



Hybrid multi-core CPU and GPU-based B&B approaches for the blocking job shop scheduling problem

Adel Dabah^{a,b,*}, Ahcène Bendjoudi^a, Abdelhakim AitZai^b, Didier El-Baz^c,
Nadia Nouali Taboudjemat^a

^a CERIST Research Center, Algiers, Algeria

^b University of Sciences and Technology Houari Boumediene (USTHB) Algiers, Algeria

^c LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

HIGHLIGHTS

- New parallel B&B schemes for heterogeneous computing nodes.
- Increasing the GPU occupation over time using Nvidia Multi-Processes Service (MPS).
- Hybrid High-level/Low-level parallel B&B schemes that exploit both the CPU-cores and the GPU simultaneously.

ARTICLE INFO

Article history:

Received 21 April 2017

Received in revised form 8 February 2018

Accepted 17 February 2018

Available online 26 February 2018

Keywords:

Job shop scheduling

Parallel B&B algorithm

Hybrid multi-core CPU and GPU B&B

Nvidia MPS

ABSTRACT

The Branch and Bound algorithm (B&B) is a well known method for solving optimally Combinatorial Optimization Problems. This method is based on intelligent enumeration of all feasible solutions which reduce considerably the search space. Nevertheless, it remains inefficient when using the sequential approach to deal with large problem instances due to its huge resolutions time. However, the execution time can be reduced considerably by using parallel computing architectures. With the huge evolution of the multi-cores CPUs and GPUs, it is quite hard to design schemes that efficiently exploit the different hardware architectures simultaneously. As a result, most of the existing works focus on exploiting one hardware architecture at a time. In this paper, we propose five parallel approaches to accelerate the B&B execution time using Multi and Many-core systems. Our goal is to solve optimally the Blocking Job Shop Scheduling problem (BJSS) which is one of the hardest scheduling problem. The first proposed approach is a *multi-search* parallelization based on master/worker paradigm, exploiting the multi-Core CPU-processors. The second and the third approaches represent a GPU-based parallelization schemes having different level of parallelism and GPU occupancy. The fourth and fifth approaches represent a hybridization between the Multi-core approach and the GPU-based parallelization approaches. The goal of this hybridization is to benefit from the power of both the CPU-cores and the GPU at the same time. This hybridization is based on concurrent kernels execution provided by Nvidia Multi process Service (MPS) that allows multiple host processes (Master and workers) to use simultaneously the GPU to launch their kernels in order to accelerate the bounding of one or several nodes at a time. Experiments using the well known Taillard instances confirm the efficiency of our proposals and show a relative speedup of 160x as compared to an optimized sequential B&B algorithm.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

The job shop scheduling problem (JSSP) consists in scheduling a set of n jobs on a set of m machines. Each job has its own sequence of crossing on machines. The execution of a job on a machine is called operation and each one uses the machine for an uninterrupted processing time. The classical JSSP assumes an

infinite storage space between machines which is not realistic. The Blocking Job Shop Scheduling problem (BJSS) is a version of the classical JSSP with no storage space between machines, where a job has to wait on the current machine until the next one becomes available. Our goal is to minimize the completion time of all jobs (*Makespan*). The classical JSSP is known to be NP-hard in the strong sense and its search space is equal to $(n!)^m$ [16]. The BJSS problem which is an extension of the JSSP, appears to be even more difficult to solve [17]. This problem has several application areas such as manufacturing systems with no storage space, trains scheduling, hospital resource scheduling, etc.

* Corresponding author at: CERIST Research Center, Algiers, Algeria.

E-mail addresses: adabah@cerist.dz, adabah@usthb.dz (A. Dabah),

abendjoudi@cerist.dz (A. Bendjoudi), h.aitzai@usthb.dz (A. AitZai), elbaz@laas.fr

(D. El-Baz), nnouali@cerist.dz (N.N. Taboudjemat).

The B&B algorithm is a well-known method to solve optimally the combinatorial optimization problems. It is based on an intelligent enumeration of all feasible solutions, which reduce considerably the search space. Nevertheless, its sequential version requires a huge execution time to solve small instances and it remains inefficient when dealing with large or real-world instances. Therefore, using parallel computing architectures such as the Multi-core CPUs and the GPUs is unavoidable to improve the running time of this method.

Nowadays, most of the processors are based on Multi-core architectures. They are composed of two or more independent processing units (cores) that can run multiple instructions at the same time. In addition to the power provided by multi-core processors, Graphics Processing Units (GPUs) have emerged as a new popular support for massively parallel computing. The GPUs are many-core co-processor devices that provide a hierarchy of memories having different sizes and access latencies and providing a highly multi-threaded environment where the threads are scheduled and executed as warps (a group of threads) using the SIMT model [10].

With the huge evolution of the Multi-core CPUs and GPUs, it is quite hard to design schemes that efficiently exploit different hardware architectures simultaneously. As a result, most of the proposed parallel B&B algorithms in the literature [8,19,3,11,6] exploit only the CPU-core or only the GPU which may lead to an under-utilization of these resources and a loss of a significant computing power. The major contribution of this paper is the new parallelization schemes that exploit and combine different parallelization levels of the B&B algorithm using Multi and Many-core systems simultaneously. The proposed schemes are based on the Nvidia Multi-Process Service (MPS) that allows us to increase the GPU occupation over time. Therefore, achieving a high relative speedup for large instances.

The first approach (*Multi-core B&B*) is a tree-based parallelization, exploiting the Multi-core CPU-processors available in all recent PCs. The proposed approach is based on Master/Worker paradigm where the workers independently explore branches (search tree nodes) sent by the master. The performance of this approach depends on the number of used CPU-cores.

The second and the third approaches are a GPU-based parallel B&B schemes. The second approach: *Parallel Evaluation of the Bound (PEB)* is a node-based parallelization exploiting the idea that the evaluation (bounding) of each node of the search tree can be calculated in parallel on the GPU using several GPU-threads. The proposed scheme here is very important to reduce the B&B execution time since the bounding phase consumes the majority of the execution time of the algorithm. Therefore, at each iteration of this scheme, one node is evaluated in parallel on the GPU by several threads organized in one GPU block. Experiments validate our idea and show the benefit of accelerating the bounding process on GPU by achieving a speedup up to 18x. The third approach (*Parallel Evaluation of Several Bounds (PESB)*) represents a generalization of the second approach (PEB) obtained by sending at each iteration a pool of nodes to the GPU for evaluation instead of one node at a time in the PEB approach. The number of nodes sent varies according to the size of instances. This approach allows us to achieve a relative speedup of 66x.

The drawback of the previous approaches is the under-utilization of either CPU or GPU resources which represents a waste of significant computing power. To increase the GPU occupation and benefit from both the multi-core CPU and the GPU at the same time, we propose a hybridization between the CPU and GPU-based approaches. This Hybridization is based on the concurrent kernel's execution provided by Nvidia MPS. *i.e.* Multiple host processes (master and workers) can execute simultaneously their kernels on the GPU. The fourth approach, *H-PEB* represents a hybridization between the Multi-core approach and the PEB GPU-based approach

Table 1
BJSS instance with two jobs and three machines.

Job	Sequence	Processing times
J1	M1, M2, M3	5, 3, 8
J2	M2, M1, M3	8, 2, 7

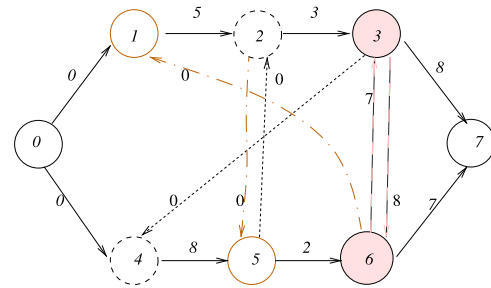


Fig. 1. Alternative graph for the BJSS instance in Table 1.

which mean that several host processes (Master and workers) can use the GPU simultaneously to accelerate the bounding of one node on the GPU according to the PEB scheme. Increasing the GPU occupation by this approach allows us to reach a relative speedup of 93x as compared to an optimized sequential B&B. In the same way, the fifth parallel approach *H-PESB* represents a hybridization between the Multi-core approach and the PESB GPU-based approach. Therefore, all host processes (Master and workers) can use simultaneously the GPU to evaluate several nodes on the GPU instead of one node in the H-PEB approach. The obtained results for this last approach show a relative speedup of 160x as compared to an optimized sequential B&B approach.

The remainder of this paper is organized as follows: Section 2 describes the blocking job shop scheduling problem and the alternative graph model. Section 3 contains a brief description of the sequential B&B algorithm and its components, in addition to a detailed related work. Section 4 presents the proposed parallelization approaches of the B&B algorithm. Section 5 discusses computational results. Finally, conclusions and perspectives are presented in Section 6.

2. Blocking job shop scheduling problem

In this section, we describe the problem treated in this paper namely The Blocking Job Shop Scheduling Problem (BJSS). Before describing formally the BJSS problem in Sections 2.2 and 2.3, we begin this section by presenting a small BJSS example of two jobs and three machines.

2.1. Example

Table 1 represents a BJSS instance with two products (jobs) and three machines. The first product (J1) has 5 min processing time on machine M1, 3 min on M2 and 8 min on machine M3. The second product (J2) has 8 min processing time on machine M2, 2 min on M1 and 7 min on machine M3.

Fig. 1 represents an alternative graph of the BJSS instance in Table 1. This graph has three alternative pairs, two between blocking operations and one between ideal operations. Both operations 2 and 4 need the same machine M2 and since M2 cannot process both operations at the same time, we associate them with an alternative pair. Since operations 2 and 4 are blocking operations the first alternative arc (3, 4) represents the choice where operation 2 must be finished before the beginning of operation 4. Its mate, arc (2, 5) represents the choice whereby operation 4 must be finished before the beginning of operation 2. We use the same process to

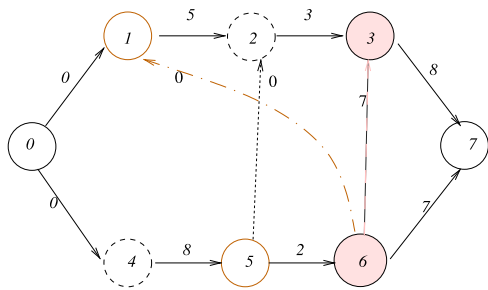


Fig. 2. Schedule for BJSS Problem in Table 1 with $C_{max} = 26$.

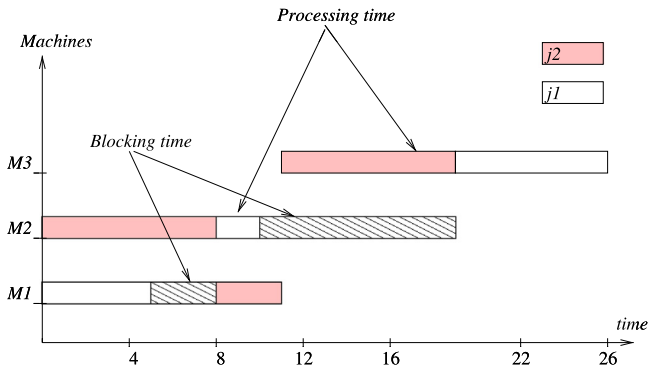


Fig. 3. Gantt chart of the schedule in Fig. 2.

generate the alternative pair $((2, 5), (6, 1))$ between operations 1 and 5. The alternative pair between operations 3 and 6 is $((3, 6), (6, 3))$ because both operations 3 and 6 are ideal.

Fig. 2 shows a alternative graph representation of a feasible schedule (solution) for the BJSS instance in Table 1. This schedule is obtained by choosing one arc from each alternative pair in Fig. 1. The *Makespan* ($C_{max} = 26$) of this schedule is the longest path in the obtained graph.

The Gantt chart in Fig. 3 represents both the processing and blocking times of the solution in Fig. 2. For example, after the end of its processing time, the job J1 blocks the machine M1 until machine M2 becomes available for processing J1.

2.2. Problem formulation

The classical JSSP can be defined by a set J of n jobs (J_1, \dots, J_n) to be processed on a set M of m machines (M_1, \dots, M_m) . Each machine can process at most one job at a given time. The execution of a job on a machine is called operation. We note by O the set of all operations $(o_1, \dots, o_{n \times m})$. Each operation o_i needs to use a machine $M(i)$ for an uninterrupted duration called processing time p_i . Each job has its own sequence of crossing on machines, which creates precedence constraints between consecutive operations of the same job. A solution (schedule) for this problem consists to assign a starting and finishing times t_i and c_i for each operation o_i ($i = 1, \dots, n \times m$); while satisfying all constraints. Our goal is to minimize the *Makespan* (C_{max}), which is the ending time of a schedule. The JSSP assumes an unlimited intermediate buffer capacity between consecutive operations of a job which is impossible in real manufacturing. The BJSS is a version of the classical JSSP with no intermediate buffers, where a job has to wait on the current machine until the next machine becomes available for processing. This problem can be presented as an alternative graph model introduced by Mascis et al. [17]. This model represents a generalization of the disjunctive graph of Roy and Sussman [28]. It can be defined as a graph $G = (N, F, A)$. N represents a set of nodes

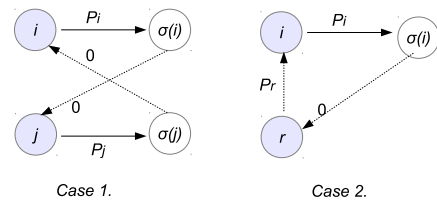


Fig. 4. Alternative pairs between blocking and ideal operations.

(operations) with two additional dummy nodes (start and finish) modeling the start and the finishing of the schedule. F represents a set of fixed arcs imposed by precedence constraints between consecutive operations of the same job and f_{qp} is the length of arc $(q, p) \in F$. Finally, A is a set of alternative pairs $((i, j), (h, k))$ representing the processing order for concurrent operations on the same machine and a_{ij} is the length of alternative arc (i, j) . Each alternative pair contains two alternative arcs and each one of them expresses the fact that one operation must be completed on the target machine before starting the processing of the other operation. Moreover, Section 2.3.1 describes the alternative arcs generation. A selection S_1 is a set of alternative arcs obtained from A by choosing at most one alternative arc from each alternative pair. We denote by $G(S_1) = (N, F \cup S_1)$ the graph representation of the selection S_1 . We note that a selection S_1 is feasible if there is no positive length cycle in $G(S_1)$ and the evaluation (*Makespan*) of this selection (S_1) is the longest path in $G(S_1)$. We say that S_1 is a complete selection if exactly one arc is chosen from each pair, therefore $|A| = |S_1|$. We define a schedule (solution of the problem) as a complete feasible selection. Finally, given a feasible selection S_1 , let $l(i, j)$ be the length of the longest path from operation i to j in $G(S_1)$. We call the last operation of each job (example o_r) an ideal operation because the machine becomes immediately available after the end of its processing time p_r . In all other cases, o_i is called a blocking operation. In this case, we denote by $\sigma(i)$ the operation immediately following o_i in the same job.

2.3. Forming the graph of all possibilities (search graph)

The search graph represents the graph of all possibilities which is used by the B&B algorithm as search tree root. This graph is obtained by generating all alternative pairs, knowing that each alternative pair represents the processing order between every two concurrent operations. Therefore, if a machine has four concurrent operations, we need six alternative pairs to show all the possibilities on how these operations are executed on the machine. The alternative pair between every two concurrent operations is generated as follow.

2.3.1. Alternative pair generation

Let us consider two blocking operations o_i, o_j and one ideal operation o_r , where $M(i) = M(j) = M(r)$. Since the three operations cannot be executed at the same time, we associate them with pairs of alternative arcs.

Case 1: the alternative pair between operations o_i and o_j (Fig. 4): The first alternative arc $(\sigma(i), j)$ having length 0 represents the situation where o_i is processed before o_j . Since o_i is a blocking operation, $M(i)$ can begin the processing of o_j only after the starting time of $\sigma(i)$ (when o_i leaves $M(i)$). The same method is followed for the other alternative arc $(\sigma(j), i)$ since o_j is a blocking operation.

Case 2: the alternative pair between operations o_i and o_r (Fig. 4): It is the same process for the alternative arc $(\sigma(i), r)$ since o_i is a blocking operation. The other alternative arc depends on the fact that o_r is an ideal operation therefore, we add the alternative arc (r, i) with length p_r .

Table 2

The description of the symbols used in our B & B algorithm.

Symbol	Description
UB, LB	Upper Bound, Lower Bound.
LIST	A set of nodes (sub-problems).
s^*	The optimal solution.
R_i	The i th successor of node R .
$LB(R_i)$	Lower bound of node R_i .

3. The branch and bound algorithm for BJSS

The B&B algorithms make an intelligent enumeration of all feasible solutions. They are mainly characterized by two operators: branching and bounding. The branching is a recursive process, which consists in replacing the search space of a given problem by a set of smaller sub-problems. The bounding operator contains two bounds: the Lower Bound (LB) and the Upper Bound (UB). The LB represents an estimation of the lowest evaluation of all feasible solutions in the considered sub-problem, while the UB represents the upper limit of the evaluation of each search tree node. Any solution to the problem can be considered as initial value for the UB which is updated as soon as a new better solution is found by the B&B algorithm. The B&B algorithm is based essentially on the bounds to make an intelligent enumeration of the search space. The elimination operator uses the bounds (LB and UB) to eliminate the sub-problems that cannot lead to improve the current best solution found by the algorithm.

The most effective B&B algorithms, for the JSSP, are based on the disjunctive graph model [4]. Our sequential B&B is based on the adaptation of this approach to the blocking case (alternative graph) [17]. Our method which consists in fixing an order between every two concurrent operations, which leads to fix the corresponding alternative pair (from A) and a set of fixed arcs represents a selection.

Algorithm 1 and Table 2 describe the used B&B algorithm.

Algorithm 1 Pseudo-code of the sequential B&B algorithm

```

LIST ← {original problem};
UB ← ∞;
while LIST ≠ ∅ do
  R ← LIST (Choose a Node R from LIST);
  Generate successors  $R_i$  from  $R \mid (i = 1, \dots, n)$ ;
  for Each successor  $R_i$  do
    if  $LB(R_i) < UB$  then
      if  $R_i$  represents one solution then
        UB =  $LB(R_i)$ ;
         $s^*$  = solution in  $R_i$ ;
      else
        LIST = LIST  $\cup$   $R_i$ ;
      end if
    end if
  end for
end while
return  $s^*$ 

```

In the following, we describe the different operators of the used B&B algorithm for the BJSS problem.

3.1. Branching

The B&B algorithm can be represented by a search tree. The tree is rooted by the original problem, *i.e.* no alternative pairs are fixed ($|S_0| = 0$). A search tree node R is characterized by (S_R, A_R) and represented by the graph $G(S_R) = (N, F \cup S_R)$. S_R denotes the set of fixed alternative arcs and A_R represents a set of unselected alternative

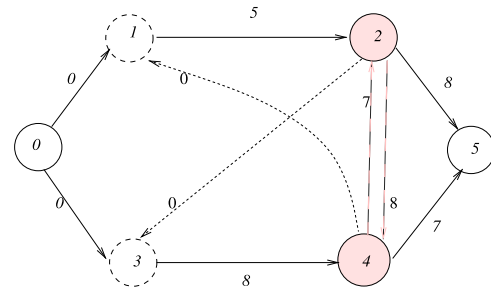


Fig. 5. Alternative graph for BJSS instance with two jobs and two machines.

pairs in this node. The branching creates two immediate successors (R_1, R_2) of R by fixing an alternative pair $((i, j), (h, k)) \in A_R$ that has a direct impact on the longest path in the graph. The node R_1 (resp. R_2) is characterized by $S_{R_1} = S_R \cup (i, j)$ (resp. $S_{R_2} = S_R \cup (h, k)$) and $A_{R_i} = A_R - \{((i, j), (h, k))\}$. Each successor represents the sub-search space related to the fixed alternative arc. After this, each successor is handled recursively in the same way until we find a complete selection, or we eliminate the sub-problem (prune the tree) if the lower bound value is bigger than the upper bound. Finally, the exploration strategy represents the way the search tree is explored. Several exploration strategies exist in the literature such as best first exploration, worst first exploration, and breadth-first exploration. The exploration strategy used in our case consists to choose the after a branching process the node with the biggest Makespan (worst first). Unlike the best-first exploration strategy, the worst-first exploration strategy has more chance to reach leaf nodes (solutions) therefore, more chance to improve the UB and eliminate a large number of branches.

Fig. 5 represents the alternative graph of a BJSS instance with 2 jobs and 2 machines. The Figure shows also the existence of two alternative pairs: the first pair $((2, 3), (4, 1))$ is between operations 1 and 3 and the second pair $((2, 4), (4, 2))$ is between operations 2 and 4.

Fig. 6 represents the search tree of the BJSS instance in Fig. 5 which contains two alternative pairs. At each level, one pair is fixed which generate two sub-problems. This process is repeated until a complete feasible selection is obtained or an infeasibility is detected.

Fig. 7 represents the alternative graph of the optimal solution obtained from the search tree in Fig. 6.

3.2. Evaluation (bounding)

Any solution of the problem can be considered as an initial value for the Upper Bound (UB) which is updated as soon as a new better solution is found. In our case, the $UB = +\infty$.

After the branching operation, the bounding process consists to calculate the Lower Bound (LB) for each sub-problem. The LB used in our case is the one used by Carlier *et al.* [5] to solve optimally the classical JSSP, it is based on the *one machine scheduling problem*. Therefore, we calculate the lower bound for each machine l using concurrent operations on it:

$$LB_l = \text{Min}r_i + \sum p_i + \text{Min}q_i / M(i) = l, l = \{1, \dots, m\}.$$

The LB for the sub-problem is equal to the maximum value of LB_l .

$$LB = \text{Max}\{LB_l\}, l = \{1, \dots, m\}.$$

To do a link with the alternative graph model, each search tree node represents an alternative graph. This lower bound cannot be calculated without the new head and tail values for each operation affected by the branching process. ($H_i = l(0, i)$, $T_i = l(i, n * m) \mid (i = 1, \dots, n * m)$). The process of adjusting the *head* and *tail*

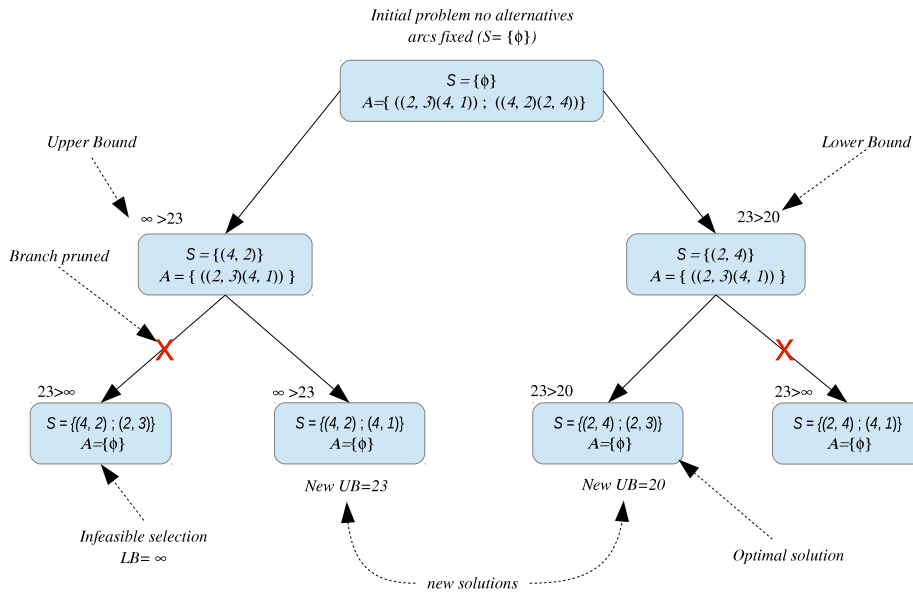


Fig. 6. Search tree for the BJSS instance in Fig. 5.

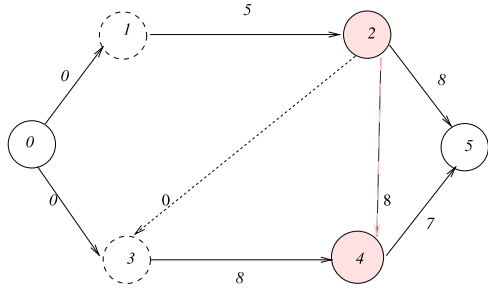


Fig. 7. Alternative graph of the optimal solution ($Cmax = 20$).

values is very expensive and consumes 70% of global execution time. This process is done sequentially for all operations affected by the change made and can be repeated several times for the same operation if there are multiple paths leading to this operation.

3.2.1. Immediate selection

The immediate selection represents several techniques allowing to accelerate the B&B algorithm by reducing the number of branching necessary to obtain the optimal solution. This process is done sequentially and costs 20% of the global processing time since there is a large number of alternative pairs (99,000 for big instances). This process uses also the head and tail values computed in the bounding process. Given a search tree node (sub-problem) R with a feasible selection S_R and a set of unselected pairs A_R . For each unselected pair $((i, j), (h, k)) \in A_R$, if $l(0, h) + a_{hk} + l(k, n) \geq UB$ then $S_R = S_R \cup (i, j)$. This rule expresses the fact that adding the arc (k, h) (resp. (i, j)) to S_R will produce a sub-problem with a lower bound greater than the upper bound. Consequently the arc (i, j) (resp. (h, k)) is added to S_R . If both alternative arcs (i, j) and (h, k) do not satisfy the condition of $LB < UB$, we prune the branch by eliminating node R , since this latter cannot contain a solution that can improve the best solution found by the B&B algorithm.

The complexity of the evaluation process depends on the number of operations $(n \times m)$ in the treated instance. Therefore, the evaluation time increases by increasing the size of the instances. The implementation of the evaluation process, requires six data structures. The matrix $Succ((n \times m) \times n)$ contains the successors of each operation, i.e. row i represents the successors of operation o_i . Similarly, the matrix $Pred((n \times m) \times n)$ contains the predecessors of

each operation. Vector S contains all selected and unselected pairs. Vector H (resp. T) contains the Head (resp. Tail) of each operation. The element $H[i] = l(0, i)$ represents the longest path from o_0 to o_i . Similarly, $T[i] = l(i, n \times m)$ the longest path from o_i to the last operation in the graph $o_{n \times m}$.

3.3. Related works

The BJSS problem can be solved using either exact or approximate methods. In the literature, several heuristics and meta-heuristics have been proposed to solve the BJSS problem as in [18,15,1,25,21,26,12]. In addition, heuristics as in [29,32] and studies on meta-heuristics based on multi-agent systems have proven their efficiency in solving other optimization problems, among them we can find [2,24,23].

In this paper, we focus on accelerating the B&B algorithm using multi-cores CPUs and GPUs simultaneously. Most of the B&B algorithms for the job shop problem are based on the *one machine scheduling problem* proposed by Carlier et al. [5].

Only a few authors tried to solve optimally the BJSS problem, the most effective one which is the base of our sequential version is proposed by Mascis et al. [17]. The authors formulate the problem by means of an alternative graph model which is a generalization of the disjunctive graph of Roy and Sussman [28]. Based on this model, they solve optimally the 10×10 benchmark instances of this problem.

Ait Zai et al. [1], proposed an original B&B method based on graph theory to solve the BJSS problem. The idea of its branching scheme relies on the implicit enumeration of all possible combinations on a given machine. The authors gave solutions for local instances only.

The B&B algorithms are not efficient when dealing with large problem instances, therefore computing accelerators like GPUs are required. Several authors have proposed to accelerate the B&B method using GPUs. This work represents an extension of our work in [13], in which we added three parallel approaches in order to fully occupy the GPU. This extension allowed us to increase the performance of our approaches by 74% to reach a speedup of 164x compared to 90x in [13]. In [14], we propose a preliminary version of the B&B algorithm dedicated to treat small instances (less than 50 jobs), the concept of the parallelization is kept for the two GPU approaches. However, to treat large instances we had to redesign the code and the GPU memory management to deal with such huge

generated data. The performance obtained in this paper boost the speedup to reach a 164X compared to 60x in [14].

Chakroun et al. [8,19], take the classical approach of sending nodes to be evaluated on GPU to solve the Flow shop scheduling problem since this step takes more than 98% of the global execution time. Therefore, each GPU thread supports the evaluation of a single node of the search tree. In [3,7] the authors extend the approach below to exploit the Multi-core CPU processors in the generalization of sub-problems that are going to be sent to the GPU for evaluation.

In [11], Alami et al. proposed a CPU–GPU based B&B applied to the knapsack problem. In this scheme, the branching and bounding can be done either on the CPU or the GPU according to the size of the search tree. This approach uses less CPU–GPU communication and better management of data-structures in the GPU memory.

In [6], Carneiro et al. applied the B&B to the traveling salesman problem where a pool of nodes is sent to the GPU for evaluation. Each GPU-thread applies the branching and bounding operators to a single node and builds its own local tree. The resulting nodes are moved back to the CPU where the promising nodes are inserted into the tree.

In [20], the authors proposed multi-core and many-core parallel B&B for big optimization problems. The authors proposed two parallel B&B implementations, the first one focuses on exploiting the traditional multi-core CPU processors, while the second one is dedicated for Intel Xeon Phi coprocessors considering both native and offload modes. The reported results show that the many-core approaches (native and offload) are twice faster as compared with the multi-core approach.

In [27], the authors proposed a parallel B&B algorithm exploiting the advantage of instance-specific computing on Field Programmable Gate Array (FPGA) which has proven to be highly efficient in term of area, energy consumption, and performance. In addition, the proposed parallelization is based on work stealing strategies to ensure dynamic load-balancing between the parallel threads. The author's approach was applied on the reconstruction of corrupted AES keys problem, the reported results show an overall speedup of 47x.

In [31], Trong and Bilel proposed parallel B&B for large-scale heterogeneous distributed platforms with several distributed CPUs and GPUs. The proposed approaches address the critical issue of how to map B&B workload with the different levels of parallelism exposed by the target compute platform. The reported results show the significant impact of the adaptive load balancing among the heterogeneous compute resources on the performance.

Most of the previously cited works focus on exploiting the GPU part and ignoring the available CPU-cores. Also, most authors use the *parallel evaluation of bounds* model in which each GPU thread supports the bounding of a single search tree node. This represents a lot of calculation and lot of resources for a single thread which may limit the performance. For this reason, we propose in this paper new parallelization schemes, that exploit and combine different parallelization levels to accelerate the B&B execution time using Nvidia MPS.

4. The proposed parallelization approaches for the B&B algorithm

The fact that each node of the B&B search-tree can be explored independently amplifies the parallelization of this algorithm. The only global information in the algorithm is the value of the upper bound.

The algorithm parallelization may depend on the architecture of the processing machine, synchronization, granularity of tasks and communication between different processes. There are several classes of the B&B parallelization. For more details, the reader may refer to [12]. In the following, we present our proposed parallel B&B approaches for the BJSS problem.

4.1. Multi-search parallel B&B on multi-core CPU

In this section, we describe the proposed parallel B&B algorithm exploiting the CPU-cores available in all recent computers. The proposed approach (see Fig. 8) is based on the master/worker paradigm. A work-pool represents a set of active sub-problems. There is a unique master work-pool owned and managed only by the master process which contains the sequential search tree and several local work-pools, empty initially, owned by the different workers. Both of the master and workers work-pools are managed in the same way. The exploration of the search-tree is done simultaneously by the master and workers since each one of them has its own instance of the B&B algorithm that uses a local work-pool. The results given by a worker can influence others. Therefore, our approach can be seen as a multi-search parallelization in which the goal is to accelerate the exploration of master search tree stored in the master pool. We call blocked worker, a worker with an empty work-pool waiting for a sub-problem to explore. The master process initializes the search by creating the root node, launches its own B&B algorithm which generates a set of active sub-problems stored in the master work-pool. After that, the master wakes up the blocked workers by sending them sub-problems from master work-pool. After that, each worker launches its own B&B algorithm. During the search, the local work-pools evolve continuously and when they become empty, the corresponding workers send a request to the master and wait for sub-problems. When the master receives a request, it satisfies the request if the global pool is not empty. When the global work-pool is empty, the master sends a request to all workers to send him back a sub-problem. Two states are then reserved for each process (blocked or active). Each time the global work-pool is empty, the master checks the state of all workers. If all the workers are blocked, then the master ends the calculation and frees the workers. The workers perform a *depth-first strategy* in order to reach quickly feasible solutions or eliminate the branches if the lower bound is greater than the upper bound. A worker which finds a better solution than the current best one broadcasts the new value to all workers via the master to ensure an efficient pruning process. An extended version of this approach that exploits the computing power provided by cluster-based architectures is presented in [11].

4.2. The proposed GPU-based B&B schemes

In this section, we describe our B&B GPU-based parallelization schemes. These schemes were firstly tested on small instances then adapted here for large Taillard instances. To handle such huge instances, we had to readjust the memory allocation and the distribution of the work-units over the threads of a block since the number of the graph operations for large instances is much bigger than the maximum number of threads in GPU block.

The GPU architectures are based on Single Instruction, Multiple Threads (SIMT) paradigm. According to this paradigm, the same program called *kernel* is executed simultaneously by a set of parallel threads with different data. The threads are organized according to a grid of thread-blocks hierarchy specified in the kernel call. The grid represents a set of thread-blocks. Threads of the same block can cooperate by using a private shared memory and barrier of synchronization. Threads can access multiple memory spaces: *constant memory* and *texture memory* are read-only cached memory accessible by all threads. The *global memory* is a read-write memory, also accessible by all threads. Unlike the global memory the *shared memory* is a cached memory accessible only by threads in the same block [10].

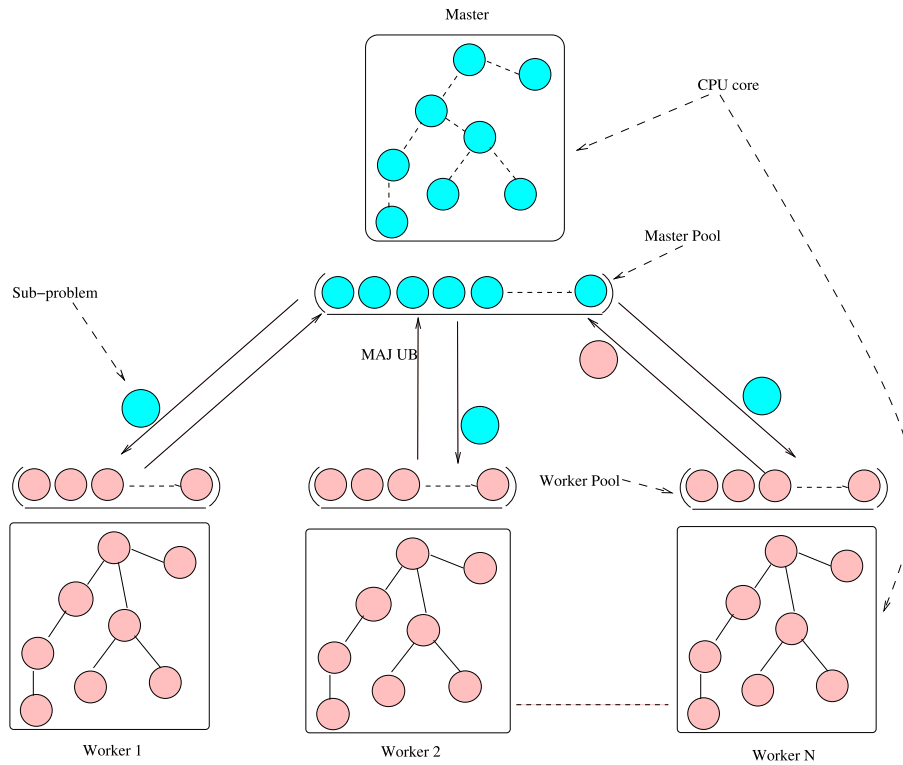


Fig. 8. Global architecture of the proposed master/worker parallelization.

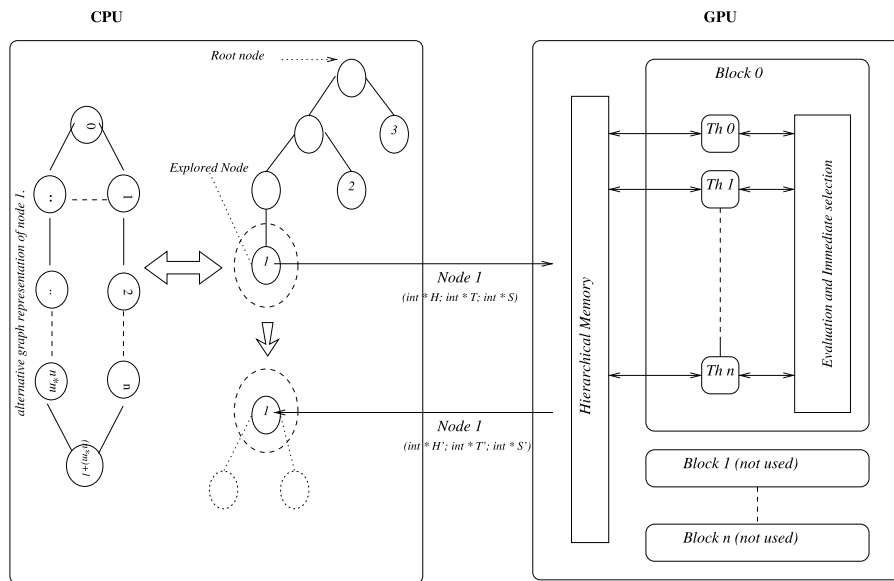


Fig. 9. GPU Evaluation of a single node.

4.2.1. Parallel evaluation of the bound (PEB)

We have seen in Section 3 that the evaluation process and the immediate selection consume together more than 90% of the global execution time, therefore, it is crucial to accelerate this phase in order to reduce the B&B execution time.

In the following, we present our proposed node-based parallelization scheme for the B&B algorithm exploiting the GPU-based architectures. The proposed scheme referred to as *Parallel Evaluation of the Bound (PEB)*, exploits the idea that the evaluation and immediate-selection for each node can be done in parallel using several threads.

As shown in Figs. 9 and 10, this approach uses the same design as the sequential B&B algorithm except that the evaluation (bounding) of each node is done in parallel on GPU. As already presented, each node of the search tree represents a graph of $n \times m$ operations. The bounding process consists in updating the head and tail values for each operation in the graph. The parallel PEB scheme is based on the idea that each GPU-thread updates the head and tail values for a single operation in the graph. This scheme exploits the fact that the updating process can be done independently for each operation. Therefore, the GPU block size is equal to the number of operations in the graph ($n \times m$). As shown in Fig. 9, at each iteration, only

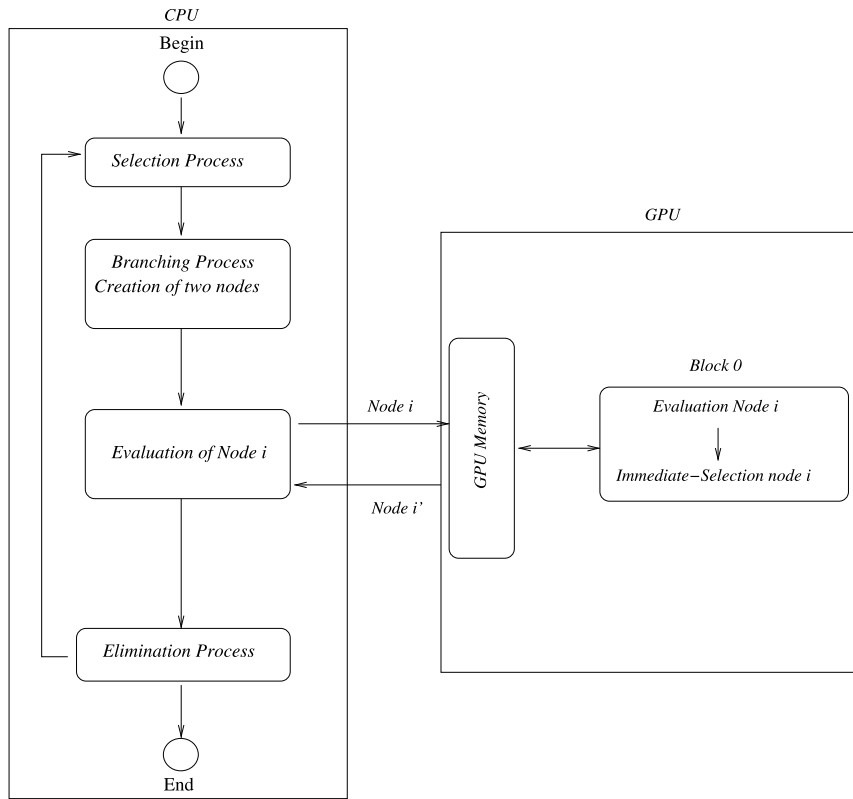


Fig. 10. Parallel evaluation of one bound (PEB scheme).

one node is sent to the GPU for evaluation and immediate selection using one thread-block. The bounding begins by copying the head and tail vectors to the shared memory, *i.e.* each thread copies the head and tail values relative to its *id*. After that, each thread updates the head and tail values ($H[i]$, $T[i]$) for the operation relative to its *id* (i) using respectively the head of its predecessors and the tail of its successors.

$$H[i] = \text{Max} \{H[r] + p_{ri}\} / r \in \text{Pred}[i].$$

$$T[i] = \text{Max} \{T[r] + p_{ir}\} / r \in \text{Succ}[i].$$

At the end of this computation, each thread waits for the other threads of the block using a barrier of synchronization to ensure the visibility of the new head and tail values to all threads of the block which is important to have a valid update process.

The work is repeated several times until there is no update of the head or the tail values for all threads or an infeasibility is detected.

After the end of the bounding process, each thread computes the immediate selection for a set of unselected alternative pairs using the new head and tail values. Finally, the new results are sent back to the CPU to be used by the branching and elimination process.

As can be seen in Fig. 9, a single block is used on the GPU to evaluate one node while the other blocks are idle. The weakness of this solution lies in the under-utilization of the GPU capacity and thus a waste of a significant computing power. To overcome this drawback, we propose a second level of parallelization.

4.2.2. Parallel evaluation of several bounds (PESB)

We propose in this section a second level of parallelization allowing to increase the occupation of the GPU. This level represents a generalization of the first scheme (PEB) called Parallel Evaluation of Several Bounds (PESB). The goal here is to increase the GPU

occupation by generalizing the idea of the first level (Bounding is faster) to exploit more efficiently the GPU computing power. Therefore, at each iteration of the B&B algorithm, a pool of nodes is sent to the GPU for the evaluation and immediate-selection instead of one. *i.e.* each GPU-block supports the evaluation of a single node. Then, the new results of nodes are sent back to the CPU to be used by the selection, branching, and elimination processes. As shown in Fig. 11, this approach has only one instance of the B&B algorithm running on the CPU. The nodes sent to the GPU for evaluation are chosen among the nodes recently added in the B&B work-pool which allows to avoid memory saturation by exploring the search tree in depth-first order.

As we have already seen, five data structures are used for the bounding of each node on the GPU. The vectors Head ($H [m \times n]$), Tail ($T [n \times m]$) and alternative pairs ($S [nbpair]$) are sent from the CPU to the GPU. Therefore, they are stored in the global memory of the GPU. The matrices *Succ* and *Pred* are also stored in the GPU global memory. These two matrices are calculated on the GPU using the vector *S* as an initialization for the bounding. To accelerate this initialization phase the calculation is divided across all the threads of the block.

The access to the global memory is much longer than the shared memory, but the latter is smaller compared to the global memory. The number of blocks that can run in parallel on each Streaming Multiprocessor depends on the amount of shared memory used by each block. Therefore, we use the shared memory only for the head and tail vectors in order to have a large number of blocks running in parallel and since there is a high number of accesses to these vectors.

The weakness of the previous approaches lies in the under-utilization of the GPU capacity in addition to the multi-core CPU processors and thus a waste of significant computing power. To overcome this drawback and increase the GPU occupation, we propose a hybridization between the Multi-core CPUs and the GPU-based approaches.

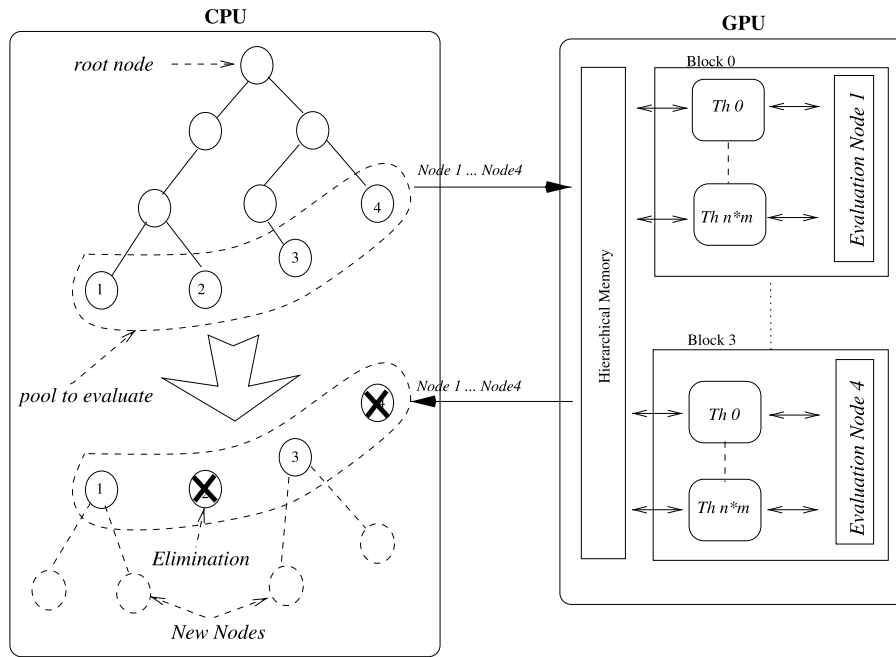


Fig. 11. GPU evaluation of several nodes.

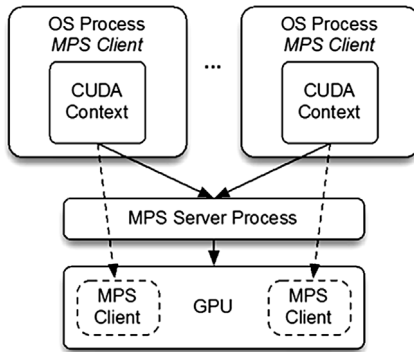


Fig. 12. MPS components.

```

mkdir /tmp/mps /tmp/mps-log
export CUDA_VISIBLE_DEVICES=0 # SELECT GPU 0.
export CUDA_MPS_PIPE_DIRECTORY=/tmp/mps # NAMED PIPES
export CUDA_MPS_LOG_DIRECTORY=/tmp/mps-log # LOGFILES
nvidia-cuda-mps-control -d # START THE DAEMON
unset CUDA_VISIBLE_DEVICES
mpirun -x CUDA_MPS_PIPE_DIRECTORY=/tmp/mps -np 35 ./BB
export CUDA_MPS_PIPE_DIRECTORY=/tmp/mps # SELECT THE LOCATION OF MPS DAEMON
echo quit | nvidia-cuda-mps-control # STOP MPS DAEMON
rm -rf /tmp/mps /tmp/mps-log
    
```

Fig. 13. Running MPI application using MPS.

4.3. Hybrid master-worker/GPU-based parallelization

In this section, we present two hybrid approaches based on the Nvidia Multi Processes Service (MPS). This tool is a client-server runtime implementation of the CUDA API used to increase the overall GPU utilization. Without MPS, only one host process can use the GPU at a given time. Therefore, it potentially may underutilize the GPU resources. To overcome this problem, Nvidia provides the MPS to enable multiple host processes, like MPI processes, to use the Hyper-Q capability on the Nvidia Kepler GPUs. Hyper-Q allows a single host process to execute multiple CUDA kernels concurrently on the same GPU. As we can see in Fig. 12, the MPS consists of several components: the Control Daemon Process is responsible for starting and stopping the MPS server, as well as coordinating connections between clients and the server [9].

The server process provides the connection between clients and the GPU which allows concurrency. Each process (client) has its own CUDA context for its GPU operations. When the MPS client connects to the control daemon, the later creates an MPS server if no server is active, then the client proceeds to connect with the server [9]. Note that all communications between MPS clients/server and MPS control daemon are done using a named

Pipe. Furthermore, Fig. 13. shows how to use the Multi Processes Service (MPS) to run MPI applications.

4.3.1. Hybrid parallel evaluation of the bound (H-PEB)

We propose in this section a hybridization of the first two approaches (Multi-core and GPU node based) to increase the GPU occupation and then improve the runtime. This version generalizes the idea of the PEB approach to exploit the advantages of both the Multi-core CPU processors and the GPU at the same time. The hybrid approach is based on concurrent kernels execution provided by Nvidia in devices with a compute capability of 2.x and higher. The maximum number of kernels that a device can execute concurrently varies between 16 and 32 according to device compute capability [10].

Therefore, in our proposed scheme (see Fig. 14), several CPU processes from the Multi-core approach (the Master or the workers) launch their kernels simultaneously on the GPU in order to accelerate the bounding of one node at a time according to the PEB scheme. Each host (MPI) process launches its own kernel in the default stream and the MPS server manages to execute the kernels in parallel by using different CUDA-Streams. The advantage of our hybrid approach based on concurrent kernel's execution is the occupation of the GPU over time. i.e. at each moment, our

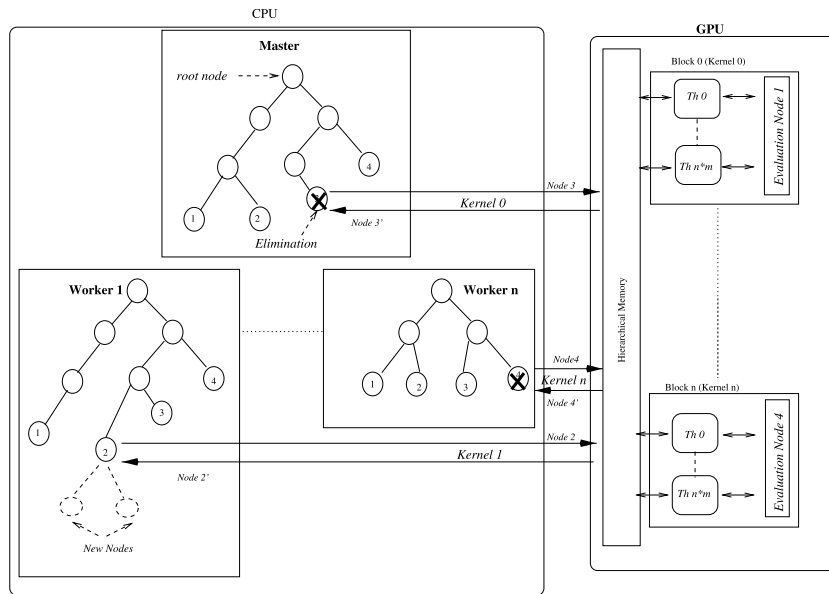


Fig. 14. Hybrid multi-core CPU/GPU approach.

hybrid approach can have simultaneously several workers executing instructions on the GPU while others perform data-transfer to/from the GPU and yet others apply the selection and elimination operators on the CPU. Using this approach, we have been able to increase the occupation of the GPU but this later still not yet fully occupied because of the limited number of parallel processes and the number of nodes sent by each process. For this reason, we propose in the following another hybrid parallel approach to fully occupy the GPU.

4.3.2. Hybrid parallel evaluation of several bounds (H-PESB)

In order to fully occupy the GPU, we propose in the following the last parallel approach called Hybrid Parallel Evaluation of Several Bounds (H-PESB). This approach represents a hybridization between the multi-core approach and the PESB approach using the Nvidia MPS allowing several MPI-processes to use the GPU at the same time. Each mpi-process (master, workers) has its own B&B algorithm and use the GPU to evaluate a pool of node according to the PESB scheme. Thanks to MPS-server, each MPI-process uses the GPU like it is the only one using it. Therefore, the communication between the CPU and the GPU is done exactly in the same way as in one host process case. This approach is also similar to the H-PEB scheme, except that at each iteration, each host process sends a pool of nodes to the GPU for evaluation and immediate-selection instead of one node at a time. *i.e.* several GPU-blocks are allocated for each MPI-process and each one of them supports the evaluation of a single node. Therefore, each thread of the block update the head and tail values for one or several operations. The nodes sent by each mpi-process to the GPU for evaluation are chosen among the nodes recently added in the mpi-process B&B work-pool. At the end, the bounding results of nodes are sent back to the CPU to be used by the corresponding mpi-process for the selection, branching and elimination operations. As shown in Fig. 15, the search tree is explored simultaneously by several B&B instances and each instance uses the GPU according to the PESB model. The master divides the search space between the workers. After that, each worker explores its search space independently from the others and uses the GPU to evaluate several nodes at a time. As explained previously, The master and workers use the GPU simultaneously thanks to the Nvidia MPS that allows concurrent kernels execution while respecting the availability of GPU resources.

5. Experimentations

In this section, computational results are given using benchmarks obtained from the well known classical job shop instances by replacing the infinite buffer capacity by a zero buffer capacity constraint. We tested our approaches using the benchmark instances proposed by Taillard's [30]. The different instances are denoted by $n \times m$, where n and m represent respectively the number of jobs and the number of machines. The size of the Taillard's instances for the job shop problem varies between 15×15 and 100×20 . The experiments have been carried out using Intel Xeon E5640 CPU with four CPU-cores, 2.67 GHz clock speed each and Nvidia Tesla K40 with CUDA cores and 12 GB GDDR5 of global memory. The approaches have been implemented using C-CUDA 7.0, C++ and MPI [22] as a communication tool between processes. All reported times in this paper represent the average time to explore an equal number of nodes for each benchmark size. In our case, this number is fixed to 700,000 nodes, which is acceptable since an optimized sequential B&B algorithm takes 19 h to explore them all.

For the 100×20 benchmark instances there are 2002 operations. Since the GPU hardware limit is 1024 threads per block, we adapt the PEB approach to enable each thread to treat 2 operations instead of one which enables us to treat such big instances.

Fig. 16 shows the execution time needed for the Multi-core CPU and H-PEB approaches to explore 700,000 nodes using a different number of MPI-processes. For the Multi-core approach, the best time is reached for five MPI-processes. After that, we notice an increase in execution time when increasing the number of parallel MPI-processes. This can be explained by the limited number of CPU-cores available in our workstation (four cores). Therefore, the worker's tasks are executed sequentially when the number of workers is above four. For the Hybrid H-PEB approach the best time is reached for 35 processes which is the maximum supported since the Nvidia MPS support up to 35 connections to the MPS server. This hybrid version supports a large number of workers compared to the Multi-core version since each worker has less than 15% of its execution time on the CPU.

For the H-PESB approach, there is no easy way to find the best configuration. For each instance, we have to test several configurations and take the best one among them. Each configuration is defined as the number of MPI-processes and the number of nodes

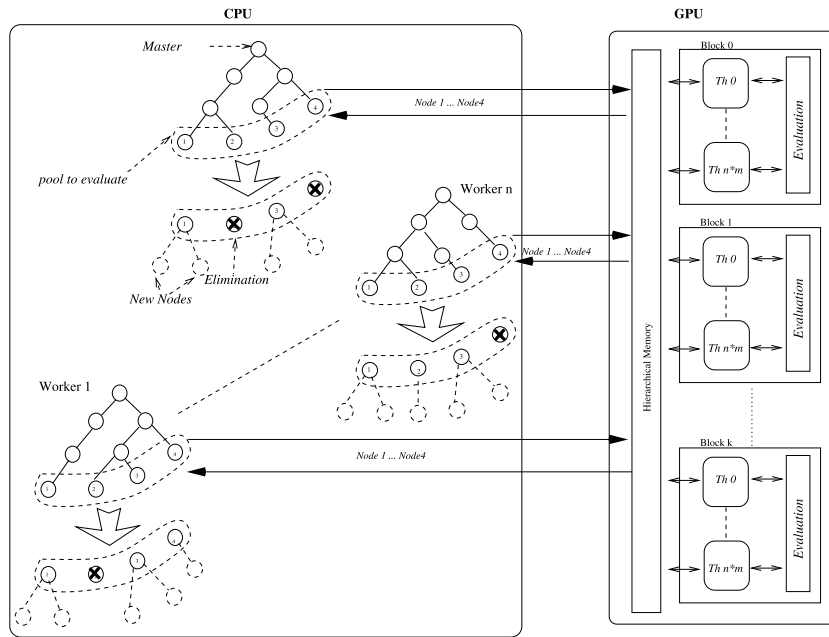


Fig. 15. Hybrid parallel evaluation of several bounds (*H-PESB*).

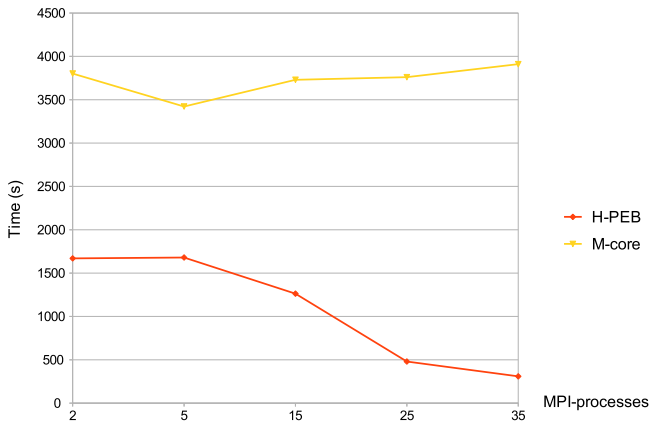


Fig. 16. Impact of using different number of MPI-processes to explore 700,000 nodes for Tai61 instance.

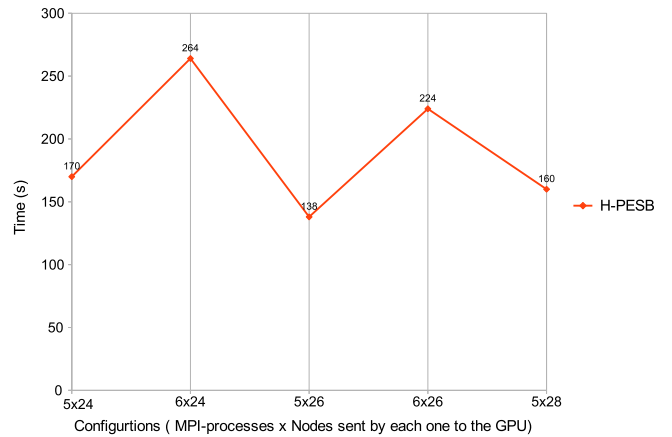


Fig. 17. *H-PESB* execution time using different configurations to explore 700,000 nodes for Tai61 instance.

sent to the GPU by each mpi-process at each iteration. As illustrated in Fig. 17 and unlike the H-PEB approach, the best performance is reached for low number of MPI-processes, because the MPS server can manage simultaneously a limited number of physical contexts (one for each parallel process) because of the large amount of virtual memory allocated to each process matching the size of nodes sent to the GPU. For the H-PEB approach, we can simultaneously manage a large number of physical context (35) because of the low virtual memory allocated to each process matching the size of one node only. We notice also from Fig. 17 that it is more benefit to increase the number of nodes sent to the GPU instead of increasing the number of MPI-process. Therefore for each instance, we fixed the number of MPI-processes to 5 and we increase the number of nodes sent to the GPU while it reduces the execution time needed to explore 700,000 nodes.

Table 3 reports the average execution times for each approach to explore 700,000 nodes. The first column (*Size*) reports the size of the benchmark instances. Column $B\&B_{Seq}$ reports the average execution time of an optimized sequential B&B algorithm. Column $B\&B_{Mcore}$ gives the execution time obtained by our Master/worker

approach exploiting only the Multi-cores CPU using 4 workers. For the all other approaches columns *Time* and *nb-nodes* report respectively the average execution time needed by each approach to explore 700,000 nodes and the number of nodes sent to the GPU at each iteration. For each column the parameter *nb_pr* indicates the number of parallel processes running simultaneously on the CPU according to the master/worker paradigm.

Column *PEB* reports the results of our GPU node-based approach obtained by sending one node at a time for parallel evaluation on the GPU. Column *PESB* reports the results of our second GPU-based approach obtained by sending several nodes to the GPU to be evaluated simultaneously and each one of them is evaluated in parallel using one block of GPU threads.

Columns *H-PEB* and *H-PESB* reports the results of the hybridization between the Multi-core approach and *PEB*, *PESB* GPU approaches using Nvidia MPS i.e. both master and workers use simultaneously the GPU to accelerate they respective bounding processes.

As mentioned before, 35 MPI-processes are used in the *H-PEB* approach and each one uses the default CUDA Stream to launch its

Table 3
Average execution time (in seconds) of the proposed approaches to explore 700,000 nodes. nb-pr: The number of parallel (host) processes running simultaneously. nb-nodes: The number of nodes evaluated simultaneously on the GPU by each parallel process.

Size	B & $B_{Seq.}$	B & B_{Mcore}	PEB (nb_pr = 1)		$H-PEB$ (nb_pr = 35)		$PESB$ (nb_pr = 1)		$H-PESB$ (nb_pr = 5)	
			nb-nodes	Time	nb-nodes	Time	nb-nodes	Time	nb-nodes	Time
15 × 15	188	52	1	603	1	162	240	37	90	60
20 × 15	384	113	1	653	1	164	240	53	90	59
20 × 20	393	120	1	736	1	173	240	71	86	58
30 × 15	1,076	375	1	795	1	180	240	104	54	60
30 × 20	1,127	447	1	955	1	209	140	156	46	71
50 × 15	4,246	1,454	1	1162	1	270	80	280	34	103
50 × 20	10,546	3,728	1	1530	1	340	30	396	26	145
100 × 20	69,300	19,200	1	3760	1	741	20	1050	20	418

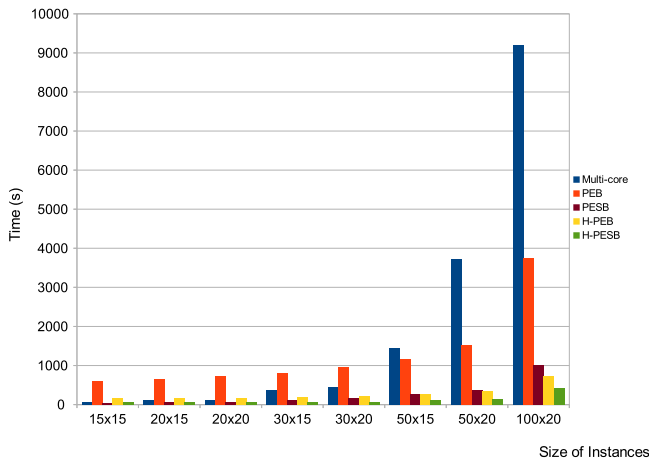


Fig. 18. Execution time of the proposed approaches.

kernel to evaluate one node at a time. For the $H-PESB$ approach, we fixed the number of parallel processes to 5 due to the huge amount of virtual memory matching the number of nodes sent to the GPU.

We notice from Table 3 that the complexity and the execution time increase when increasing the size of instances. Therefore, the need for parallelization is crucial.

Fig. 18 shows the histogram representation of the execution time for the different approaches. The first result from Table 3 and Fig. 18 is the positive impact of using parallel architectures to reduce the execution time needed to solve the BJSS problem.

The improvement obtained with the Multi-core version is low which is expected since our workstation contains only four CPU-cores. Therefore, increasing the number of workers above four reduces further the obtained performances. For the PEB version, we notice a low performance for small instances ($15 \times 15 - 30 \times 20$) against the Multi-core and sequential approaches. This can be explained by the high ratio of communication to computing time on the GPU *i.e.* the approach passes more time in sending data and recovering results to/from the GPU. By increasing the size of instances, we notice a significant improvement in execution time as compared with the sequential and multi-core cases. In addition to the efficiency in reducing the execution time for large instances, the PEB approach does not depend on the GPU capacity since it uses a small amount of GPU resources. However, only one block is used and the other blocks remain idle. Therefore, this approach cannot benefit from the entire GPU capacity.

The performance of the $PESB$ approach depends on the number of nodes that the GPU can evaluate simultaneously which is determined by the amount of the shared memory used by each GPU block to evaluate a node. The number of nodes evaluated simultaneously for small instances is equal to the maximum number of block that our GPU can run simultaneously (240). By increasing the

Table 4
The number of GPU communications needed for each approach to explore 700,000 nodes.

Approaches	#processes	#nodes sent	GPU communications
<i>Multi-core</i>	5	0	0
<i>PEB</i>	1	1	1400,000
<i>H-PEB</i>	35	1	40,000
<i>H-PESB</i>	5	20	14,000

size of instances, the number of nodes evaluated simultaneously matching the number of parallel blocks that a GPU can handle decreases. This behavior can be explained by the huge amount of shared memory needed to each block for large instances. Since the amount of shared memory is fixed, the number of parallel block decreases by increasing the shared memory used by each block matching the generated data for the handled instance used for synchronization. By sending several nodes at a time instead of one in the PEB approach, we have been able to reduce the execution time by a factor of 3 as compared with the PEB approach and a factor of 18x as compared to the multi-core version.

The hybrid approach ($H-PEB$) reduces considerably the execution time even for small instances against the sequential approach. This performance represents the results of exploiting both the CPU-cores and the GPU at the same time by using concurrent kernels execution provided by Nvidia MPS which allows us to increase the GPU occupation. Furthermore, the wasted time in CPU/GPU communications is covered by the concurrent access to the GPU where several workers execute their bounding operation at the same time. This hybridization allows us to reduce the execution time by a factor of 5x as compared with the PEB approach and a factor of 26x as compared with the multi-core version. Unlike smaller instances, the $H-PEB$ approach outperforms the results of the $PESB$ approach for large instances due to the limited number of nodes handled simultaneously by this later.

Our last parallel approach ($H-PESB$) is also based on the Nvidia MPS, it represents a hybridization between the multi-core approach and the $PESB$ approach. This approach fully occupies the GPU which explains the good obtained performance even for the smaller instances. The result of this approach is 2 times faster as compared to the $PESB$ approach and 46 times faster as compared with the multi-core version. The results of our two hybrid approaches show clearly the benefit of using both the multi-core CPU and the GPU at the same time as compared with approaches exploiting only the Multi-core CPU or only the GPU.

We notice for our experience of using Nvidia MPS that the best performance is obtained for a large number of parallel host processes as we can notice from the $H-PEB$ results in Table 3 and Fig. 16. But, This depends on the MPS server resources *i.e.* if he has enough virtual memory space for the parallel processes which is not the case in the $H-PESB$ approach.

Table 4 shows the number of GPU communications needed for each approach to explore 700,000 nodes. For each approach, column processes reports the number of used parallel processes.

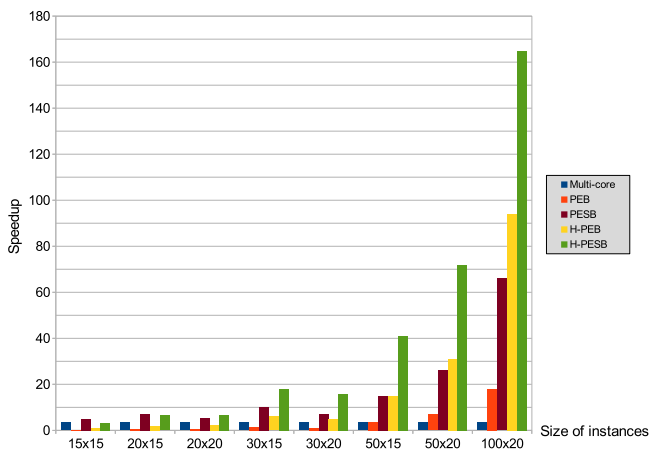


Fig. 19. The speedup of the proposed approaches.

Column nodes reports the number of nodes sent by each process. Finally, column GPU communications reports the number communications between the CPU and the GPU. The PEB approach has a huge number of GPU communications (1400,000) because, at each iteration, only one node is sent to the GPU then the results are moved back to the CPU. For the H-PEB approach we have 40,000 communications since, at each iteration, we can have 35 connections at the same time to the GPU without blocking. The same for the H-PESB approach, the later reduces more the number of communications to the GPU which explains the obtained performances.

Fig. 19 shows the relative speedup of our proposed approaches for different problem sizes. The speedup of our Multi-core version is around four, for all sizes which is expected since it depends on the number of CPU-cores available in our workstation. The speedup of the other approaches is proportional to the size of instances. Therefore, the maximum speed-up is obtained for the 100×20 instances. This is logical because the speedup of these approaches depends on the amount of the computation on the GPU. The idea used in PEB approach to accelerate the bounding of one node at a time on GPU using several threads organized into one GPU block gave good results (18 times faster) compared to the sequential version.

The PESB approach gave the best performance against all approaches for small instances, however, this is not the case for large instances as compared with the H-PEB and H-PESB approaches due to the limited number of nodes evaluated simultaneously on the GPU. This can be explained by the limited amount of shared memory available in the device and a large amount of this memory needed by each GPU block for synchronization.

The speedup obtained by the hybrid H-PEB approach is around 90 times faster which confirms the efficiency and the benefit of using both CPU-cores and GPU at the same time.

The H-PESB approach that fully occupies the GPU has achieved the best performances for almost all sizes against all other approaches. It has achieved an impressive speedup, especially for the largest instances where it is up to 160 times faster than a sequential B&B algorithm. In addition, the speedup of hybrid approaches grows according to the size of instances. This is due to the ratio of computing to communication time that increases by increasing the size of instances. This proves that the hybrid approaches are scalable and can easily deal with large instances. These last two approaches are based on the concurrent kernel's execution via Nvidia Multi Processes Service (MPS) which is rarely exploited in scientific computing. The performance (speedup) of the hybrid approaches is the result of:

- 1- Using the PEB scheme which is 18 times faster as the basis of our hybrid approaches.
- 2- Exploiting both the power of the CPU-cores and the GPU at the same time using Nvidia MPS.
- 3- The occupation of the GPU over time *i.e.* several workers run instructions on the GPU while others perform data-transfer from/to the GPU and yet others apply elimination and branching operators on the CPU.

6. Conclusions

This paper investigates the acceleration of the B&B method using Multi and Many-core systems in order to solve optimally the NP-hard Blocking Job Shop Scheduling problem. This problem represents a version of the classical job shop problem with no intermediate buffer between machines. In this paper, five approaches have been proposed. The first approach exploits only the CPU-core of our machine. The second and the third approaches, namely PEB and PESB respectively, are a GPU-based parallelization. Finally, the last two approaches (H-PEB and H-PESB) are hybrid. They exploit the Multi-core CPU and the GPU at the same time by combining the approaches above using concurrent Kernels execution provided by Nvidia MPS. The obtained results confirm the efficiency of our proposals and the positive impact of using computing accelerators like GPUs to solve this problem. The results show the advantage of increasing the GPU occupation over time by using Hybrid approaches based on the concurrent kernel's execution provided by Nvidia MPS which allowed us to achieve an impressive relative speedup.

As a future perspective, we plan to act on the granularity of tasks assigned to each thread and explore more heterogeneous architectures like Intel Xeon Phi.

Acknowledgments


We would like to thank Dr. Djamel Belazzougui for his assistance and comments that greatly improved the presentation of this manuscript.

Dr. Didier El Baz gratefully acknowledges the support of NVIDIA Corporation with the donation of the Tesla K40 GPU used for this research work.

References

- [1] Abdelhakim AitZai, Brahim Benmedjdoub, Mourad Boudhar, A branch and bound and parallel genetic algorithm for the job shop scheduling problem with blocking, *Int. J. Oper. Res.* 14 (3) (2012) 343–365.
- [2] Leila Asadzadeh, A local search genetic algorithm for the job shop scheduling problem with intelligent agents, *Comput. Ind. Eng.* 85 (2015) 376–383.
- [3] Ahcène Bendjoudi, Mehdi Chekini, Makhlof Gharbi, Malika Mehdi, Karima Benatchba, Fatima Sitayeb-Benbouzid, Nouredine Melab, Parallel B&B algorithm for hybrid multi-core/GPU architectures, in: *High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing (HPCC_EUC)*, 2013 IEEE 10th International Conference on, IEEE, 2013, pp. 914–921.
- [4] Peter Brucker, P. Brucker, *Scheduling Algorithms*, Springer, 2007.
- [5] Jacques Carlier, Eric Pinson, Adjustment of heads and tails for the job-shop problem, *European J. Oper. Res.* 78 (2) (1994) 146–161.
- [6] Tiago Carneiro, Albert Einstein Muritiba, Marcos Negreiros, Gustavo Augusto Lima de Campos, A new parallel schema for branch-and-bound algorithms using GPGPU, in: *Computer Architecture and High Performance Computing (SBAC-PAD)*, 2011 23rd International Symposium on, IEEE, 2011, pp. 41–47.
- [7] Imen Chakroun, Nordine Melab, Mohand Mezma, Daniel Tuytens, Combining multi-core and GPU computing for solving combinatorial optimization problems, *J. Parallel Distrib. Comput.* 73 (12) (2013) 1563–1577.
- [8] Imen Chakroun, Mohand Mezma, Nouredine Melab, Ahcène Bendjoudi, Reducing thread divergence in a GPU-accelerated branch-and-bound algorithm, *Concurr. Comput.: Pract. Exper.* 25 (8) (2013) 1121–1136.
- [9] NVIDIA Corporation, Multi-Process service, 2012, https://docs.nvidia.com/deploy/pdf/CUDA_Multi_Process_Service_Overview.pdf.

- [10] C. Cuda, Programming guide, 2012, <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>.
- [11] Adel Dabah, Ahcene Bendjoudi, Abdelhakim AitZai, Efficient parallel B&B method for the blocking job shop scheduling problem, in: High Performance Computing & Simulation (HPCS), 2016 International Conference on, IEEE, 2016, pp. 784–791.
- [12] Adel Dabah, Ahcene Bendjoudi, Abdelhakim AitZai, An efficient tabu search neighborhood based on reconstruction strategy to solve the blocking job shop scheduling problem, J. Indust. Manage. Optim. 13 (4) (2017) 2015–2031.
- [13] Adel Dabah, Ahcene Bendjoudi, Abdelhakim AitZai, Didier El-Baz, Nadia Nouali Taboudjemmat, Multi and many-core parallel B&B approaches for the blocking job shop scheduling problem, in: High Performance Computing & Simulation (HPCS), 2016 International Conference on, IEEE, 2016, pp. 705–712.
- [14] Adel Dabah, Ahcene Bendjoudi, Didier El-Baz, Abdelhakim AitZai, GPU-based two level parallel B&B for the blocking job shop scheduling problem, in: Parallel and Distributed Processing Symposium Workshops, 2016 IEEE International, IEEE, 2016, pp. 747–755.
- [15] Heinz Gröflin, Andreas Klinkert, A new neighborhood and tabu search for the blocking job shop, Discrete Appl. Math. 157 (17) (2009) 3643–3655.
- [16] Nicholas G. Hall, Chelliah Sriskandarajah, A survey of machine scheduling problems with blocking and no-wait in process, Oper. Res. 44 (3) (1996) 510–525.
- [17] Alessandro Mascis, Dario Pacciarelli, Job-shop scheduling with blocking and no-wait constraints, European J. Oper. Res. 143 (3) (2002) 498–517.
- [18] Yazid Mati, Nidhal Rezg, Xiaolan Xie, A taboo search approach for deadlock-free scheduling of automated manufacturing systems, J. Intell. Manuf. 12 (5–6) (2001) 535–552.
- [19] Nouredine Melab, Imen Chakroun, Ahcene Bendjoudi, Graphics processing unit-accelerated bounding for branch-and-bound applied to a permutation problem using data access optimization, Concurr. Comput.: Pract. Exper. 26 (16) (2014) 2667–2683.
- [20] Nouredine Melab, Jan Gmys, Mohand Mezma, Daniel Tuytens, Multi-core versus many-core computing for many-task Branch-and-Bound applied to big optimization problems, Future Gener. Comput. Syst. (2017).
- [21] Carlo Meloni, Dario Pacciarelli, Marco Pranzo, A rollout metaheuristic for job shop scheduling problems, Ann. Oper. Res. 131 (1–4) (2004) 215–235.
- [22] MPI Message passing interface forum, MPI: A Message Passing Interface Standard, Version 3.0, 2012.
- [23] Housseem Eddine Nouri, Olfa Belkahlia Driss, Khaled Ghédira, Hybrid metaheuristics within a holonic multiagent model for the flexible job shop problem, Procedia Comput. Sci. 60 (2015) 83–92.
- [24] Housseem Eddine Nouri, Olfa Belkahlia Driss, Khaled Ghédira, Hybrid metaheuristics for scheduling of machines and transport robots in job shop environment, Appl. Intell. 45 (3) (2016) 808–828.
- [25] Angelo Oddi, Riccardo Rasconi, Amedeo Cesta, Stephen F. Smith, Iterative improvement algorithms for the blocking job shop, in: ICAPS, 2012.
- [26] Marco Pranzo, Dario Pacciarelli, An iterated greedy metaheuristic for the blocking job shop scheduling problem, J. Heuristics 22 (4) (2016) 587–611.
- [27] Heinrich Riebler, Michael Lass, Robert Mittendorf, Thomas Lücke, Christian Plessl, Efficient branch and bound on fpgas using work stealing and instance-specific designs, ACM Trans. Reconfigur. Technol. Syst. (TRETS) 10 (3) (2017) 24.
- [28] Bernard Roy, B. Sussmann, Les problemes d'ordonnement avec contraintes disjonctives, Note ds 9 (1964).
- [29] Danilo S. Souza, Haroldo G. Santos, Igor M. Coelho, A hybrid heuristic in GPU-CPU based on scatter search for the generalized assignment problem, Procedia Comput. Sci. 108 (2017) 1404–1413.
- [30] Eric Taillard, Benchmarks for basic scheduling problems, Eur. J. Oper. Res. 64 (2) (1993) 278–285.
- [31] Trong-Tuan Vu, Bilel Derbel, Parallel branch-and-bound in multi-core multi-cpu multi-gpu heterogeneous environments, Future Gener. Comput. Syst. 56 (2016) 95–109.
- [32] Wangdong Yang, Kenli Li, Keqin Li, A hybrid computing method of SpMV on CPU-GPU heterogeneous computing systems, J. Parallel Distrib. Comput. 104 (2017) 49–60.



Ahcene Bendjoudi received his Master and Ph.D. degrees in Computer Science from the University of Bejaia, Algeria, and his HDR and state engineering degrees in Computer Science from the High School of Computer Science (ESI). He is currently full-time senior researcher at the Research Centre in Scientific and Technical Information (CERIST, Algiers) within Theory and Engineering of Computer Systems division (DTISI) since 2012. He is the founder and head of the Parallel Computing and Applications (CAPA) research group, and adjunct professor at the University of Science and Technology USTHB, Algiers. His major research interests include Parallel Combinatorial Optimization, Cluster, Multi-core and GPU-Computing, and recently, Big Data, Big Graphs, Intelligent Transportation Systems, and Smart Cities. He conducted and he was involved in several research projects on these topics. He has more than 35 international publications including journal papers, book chapters and conference proceedings.



Abdelhakim AitZai was born in Algeria in 1970. He is currently a lecturer and researcher at the Department of Computer Sciences, Houari Boumediene University of Algiers (USTHB) and member of RECITS Laboratory, Faculty of Mathematics (USTHB). He received his engineer and master degrees in mathematics in 1994 and 1998, respectively from USTHB University. He received also his Ph.D. and HDR in computer sciences from USTHB University in 2012 and 2014 respectively. His research interests include, manufacturing systems, scheduling and optimization and parallel and distributed computing. He has published several research papers in national and international journals and conference proceedings.



Didier El Baz was born in Toulouse, France in 1958, he received the Engineer degree in Electrical Engineering and Computer Science from National Institute of Applied Sciences in Toulouse (Institut National des Sciences Appliquées, INSA) in 1981 and the Doctor Engineer degree in Control Theory from INSA Toulouse in January 1984. Dr. El Baz was visiting scientist in the Laboratory for Information and Decision Systems, MIT Cambridge Massachusetts, USA, in 1984. He received the HDR in Computer Sciences from Institut National Polytechnique of Toulouse (INP) in 1998. He is the founder and head of the team Distributed Computing and Asynchronism at the Laboratory of Analysis and Systems Architecture (LAAS-CNRS). Dr El Baz is the author of 40 papers in referred international journals and 80 papers in referred International Conference. His fields of interest are in parallel and distributed computing, GPU computing, optimization, scheduling and Internet of Things. Dr El Baz was the coordinator of ANR project CIP that dealt with high performance peer-to-peer computing. Dr El Baz was the advisor of 11 Ph.D. students and 21 Master students. Dr El Baz was Program Chairman and Organizing Chairman of the 16th International Conference on Parallel, Distributed and networked-based Processing, PDP 2008, Toulouse, France. Program Chairman of the 17th International Conference on Parallel, Distributed and networked-based Processing, PDP 2009, Weimar, Germany, General Cochairman of IEEE international Conference iThings 2013, Program Chairman of IEEE international Conference CSE 2014, Chengdu China, General Chair of IEEE CSE 2015, Porto, Portugal, Executive Chair of the 15th IEEE International Conference on Scalable Computing and Communications (ScalCom 2015) Beijing China and General Chair of IEEE ScalCom 2016 as well as General Chair of the 13th IEEE International Conference on Ubiquitous Intelligence and Computing (UIC 2016) Toulouse. Dr El Baz has been invited to give Talks at China University of Petroleum Qingdao, China University of Geosciences Beijing, CERIST Algiers Algeria, as well as several University and Institute in France like University of Paris XIII, University of Perpignan, CEA, CNES and CNRS. Dr. El Baz has also given courses on High Performance Computing at several Universities like University of Sciences and Technology Beijing, University of Sichuan and University Houari Boumediene.



Nadia Nouali-Taboudjemmat currently a Director of Research and the head of the Theory and Engineering of Computer Systems and of the Ubiquitous Computing Group at Research Centre in Scientific and Technical Information (CERIST) in Algeria. She graduated in C.S. from the Houari Boumediene University (USTHB, Algiers) and she obtained her Magister Degree from the Advanced Technologies Research Centre (Algiers). She also obtained her Doctorate Degree from USTHB. Her research interests include pervasive/ubiquitous computing, distributed and mobile computing, wireless networks, etc. Her recent researches include big data management, cloud computing and she is particularly interested by the application domain of ICT-based disaster management and smart cities.



Adel Dabah is Ph.D. student and full time researcher in the CERIST Research Center Algiers, Algeria since 2014. He received his master and Magister degrees in Computer Science from the University of Sciences and Technology Houari Boumediene (USTHB), Algiers, Algeria, in 2010 and 2013 respectively. His research interests include, manufacturing systems, scheduling and optimization, and parallel and distributed computing.