



Libraries and Learning Services

University of Auckland Research Repository, ResearchSpace

Version

This is the Accepted Manuscript version. This version is defined in the NISO recommended practice RP-8-2008 <http://www.niso.org/publications/rp/>

Suggested Reference

Sinnen, O., Sousa, L. A., & Sandnes, F. E. (2006). Toward a Realistic Task Scheduling Model. *IEEE Transactions on Parallel and Distributed Systems*, 17(3), 263-275. doi: [10.1109/TPDS.2006.40](https://doi.org/10.1109/TPDS.2006.40)

Copyright

Items in ResearchSpace are protected by copyright, with all rights reserved, unless otherwise indicated. Previously published items are made available in accordance with the copyright policy of the publisher.

© 2006 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

For more information, see [General copyright](#), [Publisher copyright](#), [SHERPA/RoMEO](#).

Towards a Realistic Task Scheduling Model

Oliver Sinnen, Leonel Sousa, Frode Eika Sandnes

Abstract

Task scheduling is an important aspect of parallel programming. Most of the heuristics for this NP-hard problem are based on a very simple system model of the target parallel system. Experiments revealed the inappropriateness of this classic model to obtain accurate and efficient schedules for real systems. In order to overcome this shortcoming, a new scheduling model was proposed that considers the contention for communication resources.

Even though the accuracy and efficiency improved with the consideration of contention, the new contention model is still not good enough. The crucial aspect is the involvement of the processor in communication. This paper investigates the involvement of the processor in communication and its impact on task scheduling. A new system model is proposed based on the contention model that is aware of the processor involvement. The challenges for the scheduling techniques are analysed and two scheduling algorithms are proposed. Experiments on real parallel systems show the significantly improved accuracy and efficiency of the new model and algorithms.

Index Terms

Parallel processing, concurrent programming, scheduling and task partitioning, processor involvement, heterogeneous system model

I. INTRODUCTION

Scheduling is an important aspect of efficient parallel computer utilisation. In task scheduling, the program is represented by a task graph, or Directed Acyclic Graph (DAG), where the nodes represent the tasks of the program and the edges the communications between the tasks. The scheduling problem is to find the spatial and temporal assignments of the tasks onto the processors of the target system which results in the shortest possible execution time of the program. In its general form, this problem is NP-hard [28], [35]. Many heuristics have been proposed for the near optimal solution of the problem [3], [4], [12], [14], [19], [22], [26], [37], [38].

Oliver Sinnen is with Dept. of Electrical and Computer Engineering, University of Auckland, Private Bag 92019, Auckland, New Zealand; Email: o.sinnen@auckland.ac.nz

Leonel Sousa, *Senior Member, IEEE*, is with INESC-ID, Instituto Superior Técnico, Technical University of Lisbon, Rua Alves Redol 9, P-1000-029 Lisboa, Portugal; Email: las@inesc-id.pt

Frode Eika Sandnes is with Dept. Computer Science, Faculty of Engineering, Oslo University College, P.O. Box 4, St. Olav Plass, N-0130 Oslo, Norway. Email: frode-eika.sandnes@iu.hio.no

We would like to thank the anonymous reviewers for their valuable comments and suggestions. This work was partially supported by Portuguese Foundation for Science and Technology (FCT) through Program FEDER, by FCT and FSE (Fundo Social Europeu) in the context of III European Framework of Support, and by the European Commission through ARI grant HPRI-CT-1999-00026 and TMR grant ERB FMGE CT950051 (TRACS Programme at EPCC), which we gratefully acknowledge.

Unfortunately, most of the algorithms are based on a very simple system model, which does not accurately reflect real parallel systems. The main problematic assumptions are: i) a dedicated subsystem for the interprocessor communication; ii) completely concurrent communication; iii) a fully connected communication network. The last two assumptions avoid the consideration of contention for communication resources in task scheduling. Experiments [24], [31] showed that the consideration of contention is essential for the generation of accurate and efficient schedules. A contention aware task scheduling strategy that captures end-point and network contention has been proposed in [34].

Even though the accuracy and efficiency of scheduling was significantly improved through the consideration of contention in task scheduling, the experiments also demonstrated that there is still one aspect regarding communication, which is not sufficiently addressed by the new contention model, namely the involvement of the processors in communication, which is in opposition to the assumption of a dedicated subsystem for the interprocessor communication.

This paper investigates the involvement of the processor in communication, its impact on task scheduling and how it can be considered in task scheduling. A new system model for scheduling that considers the involvement of the processor in communication is proposed based on the contention model. As a result, the new model is general and unifies the existing scheduling models. The concept of edge scheduling, used in contention aware scheduling, is extended to the scheduling of the edges on the processors in order to reflect the processors' involvement. Since scheduling under the new model requires the adjustment of the existing techniques, it is shown how this can be done for list scheduling and a new genetic algorithm based heuristic is proposed. Experiments are performed which demonstrate the greatly improved accuracy and efficiency of the schedules produced under the new involvement-contention model.

The remainder of this paper is organised as follows. Section II establishes the background and definitions of task scheduling under the classic and the contention model. Section III generically analyses the processor involvement in communication. Based on this analysis, Section IV investigates the integration of the awareness for the processor involvement in task scheduling. Section V discusses scheduling heuristics based on the new model. Experimental results are

presented in Section VI and this paper concludes with Section VII.

II. TASK SCHEDULING

In task scheduling the program to be scheduled is represented by a directed acyclic graph.

Definition 1 (Directed Acyclic Graph (DAG)): A DAG is a directed acyclic graph $G = (\mathbf{V}, \mathbf{E}, w, c)$ representing a program \mathcal{P} . The nodes in \mathbf{V} represent the tasks of \mathcal{P} and the edges in \mathbf{E} the communications between the tasks. An edge $e_{ij} \in \mathbf{E}$ represents the communication from node n_i to node n_j . The positive weight $w(n)$ associated with node $n \in \mathbf{V}$ represents its computation cost and the non-negative weight $c(e_{ij})$ associated with edge $e_{ij} \in \mathbf{E}$ represents its communication cost.

All instructions, or operations, of one task are executed in sequential order; there is no parallelism within a task. The nodes are strict with respect to both their inputs and their outputs: that is, a node cannot begin execution until all its inputs have arrived, and no output is available until the computation has finished and at that time all outputs are available for communication simultaneously.

The set $\{n_x \in \mathbf{V} : e_{xi} \in \mathbf{E}\}$ of all direct **predecessors** of n_i is denoted by $\text{pred}(n_i)$ and the set $\{n_x \in \mathbf{V} : e_{ix} \in \mathbf{E}\}$ of all direct **successors** of n_i , is denoted by $\text{succ}(n_i)$. A node $n \in \mathbf{V}$ without predecessors, $\text{pred}(n) = \emptyset$, is named **source** node and if it is without successors, $\text{succ}(n) = \emptyset$, it is named **sink** node.

A. Basis

A schedule of a DAG is the association of a start time and a processor with every node of the DAG. To describe a schedule \mathcal{S} of a DAG $G = (\mathbf{V}, \mathbf{E}, w, c)$ on a target system consisting of a set \mathbf{P} of **dedicated** processors the following terms are defined: $t_s(n, P)$ denotes the **start time** and $\omega(n, P)$ the **execution time** of node $n \in \mathbf{V}$ on processor $P \in \mathbf{P}$. Thus, the node's **finish time** is given by $t_f(n, P) = t_s(n, P) + \omega(n, P)$. In a homogeneous system the execution time is equivalent to the computation cost of the node, thus $\omega(n, P) = w(n)$. In a heterogeneous system the computation cost $w(n)$ of node n describes its *average* computation cost. The processor to which n is allocated is denoted by $\text{proc}(n)$. Further, let $t_f(P) = \max_{n \in \mathbf{V} : \text{proc}(n)=P} \{t_f(n)\}$ be

the **processor finish time** of P and let $sl(\mathcal{S}) = \max_{n \in \mathbf{V}} \{t_f(n)\}$ be the **schedule length** (or makespan) of \mathcal{S} , assuming $\min_{n \in \mathbf{V}} \{t_s(n)\} = 0$.

For such a schedule to be feasible, the following two conditions must be fulfilled for all nodes in G .

Condition 1 (Processor Constraint (dedicated processor)): For any two nodes $n_i, n_j \in \mathbf{V}$

$$proc(n_i) = proc(n_j) = P \Rightarrow \begin{cases} t_f(n_i, P) \leq t_s(n_j, P) \\ \text{or } t_f(n_j, P) \leq t_s(n_i, P) \end{cases} \quad (1)$$

Condition 2 (Precedence Constraint (node strictness)): For $n_i, n_j \in \mathbf{V}$, $e_{ij} \in \mathbf{E}$, $P \in \mathbf{P}$,

$$t_s(n_j, P) \geq t_f(e_{ij}), \quad (2)$$

where $t_f(e_{ij})$ is the edge finish time of the communication associated with e_{ij} , which is defined later, depending on the model.

The earliest time a node $n_j \in \mathbf{V}$ can start execution on processor $P \in \mathbf{P}$, which is constrained by n_j 's entering edges (2), is called the **Data Ready Time (DRT)** t_{dr} with

$$t_{dr}(n_j, P) = \max_{e_{ij} \in \mathbf{E}, n_i \in \text{pred}(n_j)} \{t_f(e_{ij})\} \quad (3)$$

and hence

$$t_s(n, P) \geq t_{dr}(n, P) \quad (4)$$

for all $n \in \mathbf{V}$. If $\text{pred}(n) = \emptyset$, i.e. n is a source node, $t_{dr}(n) = t_{dr}(n, P) = 0$, for all $P \in \mathbf{P}$.

B. Classic scheduling

Most scheduling algorithms employ a strongly idealised model of the target parallel system [3], [4], [12], [14], [19], [22], [26], [37], [38]. This model, which shall be referred to as the classic model, is defined in the following, including a generalisation towards heterogeneous processors.

Definition 2 (Classic System Model): A parallel system $M_{classic} = (\mathbf{P}, \omega)$ consists of a finite set of dedicated processors \mathbf{P} connected by a communication network. The processor heterogeneity, in terms of processing speed, is described by the execution time function ω .

This dedicated system has the following properties: i) local communication has zero costs; ii) communication is performed by a communication subsystem; iii) communication can be performed concurrently; iv) the communication network is fully connected.

Based on this system model, the edge finish time only depends on the finish time of the origin node and the communication time.

Definition 3 (Edge Finish Time): The edge finish time of $e_{ij} \in \mathbf{E}$ is given by

$$t_f(e_{ij}) = t_f(n_i) + \begin{cases} 0 & \text{if } \text{proc}(n_i) = \text{proc}(n_j) \\ c(e_{ij}) & \text{otherwise} \end{cases} \quad (5)$$

Thus, communication can overlap with the computation of other nodes, an unlimited number of communications can be performed at the same time, and communication has the same cost $c(e_{ij})$, regardless of the origin and the destination processor, unless the communication is local.

C. Contention aware scheduling

The classic scheduling model (Definition 2) does not consider any kind of contention for communication resources. To make task scheduling contention aware, and thereby more realistic, the communication network is modelled by a graph, where processors are represented by vertices and the edges reflect the communication links. The awareness for contention is achieved by edge scheduling [29], i.e. the scheduling of the edges of the DAG onto the links of the network graph, in a very similar manner to how the nodes are scheduled on the processors.

The details of contention aware scheduling are out of the scope of this paper, the interested reader should refer to [34]. The network model proposed in [34] captures network [29], [32] as well as end-point contention [6], [17] and can represent heterogeneous communication links. This is achieved by using different types of edges and by using switch vertices in addition to processor vertices. Here, it suffices to define the topology network graph to be $TG = (\mathbf{P}, \mathbf{L})$, where \mathbf{P} is a set of vertices representing the processors and \mathbf{L} is a set of edges representing the communication links. The system model is then defined as follows.

Definition 4 (Target Parallel System – Contention Model): A target parallel system $M_{TG} = (TG, \omega)$ consists of a set of possibly heterogeneous processors \mathbf{P} connected by the commu-

nication network $TG = (\mathbf{P}, \mathbf{L})$. This dedicated system has the following properties: i) local communications have zero costs; ii) communication is performed by a communication subsystem.

The notions of concurrent communication and a fully connected network found in the classic model (Definition 2) are substituted by the notion of scheduling the edges \mathbf{E} on the communication links \mathbf{L} . Corresponding to the scheduling of the nodes, $t_s(e, L)$ and $t_f(e, L)$ denote the **start** and **finish time** of edge $e \in \mathbf{E}$ on link $L \in \mathbf{L}$, respectively.

When a communication, represented by the edge e , is performed between two distinct processors P_{src} and P_{dst} , the routing algorithm of TG returns a **route** from P_{src} to P_{dst} : $R = \langle L_1, L_2, \dots, L_l \rangle$, $L_i \in \mathbf{L}$ for $i = 1, \dots, l$. The edge e is scheduled on each link of the route.

This only affects the scheduling of the nodes with an altered definition of the edge finish time (Definition 3).

Definition 5 (Edge Finish Time – Contention Model): Let $G = (\mathbf{V}, \mathbf{E}, w, c)$ be a DAG and $M_{TG} = ((\mathbf{P}, \mathbf{L}), \omega)$ a parallel system. Let $R = \langle L_1, L_2, \dots, L_l \rangle$ be the route for the communication of $e_{ij} \in \mathbf{E}$, $n_i, n_j \in \mathbf{V}$, if $proc(n_i) \neq proc(n_j)$. The finish time of e_{ij} is

$$t_f(e_{ij}) = \begin{cases} t_f(n_i) & \text{if } proc(n_i) = proc(n_j) \\ t_f(e_{ij}, L_l) & \text{otherwise} \end{cases} \quad (6)$$

Thus, the edge finish time $t_f(e_{ij})$ is now the finish time of e_{ij} on the last link of the route, L_l , unless the communication is local.

D. Scheduling heuristics

The scheduling problem is to find a schedule with minimal length. As this problem is NP-hard, under the classic model [28], [35] as well as under the contention model [34], many heuristics have been proposed for its solution.

A heuristic must schedule a node on a processor so that it adheres to all resource (1) and precedence constraints (2). The following Condition 3 ensures this. The notion of a free node, used in the condition, is a node whose predecessors have already been scheduled, which is a requisite for the calculation of the DRT.

Condition 3 (Scheduling Condition): Let $G = (\mathbf{V}, \mathbf{E}, w, c)$ be a DAG and $\mathbf{M}_{classic} = (\mathbf{P}, \omega)$ a parallel system. Let $[A, B]$, $A, B \in [0, \infty]$, be an idle time interval on $P \in \mathbf{P}$, i.e. an interval in which no node is executed. A free node $n \in \mathbf{V}$ can be scheduled on P within $[A, B]$ if

$$\max\{A, t_{dr}(n, P)\} + \omega(n, P) \leq B. \quad (7)$$

So, Condition 3 allows node n to be scheduled between already scheduled nodes (**insertion technique**) [18], i.e. $[A, B] = [t_f(n_{P_l}, P), t_s(n_{P_{l+1}}, P)]$, or after the finish time of processor P (**end technique**) [1], i.e. $[A, B] = [t_f(P), \infty]$.

A similar condition is formulated for the scheduling of the edges on the links [34], additionally considering routing aspects and the causality of the communication along the route.

III. PROCESSOR INVOLVEMENT IN COMMUNICATION

Experimental results demonstrated that the utilisation of the contention aware model in scheduling heuristics significantly improves the accuracy and efficiency of the produced schedules [33]. Yet, the experiments also showed that the contention model is still not sufficiently realistic in terms of communication [31].

The contention model (Definition 4) supposes, as does the classic model (Definition 2), a dedicated communication subsystem to be present in the target system. With the assumed subsystem, computations can overlap with communications, because the processor is not involved in communication. However, many parallel systems do not possess such a subsystem [10]. Therefore, in many systems the processors are involved, in one way or the other, in interprocessor communication. Furthermore, the involvement of the processor also serialises communication, even if the network interfaces were capable of performing multiple message transfers at the same time, since a processor can only handle one communication at a time. For example a processor can usually only perform one memory copy at a time. Thus, considering the processors' involvement in task scheduling is of utmost importance as it serialises the communication and, more importantly, prevents the overlap of computation and communication.

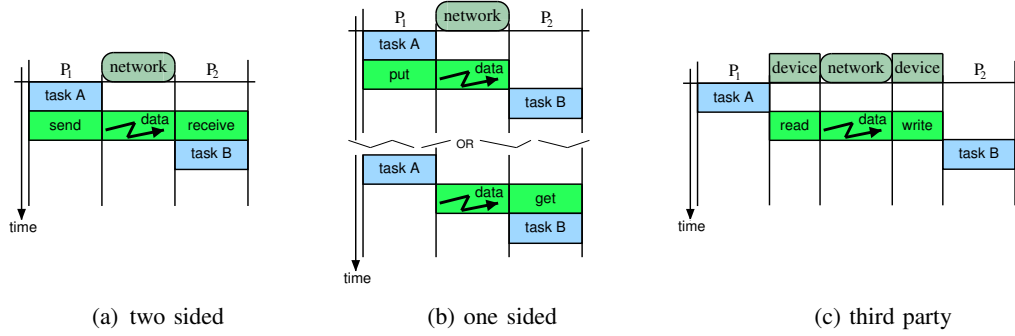


Fig. 1. The three types of interprocessor communication

A. Involvement types

In the context of the processor involvement, interprocessor communication can be divided into three basic types: **two sided**, **one sided** and **third party**, as illustrated in Figure 1.

Two sided – In two sided interprocessor communication both the source and the destination processor are involved in the communication (Figure 1(a)). For example in a PC cluster, the TCP/IP based communication over the LAN involves both processors. The sending processor must divide a message into packages and wrap them into TCP/IP protocol envelopes before setting up the network card for the transfer. On the receiving side, the processor is involved in the unwrapping and assembling of the packages into the original message [10].

One sided – Communication is said to be one sided, if only one of the two participating processors is involved (Figure 1(b)). Thus, this type of communication is limited to shared memory systems: either the source processor writes the data into its destination location (shared memory “put”) or the destination processor reads the data from its source location (shared memory “get”). For example, on the Cray T3E a processor can read from or write into remote memory using special registers [10].

Third party – In third party interprocessor communication, the data transfer is performed by dedicated communication devices, as illustrated in Figure 1(c). The processor only informs the communication device of the memory area it wants transferred and the rest of the communication is performed by the device, usually employing some kind of Direct Memory Access (DMA). Examples for machines that possess such subsystems are Meiko CS-2 [5] or IBM SP-2 [15]. Task scheduling, both under the classic and under the contention model, assumes the third-party

type of interprocessor communication (Definitions 2 and 4).

It is important to note that the software layer employed in parallel programming strongly affects the type of communication used. For example, in a shared memory system, communication can be one sided, but the software layer might use a common buffer (one processor writes, the other reads) which turns it into two sided communication. This effect is not uncommon, as shown by the analysis of MPI implementations on common parallel systems in [16]. Task scheduling should of course reflect the effective type of involvement, taking into account the software layer employed in the code generation.

B. Involvement characteristics

The classification of interprocessor communication into three different types can be refined with the notions of overhead and direct involvement.

Overhead – The first component of the processor’s involvement is the communication setup or *overhead*. From the initiation of the communication process until the first data item is put onto the network, the processor is engaged in preparing the communication. An overhead is in general also imposed on the destination processor from after the data has arrived until it can be used by the receiving task. The overhead consists mainly of the path through the communication layers (e.g. MPI [25], Active Messages [36], TCP/IP) and hence is usually of constant time on both processors. In some environments, however, data might be copied into and from a buffer, which is an operation taking time proportional to the data size. Examples are some MPI implementations as described in [16]. Note, the overhead does not arise for communication between tasks executed on the same processor. Therefore, it cannot be made part of the computation reflected by the origin and the destination tasks of the communication.

Direct involvement – After the communication has been prepared by the processor during the overhead, any further participation in communication is the direct involvement of the processor. It is characterised by the circumstance that data is already in transit on the communication network. Figure 2 features both the overhead and the direct involvement of the processor and thereby illustrates their differences. From now on, the term *involvement* means the direct involvement of the processor and the term *overhead* is used for the pre and post phases discussed above.




Fig. 2. The decomposing of the processor's participation in communication into overhead and direct involvement

The distinction between communication types becomes obsolete if the communication is described in terms of overhead and involvement, even when the processors' participation is only partial. Therefore it is assumed that overhead and involvement are imposed on both the source and the destination processor. Then, one communication type differs from another merely by the length of the two components. Logically, the involvement time is zero on some processors, namely on both in third party communication and on one – either the sending or the receiving processor – in one sided communication. In contrast, the overhead is a common part of interprocessor communication, independent of the type of communication. Thus, third party communication is affected just as much as two and one sided communication, in terms of overhead.

The separation into overhead and involvement is also more general than the approach taken by the LogP model [9] and the computational overhead of communication considered in [28], which is discussed in detail in [30].

IV. INVOLVEMENT SCHEDULING

The notions of overhead and involvement discussed in the last section are the key concept to enhance task scheduling towards the awareness of the processor involvement in communication. In the first step, a new target system model is defined.

Definition 6 (Target Parallel System – Involvement-Contention Model): A target parallel system $M = (TG, \omega, o, i)$ consists of a set of possibly heterogeneous processors \mathbf{P} connected by the communication network $TG = (\mathbf{P}, \mathbf{L})$. This dedicated system has the following property: i) local communication has zero costs.

So, in comparison with the contention model (Definition 4), the involvement-contention model departs from the assumption of a dedicated communication subsystem. Instead, the role of the

processors in communication is described by the new components o – for overhead – and i – for (direct) involvement.

Let $R = \langle L_1, L_2, \dots, L_l \rangle$ be the route for the communication of edge $e \in \mathbf{E}$ from $P_{src} \in \mathbf{P}$ to $P_{dst} \in \mathbf{P}$, $P_{src} \neq P_{dst}$.

Overhead – $o_s(e, P_{src})$ is the computational overhead, i.e. the execution time, incurred by processor P_{src} for preparing the transfer of the communication associated with edge e and $o_r(e, P_{dst})$ is the overhead incurred by processor P_{dst} after receiving e .

Involvement – $i_s(e, L_1)$ is the computational involvement, i.e. execution time, incurred by processor P_{src} during the transfer of edge e and $i_r(e, L_l)$ is the computational involvement incurred by P_{dst} during the transfer of e .

This is the general definition of overhead and involvement for heterogeneous arbitrary systems. Therefore, the overhead is made a function of the processor and the involvement a function of the utilised communication link. As discussed in the previous section, the overhead depends largely on the employed communication environment and is thereby normally unaffected by the utilised communication resources. In contrast, the involvement depends to a large extent on the capabilities of the utilised communication resources. Hence, the processor involvement is characterised by the outgoing or incoming link utilised for a communication.

With the distinction between the sending (o_s, i_s) and the receiving side (o_r, i_r) of communication, all three types of communication – third party, one sided, two sided – can be precisely represented. The corresponding functions are simply defined accordingly, e.g. $i_s(e, L) = i_r(e, L) = 0$ for involvement-free third party communication.

Note, for homogeneous systems or systems that have homogeneous parts, the definition of overhead and involvement can be significantly simplified. For example, in a systems with a homogenous network, the involvement functions can be defined globally, i.e. $i_{s,r}(e, L) = i_{s,r}(e)$.

A. Scheduling edges on the processors

Incorporating overhead and involvement into contention aware task scheduling is accomplished by extending edge scheduling so that edges are not only scheduled on the links but also on the processors.




Fig. 3. Scheduling under the involvement-contention model: edges are also scheduled on the processors; (S – switch or other processor)

So, the **start time** of an edge $e \in \mathbf{E}$ on a processor $P \in \mathbf{P}$ is denoted by $t_s(e, P)$. Let $R = \langle L_1, L_2, \dots, L_l \rangle$ be the route for the communication of edge $e \in \mathbf{E}$ from $P_{src} \in \mathbf{P}$ to $P_{dst} \in \mathbf{P}$, $P_{src} \neq P_{dst}$. The **finish time** of e on P_{src} is

$$t_f(e, P_{src}) = t_s(e, P_{src}) + o_s(e, P_{src}) + i_s(e, L_1) \quad (8)$$

and on P_{dst} it is

$$t_f(e, P_{dst}) = t_s(e, P_{dst}) + o_s(e, P_{dst}) + i_s(e, L_l). \quad (9)$$

Figure 3 illustrates scheduling under the involvement-contention model. The execution time of an edge on a processor is the sum of the overhead and the involvement (see (8) and (9)).

As an edge scheduled on a processor represents computation, precisely the computation necessary for the communication of the edge, its scheduling must fulfil the processor constraint as formulated in Condition 1. For a meaningful and feasible schedule, the scheduling of the edges on the processors must obey this condition:

Condition 4 (Causality in Involvement Scheduling): Let $R = \langle L_1, L_2, \dots, L_l \rangle$ be the route for the communication of edge $e_{ij} \in \mathbf{E}$, $n_i, n_j \in \mathbf{V}$, from $P_{src} \in \mathbf{P}$ to $P_{dst} \in \mathbf{P}$, $P_{src} \neq P_{dst}$.

To assure the node strictness of n_i

$$t_s(e_{ij}, P_{src}) \geq t_f(n_i, P_{src}). \quad (10)$$

Edge e_{ij} can be transferred on the first link L_1 only after the overhead completed on the source

processor P_{src} :

$$t_s(e_{ij}, L_1) \geq t_s(e_{ij}, P_{src}) + o_s(e_{ij}, P_{src}) \quad (11)$$

To assure the causality of the direct involvement on the destination processor P_{dst}

$$t_s(e_{ij}, P_{dst}) \geq t_f(e_{ij}, L_l) - i_r(e_{ij}, L_l). \quad (12)$$

The three inequalities can be observed in effect in Figure 3. Edge e_{AB} starts on P_1 after the origin node A finishes (inequality (10)). On the first link L_1 , e_{AB} starts after the overhead finishes on P_1 (inequality (11)), at which time the involvement of P_1 begins. And last, e_{AB} starts on P_2 so that the involvement finishes at the same time as e_{ij} on L_2 (inequality (12)).

Condition 4 leaves scheduling algorithms some freedom for the node-edge order on the processor. An edge leaving a processor does not have to start immediately after the processor (10) – other edges or nodes can be executed before. The same principle holds on the destination processor (12).

1) Scheduling: Just as for the scheduling of the nodes on the processors (Section II-D, Condition 3), a scheduling condition is formulated for the correct choice of an idle time interval into which an edge can be scheduled on a processor, with either the end or insertion scheduling technique.

Condition 5 (Edge Scheduling Condition on a Processor): Let $G = (\mathbf{V}, \mathbf{E}, w, c)$ be a DAG, $M = ((\mathbf{P}, \mathbf{L}), \omega, o, i)$ a parallel system and $R = \langle L_1, L_2, \dots, L_l \rangle$ the route for the communication of edge $e_{ij} \in \mathbf{E}$, $n_i, n_j \in \mathbf{V}$, from $P_{src} \in \mathbf{P}$ to $P_{dst} \in \mathbf{P}$, $P_{src} \neq P_{dst}$. Let $[A, B]$, $A, B \in [0, \infty]$, be an idle time interval on P , $P \in \{P_{src}, P_{dst}\}$. Edge e_{ij} can be scheduled on P within $[A, B]$ if

$$B - A \geq \begin{cases} o_s(e_{ij}, P_{src}) + i_s(e_{ij}, L_1) & \text{if } P = P_{src} \\ o_r(e_{ij}, P_{dst}) + i_r(e_{ij}, L_l) & \text{if } P = P_{dst} \end{cases} \quad (13)$$

$$B \geq \begin{cases} t_f(n_i, P_{src}) + o_s(e_{ij}, P_{src}) + i_s(e_{ij}, L_1) & \text{if } P = P_{src} \\ t_f(e_{ij}, L_l) + o_r(e_{ij}, P_{dst}) & \text{if } P = P_{dst} \end{cases} \quad (14)$$

This condition ensures that the time interval $[A, B]$ adheres to the inequalities (10) and (12) of the causality Condition 4. For an idle time interval $[A, B]$ adhering to Condition 5, the start

time of e_{ij} on P_{src} and P_{dst} is

$$t_s(e_{ij}, P) = \begin{cases} \max\{A, t_f(n_i)\} & \text{if } P = P_{src} \\ \max\{A, t_f(e_{ij}, L_l) - i_r(e_{ij}, L_l)\} & \text{if } P = P_{dst} \end{cases} \quad (15)$$

So, the edge is scheduled as early as possible within the limits of the interval. Of course, the choice of the interval should follow the same policy on the links and on the processors, i.e. either end or insertion scheduling should be used.

Note, the size of the involvement does not depend on the contention in the network. The assumption is that if the processor has to wait to send or receive a communication due to contention, this wait is passive or non-blocking, which means it can perform other operations in the meantime.

B. Node and edge scheduling

Few alterations are imposed by the new model on the edge scheduling on the links and on the scheduling of the nodes.

Edge scheduling on links – The Causality Condition 4 of Involvement Scheduling only imposes a constraint on the scheduling of an edge on the first link of its communication route. This is formulated in (11), which requires that edge e_{ij} must not start on the first link L_1 of its route until after the overhead has finished on the source processor P_{src} , $t_s(e_{ij}, L_1) \geq t_s(e_{ij}, P_{src}) + o_s(e_{ij}, P_{src})$. In comparison, under the contention model edge e_{ij} can start on the first link L_1 immediately after its origin node n_i has finished, $t_s(e_{ij}, L_1) \geq t_f(n_i)$.

Note, the rest of the edge scheduling procedure is completely unaffected by the scheduling of the edges on the processors and remains unchanged.

Node scheduling – To adapt the scheduling of the nodes to the new model, it is only necessary to redefine the finish time of the edge.

Definition 7 (Edge Finish Time): Let $G = (\mathbf{V}, \mathbf{E}, w, c)$ be a DAG and $M = ((\mathbf{P}, \mathbf{L}), \omega, o, i)$ a parallel system. The finish time of $e_{ij} \in \mathbf{E}$, $n_i, n_j \in \mathbf{V}$ is

$$t_f(e_{ij}) = \begin{cases} t_f(n_i) & \text{if } \text{proc}(n_i) = \text{proc}(n_j) \\ t_f(e_{ij}, \text{proc}(n_j)) & \text{otherwise} \end{cases} \quad (16)$$

C. NP-completeness

Scheduling under the involvement-contention model remains an NP-hard problem. This is easy to see, as the involvement model is based on the contention model, which is NP-hard. It is proven with the straightforward reduction from the NP-complete decision problem C-SCHED(G, M_{TG}) associated with the scheduling under the contention model [30].

Theorem 1 (NP-Completeness – Involvement-Contention Model): Let $G = (\mathbf{V}, \mathbf{E}, w, c)$ be a DAG and $M = ((\mathbf{P}, \mathbf{L}), \omega, o, i)$ a parallel system. The decision problem IC-SCHED(G, M) associated with the scheduling problem is as follows: Is there a schedule \mathcal{S} for G on M with length $sl(\mathcal{S}) \leq T, T \in \mathbb{Q}^+$? IC-SCHED(G, M) is NP-complete.

The formal proof is given in [30].

V. SCHEDULING ALGORITHMS

In contrast to scheduling on the links, the scheduling of the edges on the processors, which seems at first sight a simple extension, has a strong impact on the operating mode of scheduling algorithms. Essentially, the problem is that at the time a free node n is scheduled, it is generally unknown to where its successor nodes will be scheduled. It is not even known if the corresponding outgoing communications will be local or remote. Thus, no decision can be taken whether to schedule n 's leaving edges on its processor or not. Later, at the time a successor is scheduled, the period of time directly after node n might have been occupied with other nodes. Hence, there is no space left for the scheduling of the corresponding edge. Scheduling under the LogP model faces the same problem with the scheduling of o for each communication [17].

Two approaches to scheduling under the involvement-contention model are proposed, namely i) direct scheduling and ii) scheduling based on a given processor allocation. Both approaches are investigated in the next sections, followed by the proposal of two scheduling heuristics, one for each approach.

A. Direct scheduling

Direct scheduling means that the processor allocation and the start/finish time attribution of a node are done in one single step. The application of the scheduling method from contention

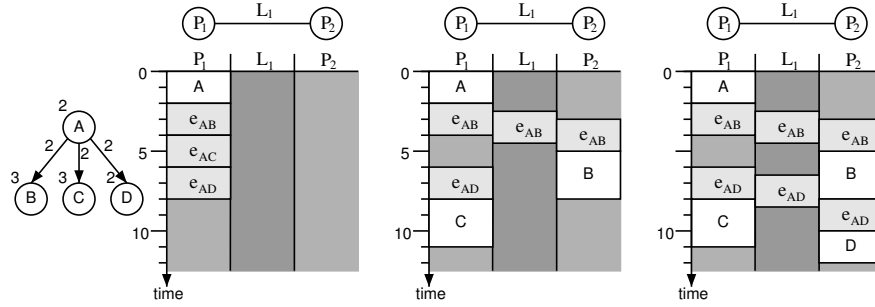


Fig. 4. Direct scheduling: edges are provisionally scheduled on source processor; left and centre chart illustrate intermediate schedules, right chart the final schedule of the depicted DAG; $o_{s,r}(e, P) = 0.5$, $i_{s,r}(e, L) = 0.75 \cdot (t_f(e, L) - t_s(e, L))$

scheduling is inadequate under the new model, since the decision whether a communication is remote or local is made too late. Consequently, it is necessary to investigate how edges can be scheduled earlier.

The most viable solution is to reserve an appropriate time interval after a node for the postponed scheduling of the leaving edges. This must be done in a worst case fashion, which means the interval must be large enough to accommodate all leaving edges. A straightforward manner to do so, is to schedule all leaving edges on the source processor, directly after the origin node. The scheduling of the edges on the links and the destination processors can take place when the destination node is scheduled. If the destination node is scheduled on the same processor as the origin node, the corresponding edge, which was provisionally scheduled with the origin node, is simply removed from that processor.

Figure 4 shows an example. First, A is scheduled on P_1 , together with its three leaving edges (shown in the left chart), hence the worst case that B , C and D are going to be scheduled on P_2 is assumed. Indeed, node B is scheduled on P_2 , which includes the preceding scheduling of e_{AB} on the link and on P_2 . Next, C is scheduled on P_1 , hence e_{AC} is removed from P_2 (centre chart). Finally, D is scheduled on P_2 with the respective scheduling of e_{AD} on the link and P_2 (right chart).

On heterogeneous systems, provisional scheduling of an edge on its source processor must take the fact that the involvement depends on the first link of the utilised route into consideration. Again, as the route is unknown during the scheduling, the worst case must be assumed. So, the

provisional finish time of edge $e_{ij} \in \mathbf{E}$ on its source processor $P = \text{proc}(n_i)$, $P \in \mathbf{P}$, is

$$t_f(e_{ij}, P) = t_s(e_{ij}, P) + o_s(e_{ij}, P) + i_{s,max}(e_{ij}, P), \quad (17)$$

where $i_{s,max}(e_{ij}, P) = \max_{L \in \mathbf{L}: L \text{ leaving } P} \{i_s(e_{ij}, L)\}$. When the destination node n_j is scheduled, the finish time must be reduced, if applicable, to the correct value.

With the reservation of a time interval for the outgoing edges on the processor, the remaining scheduling can be performed as under the contention model. The disadvantage of this approach is the gaps left behind by removed edges, which can make a schedule less efficient. This shortcoming can be relieved using the insertion technique in scheduling. Inserting a node or an edge into a gap, is very likely to separate edges from their origin or destination node, but this is allowed by the causality Condition 4. Further, the gaps can be eliminated, i.e. the schedule is compacted, by repeating the scheduling. In this rescheduling the nodes and their edges must be scheduled in the exact same order as in the first run, because as under the contention model the scheduling order of the edges matters. Then, with the processor allocation taken from the completed schedule, the provisional scheduling of edges becomes obsolete and the gaps are avoided.

B. Scheduling with given processor allocation

The second approach to involvement scheduling assumes a given processor allocation, also referred to as a mapping, of the nodes to the processors as a prerequisite to the scheduling procedure. Hence, for every node it is known to where its successors will be scheduled. With a given mapping, the schedule is constructed with a list scheduling heuristic (Section V-C), where the processor choice is simply a lookup from the given mapping.

The scheduling of an edge e_{ij} can be divided into three parts: scheduling on the source processor P_{src} , on the links of the route R , and on the destination processor P_{dst} . On the source processor, an edge must be scheduled together with its origin node n_i , as the foregoing considerations in the context of the direct scheduling showed. The scheduling on the links and on the destination processor can take place with either the origin node n_i or the destination node n_j . Hence, there are three alternatives for the scheduling of an edge e_{ij} .

1. – The first alternative is identical to the approach of direct scheduling, where the edge e_{ij} is scheduled on the links of R and on the destination processor P_{dst} , when its destination node n_j is scheduled.

2. – In this alternative, e_{ij} is not only scheduled on P_{src} , but also on the links, when n_i is scheduled. This way, the edges are scheduled on the links of R in the scheduling order of their origin nodes, while in the first alternative the edges are scheduled on the links in the order of their destination nodes. There is no clear advantage of the first over the second alternative or vice versa. Which one is better, i.e. more realistic, depends on the way the communication is realised in the target parallel system, whether it is initiated by the receiving (first alternative) or by the sending side (second alternative).

3. – Edge e_{ij} is completely scheduled (i.e. on P_{src} , R and P_{dst}), when its origin node n_i is scheduled. This alternative is likely to produce the best scheduling alignment of the edge on the source processor, the links and the destination processor, as the scheduling is done at once. Unfortunately, it has the disadvantage that the scheduling of the edges on their destination processors might prevent the efficient scheduling of the nodes. The early scheduling of the edges on their destination processors rapidly increases their finish times, leaving large idle gaps. Therefore, the conjoint scheduling of an edge on all resources is only sensible with the insertion scheduling technique.

The processor allocation can originate from any heuristic or can be extracted from a given schedule. For example, a schedule produced under the classic or contention model might serve as the input. In Section V-D a genetic algorithm is proposed for the determination of the processor allocation.

C. List scheduling

In this section list scheduling [1] is adapted for the involvement-contention model, using the direct scheduling approach (Section V-A). In the simple, but common, variant of list scheduling (Algorithm 1) the nodes are ordered according to a priority in the first part of the algorithm.

The schedule order of the nodes is important for the schedule length and many different priority schemes have been proposed, e.g. [1], [32], [37]. A common and usually good priority is the

Algorithm 1 List scheduling

- 1: \triangleright *I. Part:*
 - 2: Sort nodes $n \in \mathbf{V}$ into list L , according to priority scheme and precedence constraints.
 - 3: \triangleright *2. Part:*
 - 4: **for** each $n \in L$ **do**
 - 5: Find processor $P \in \mathbf{P}$ that allows earliest finish time of n .
 - 6: Schedule n on P .
-

Algorithm 2 Scheduling of node n_j on processor P in involvement-contention model

- 1: **for** each $n_i \in \text{pred}(n_j)$ **do**
 - 2: **if** $\text{proc}(n_i) = P$ **then**
 - 3: remove e_{ij} from P
 - 4: **for** each $n_i \in \text{pred}(n_j)$ in a definite order **do**
 - 5: **if** $\text{proc}(n_i) \neq P$ **then**
 - 6: determine route $R = \langle L_1, L_2, \dots, L_l \rangle$ from $\text{proc}(n_i)$ to P
 - 7: correct $t_f(e_{ij}, \text{proc}(n_i))$
 - 8: schedule e_{ij} on R
 - 9: schedule e_{ij} on P
 - 10: schedule n_j on P
 - 11: **for** each $n_k \in \text{succ}(n_j)$ in a definite order **do** \triangleright *reserve space for leaving edges*
 - 12: schedule e_{jk} on P with worst case finish time
-

node's **bottom level** bl , which is the length of the longest path leaving the node. Recursively bl is defined as

$$bl(n_i) = w(n_i) + \max_{n_j \in \text{succ}(n_i)} \{c(e_{ij}) + bl(n_j)\} \quad (18)$$

Under the involvement-contention model, and in accordance with the direct scheduling approach, the scheduling of a node (line 6 in Algorithm 1) is performed as described in Algorithm 2.

As under the contention model, finding the processor that allows the earliest finish time of a node involves the tentative scheduling on every processor (line 5 of Algorithm 1), including the incoming edges on the links and the destination processor. In this way it is possible to consider the communication contention and the processor involvement in the scheduling decisions.

To determine the start time (i.e. the “schedule” lines in Algorithm 2) of a node or edge on a processor or link, both the end technique and insertion technique (Section II-D) can be employed with list scheduling. Under the involvement-contention model, the insertion technique is more

indicated, since the removing of provisionally scheduled edges leaves gaps (Section V-A), which should be filled by other nodes or edges.

Compared to contention aware list scheduling, the time complexity under the involvement-contention model does not increase. The complexity of the second part of list scheduling is $O(\mathbf{P}(\mathbf{V} + \mathbf{E}O(\text{routing})))$ (end technique) or $O(\mathbf{V}^2 + \mathbf{P}\mathbf{E}^2O(\text{routing}))$ (insertion technique) [30]. $O(\text{routing})$ is the complexity of determining the communication route and scheduling an edge on this route (lines 6 and 8 of Algorithm 2). In regular networks determining the route is often $O(\mathbf{P})$ or even $O(1)$ (e.g. central switch or fully connected network). If the route can be determined in $O(1)$ time (calculated or looked up), then $O(\text{routing})$ is just the complexity of the length of the route (the edge must be scheduled on each link of the route). For comparison, the complexity expressions under the classic model are $O(\mathbf{P}(\mathbf{V} + \mathbf{E}))$ (end technique) and $O(\mathbf{V}^2 + \mathbf{P}\mathbf{E})$ (insertion technique).

D. A genetic scheduling algorithm

In this section a genetic algorithm (GA) based scheduling heuristic is proposed that follows the approach of using a given processor allocation (Section V-B). Genetic algorithms have been successfully employed for the scheduling problem, e.g. [2], [7], [13], [21], [27]. A GA is a search algorithm which is inspired by the principles of evolution and natural genetics [11]. It begins with an initial population (a set of chromosomes) and then operates through a simple cycle of stages: evaluation of the chromosomes of the population, stochastic selection of the best chromosomes and reproduction to create a new population, using crossover and mutation operators. This process is repeated and terminates after a specified number of generations or when a time limit is exceeded.

The idea of the proposed heuristic, referred to as Genetic Involvement-Contention Scheduling (GICS), is that the genetic algorithm searches for an efficient processor allocation, while the actual scheduling is performed with a list-scheduling based heuristic, where the decision for a processor is taken from the calculated processor allocation. Below, the components of GICS are described.

Chromosome – The chromosome, encoded as an array of size $|V|$, represents the processor allocation [7]. Hence, the value of array element i is the index of the processor assigned to node $n_i \in V$.

Initial population – Randomly created chromosomes make the largest part of the initial pool of chromosomes. In order to avoid the creation of schedules that are slower than the sequential execution, sequential processor allocations, i.e. all nodes are allocated to one single processor, are also included. The pool is completed by one allocation extracted from a schedule produced with a list scheduling heuristic [21]. The chosen list scheduling uses finish time minimisation, insertion scheduling and the nodes are in bottom-level order (Section V-C).

Evaluation – At each iteration, a schedule is produced for every chromosome, i.e. processor allocation, with a heuristic based on the list scheduling structure (the decision for the processor is taken from the chromosome). The heuristic employs the third alternative for the scheduling of the edges (Section V-B), where an edge is completely scheduled with its origin node. This kind of edge scheduling requires the insertion technique. To reduce the running time of the evaluation, the node order is determined only once at the beginning of the algorithm, namely according to their bottom-levels. The **fitness value** of each chromosome is the length of its schedule.

Selection – Selection is performed by randomly picking two chromosomes out of the current pool from which the fitter one goes into the next generation, i.e. so called tournament selection. This process is repeated until the pool of the next generation has the specified population size. It is also guaranteed that the fittest chromosome, the one with the shortest schedule, survives.

Crossover – Two new chromosomes, nc_1 , nc_2 are generated from two randomly selected chromosomes, c_1 , c_2 , by swapping the array values of a randomly determined element interval. Let the interval range from i to j . Outside this interval, i.e. $[1, i - 1]$ and $[j + 1, |V|]$, the elements of nc_1 have the values of c_1 and inside those of c_2 . For nc_2 it is converse.

Mutation – The mutation operator creates a new chromosome by copying a randomly picked chromosome and swapping the values of two randomly determined array elements. This operation helps to balance the load, as the number of nodes per processor is not changed. For DAGs with a high CCR, i.e. with a lot of communication, load balancing is not very promising. The CCR is

the communication to computation ratio defined as the sum of all communication costs divided by the sum of all computation costs, $CCR = \frac{\sum_{e \in \mathbf{E}} c(e)}{\sum_{n \in \mathbf{V}} w(n)}$. Therefore, if the CCR is high (currently > 5), the swapping is substituted with an increase or decrease (randomly decided) by one of an array element.

The complexity of GICS is $O(\text{population_size} \times \text{number_of_generations} \times (\mathbf{V} + \mathbf{EO}(\text{routing})))$, where the last term is the complexity of list scheduling without the processor choice (hence without factor \mathbf{P} , see Section V-C).

The edge scheduling of GICS is different to the edge scheduling of list scheduling as proposed in the previous section. Consequently, the comparison of the two heuristics in the following experiments will also provide some insights regarding the quality of the different edge scheduling approaches.

Finally, the control parameters which delivered good results and were used in the experiments are: population size = 100, number of generations = 40 (60 for DAGs with high CCR, see mutation operator), probability of a chromosome to participate in a crossover = 0.4, probability of mutation = 0.15.

VI. EXPERIMENTAL RESULTS

For the evaluation of the new involvement-contention model and the two proposed heuristics, the experimental methodology proposed in [33] is employed. A large set of graphs is generated and scheduled by algorithms under the different models to several target systems. Code is generated from the produced schedules, using C and MPI, and executed on the real parallel systems. The execution times of these codes directly show which algorithms and models produce the best schedules.

The evaluation is divided into two parts: accuracy and execution time. In the following, only the most important results are presented, but more experiments and details can be found in [30].

A. Improved accuracy

In [31], the accuracy of the classic and the contention model are examined using the mentioned methodology. To evaluate the accuracy of the new involvement-contention model, the schedules

produced in those experiments under classic and the contention model are rescheduled, but now under the involvement-contention model. The two original heuristics are **Cl-LS(bl)** – list scheduling with bottom-level order – as a **classic** model heuristic, and **LS(dls)** – list scheduling with DLS’s node order [29] – as a **contention** model heuristic. This rescheduling allocates the nodes to the same processors in the same local order as in the original schedule. Consequently, the code generated for the schedule under the involvement-contention model would be identical to the one generated for the original schedule, under the contention model. Hence the execution time of that code would be identical to the execution time of the original code, which was already obtained experimentally. By comparing this execution time with the prediction under the involvement-contention model, the new model’s accuracy is determined.

Three target systems were employed in the experiments: a cluster (BOBCAT) of 16 PCs, modelled as a switched network; a shared memory multiprocessor system Sun E3500 with 8 processors, modelled as a bus network; a massively parallel system Cray T3E-900 with a total of 344 processors, modelled as a fully connected network.

Due to the lack of a profound insight into the target systems’ communication mechanisms and their MPI implementations, 100% involvement is assumed, i.e. the source and destination processors are involved during the entire communication time on the first and last link, respectively: $i_s(e, L_1) = t_f(e, L_1) - t_s(e, L_1)$ and $i_r(e, L_l) = t_f(e, L_l) - t_s(e, L_l)$. The overhead is intuitively set to an experimentally measured setup time: $o_s(e, P) = o_r(e, P) = \text{setup_time}$. While it is clear that this definition of the overhead and the involvement is probably not an accurate description of the target systems’ communication behaviour, it is very simple. The idea is to demonstrate that accuracy and efficiency of scheduling can be improved even with a rough but simple estimate of the overhead and involvement functions.

1) *Results*: Figure 5 visualises the average accuracy deviations $\Delta acc(\mathcal{S})$ with

$$\Delta acc(\mathcal{S}) = \begin{cases} acc(\mathcal{S}) - 1 & \text{if } acc(\mathcal{S}) \geq 1 \\ 1/acc(\mathcal{S}) - 1 & \text{if } acc(\mathcal{S}) < 1 \end{cases}, \quad (19)$$

where $acc(\mathcal{S})$ is the ratio of the execution time of the code produced for schedule \mathcal{S} to its schedule length $sl(\mathcal{S})$.

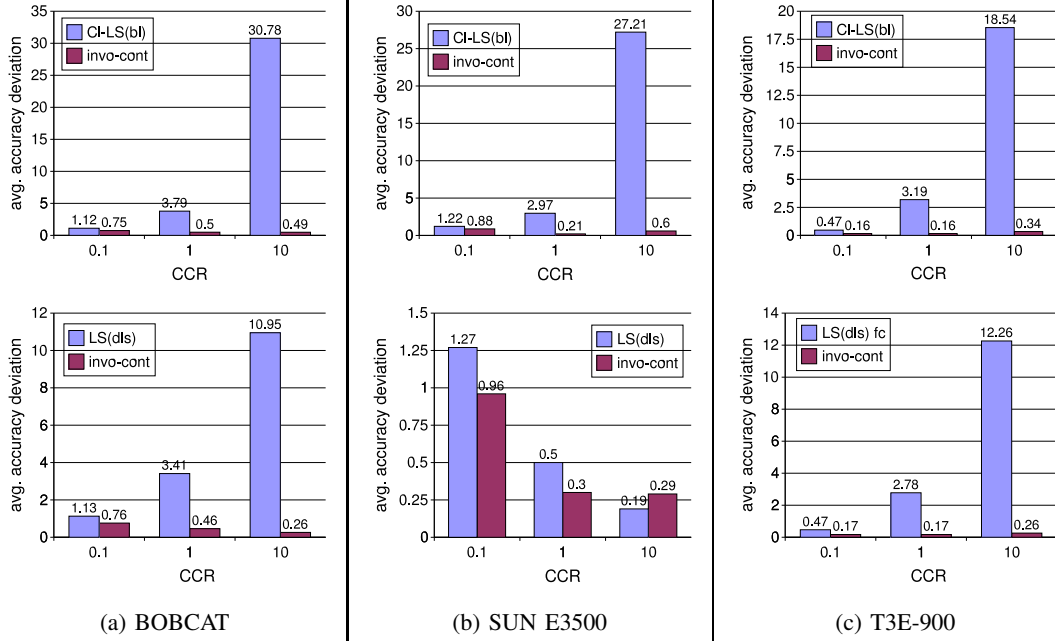


Fig. 5. Avg. accuracy deviation of CI-LS(bl) (classic model) [top] and LS(dls) (contention model) [bottom] compared with rescheduling under the involvement-contention model (invo-cont)

It is immediately apparent from Figure 5 that the accuracy profoundly improved under the new involvement-contention model. While this improvement is already considerable for low communication ($CCR = 0.1$), it is more significant for medium ($CCR = 1$) and especially so for high communication ($CCR = 10$). The length of a schedule is now in a region, where it can be seriously considered an estimation of the real execution time.

Generally, the difference between the reschedules from the classic model (CI-LS(bl)) and the contention model (LS(dls)) on the same system is small. This is desirable, as the origin of the original schedule should not affect the accuracy under the involvement-contention model.

The scheduling accuracy under the involvement-contention model is still not perfect, especially for low communication ($CCR = 0.1$). A possible explanation might be the blocking communication mechanisms used in MPI implementations [16], which does not match the assumption of non-blocking communication made in the involvement contention model. Further, the employed overhead and involvement functions are very rough estimates, a better approximation of these functions might improve the accuracy. In any case, it is in the nature of any model that there is a difference between prediction and reality. Under this perspective, the results obtained for the T3E are very satisfying, which is probably due to the fact that the T3E-900, being a massively

structure	# nodes	avg. in/out-degree $\bar{\delta}$	weights
random in-tree	175, 406, 233 ($CCR = 0.1, 1, 10$)	≈ 1.0	random
random out-tree	176, 648, 319 ($CCR = 0.1, 1, 10$)	≈ 1.0	random
random multiple fork-join	500	1.94	random
Gaussian elimination (e.g. [8])	495 ($n = 31$)	1.88	regular
stencil (e.g. [23])	400 (20×20)	2.76	unit
pipeline (e.g. [20])	500 (25×20)	1.86	unit
random	600	≈ 2.0	random

TABLE I
WORKLOAD DAGS

parallel system specifically designed for parallel processing, is the most predictable among the target systems.

B. Improved execution time

To determine if the new model also produces schedules, which lead to shorter execution times, new experiments were performed using the mentioned methodology of [33]. As the intention here is to compare scheduling models and not algorithms, the same heuristic is employed for each model analysed.

Seven different graph types constitute the DAG workload, listed in Table I. The average in-degree or the average out-degree (they are identical) of the nodes ($\bar{\delta}$), i.e. the average number of incoming or outgoing edges, characterises the density of a graph. Three instances of each of the seven DAGs are generated, one instance with a CCR of 0.1, one with a CCR of 1 and one with a CCR of 10. To achieve the CCR values in the DAGs with regular weights, the node and edge weights are scaled accordingly.

These graphs are scheduled by a list scheduling heuristic with the insertion technique, where the nodes are ordered according to their bottom levels. This algorithm is applied under the classic, the contention and the involvement-contention model, denoted by “classic”, “cont”, “invo-cont”, respectively. Using the same algorithm for each model allows the impact of the model on the quality of the produced schedules to be analysed and compared, without the influence of different scheduling techniques. Further, in order to evaluate scheduling with a given processor allocation under the involvement-contention model, the genetic algorithm “GICS” (Section V-D) is utilised.




Fig. 6. BOBCAT: avg. speedups under the three models (speedups are normalised to values under classic model)

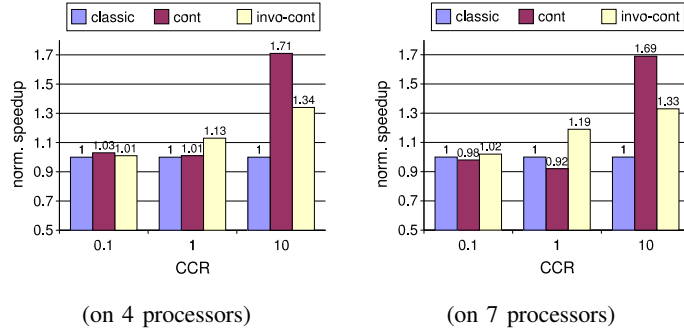


Fig. 7. Sun E3500: avg. normalised speedups under the three models (speedups are normalised to values under classic model)

Code is generated for the obtained schedules and executed on two different target systems: Sun E3500 and BOBCAT (Cray T3E-900 was not used, as it simply was taken out of service at the Edinburgh Parallel Computing Centre before the realisation of the experiments). Both systems were modelled as in the previous experiments of Section VI-A.

1) Results: Figures 6 and 7 visualise the average normalised speedups of the three different models on different configurations of the two target systems. Since GICS is based on a different algorithm (not list scheduling), GICS is treated separately below. The speedup is defined as the ratio of the sequential time $seq(G) = \sum_{n \in \mathbf{V}} w(n)$ (local communication has zero costs) to the execution time of the code produced for schedule \mathcal{S} .

Low communication ($CCR = 0.1$) – For a CCR of 0.1, communication is of low significance compared to the computation costs. Nevertheless, the involvement-contention model noticeably reduces the execution times – compared to the classic model – on BOBCAT with 16 processors (average speedup is improved by 9%). With more processors, communication becomes more important, because it is probable that more communications are performed remotely.

Medium communication ($CCR = 1$) – The situation changes for DAGs with medium communication. Schedules produced under the involvement-contention model have significantly shorter execution times, with speedup improvements of up to 82% (Gauss elimination). Again, the improvement increases with the number of the utilised processors.

High communication ($CCR = 10$) – The highest reduction in the execution time is apparent in the results of the involvement-contention for DAGs with high communication. Unfortunately, this reduction is irrelevant: the absolute speedup of the schedules of all models is below 1. Hence, a parallelisation is meaningless and the sequential execution is indicated.

GICS – Figures 8 and 9 summarise the average normalised speedups of GICS compared to invo-cont. Not regarding high communication, the execution times of the schedules by GICS are almost identical to those by invo-cont. This is surprising, as at least a slight improvement was expected, since one of GICS’s initial chromosomes is the processor allocation of the invo-cont schedule (Section V-D). Yet, invo-cont and GICS differ in the way they schedule edges: invo-cont employs the direct scheduling approach (V-A), which corresponds to the first alternative of the edge scheduling approaches (Section V-B). In contrast, GICS uses the third alternative, which apparently leads to no improvement by GICS. Hence, the first alternative should be preferred over the third in scheduling under the involvement-contention model.

GICS’s good results for high communication are due to the fact that sequential processor allocations are part of its initial population. Indeed, nearly all schedules by GICS for $CCR = 10$ are (almost) sequential, with an absolute speedup of 1. Thus, neither GICS is able to efficiently parallelise the DAGs with $CCR = 10$. However, the accuracy of the involvement-contention model allows GICS to know that.

2) *Discussion:* In the experiments conducted, the involvement-contention model clearly demonstrated its ability to produce schedules with significantly reduced execution times.

Despite the very good results, the efficiency improvement lags behind the accuracy improvement demonstrated in the previous subsection. A possible explanation lies in the employed heuristic. List scheduling is a greedy algorithm, which tries to reduce the finish time of each node to be scheduled. Thereby it does not consider the leaving communications of a node, which




Fig. 8. BOBCAT: avg. normalised speedups of cont-invo and GICS (speedups are normalised to values under classic model)




Fig. 9. Sun E3500: avg. normalised speedups of cont-invo and GICS (speedups are normalised to values under classic model)

may impede an early start of following nodes. The high importance of communication under the involvement-contention model seems to demand the research of more sophisticated algorithms in order to exploit the full potential of this new model.

VII. CONCLUSIONS

This paper investigated the involvement of the processor in communication, its impact on task scheduling and how to make task scheduling aware of it. First, the different types of processor involvement and their characteristics were investigated. A new system model was proposed, which extends the scheduling of the edges on the links to their scheduling on the processors. This technique can reflect the three types of processor involvement and the distinction between overhead and direct involvement. Scheduling the edges on the processors has an impact on the operating techniques of scheduling heuristics. This challenge was investigated and two solutions were proposed: provisional scheduling and using a given processor allocation. Based on these solutions, two scheduling algorithms were proposed for the new involvement-contention model: an adapted list scheduling and a genetic algorithm (GICS). Extensive experiments demonstrated

that the involvement-contention model significantly improves the accuracy and the execution time of the produced schedules. The improved accuracy now allows for a useful estimation of the execution time. In order to achieve further improvements, research into algorithms that are better at exploiting the new model is needed.

REFERENCES

- [1] T. L. Adam, K. M. Chandy, and J. R. Dickson. A comparison of list schedules for parallel processing systems. *Communications of the ACM*, 17:685–689, 1974.
- [2] I. Ahmad and M. K. Dhodhi. Multiprocessor scheduling in a genetic paradigm. *Parallel Computing*, 22:395–406, 1996.
- [3] I. Ahmad and Y.-K. Kwok. On exploiting task duplication in parallel program scheduling. *IEEE Transactions on Parallel and Distributed Systems*, 9(8):872–892, September 1998.
- [4] I. Ahmad, Y.-K. Kwok, and M.-Y. Wu. Analysis, evaluation, and comparison of algorithms for scheduling task graphs on parallel processors. In G.-J. Li, D. F. Hsu, S. Horiguchi, and B. Maggs, editors, *Second International Symposium on Parallel Architectures, Algorithms, and Networks, 1996*, pages 207 – 213, June 1996.
- [5] A. Alexandrov, M. Ionescu, K. E. Schauser, and C. Scheimann. LogGP: Incorporating long messages into the LogP-model - one step closer towards a realistic model for parallel computation. In *7th Annual Symposium on Parallel Algorithms and Architectures*, pages 95–105. ACM press, 1995.
- [6] O. Beaumont, V. Boudet, and Y. Robert. A realistic model and an efficient heuristic for scheduling with heterogeneous processors. In *HCW'2002, the 11th Heterogeneous Computing Workshop*. IEEE Computer Society Press, 2002.
- [7] M. S. T. Benten and S. M. Sait. Genetic scheduling of task graphs. *Int. Journal of Electronics*, 77(4):401–415, 1994.
- [8] M. Cosnard and D. Trystram. *Parallel Algorithms and Architectures*. Int. Thomson Computer Press, London, UK, 1995.
- [9] D. E. Culler, R. M. Karp, D. A. Patterson, A. Sahay, K. E. Schauser, E. Santos, R. Subramonian, and T. von Eicken. LogP: towards a realistic model of parallel computation. *ACM SIGPLAN Notices, Proc. Symposium on Principles and practice of parallel programming*, 28(7):1–12, July 1993.
- [10] D. E. Culler and J. P. Singh. *Parallel Computer Architecture*. Morgan Kaufmann Publishers, 1999.
- [11] L. Davis. *Handbook of Genetic Algorithms*. Van Nostrand-Reinhold, New York, USA, 1991.
- [12] A. Gerasoulis and T. Yang. A comparison of clustering heuristics for scheduling DAGs on multiprocessors. *Journal of Parallel and Distributed Computing*, 16(4):276–291, December 1992.
- [13] E. S. H. Hou, N. Ansari, and H. Ren. Genetic algorithm for multiprocessor scheduling. *IEEE Transactions on Parallel and Distributed Systems*, 5(2):113–120, February 1994.
- [14] J. J. Hwang, Y. C. Chow, F. D. Anger, and C. Y. Lee. Scheduling precedence graphs in systems with interprocessor communication times. *SIAM Journal of Computing*, 18(2):244–257, April 1989.
- [15] IBM. *SP Switch2 Technology and Architecture*, March 2001. http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/sp_switch2.pdf.
- [16] J. White III and S. Bova. Where's the overlap? An analysis of popular MPI implementations. In *Proc. of MPIDC 1999*, 1999.
- [17] T. Kalinowski, I. Kort, and D. Trystram. List scheduling of general task graphs under LogP. *Parallel Computing*, 26:1109–1128, 2000.
- [18] B. Kruatrachue. *Static Task Scheduling and Grain Packing in Parallel Processing Systems*. PhD thesis, Oregon State University, USA, 1987.
- [19] B. Kruatrachue and T. Lewis. Grain size determination for parallel processing. *IEEE Software*, 5(1):23–32, January 1988.
- [20] S. Y. Kung. *VLSI Array Processors*. Information and System Sciences Series. Prentice Hall, 1988.
- [21] Y.-K. Kwok and I. Ahmad. Efficient scheduling of arbitrary task graphs to multiprocessors using a parallel genetic algorithm. *Journal of Parallel and Distributed Computing*, 47(1):58–77, November 1997.
- [22] Y.-K. Kwok and I. Ahmad. Benchmarking the task graph scheduling algorithms. In *Proc. of Int. Par. Processing Symposium/Symposium on Par. and Distributed Processing (IPPS/SPDP-98)*, pages 531–537, Orlando, Florida, USA, April 1998.

- [23] W. Löwe, W. Zimmermann, S. Dickert, and J. Eisenbiegler. Source code and task graphs in program optimization. In *High Performance Computing and Networking*, volume 2110 of *Lecture Notes in Computer Science*, pages 273–282. Springer-Verlag, 2001.
- [24] B. S. Macey and A. Y. Zomaya. A performance evaluation of CP list scheduling heuristics for communication intensive task graphs. In *Parallel Processing Symposium, 1998. Proc. of IPPS/SPDP 1998*, pages 538–541, 1998.
- [25] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard*, June 1995. <http://www.mpi-forum.org/docs/docs.html>.
- [26] P. Rebreyend, F. E. Sandnes, and G. M. Megson. Static multiprocessor task graph scheduling in the genetic paradigm: A comparison of genotype representations. Research Report 98-25, Ecole Normale Supérieure de Lyon, Laboratoire de Informatique du Parallelisme, Lyon, France, 1998.
- [27] F. E. Sandnes and G. M. Megson. An evolutionary approach to static taskgraph scheduling with task duplication for minimised interprocessor traffic. In *Proc. Int. Conf. on Parallel and Distributed Computing, Applications and Technologies (PDCAT 2001)*, pages 101–108, Taipei, Taiwan, July 2001. Tamkang University Press.
- [28] V. Sarkar. *Partitioning and Scheduling Parallel Programs for Execution on Multiprocessors*. MIT Press, 1989.
- [29] G. C. Sih and E. A. Lee. A compile-time scheduling heuristic for interconnection-constrained heterogeneous processor architectures. *IEEE Transactions on Parallel and Distributed Systems*, 4(2):175–186, February 1993.
- [30] O. Sinnen. *Accurate Task Scheduling for Parallel Systems*. PhD thesis, Instituto Superior Técnico, Technical University of Lisbon, Portugal, April 2003.
- [31] O. Sinnen and L. Sousa. Experimental evaluation of task scheduling accuracy: Implications for the scheduling model. *IEICE Transactions on Information and Systems*, E86-D(9):1620–1627, September 2003.
- [32] O. Sinnen and L. Sousa. List scheduling: Extension for contention awareness and evaluation of node priorities for heterogeneous cluster architectures. *Parallel Computing*, 30(1):81–101, January 2004.
- [33] O. Sinnen and L. Sousa. On task scheduling accuracy: Evaluation methodology and results. *The Journal of Supercomputing*, 27(2):177–194, February 2004.
- [34] O. Sinnen and L. Sousa. Communication contention in task scheduling. *IEEE Transactions on Parallel and Distributed Systems*, 16(6):503–515, June 2005.
- [35] J. D. Ullman. NP-complete scheduling problems. *Journal of Computing System Science*, 10:384–393, 1975.
- [36] T. von Eicken, D. E. Culler, S. C. Goldstein, and K. E. Schauer. Active Messages: A mechanism for integrated communication and computation. In *Proc. 19th Annual Int. Symposium on Computer Architecture*, pages 256–266, Gold Coast, Australia, May 1992.
- [37] M. Y. Wu and D. D. Gajski. Hypertool: A programming aid for message-passing systems. *IEEE Transactions on Parallel and Distributed Systems*, 1(3):330–343, July 1990.
- [38] T. Yang and A. Gerasoulis. PYRROS: static scheduling and code generation for message passing multiprocessors. In *Proc. of 6th ACM International Conference on Supercomputing*, pages 428–437, Washington D.C, August 1992.