset is executed in each period, whereas with replication, several datasets can be executed in each period.

Communication models affect the schedule formulation. The easiest communication model is the one-port model where a machine communicates with a single other machine at a time. Therefore, in the schedule, each machine is represented by two processors, one for the computations and one for the communications. A valid schedule needs to "execute" a communication task at the same time on the communication processor of both machines involved in the data transfer. A common variation on the one-port model is to forbid communication and computation overlap. This model is used in Hary and Özgüner [1999]. In this case, there is no need for a communication processor; the communication tasks have to be scheduled on the computation processor [Benoit et al. 2007].

To deal with more than one communication at a time, a realistic model would be to split the bandwidth equally among the communications. However, such models are more complicated to analyze and are therefore not used in practice. Two ways of overcoming the problem exist. The first one is to consider the k-port model where each machine has a bandwidth B divided equally into k channels. The scheduling problem amounts to using k communication processors per machine. This model has been used in Vydyanathan et al. [2010]. When only the throughput matters (and not the latency), it is enough to ensure that no network link is overloaded. One can reconstruct a periodic schedule explicitly by using the model detailed previously, considering each network link as a processor. This approach has been used in Taura and Chien [1999].

4.3. Complexity

The goal of this section is to provide reference pointers for the complexity of the pipelined scheduling problem. Lots of works are dedicated to highlighting the frontier between polynomial and NP-hard optimization problems in pipelined scheduling.

The complexity of classical scheduling problems has been studied in Brucker [2007]. One of the main contributions was to determine some constraint changes that always make the problem harder. Some similar results are valid on pipelined scheduling. For instance, heterogeneous versions of problems are always harder than their homogeneous counterpart, since homogeneous cases can be easily represented as heterogeneous problem instances but not vice versa. A problem with an arbitrary task graph or architecture graph is always harder than the structured counterpart, and in general considering a superset of graphs makes problems harder. Also, removing communications makes the problem easier.

As seen in the previous examples, throughput optimization is always NP-hard for general mappings, but polynomial instances can be found for interval mappings. The communication model plays a key role in complexity. The optimization of latency is usually equivalent to the optimization of the makespan in classical DAG scheduling [Kwok and Ahmad 1999b].

The complexity of multi-objective problem instances relates to three different types of questions. First, the decision problems of multi-objective problems are directly related to those for mono-objective problems. A threshold value is given for all the objectives and the problem is to decide whether a solution exist that matches all these thresholds. Multi-objective decision problems are obviously harder than their mono-objective counterpart. Second, the counting problem consists in computing how many Pareto-optimal solutions a multi-objective problem has; a Pareto-optimal solution is such that no other solution is strictly better than it. Finally, the enumeration problem consists in enumerating all the Pareto-optimal solutions of an instance. The enumeration problem is obviously harder than the decision problem and the counting problem, since it is possible to count the number of solutions with an enumeration algorithm and to decide whether given thresholds are feasible. A complete discussion of these problems can be found in T'kindt and Billaut [2007].

The complexity class of enumeration problems expresses the complexity of the problem as a function of both the size of the input of the problem and the number of Pareto-optimal solutions leading to classes EP (Enumeration Polynomial) and ENP (for Enumeration Nondeterministic Polynomial) [T'kindt et al. 2007]. Therefore, the decision version of a multi-objective problem might be NP-complete, but since it has an exponential number of Pareto optimal solution, its enumeration version is in EP (the problem 1 || $\sum C_i^A$; $\sum C_i^B$ of Agnetis et al. [2004] is one of the many examples that exhibit this property). Therefore, the approaches based on exhaustive enumeration can take a long time. However, Papadimitriou and Yannakakis [2000] show that most multi-objective problems admit an approximate set of Pareto-optimal solutions whose cardinality is polynomial in the size of the instance (but it is exponential in the number of objectives and in the quality of the approximation). It was also shown in Papadimitriou and Yannakakis [2000] that an approximation algorithm for the decision problem can be used to derive an approximation of the Pareto set in polynomial time. These results motivate the investigation of algorithms that enforce thresholds on some objectives and optimize the other ones.

Mapping rule	Platform type				
	Fully Hom.	Comm. Hom.	Hetero.		
one-to-one	polynomial	polynomial, NP-hard (rep.)	NP-hard		
interval	polynomial	NP-hard	NP-hard		
general	NP-hard, polynomial (rep.)	NP-hard			

Table II. Summary of Complexity Results for Period Minimization of a Linear Task Graph

Table III. Summary of Complexity Results for Latency Minimization of a Linear Task Graph

Mapping rule	Platform type				
	Fully Hom.	Comm. Hom.	Hetero.		
one-to-one	polynomial [Benoit and Robert 2009]		NP-hard [Benoit et al. 2008]		
interval	polynomial [Benoit and Robert 2009]		NP-hard [Benoit et al. 2009c]		
general	polynomial [Benoit et al. 2008]				

The complexity of linear graph problems has been widely studied since it roots the general DAG case and most of the structured graph ones [Benoit and Robert 2008, 2009, 2010; Benoit et al. 2007, 2008, 2009c; Agrawal et al. 2008, 2010]. The large number of variants for these scheduling problems makes complexity results very difficult to apprehend. An exhaustive list of such results can be found in Benoit [2009]. We provide in Tables II and III a summary of complexity results for period and latency optimization problems, which hold for all communication models. *Fully Hom.* platforms refer to homogeneous computations and communications. *Comm. Hom.* platforms add one level of heterogeneity (heterogeneous-related processors). Finally, *Hetero.* platforms are fully heterogeneous (heterogeneous-related processors and heterogeneous communication links).

For the period minimization problem, the reader can refer to Benoit and Robert [2008] for the variant with no replication, and to Benoit and Robert [2010] otherwise (results denoted with *(rep.)*). For the latency minimization problem, we report here results with no data parallelism; otherwise the problem becomes NP-hard as soon as processors have different speeds (related model), with no communication costs [Benoit and Robert 2010].

5. SOLVING PROBLEMS

The goal of this section is to give methods to solve the pipelined scheduling problem using exact algorithms or heuristic techniques.

5.1. Scheduling a Chain of Tasks with Interval Mappings

The first problem that we consider has been presented in Section 4.2.1. It consists in scheduling a chain of n tasks onto m identical (homogeneous) processors, without communication, and enforcing the interval mapping constraint. Section 4.2.1 states that the latency of such schedules is constant, however, the throughput can be optimized by minimizing the length of the longest interval.

The optimization of the throughput problem is the same combinatorial problem as the well-known chains-on-chains partitioning problem, which has been solved by a polynomial algorithm in Bokhari [1988], and then refined to reach lower complexity in Iqbal [1992], Nicol [1994], Manne and Olstad [1995]. For very large problems, some heuristics have also been designed to reduce the scheduling times even further (see Pinar and Aykanat [2004] for a survey). The first algorithm was based on a shortest path algorithm in an assignment graph. The approach that follows has a lower complexity and is easier to understand.

50:24

The core of the technique is the Probe function that takes as a parameter the processing time of the tasks and the length of the longest interval \mathcal{P} . It constructs intervals $I = \{I_1, \ldots, I_m\}$ such that $\max_{1 \le j \le m} L_j \le \mathcal{P}$, or shows that no such intervals exist (remember that $L_j = \sum_{i \in I_j} p_i$, where p_i is the processing time of task t_i). Probe recursively allocates the first x tasks of the chain to the first processor so that $\sum_{i \leq x} p_i \leq \mathcal{P}$ and $\sum_{i \leq x+1} p_i > \mathcal{P}$ until no task remains, and then returns the schedule. If the number of intervals is less than the number of processors, this function builds a schedule having no interval of length exceeding \mathcal{P} . Otherwise, no schedule of maximal interval length less than \mathcal{P} exists with *m* processors. It can be easily shown that the schedules constructed are dominant, that is, if a schedule exists, then there is one respecting this construction. The last problem is to choose the optimal value for the threshold \mathcal{P} . The optimal value is obtained by using a binary search on the possible values of \mathcal{P} , which are tested using the Probe function. This construction is polynomial but has a quite high complexity. It is possible to reduce the complexity of the Probe function using prefix sum arrays and binary search so that fewer values of \mathcal{P} can be tested by analyzing the processing time values. In the general case, the lowest complexity is reached by using Nicol's algorithm [Nicol 1994] with the algorithm for the Probe function described in Han et al. [1992], leading to a total complexity of $O(n + m^2 \log(n) \log(n/m))$ (see Pinar and Aykanat [2004] for details).

The same idea can be used to deal with different problems, for instance, with nonoverlapping communications following the one-port model [Iqbal 1992]. Since there is no overlapping, the communication time is included in the computation of the throughput of both processors involved in the communication. Then, setting the boundary of the interval on a communication larger than the length of the previous task plus its in-bound communication is never efficient: the two tasks can be merged without losing optimality. (The same argument applies towards the next task.) Detecting these tasks and merging them can be performed in linear time. The resulting chain is said to be monotonic. Then, an algorithm similar to the partitioning of a chain in intervals without communication can be applied, leading to the optimal solution [Iqbal 1992].

The same algorithm can also be used to optimally solve the case with related processor heterogeneity (processor speeds differ) if the order in which a dataset goes through the processors is known. This is the case on dedicated hardware where the processor network forces the order of execution between the processors. However, if this order is not known, the problem is NP-complete in the strong sense [Benoit et al. 2007], even without taking communication costs into account. There are too many permutations to try, but the Probe algorithm sets a solid ground to build heuristics upon.

Benoit and Robert [2008] propose three heuristics to build interval mappings for optimizing the throughput on heterogeneous processors. The first one, called SPL, starts by assigning all the tasks to the fastest processor and then greedily splits the largest interval by unloading work to the fastest available processor. The splitting point is chosen so as to minimize the period of the new solution. The two other heuristics BSL and BSC use a binary search on the period of the solution. This period is used as a goal in the greedy allocation of tasks to processors. BSL allocates the beginning of the chain of tasks to the processor that will execute the most computations while respecting the threshold. On the other hand, BSC chooses the allocation that is the closest to the period. Note that Pinar et al. [2008] survey chains-on-chains partitioning problems with heterogeneous resources, and propose several heuristics that can be transposed to throughput optimization without communication costs.

Kijsipongse and Ngamsuriyaroj [2010] propose a heuristic algorithm called ThroughputPipeline to schedule a chain using interval mappings on Grid computing systems (related processors, bounded multiport, communication and computation overlapping, no replication), considering routing through intermediate nodes. ThroughputPipeline is based on Dijkstra's shortest path algorithm. Simulation shows that its performance is within 30% of the general mapping optimal solution (obtained by an Integer Linear Programming formulation).

Solving the problem of optimizing both the latency and the throughput of a linear pipeline application has been considered in Subhlok and Vondran [1996] and Benoit et al. [2007, 2009b]. Bi-objective optimization problems are usually solved by providing a set of efficient solutions. This set of solutions is generated by using an algorithm that targets values of one objective while optimizing the other one. The solution space is explored by executing this algorithm with different values of the threshold, hence covering the efficient part of the solution space.

Subhlok and Vondran [1995] address the problem of scheduling a chain of tasks on homogeneous processors to optimize the throughput of the application without computation and communication overlap. It covers a large range of problems since it addresses moldable tasks with dedicated communication functions and replicable tasks. The network is supposed to be homogeneous but the details of the communication model are abstracted by explicitly giving the communication time in the instance of the problem. The technique used is based on dynamic programming and leads to an optimal polynomial algorithm. This result can be extended by adding a latency dimension in the dynamic program to allow the optimization of the latency under a throughput constraint [Subhlok and Vondran 1996].

Benoit et al. [2007] propose heuristics that optimize the throughput and latency when link bandwidths are identical but processors have different speeds (one-port communications without overlap). Six heuristics are presented, enforcing a constraint on either the throughput or the latency. All six heuristics are similar to SPL in that they start by allocating all the tasks to the fastest processor and split the interval iteratively. The differences are that each interval may be split in two to use the fastest available processor, or split in three to use the fastest two processors available. The other differences are about the solution chosen; it could be the one that maximizes one objective or a ratio of improvement. Benoit et al. [2009b] propose an integer linear programming formulation to solve the problem optimally (and with heterogeneous bandwidths). The solving procedure takes a long time even on a simple instance of 7 tasks and 10 processors (a few hours on a modern computer) but allows assessment of the absolute performance of the previously proposed heuristics.

5.2. Scheduling a Chain of Tasks with General Mappings

Using general mappings instead of restricting to interval mappings leads to better throughput. Without replication and without communication costs, the optimization of the throughput on homogeneous processors is NP-complete by reduction to 3-PARTITION. In fact, the mathematical problem is to partition *n* integers p_1, \ldots, p_n into *m* sets A_1, \ldots, A_m so that the length of the largest set $\max_{1 \le j \le m} \sum_{i \in A_j} p_i$ is minimized. This formulation corresponds exactly to the problem of scheduling independent tasks on identical processors to minimize the makespan, that has been studied for a long time Graham [1966, 1969] and considered theoretically solved since Hochbaum and Shmoys [1987].

On homogeneous processors, the classical List Scheduling algorithm schedules tasks greedily on the least loaded processor, and it is a 2-approximation [Graham 1966], that is, the value of the obtained solution is at most twice the optimal value. Sorting tasks by nonincreasing processing times leads to the Largest Processing Time (LPT) algorithm, which is known to be a 4/3-approximation algorithm [Graham 1969]. An approximation

algorithm with arbitrary precision that is polynomial in the size of the problem but exponential in the inverse of the precision (also known as PTAS) based on binary search and dynamic programming has been proposed in Hochbaum and Shmoys [1987]. The question of whether an algorithm with arbitrary precision that is polynomial in the size of the problem and in the inverse of the precision (also known as FPTAS) arises. However, the problem is NP-complete in the strong sense and FPTAS do not exist for this problem unless P = NP (see Hochbaum [1997] for details on complexity classes based on approximation properties).

With heterogeneous processors, there is still a link with the classical makespan optimization problem. If processors are heterogeneous related (computing at different speeds), the throughput optimization problem is the same as scheduling independent tasks on heterogeneous-related processors. This problem admits a 2-approximation algorithm similar to LPT [Gonzalez et al. 1977]. Hochbaum and Shmoys [1988] provide an elaborate approximation scheme with very high runtime complexity as well as a simple 3/2-approximation algorithm. If processors are unrelated (i.e., their speeds depend on the task they are handling), the throughput optimization problem is the same as scheduling independent tasks on unrelated processors to minimize the makespan. It can be shown that there exists no approximation algorithm with a ratio better than 3/2 [Lenstra et al. 1990]. Moreover, a 2-approximation algorithm based on binary search and linear programming has been proposed in Lenstra et al. [1990] and recently made simpler and faster by Gairing et al. [2005].

The results on classical scheduling problems remain valid even if the graph is not linear as long as the performance index is the throughput and there are no communications. However, they still provide an interesting baseline to assess the impact of communications.

Kijsipongse and Ngamsuriyaroj [2010] consider the problem of scheduling a chain on a Grid computer with routing to optimize the throughput. Processors are heterogeneous (related) and communications follow the bounded multiport model with overlapping. It first considers the case with replication (called multipath in this terminology). An optimal general mapping is constructed in polynomial time by LPSAG, a flow-based linear programming formulation (somehow similar to LPsched [do Nascimento et al. 2005]). Finding the optimal general mapping without replication (single path) is investigated and shown to be NP-complete. LPSAG is extended into an Integer Linear Program that finds the optimal general mapping without replication (but possibly in exponential time). A polynomial heuristic based on Dijkstra's shortest path algorithm that only constructs interval mappings is proposed. Experiments show that the heuristic is within 30% of the Integer Linear Programming solution.

Benoit et al. [2009c] provide a polynomial algorithm to optimize the latency of a pipelined chain on heterogeneous (related) networks of processors under the one-port model. The algorithm is based on a dynamic programming formulation.

5.3. Structured Application Graphs

Choudhary et al. [1994] tackle the problem of scheduling pipelined series-parallel graphs of moldable tasks. (This problem has been presented in Section 4.2.4.) A bicriteria problem is solved optimally by optimizing the latency under throughput constraint, under the assumption that a processor executes at most one task. The latency is optimized by computing, for each node of the binary decomposition tree, the optimal latency achievable for any number of processors using dynamic programming. The latency of the series node S(l,r) using m processors is obtained by evaluating $\mathcal{L}(S(l,r),m) = \min_{1 \le j \le m-1}(\mathcal{L}(l,j) + \mathcal{L}(r,m-j))$. The latency of the parallel node P(l,r) on m processors is obtained by evaluating $\mathcal{L}(P(l,r),m) = \min_{1 \le j \le m-1}(\mathcal{L}(l,j),\mathcal{L}(r,m-j))$. The leaves of the binary decomposition tree are the tasks of the application and

their latency is given as an input of the problem. The throughput constraint is ensured by setting the latency of the tasks to infinity on processor allocations that would not respect the throughput constraint. Evaluating the latency of a node for a given number of processors requires O(m) computations and there are $2n-1 \in O(n)$ nodes to estimate in the tree for *m* different values of the number of processors. The overall complexity of the algorithm is $O(nm^2)$. Moreover, the authors remark that if both $\mathcal{L}(l, j) - \mathcal{L}(l, j-1)$ and $\mathcal{L}(r, j) - \mathcal{L}(r, j-1)$ decrease when *j* increases, then $\mathcal{L}(S(l, r), j)$ has the same property. Using that property appropriately enables to obtain the optimal latency for chains (series graph) in $O(m \log n)$ by considering the whole chain at once (instead of using the Binary Decomposition Tree form). If a graph is only composed of parallel tasks, considering the whole graph at once leads to an algorithm of similar complexity. Finally, Choudhary et al. [1994] show how to efficiently solve the problem of optimizing the throughput under latency constraint.

Hasan and Motwani [1994] and its refinement [Chekuri et al. 1995] are interested in optimizing the throughput of pipelined trees for database applications. Homogeneous processors are used, the communications do not overlap with computations, and the communications follow the bandwidth bounded multiport model. Therefore the load of a processor is the sum of the weights of the nodes executed by this processor plus the weights of the edges to other processors. Since latency is not a concern here, there is no fine-grain scheduling of tasks and datasets, but only a flow-like solution where each processor has a large buffer of datasets to execute. This solution is very similar to the notion of core schedule described in Section 4.2.2. The main contribution of Hasan and Motwani [1994] is the definition of a monotone tree, which is a modified version of a tree where two nodes linked by too high communication edge are merged. This technique is an extension of the monotonic chains used in Iqbal [1992]. It is shown that such a modification is optimal.

Chekuri et al. [1995] present two approximation algorithms for the previous problem. Both are based on a two-phase decomposition: first the tree is decomposed into a forest by removing some edges; then the trees are allocated to processors using LPT. Removing an edge incurs communication costs to both extremities of the edge. It is shown that if the obtained forest does not have too large trees and the load is kept reasonable, then LPT will generate an approximation of the optimal solution. Two tree decomposition algorithms follow. The first one is a simple greedy algorithm with approximation ratio 3.56, the second one is a more complex greedy algorithm with approximation ratio 2.87.

5.4. Scheduling with Replication

Subhlok and Vondran [1995] address the problem of scheduling a chain of moldable tasks to optimize the throughput using replicated interval mappings: if a task is replicated, its whole interval is replicated too. The algorithm uses dynamic programming to find the intervals $I = \{I_1, \ldots, I_k\}$, and for interval $I_j \in I$, the number of processors m_{I_j} and the number of replications r_{I_j} for this interval, so that the period $\mathcal{P} = \max_{I_j \in I} \frac{p(I_j, m_{I_j})}{r_{I_j}}$ is minimized. Note that $p(I_j, m_{I_j})$ is the period of interval $I_j \in I$ executed on m_{I_j} processors. However, the algorithm does not return the periodic schedule that the system should follow. It just states where the tasks should be executed and it relies on a demand-driven middleware to execute them correctly. A periodic schedule reaching the same throughput can be computed from the intervals, and the number of a number of datasets equal to the least common multiple of the number of replication of all the intervals to completely provide the schedule. Indeed, if one task is replicated two times and another is replicated three times, the execution of six datasets must be unrolled for the schedule to be periodic.

Subhlok and Vondran [1996] add the computation of the latency to Subhlok and Vondran [1995]. Since the graph is a chain and all the intervals are executed independently, it is possible to build a schedule that reaches the optimal latency for a given processor allocation $\mathcal{L} = \sum_{I_j \in I} p(I_j, m_{I_j})$. The interval that constrains the throughput must be executed without idle time, the preceding tasks are scheduled as late as possible, and the following tasks are scheduled as soon as possible. The optimization of the latency under a throughput constraint is obtained using a dynamic programming algorithm, by forbidding the numbers of processors and numbers of replications for each interval that violate the throughput constraint.

5.5. General Method to Optimize the Throughput

Banerjee et al. [1995] deal with executing a signal processing application on heterogeneous machines, where not all tasks can be executed on all type of processors. They schedule a precedence task graph of sequential monolithic tasks. Communications follow the bandwidth bounded multiport model with latency, and they overlap with computations. The proposed algorithm relies on the notion of processor-ordered schedule. (Recall that a schedule is processor ordered if the graph of the communications between the processors is acyclic.) This property helps ensuring that precedence constraints are respected. The algorithm first builds a schedule by clustering some tasks together to reduce the size of the graph. Then, an exponential algorithm finds the optimal processor-ordered schedule of the clustered task graph. Finally, tasks are greedily moved from one processor to the other. The last two steps are alternatively executed as long as the throughput improves.

Beynon [2001] deals with scheduling pipelined task graphs on the Grid. The resources of the Grid are exploited using replication. The author proposes the Filter Copy Pipeline (FCP) algorithm. FCP considers the application graph in a topological order, and chooses the number of copies for each task, so that it can process the data it receives without getting a large backlog. In other words, if the predecessor of a task handles x datasets per time unit, FCP replicates this task to handle x datasets per time unit. Those replicas are allocated to processors using the earliest completion time rule. To improve this approach, Spencer et al. [2002] propose Balanced Filter Copies that allocates a processor to a single task. In order to ensure the balance, it keeps track of the network bandwidth used while computing the schedule.

Taura and Chien [1999] are concerned with scheduling a pipelined task graph on a heterogeneous network of heterogeneous (related) processors with computation and communication overlap. Since the authors are only interested in the throughput, the problem reduces to a mapping of tasks to processors, and the throughput of the solution is given by the most loaded processor or communication link. The algorithm starts by ordering the tasks in depth-first traversal of a clustering tree. Then tasks are mapped to the processors using the following algorithm. The processors are, one after the other, loaded with the first unallocated tasks that minimize the maximum of three quantities: the current load, the perfectly balanced load of the unallocated tasks on the unallocated processors, and the communication volume between the tasks currently allocated to the processor. Finally, the obtained schedule is iteratively improved by unscheduling some of the tasks on the most loaded processors and/or links, and scheduling them again.

Yang et al. [2003] deal with scheduling arbitrary precedence task graphs on a Network Of Workstations (NOW). The processors are heterogeneous (related) and allow for communication and computation overlap. The communications are modeled using the delay model where the delay is computed using a per-link latency and bandwidth. Two objectives are optimized: the throughput and number of machines used from the NOW. The throughput is given by the user and then the execution is cut in stages whose lengths are given by the throughput. A processor used in one stage is not reused in the next one, so that the throughput can be guaranteed. The tasks are allocated using the earliest-task-first heuristic. The authors also propose some techniques to compact the schedule, reducing the number of processors used.

5.6. General Method to Optimize Throughput and Latency

Guirado et al. [2005] are interested in scheduling a pipelined precedence task graph on a homogeneous cluster with communication and computation overlap to optimize both latency and throughput. The network is assumed completely connected and the delay model is used. The delay between two tasks scheduled on two processors is computed using a latency plus bandwidth model. The EXPERT algorithm optimizes the latency under a throughput constraint. Given a throughput goal $\mathcal{T} = 1/\mathcal{P}$, all tasks are partitioned into stages. Each stage is identified by an integer. Task t is allocated to the minimum stage k such that $topLevel(t) \leq k \times \mathcal{P}$, where topLevel(t)is the critical path from the root of the graph to the completion of t. Then, EXPERT considers all the paths of the graph from root to sink in decreasing order of length, including communication delays, and for each edge of each path it applies the following greedy rule: if the two clusters linked by the edge belong to the same stage, and if the sum of the processing times of the tasks in these clusters is smaller than \mathcal{P} , then the two clusters are merged. Finally, inter-stage clusters are merged as long as the sum of the processing times of the tasks of the resulting cluster is less than \mathcal{P} . Each cluster is assigned to a different processor (and the throughput goal is declared infeasible if there are not enough processors). Communication between the processors are grouped at the end of the execution of the cluster to which they are assigned.

Hary and Ozgüner [1999] deal with arbitrary application graphs and homogeneous processors and network links. The technique was originally designed for hypercube networks, but can be adapted to arbitrary networks. It assumes communication and computation overlap. The authors are interested in optimizing both latency and throughput. The proposed algorithm only provides a processor allocation, and the periodic schedule is reconstructed using a technique similar to the one presented in Section 4.2.5: within a period, the tasks are ordered in topological order. If a task precedence constraint is not satisfied inside the current period, it enforces the dependency from the previous period. Given a target period, and therefore throughput, the proposed method has three phases. The first phase groups the tasks in as many clusters as there are processors. This phase orders communications in nonincreasing order of their size, and greedily considers grouping the tasks at both ends of the communication inside the same cluster, unless both tasks are already assigned to a cluster, or the assignment would make the sum of the processing times of the tasks in the cluster larger than the period. At the end of this phase, tasks that are not assigned to a cluster are assigned using a first-fit rule. Then, in the second phase, the clusters are mapped to computation nodes to minimize the amount of communication. This is done by first mapping the clusters randomly to the nodes. Then the processor set is cut in two equals parts and clusters are pairwise exchanged to decrease communications on the processor cut. The communications in each part of the processor set are optimized by applying the "cut in two parts and exchange clusters" procedure recursively. Finally, in the third phase, the solution is improved iteratively by moving tasks between processors to decrease the load of the most loaded link.

Vydyanathan et al. [2010] consider optimizing the latency and throughput of arbitrary DAGs on homogeneous processors linked by a network of different bandwidths, with communication/computation overlap, and using replication and duplication. Communications are explicitly scheduled according to the k-port model. The algorithm

50:29

operates in three phases and takes a throughput constraint as a parameter. The first phase groups tasks together in as many clusters as necessary to match the throughput constraint. This is achieved by considering the replication of each task to deal with computational bottlenecks, and the duplication of each task and merging the clusters of the tasks at both ends of an edge to decrease communication bottlenecks. For each alteration of the mapping considered, the throughput is evaluated by scheduling the computations and the communications using a greedy algorithm. In a second phase, the number of clusters is reduced to the number of processors in the system by merging clusters to minimize processor idle times. Finally, in the last phase, the latency is minimized by considering for each task of the critical path its duplication, and merging to its predecessor cluster or successor cluster.

6. CONCLUSION AND FUTURE WORK

50:30

In this survey, we have presented an overview of pipelined workflow scheduling, a problem that asks for an efficient execution of a streaming application that operates on a set of consecutive datasets. We described the components of application and platform models, and how a scheduling problem can be formulated for a given application. We presented a brief summary of the solution methods for specific problems, highlighting the frontier between polynomial and NP-hard optimization problems.

Although there is a significant body of literature for this complex problem, realistic application scenarios still call for more work in the area, both theoretical and practical. When developing solutions for real-life applications, one has to consider all the ingredients of the schedule as a whole, including detailed communication models and memory requirements (especially when more than one dataset is processed in a single period). Such additional constraints make the development of efficient scheduling methods even more difficult.

As the literature shows, having structure either in the application graph or in the execution platform graph dramatically helps for deriving effective solutions. We think that extending this concept to the schedule could be useful too. For example, for scheduling arbitrary DAGs, developing structured schedules, such as convex clustered schedules, has a potential for yielding new results in this area.

Finally, as the domain evolves, new optimization criteria must be introduced. In this article, we have mainly dealt with throughput and latency. Other performance-related objectives arise with the advent of very large-scale platforms, such as increasing the reliability of the schedule (e.g., through task duplication). Environmental and economic criteria, such as the energy dissipated throughout the execution, or the rental cost of the platform, are also likely to play an increasing role. Altogether, we believe that future research will be devoted to optimizing several performance-oriented and environmental criteria simultaneously. Achieving a reasonable trade-off between all these multiple and antagonistic objectives will prove a very interesting algorithmic challenge.

ACKNOWLEDGMENTS

We would like to wholeheartedly thank the three reviewers, whose comments and suggestions greatly helped us to improve the final version of the article.

REFERENCES

AGNETIS, A., MIRCHANDANI, P. B., PACCIARELLI, D., AND PACIFICI, A. 2004. Scheduling problems with two competing agents. *Oper. Res.* 52, 2, 229–242.

AGRAWAL, K., BENOIT, A., MAGNAN, L., AND ROBERT, Y. 2010. Scheduling algorithms for linear work-flow optimization. In Proceedings of the 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS'10). IEEE Computer Society Press.

- AGRAWAL, K., BENOIT, A., AND ROBERT, Y. 2008. Mapping linear workflows with computation/communication overlap. In Proceedings of the 14th IEEE International Conference on Parallel and Distributed Systems (ICPADS'08). IEEE.
- AHMAD, I. AND KWOK, Y.-K. 1998. On exploiting task duplication in parallel program scheduling. IEEE Trans. Parallel Distrib. Syst. 9, 9, 872–892.
- ALEXANDROV, A., IONESCU, M. F., SCHAUSER, K. E., AND SCHEIMAN, C. 1995. LogGP: Incorporating long messages into the LogP model - One step closer towards a realistic model for parallel computation. In Proceedings of the 7th Annual Symposium on Parallelism in Algorithms and Architectures (SPAA'95).
- ALLAN, V. H., JONES, R. B., LEE, R. M., AND ALLAN, S. J. 1995. Software pipelining. ACM Comput. Surv. 27, 3, 367–432.
- AMDAHL, G. M. 1967. Validity of the single processor approach to achieving large scale computing capabilities. In Proceedings of the Spring Joint Computer Conference (AFIPS'67). ACM Press, New York, 483–485.
- BANERJEE, S., HAMADA, T., CHAU, P. M., AND FELLMAN, R. D. 1995. Macro pipelining basedscheduling on high performance heterogeneous multiprocessor systems. *IEEE Trans. Signal Process.* 43, 6, 1468–1484.
- BANSAL, N., KIMBREL, T., AND PRUHS, K. 2007. Speed scaling to manage energy and temperature. J. ACM 54, 1, 1–39.
- BEAUMONT, O., LEGRAND, A., MARCHAL, L., AND ROBERT, Y. 2004. Assessing the impact and limits of steadystate scheduling for mixed task and data parallelism on heterogeneous platforms. In Proceedings of the 3rd International Symposium on Parallel and Distributed Computing / 3rd International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks (ISPDC'04). IEEE Computer Society, 296–302.
- BENOIT, A. 2009. Scheduling pipelined applications: Models, algorithms and complexity. Habilitation a diriger des recherches. Tech. rep., Ecole Normale Superieure de Lyon.
- BENOIT, A., GAUJAL, B., GALLET, M., AND ROBERT, Y. 2009a. Computing the throughput of replicated workflows on heterogeneous platforms. In *Proceedings of the 38th International Conference on Parallel Processing* (ICPP'09). IEEE Computer Society Press.
- BENOIT, A., KOSCH, H., REHN-SONIGO, V., AND ROBERT, Y. 2009b. Multi-criteria scheduling of pipeline workflows (and application to the JPEG encoder). Int. J. High Perform. Comput. Appl. 23, 2, 171–187.
- BENOIT, A., ROBERT, Y., AND THIERRY, E. 2009c. On the complexity of mapping linear chain applications onto heterogeneous platforms. *Parallel Process. Lett.* 19, 3, 383–397.
- BENOIT, A., REHN-SONIGO, V., AND ROBERT, Y. 2007. Multi-criteria scheduling of pipeline workflows. In Proceedings of the 6th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks (HeteroPar'07).
- BENOIT, A., REHN-SONIGO, V., AND ROBERT, Y. 2008. Optimizing latency and reliability of pipeline workflow applications. In *Proceedings of the 17th International Heterogeneity in Computing Workshop (HCW'08)*. IEEE.
- BENOIT, A., RENAUD-GOUD, P., AND ROBERT, Y. 2010. Performance and energy optimization of concurrent pipelined applications. In Proceedings of the 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS'10). IEEE Computer Society Press.
- BENOIT, A. AND ROBERT, Y. 2008. Mapping pipeline skeletons onto heterogeneous platforms. J. Parallel Distrib. Comput. 68, 6, 790–808.
- BENOIT, A. AND ROBERT, Y. 2009. Multi-criteria mapping techniques for pipeline workflows on heterogeneous platforms. In *Recent Developments in Grid Technology and Applications*. G. A. Gravvanis, J. P. Morrison, H. R. Arabnia, and D. A. Power, Eds., Nova Science Publishers, 65–99.
- BENOIT, A. AND ROBERT, Y. 2010. Complexity results for throughput and latency optimization of replicated and data-parallel workflows. *Algorithmica* 57, 4, 689–724.
- BERMAN, F., CHIEN, A., COOPER, K., DONGARRA, J., FOSTER, I., GANNON, D., JOHNSSON, L., KENNEDY, K., KESSELMAN, C., MELLOR-CRUMME, J., REED, D., TORCZON, L., AND WOLSKI, R. 2001. The GrADS project: Software support for high-level grid application development. Int. J. High Perform. Comput. Appl. 15, 4, 327–344.
- BESSERON, X., BOUGUERRA, S., GAUTIER, T., SAULE, E., AND TRYSTRAM, D. 2009. Fault tolerance and availability awareness in computational grids. In *Fundamentals of Grid Computing (Numerical Analysis and Scientific Computing)*, F. Magoules, Ed., Chapman and Hall/CRC Press.
- BEYNON, M. D. 2001. Supporting data intensive applications in a heterogeneous environment. Ph.D. dissertation, University of Maryland.
- BEYNON, M. D., KURC, T., CATALYUREK, U. V., CHANG, C., SUSSMAN, A., AND SALTZ, J. 2001. Distributed processing of very large datasets with datacutter. *Parallel Comput.* 27, 11, 1457–1478.
- BHAT, P. B., RAGHAVENDRA, C.S., AND PRASANNA, V. K. 2003. Efficient collective communication in distributed heterogeneous systems. J. Parallel Distrib. Comput. 63, 3, 251–263.

- BLELLOCH, G. E., HARDWICK, J. C., SIPELSTEIN, J., ZAGHA, M., AND CHATTERJEE, S. 1994. Implementation of a portable nested data-parallel language. J. Parallel Distrib. Comput. 21, 1, 4–14.
- BLIKBERG, R. AND SOREVIK, T. 2005. Load balancing and OpenMP implementation of nested parallelism. Parallel Comput. 31, 10-12, 984-998.
- BOKHARI, S. H. 1988. Partitioning problems in parallel, pipeline, and distributed computing. *IEEE Trans. Comput.* 37, 1, 48–57.
- BOWERS, S., MCPHILLIPS, T. M., LUDASCHER, B., COHEN, S., AND DAVIDSON, S. B. 2006. A model for user-oriented data provenance in pipelined scientific workflows. In Proceedings of the Provenance and Annotation of Data, International Provenance and Annotation Workshop (IPAW'06). 133–147.
- BRUCKER, P. 2007. Scheduling Algorithms 5th Ed. Springer.
- CHEKURI, C., HASAN, W., AND MOTWANI, R. 1995. Scheduling problems in parallel query optimization. In Proceedings of the 14th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'95). ACM Press, New York, 255–265.
- CHOUDHARY, A., LIAO, W.-K., WEINER, D., VARSHNEY, P., LINDERMAN, R., LINDERMAN, M., AND BROWN, R. 2000. Design, implementation and evaluation of parallel pipelined STAP on parallel computers. *IEEE Trans. Aerospace Electron. Syst.* 36, 2, 655–662.
- CHOUDHARY, A., NARAHARI, B., NICOL, D., AND SIMHA, R. 1994. Optimal processor assignment for a class of pipeline computations. *IEEE Trans. Parallel Distrib. Syst.* 5, 4, 439–443.
- COLE, M. 2004. Bringing skeletons out of the closet: A pragmatic manifesto for skeletal parallel programming. Parallel Comput. 30, 3, 389–406.
- CULLER, D., KARP, R., PATTERSON, D., SAHAY, A., SCHAUSER, K. E., SANTOS, E., SUBRAMONIAN, R., AND VON EICKEN, T. 1993. LogP: Towards a realistic model of parallel computation. In Proceedings of the 4th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'93). ACM, New York, 1–12.
- DARTE, A., ROBERT, Y., AND VIVIEN, F. 2000. Scheduling and Automatic Parallelization. Birkhauser.
- DAVIS, A. L. 1978. Data driven nets: A maximally concurrent, procedural, parallel process representation for distributed control systems. Tech. rep., Department of Computer Science, University of Utah, Salt Lake City, UT.
- DEELMAN, E., BLYTHE, J., GIL, Y., AND KESSELMAN, C. 2003. Workflow management in GriPhyN. In *Grid Resource Management*, Springer.
- DEELMAN, E., SINGH, G., SU, M. H., BLYTHE, J., GIL, A., KESSELMAN, C., MEHTA, G., VAHI, K., BERRIMAN, G. B., GOOD, J., LAITY, A., JACOB, J. C., AND KATZ, D. S. 2005. Pegasus: A framework for mapping complex scientific workflows onto distributed systems. Sci. Program. J. 13, 219–237.
- DENNIS, J. B. 1974. First version of a data flow procedure language. In Proceedings of the Symposium on Programming. 362–376.
- DENNIS, J. B. 1980. Data flow supercomputers. Comput. 13, 11, 48-56.
- MAHESWARI, U. AND DEVI, C. 2009. Scheduling recurrent precedence-constrained task graphs on a symmetric shared-memory multiprocessor. In *Proceedings of the European Conference on Parallel Processing* (*EuroPar'09*). Lecture Notes in Computer Science, vol. 5704, Springer, 265–280.
- NASCIMENTO, L. T. D., FERREIRA, R. A., MEIRA, W. JR., AND GUEDES, D. 2005. Scheduling data flow applications using linear programming. In Proceedings of the 34th International Conference on Parallel Processing (ICPP'05). IEEE Computer Society, 638–645.
- DUTOT, P.-F., RZADCA, K., SAULE, E., AND TRYSTRAM, D. 2009. Multi-objective scheduling. In *Introduction to Scheduling*, Y. Robert and F. Vivien, Eds., CRC Press, Boca Raton, FL.
- FAHRINGER, T., JUGRAVU, A., PLLANA, S., PRODAN, R., SERAGIOTTO, C., JR., AND TRUONG, H.-L. 2005. ASKALON: A tool set for cluster and Grid computing: Research articles. Concurr. Comput. Pract. Exper. 17, 2–4, 143–169.
- FAHRINGER, T., PLLANA, S., AND TESTORI, J. 2004. Teuta: Tool support for performance modeling of distributed and parallel applications. In *Proceedings of the International Conference on Computational Science, Tools for Program Development and Analysis in Computational Science.*
- FEITELSON, D. G., RUDOLPH, L., SCHWIEGELSHOHN, U., SEVCIK, K. C., AND WONG, P. 1997. Theory and practice in parallel job scheduling. In Proceedings of the Conference on Job Scheduling Strategies for Parallel Processing. Springer, 1–34.
- FOSTER, I., KESSELMAN, C., AND TUECKE, S. 2001. The anatomy of the grid: Enabling scalable virtual organizations. Int. J. High Perform. Comput. Appl. 15, 3, 200–222.
- GAIRING, M., MONIEN, B., AND WOCLAW, A. 2005. A faster combinatorial approximation algorithm for scheduling unrelated parallel machines. In Automata, Languages and Programming, vol. 3580, Springer, 828–839.
- GAREY, M. R. AND JOHNSON, D. S. 1979. Computers and Intractability. Freeman, San Francisco.

- GIRAULT, A., SAULE, E., AND TRYSTRAM, D. 2009. Reliability versus performance for critical applications. J. Parallel Distrib. Comput. 69, 3, 326–336.
- GONZALEZ, T. F., IBARRA, O. H., AND SAHNI, S. 1977. Bounds for LPT schedules on uniform processors. SIAM J. Comput. 6, 155–166.
- GRAHAM, R. L. 1966. Bounds for certain multiprocessing anomalies. Bell Syst. Tech. J. 45, 1563–1581.
- GRAHAM, R. L. 1969. Bounds on multiprocessing timing anomalies. SIAM J. Appl. Math. 17, 2, 416-429.
- GUINAND, F., MOUKRIM, A., AND SANLAVILLE, E. 2004. Sensitivity analysis of tree scheduling on two machines with communication delays. *Parallel Comput.* 30, 103–120.
- GUIRADO, F., RIPOLL, A., ROIG, C., HERNANDEZ, A., AND LUQUE, E. 2006. Exploiting throughput for pipeline execution in streaming image processing applications. In *Proceedings of the European Conference on Parallel Processing (EuroPar'06)*. Lecture Notes in Computer Science, vol. 4128, Springer, 1095–1105.
- GUIRADO, F., RIPOLL, A., ROIG, C., AND LUQUE, E. 2005. Optimizing latency under throughput requirements for streaming applications on cluster execution. In Proceedings of the IEEE International Conference on Cluster Computing. IEEE, 1–10.
- HA, S. AND LEE, E. A. 1997. Compile-time scheduling of dynamic constructs in dataflow program graphs. IEEE Trans. Comput. 46, 7, 768–778.
- HAN, Y., NARAHARI, B., AND CHOI, H.-A. 1992. Mapping a chain task to chained processors. *Inf. Process. Lett.* 44, 141–148.
- HARTLEY, T. D. R. AND CATALYUREK, U. V. 2009. A component-based framework for the cell broadband engine. In Proceedings of 23rd International Parallel and Distributed Processing Symposium, The 18th Heterogeneous Computing Workshop (HCW'09).
- HARTLEY, T. D. R., CATALYUREK, U. V., RUIZ, A., IGUAL, F., MAYO, R., AND UJALDON, M. 2008. Biomedical image analysis on a cooperative cluster of GPUs and multicores. In Proceedings of the 22nd Annual International Conference on Supercomputing (ICS'08). 15–25.
- HARTLEY, T. D. R., FASIH, A. R., BERDANIER, C. A., OZGUNER, F., AND CATALYUREK, U. V. 2009. Investigating the use of GPU-accelerated nodes for SAR image formation. In Proceedings of the IEEE International Conference on Cluster Computing, Workshop on Parallel Programming on Accelerator Clusters (PPAC'09).
- HARY, S. L. AND OZGUNER, F. 1999. Precedence-constrained task allocation onto point-to-point networks for pipelined execution. IEEE Trans. Parallel Distrib. Syst. 10, 8, 838–851.
- HASAN, W. AND MOTWANI, R. 1994. Optimization algorithms for exploiting the parallelism-communication trade-off in pipelined parallelism. In Proceedings of the 20th International Conference on Very Large Databases (VLDB'94). 36–47.
- HOCHBAUM, D. S., ED. 1997. Approximation Algorithms for NP-Hard Problems. PWS Publishing.
- HOCHBAUM, D. S. AND SHMOYS, D. B. 1987. Using dual approximation algorithms for scheduling problems: Practical and theoretical results. J. ACM 34, 144–162.
- HOCHBAUM, D. S. AND SHMOYS, D. B. 1988. A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach. SIAM J. Comput. 17, 3, 539–551.
- HONG, B. AND PRASANNA, V. K. 2003. Bandwidth-aware resource allocation for heterogeneous computing systems to maximize throughput. In Proceedings of the 32th International Conference on Parallel Processing (ICPP'03). IEEE Computer Society Press.
- IQBAL, M. A. 1992. Approximate algorithms for partitioning problems. Int. J. Parallel Program. 20, 5, 341–361.
- ISARD, M., BUDIU, M., YU, Y., BIRRELL, A., AND FETTERLY, D. 2007. Dryad: Distributed data-parallel programs from sequential building blocks. In Proceedings of the 2nd ACM SIGOPS European Conference on Computer Systems (EuroSys'07). ACM, New York, 59–72.
- JEJURIKAR, R., PEREIRA, C., AND GUPTA, R. 2004. Leakage aware dynamic voltage scaling for real-time embedded systems. In Proceedings of the 41st Annual Design Automation Conference (DAC'04). ACM, New York, 275–280.
- JONSSON, J. AND VASELL, J. 1996. Real-time scheduling for pipelined execution of data flow graphs on a realistic multiprocessor architecture. In Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'96). Vol. 6, IEEE, 3314–3317.
- KAHN, G. 1974. The semantics of simple language for parallel programming. In Proceedings of the IFIP Congress. 471–475.
- KENNEDY, K. AND ALLEN, J. R. 2002. Optimizing Compilers for Modern Architectures: A Dependence-Based Approach. Morgan Kaufmann, San Fransisco.
- KLISSIPONGSE, E. AND NGAMSURIYAROJ, S. 2010. Placing pipeline stages on a grid: Single path and multipath pipeline execution. Future Generat. Comput. Syst. 26, 1, 50–62.

- KIM, J., GIL, Y., AND SPRARAGEN, M. 2004. A knowledge-based approach to interactive workflow composition. In Proceedings of the 14th International Conference on Automatic Planning and Scheduling (ICAPS 04).
- KNOBE, K., REHG, J. M., CHAUHAN, A., NIKHIL, R. S., AND RAMACHANDRAN, U. 1999. Scheduling constrained dynamic applications on clusters. In Proceedings of the ACM/IEEE Conference on Supercomputing. ACM, New York, 46.
- KWOK, Y.-K. AND AHMAD, I. 1999a. Benchmarking and comparison of the task graph scheduling algorithms. J. Parallel Distrib. Comput. 59, 3, 381–422.
- KWOK, Y.-K. AND AHMAD, I. 1999b. Static scheduling algorithms for allocating directed task graphs to multiprocessors. ACM Comput. Surv. 31, 4, 406–471.
- LEE, E. A. AND PARKS, T. M. 1995. Dataflow process networks. Proc. IEEE 83, 5, 773-801.
- LEE, M., LIU, W., AND PRASANNA, V. K. 1998. A mapping methodology for designing software task pipelines for embedded signal processing. In Proceedings of the Workshop on Embedded HPC Systems and Applications of IPPS/SPDP. 937–944.
- LENSTRA, J. K., SHMOYS, D. B., AND TARDOS, E. 1990. Approximation algorithms for scheduling unrelated parallel machines. Math. Program. 46, 259–271.
- LEPERE, R. AND TRYSTRAM, D. 2002. A new clustering algorithm for large communication delays. In *Proceedings* of the International Parallel and Distributed Processing Symposium (IPDPS'02). IEEE Computer Society Press.
- LEVNER, E., KATS, V., PABLO, D. A. L. D., AND CHENG, T.C.E. 2010. Complexity of cyclic scheduling problems: A state-of-the-art survey. Comput. Industr. Engin. 59, 352–361.
- LITZKOW, M. J., LIVNY, M., AND MUTKA, M. W. 1988. Condor-A hunter of idle workstations. In Proceedings of the 8th International Conference on Distributed Computing Systems. 104–111.
- MACKENZIE-GRAHAM, A., PAYAN, A., DINOV, I. D., HORN, J. D. V., AND TOGA, A. W. 2008. Neuroimaging data provenance using the LONI pipeline workflow environment. In Proceedings of the Provenance and Annotation of Data, International Provenance and Annotation Workshop (IPAW'08). 208–220.
- MANNE, F. AND OLSTAD, B. 1995. Efficient partitioning of sequences. *IEEE Trans. Comput.* 44, 11, 1322–1326. MICROSOFT. 2009. AXUM webpage. http://msdn.microsoft.com/en-us/devlabs/dd795202.aspx.
- MILLS, M. P. 1999. The Internet Begins with Coal: A Preliminary Exploration of the Impact of the Internet on Electricity Consumption: A Green Policy Paper for the Greening Earth Society. Mills-McCarthy & Associates.
- MORENO, A., CESAR, E., GUEVARA, A., SORRIBES, J., MARGALEF, T., AND LUQUE, E. 2008. Dynamic pipeline mapping (DPM). http://www.sciencedirect.com/science/article/pii/S0167819111001566.
- NICOL, D. 1994. Rectilinear partitioning of irregular data parallel computations. J. Parallel Distrib. Comput. 23, 119–134.
- NIKOLOV, H., THOMPSON, M., STEFANOV, T., PIMENTEL, A. D., POLSTRA, S., BOSE, R., ZISSULESCU, C., AND DEPRETTERE, E. F. 2008. Daedalus: Toward composable multimedia MP-SoC design. In Proceedings of the 45th Annual Design Automation Conference (DAC'08). ACM, New York, 574–579.
- OINN, T., GREENWOOD, M., ADDIS, M., ALPDEMIR, N., FERRIS, J., GLOVER, K., GOBLE, C., GODERIS, A., HULL, D., MARVIN, D., LI, P., LORD, P., POCOCK, M., SENGER, M., STEVENS, R., WIPAT, A., AND WROE, C. 2006. Taverna: Lessons in creating a work- flow environment for the life sciences. *Concurr. Comput. Pract. Exper.* 18, 10, 1067–1100.
- OKUMA, T., YASUURA, H., AND ISHIHARA, T. 2001. Software energy reduction techniques for variable-voltage processors. *IEEE Des. Test Comput.* 18, 2, 31–41.
- PAPADIMITRIOU, C. H. AND YANNAKAKIS, M. 2000. On the approximability of trade-offs and optimal access of web sources. In Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS'00). 86–92.
- PECERO-SANCHEZ, J. E. AND TRYSTRAM, D. 2005. A new genetic convex clustering algorithm for parallel time minimization with large communication delays. In *PARCO* (John von Neumann Institute for Computing Series), G. R. Joubert, W. E. Nagel, F. J. Peters, O. G. Plata, P. Tirado, and E. L. Zapata, Eds., vol. 33, Central Institute for Applied Mathematics, Julich, Germany, 709–716.
- PINAR, A. AND AYKANAT, C. 2004. Fast optimal load balancing algorithms for 1D partitioning. J. Parallel Distrib. Comput. 64, 8, 974–996.
- PINAR, A., TABAK, E. K., AND AYKANAT, C. 2008. One-dimensional partitioning for heterogeneous systems: Theory and practice. J. Parallel Distrib. Comput. 68, 1473–1486.
- PRATHIPATI, R. B. 2004. Energy efficient scheduling techniques for real-time embedded systems. Master's thesis, Texas A&M University.

- RANAWEERA, S. AND AGRAWAL, D. P. 2001. Scheduling of periodic time critical applications for pipelined execution on heterogeneous systems. In Proceedings of the International Conference on Parallel Processing (ICPP'01). IEEE Computer Society, 131–140.
- RAYWARD-SMITH, V. J. 1987. UET scheduling with interprocessor communication delays. *Discr. Appl. Math.* 18, 55–71.
- RAYWARD-SMITH, V. J., BURTON, F. W., AND JANACEK, G. J. 1995. Scheduling parallel program assuming preallocation. In Scheduling Theory and its Applications, P. Chretienne, E. G. Coffman Jr., J. K. Lenstra, and Z. Liu, Eds., Wiley, 146–165.
- REINDERS, J. 2007. Intel Threading Building Blocks. O' Reilly.
- Rowe, A., KALAITZOPOULOS, D., OSMOND, M., GHANEM, M., AND GUO, Y. 2003. The discovery net system for high throughput bioinformatics. *Bioinf*, 19, 1, 225–231.
- SAIF, T. AND PARASHAR, M. 2004. Understanding the behavior and performance of non-blocking communications in MPI. In Proceedings of the European Conference on Parallel Processing (EuroPar'04). Lecture Notes in Computer Science, vol. 3149, Springer, 173–182.
- SERTEL, O., KONG, J., SHIMADA, H., CATALYUREK, U. V., SALTZ, J. H., AND GURCAN, M. N. 2009. Computer-aided prognosis of neuroblastoma on whole-slide images: Classification of stromal development. *Pattern Recogn.* 42, 6, 1093–1103.
- SPENCER, M., FERREIRA, R., BEYNON, M. D., KURC, T., CATALYUREK, U. V., SUSSMAN, A., AND SALTZ, J. 2002. Executing multiple pipelined data analysis operations in the grid. In *Proceedings of the ACM/IEEE Conference on Supercomputing*. IEEE Computer Society Press, 1–18.
- SUBHLOK, J. AND VONDRAN, G. 1995. Optimal mapping of sequences of data parallel tasks. In Proceedings of the 5th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'95). ACM, New York, 134–143.
- SUBHLOK, J. AND VONDRAN, G. 1996. Optimal latency-throughput tradeoffs for data parallel pipelines. In Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'96). ACM Press, New York, 62–71.
- SUHENDRA, V., RAGHAVAN, C., AND MITRA, T. 2006. Integrated scratchpad memory optimization and task scheduling for MPSoC architectures. In Proceedings of the ACM/IEEE International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES'06).
- TANNENBAUM, T., WRIGHT, D., MILLER, K., AND LIVNY, M. 2001. Condor- A distributed job scheduler. In Beowulf Cluster Computing with Linux, T. Sterling, Ed. MIT Press.
- TAURA, K. AND CHIEN, A. 1999. A heuristic algorithm for mapping communicating tasks on heterogeneous resources. In Proceedings of the Heterogeneous Computing Workshop (HCW'99). IEEE Computer Society Press, 102–115.
- TAYLOR, V., WU, X., AND STEVENS, R. 2003. Prophesy: An infrastructure for performance analysis and modeling of parallel and grid applications. SIGMETRICS Perform. Eval. Rev. 30, 4, 13–18.
- TEODORO, G., FIREMAN, D., GUEDES, D., MEIRA, W. JR., AND FERREIRA, R. 2008. Achieving multi-level parallelism in filter-labeled stream programming model. In Proceedings of the 37th International Conference on Parallel Processing (ICPP'08).
- THAIN, D., TANNENBAUM, T., AND LIVNY, M. 2002. Condor and the grid. In *Grid Computing: Making the Global Infrastructure a Reality*, F. Berman, G. Fox, and T. Hey, Eds., John Wiley & Sons.
- T'KINDT, V. AND BILLAUT, J.-C. 2007. Multicriteria Scheduling. Springer.
- T'KINDT, V., BOUIBEDE-HOCINE, K., AND ESSWEIN, C. 2007. Counting and enumeration complexity with application to multicriteria scheduling. Ann. Oper. Res. 153, 215–234.
- VALDES, J., TARJAN, R. E., AND LAWLER, E. L. 1982. The recognition of series parallel digraphs. SIAM J. Comput. 11, 2, 298–313.
- VYDYANATHAN, N., CATALYUREK, U. V., KURC, T. M., SADAYAPPAN, P., AND SALTZ, J. H. 2007. Toward optimizing latency under throughput constraints for application workflows on clusters. In Proceedings of the European Conference on Parallel Processing (EuroPar'07). 173–183.
- VYDYANATHAN, N., CATALYUREK, U. V., KURC, T. M., SADAYAPPAN, P., AND SALTZ, J. H. 2010. Optimizing latency and throughput of application workflows on clusters. *Parallel Comput.* 37, 10–11, 694–712.
- WANG, L., VON LASZEWSKI, G., DAYAL, J., AND WANG, F. 2010. Towards energy aware scheduling for precedence constrained parallel tasks in a cluster with DVFS. In Proceedings of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGrid'10). 368–377.
- WOLFE, M. 1989. Optimizing Supercompilers for Supercomputers. MIT Press, Cambridge MA.
- WOLSKI, R., SPRING, N. T., AND HAYES, J. 1999. The network weather service: A distributed resource performance forecasting service for metacomputing. *Future Gener. Comput. Syst.* 15, 5–6, 757–768.

- YANG, M.-T., KASTURI, R., AND SIVASUBRAMANIAM, A. 2003. A pipeline-based approach for scheduling video processing algorithms on NOW. IEEE Trans. Parallel Distrib. Syst. 14, 2, 119–130.
- YAO, F., DEMERS, A., AND SHENKER, S. 1995. A scheduling model for reduced CPU energy. In Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS'95). IEEE Computer Society, 374.
- YU, J. AND BUYYA, R. 2005. A taxonomy of workflow management systems for grid computing. J. Grid Comput. 3, 3, 171–200.

Received June 2010; revised September 2011; accepted August 2012