



PUC

ISSN 0103-9741

Monografias em Ciência da Computação
n° 08/09

Competitive Deterministic Heuristics for Permutation Flow Shop Scheduling

Sanjay Dominik Jena

Marcus Vinicius Soledade Poggi de Aragão

David Sotelo Pinheiro da Silva

Departamento de Informática

PONTIFÍCIA UNIVERSIDADE CATÓLICA DO RIO DE JANEIRO

RUA MARQUÊS DE SÃO VICENTE, 225 - CEP 22451-900

RIO DE JANEIRO - BRASIL

Competitive Deterministic Heuristics for Permutation Flow Shop Scheduling

Sanjay D. Jena and Marcus V. S. Poggi de Aragão and David Sotelo P. da
Silva

sjena@inf.puc-rio.br, poggi@inf.puc-rio.br, dsilva@inf.puc-rio.br

Abstract. In the last 50 years, the Permutation Flow Shop Scheduling Problem with makespan minimization (PFSP) has been a central and well-studied problem in Scheduling and Operations Research communities, known by its intractability, from theoretical and practical aspects. This work introduces polynomial time deterministic heuristics for the PFSP based on pruning techniques of the implicit enumeration tree following new extensions of the classical NEH heuristic. Experimental results attest that the proposed methods currently stand among the most effective for the PFSP.

Keywords: Permutation flow shop scheduling; Deterministic heuristics; NEH extensions.

Resumo. Durante os últimos 50 anos, o problema Permutation Flow Shop Scheduling com minimização de makespan (PFSP) tem sido um problema central e intensamente estudado pelas comunidades científicas de escalonamento e pesquisa operacional, sendo conhecido por sua intratabilidade, em aspectos tanto teóricos quanto práticos. O presente trabalho apresenta heurísticas determinísticas de tempo polinomial para o PFSP baseadas em técnicas de poda na árvore de enumeração implícita, originando novas extensões da heurística clássica NEH. Resultados experimentais comprovam que os novos métodos propostos situam-se atualmente entre os mais efetivos para o NEH.

Palavras-chave: Escalonamento permutation flow shop; Heurísticas determinísticas; Extensões para o NEH.

In charge of publications:

Rosane Teles Lins Castilho
Assessoria de Biblioteca, Documentação e Informação
PUC-Rio Departamento de Informática
Rua Marquês de São Vicente, 225 - Gávea
22451-900 Rio de Janeiro RJ Brasil
Tel. +55 21 3527-1516 Fax: +55 21 3527-1530
E-mail: bib-di@inf.puc-rio.br
Web site: <http://bib-di.inf.puc-rio.br/techreports/>

1 Introduction

The Permutation Flow Shop Scheduling Problem (PFSP) is one of the most studied scheduling problems in combinatorial optimization. Given a set J of n jobs, a set M of m machines and processing times $t_{i,j}$ for each job j on each machine i , the problem consists of scheduling all n jobs at each one of the m machines. The processing sequence of the jobs must be the same on all machines and each job j can only start its execution on a machine i if both the previous job on the same machine i as well as the same job j on the previous machine $i - 1$ have already been processed. Furthermore, the order in which a job must pass through the machine set is predefined and identical for all jobs. The objective of this problem is to determine a job ordering that minimizes the completion time of the last job in the last machine, called the *makespan* criterion C_{max} . Though Johnson et al. (1976) showed that the 2-machine problem can be solved in polynomial time, the general case with m machines is known to be strongly NP-complete. This work focuses on the PFSP with the makespan criterion, which was firstly denoted as $F_m|prmu|C_{max}$ by Graham et al. (1979). Its main objective is to introduce new deterministic and polynomial time heuristics which are partially based on the NEH-T heuristic of Taillard (1990), extended by an improvement phase. The heuristics construct k solutions and perform l iterations in order to improve them (see Section 3). These extensions lead to a total time complexity of $O(n^2 \cdot k \cdot l \cdot (m + \log k))$.

This work is organized as follows. A literature review of heuristics for the PFSP is provided in Section 1.1). In Section 2, the original NEH of Nawaz et al. (1983) and NEH-T heuristics are considered. Section 3 discusses the effect of initial job ordering and tie breaking rules on the quality of the solutions generated by the NEH heuristic, suggesting new approaches to explore the solution space of the PFSP. The construction and improvement phases of the new heuristics proposed in this work, named NEH-Delta and NEH-Alpha, are presented in Sections 4 e 5, respectively. In Section 6, experiments on Taillard (1993)'s benchmark instances are taken, showing that these new deterministic heuristics are competitive with the metaheuristics currently standing among the most effective ones for the PFSP. Final conclusions are drawn in Section 7.

1.1 Previous work review

Since the PFSP is strongly NP-complete, as shown by Graham et al. (1979), fairly much research was attracted to find good heuristic algorithms for this problem. The algorithms range from deterministic and polynomial time constructive methods to metaheuristics based on non-deterministic approaches.

Deterministic approaches appeared already four decades ago. Johnson (1954)'s pioneers work was followed by Page (1961)'s pairing algorithm, Palmer (1965)'s algorithm, Campbell et al. (1970)'s CDS algorithm, the algorithms of Gupta (1971, 1972) and Bonney and Gundry (1976), Dannenbring (1977)'s effective Rapid Access heuristic (RA) and some further variations based on RA, the famous NEH heuristic of Nawaz et al. (1983), the algorithm of Hundal and Rajgopal (1988) and the heuristic of Koulamas (1998). In 1990, Taillard (1990) showed how to bring the NEH's original runtime of $O(n^3 \cdot m)$ down to $O(n^2 \cdot m)$ (see Section 2). As the NEH is still considered one of the best heuristics for the PFSP, as demonstrated in recent works such as Ruiz and Maroto (2005) and Kalczynski and Kamburowski (2008), there can be found many NEH based works, extending or modifying

the original algorithm. Framinan et al. (2003), Kalczynski and Kamburowski (2008) and Dong et al. (2008) work with different initial orders of the jobs and special tie breaking rules to improve the performance of the original NEH. Other recent works on deterministic methods include the competitive heuristics of Farahmand et al. (2009).

Non-deterministic approaches include metaheuristics such as Tabu Search, Simulated Annealing, Local Search, Genetic Algorithms and Ant Colonies. Taillard (1990)'s tabu search was one of the first non-deterministic approaches for the PFSP. One of the today's most competitive metaheuristics is the Iterated Greedy Algorithm of Ruiz and Stützle (2007). They review many of the most effective metaheuristics and, given a specified time limit, outperform all of them in the computational experiments. Other references to metaheuristics for the PFSP can be found in numerous surveys and reviews such as Framinan et al. (2004), Ruiz and Maroto (2005) and Hejazi and Saghafian (2005) that appeared due to the large number of algorithms for this problem.

Most of the surveys concerning heuristics for the PFSP tried to address different scopes instead of giving just an up-to-date review of past works. Framinan et al. (2004) introduced a new framework to classify deterministic as well as non-deterministic approaches into three phases: index development, solution construction and solution improvement. Ruiz and Maroto (2005) made a comprehensive evaluation of deterministic heuristics and metaheuristics for the PFSP, presenting vast computational experiments regarding its comparison. Hejazi and Saghafian (2005) extended this effort to exact methods as well.

2 NEH and NEH-T

The NEH algorithm of Nawaz et al. (1983) can be described by the following steps:

1. For each job j , compute the sum S_j of its processing times on all machines:

$$S_j = \sum_{i=1}^m t_{i,j}, \quad \forall j \in J$$

2. Sort all jobs in descending order of S_j to form the sequence of jobs j_1, \dots, j_n .
3. Take the first two jobs j_1 and j_2 and order them so that the makespan of the sequence given by the ordering of these two jobs is minimized.
4. Do for each of the jobs j_3, \dots, j_n , successively: Insert the next job into the sequence at the position that results in the smallest partial makespan among all possible insertion positions.

The first step needs $O(n \cdot m)$ time to calculate the sum of processing times for each job and $O(n \log n)$ to sort the jobs. In the following insertions for the n jobs, the original NEH evaluates each insertion position by determining the new makespan in $O(n \cdot m)$. As there are j insertion positions for the j -th job to be inserted, it takes $O(n^2 \cdot m)$ to evaluate all insertion positions for one job. Hence, the original NEH has a total time complexity of $O(n^3 \cdot m)$. Taillard (1990) shows that the j insertion positions can be evaluated in $O(n \cdot m)$ time using dynamic programming and thus reduced the total running time to $O(n^2 \cdot m)$. This improved version is referred to the NEH-T heuristic.

Due to its exceptional results in fairly short running time, the NEH is up-to-date one of the most discussed and analyzed PFSP heuristics. Comprehensive computational experiments given by Ruiz and Maroto (2005) argue that, considering its low computation costs, the NEH can still be considered the best heuristic among all deterministic ones.

3 Extending NEH

The outstanding practical performance of the NEH heuristic is mainly related to two factors: first, the job grouping mechanism, based on the operation of inserting a new job at the best partial position, can be computed in $O(m)$ time. Second, the order in which jobs are taken for insertion (decreasing sum of processing times) has a considerable impact if compared with other ordering mechanisms. Framinan et al. (2003) showed that the initial order of the original NEH is the best among roughly 140 evaluated starting sequences. Recently, Kalczynski and Kamburowski (2008) presented a new initial order with special tie breaking rules which yields better results than the original NEH order.

Considering that the makespan of a partial permutation schedule is an effective criterion to decide whether this partial solution must be explored or discarded, two important points, not directly treated by the NEH, come to mind. First, how to select the partial permutation schedule to be explored in the case of ties, when two or more solutions have the same makespan. Nowicki and Smutnicki (1993) constructed a special family of instances for which the NEH algorithm have an approximation factor of $\Omega(\sqrt{m})$ due to ties on the selection of job insertion positions. Second, what is the effect of exploring not only one partial solution with minimum makespan but a set of partial solutions with low makespans.

The effect of initial job ordering and tie breaking rules. The following example illustrates how the initial job ordering and arbitrary tie breaking for the NEH have effect on the quality of the generated solutions. Consider an instance of the PFSP composed of $n = 3$ jobs and $m = 9$ machines in which the processing times matrix is defined in Table 1.

	M_1	M_2	M_3	M_4	M_5	M_6	M_7	M_8	M_9
J_1	$1+2\epsilon$	0	0	0	0	1	0	1	0
J_2	0	$1+\epsilon$	0	0	1	0	1	0	0
J_3	0	0	1	1	0	0	0	0	1

Table 1: Example of a processing times matrix

Some considerations are necessary at this point. First, zero processing times on the matrix represent extremely small values, once by definition all processing times are strictly positive on the PFSP. More precisely, the value represented by 0 is less than $\frac{1}{n+m-1}$. Consider also that $0 \leq \epsilon < 1/2$. The application of the NEH heuristic to such instance leads to an initial insertion ordering (J_1, J_2, J_3) calculated by the sums of jobs processing times. On the insertion phase, J_2 is positioned right before J_1 , minimizing the partial makespan of these two first jobs. As consequence, any insertion point for J_3 will result in a final makespan of $5+\epsilon$. However, the optimal solution $\pi = (J_3, J_1, J_2)$ generates a

makespan of $4+3\epsilon$. Making ϵ also very small, leads to a gap¹ of about 25% between the optimal solution and the one found by the NEH. This simple example illustrates how the selection of the second best partial solution (insertion of J_2 after J_1) can result in a better choice. In particular, when $\epsilon = 0$ ties can occur on the initial job ordering phase and on the job insertion phase. Similarly, tie breaking decisions can affect the quality of the generated solutions.

Many works such as Framinan et al. (2003) and Kalczynski and Kamburowski (2005, 2008) confirmed the strong influence of the tie breaking decisions and showed that both the initial order and the tie breaking decisions are crucial to the good results of the NEH. Based on such studies, Kalczynski and Kamburowski (2008) presented a new initial order and a simple tie breaking method which outperform the solutions obtained by the original NEH method.

NEH and Enumeration Trees. Consider that the construction mechanism for the enumeration tree of the PFSP selects, at the level j , the position $p \in \{1, \dots, j\}$ in which the j -th job should be inserted in the partial permutation consisting of the first $j-1$ jobs. Once position p is selected, all jobs previously allocated at positions $p' \geq p$ are shifted to positions $p' + 1$. Clearly, all possible permutation schedules can be generated from such enumeration process. As consequence, the partial permutation schedule represented by a node is the path from the root of the enumeration tree to it. The NEH heuristic defines a strategy that limits the exploration of every branch of a node on this enumeration tree. Following the insertion phase of NEH, only the node with minimum cost is explored at level j . All other nodes and its corresponding subtrees are pruned. The cost of a node in the enumeration tree is the makespan of the partial permutation schedule that it represents. Figure 1(a) illustrates such explicit enumeration process for the previous example with three jobs. The final permutation $\pi = (J_3, J_2, J_1)$ found by NEH is represented by the node $(3, 2, 1)$ on the enumeration tree, the path from root to this node denotes the job insertions carried out by NEH. Dashed edges represent pruned nodes on the enumeration tree.




Figure 1: NEH and Enumeration Trees

¹We define the gap of a solution's makespan, denoted by $makespan_{sol}$, as its deviation from the optimal makespan for this instance, denoted by $makespan_{opt}$, i.e. $gap = \frac{makespan_{sol} - makespan_{opt}}{makespan_{opt}}$.

New approaches to explore the PFSP solution space. Following the principles introduced by NEH, a natural extended strategy to explore the solution space consists of considering, at each level of the enumeration tree, not only the node with the minimum cost (makespan of its corresponding partial permutation) but a set of k nodes with costs close to the lowest one on that level. The previously introduced example illustrates the benefits that can be obtained by following such approach, as shown in Figure 1(b). In such example, considering that $k = 2$, both nodes (1, 2) and (2, 1) are kept on the second level of the enumeration tree. Consequently, the two solutions of lowest costs are selected at the third level, one of them is the optimal solution. Finally, fixing parameter k or defining it as a polynomial function on input size lead to algorithms that explore the PFSP solution space in polynomial time.

4 Construction Phase

This work proposes two extensions over the NEH heuristic, named NEH-Delta and NEH-Alpha. The common framework governing these extensions is related to the definition of a set of nodes to be explored at a level of the enumeration tree. In the classical NEH heuristic, this set has only one element. The extensions proposed in this work consider sets of nodes with larger cardinality but always limited to a polynomial function on the input size so that the final heuristic can be executed in polynomial time. Throughout this work, we refer the execution of the algorithms NEH-Delta and NEH-Alpha to the *construction phase*.

4.1 NEH-Delta

Given a positive integer k defined as a parameter, this heuristic makes use of a greedy strategy to select at most k nodes to be explored from a level of the enumeration tree. All other nodes at this level and its corresponding subtrees are pruned. The k nodes selected at a level are those of minimum cost that are *valid*. A node is called *valid* if none of its ancestral on the enumeration tree was pruned. A node u is said an ancestral of a node v on the enumeration tree if and only if u is on the unique path from the root of the enumeration tree to v . Each level of the enumeration tree corresponds to the insertion of a job, selected on the initialization phase of NEH that sorts the jobs by decreasing sum of its processing times. At the end of the execution, k solutions are obtained from which that of minimum makespan is returned. The complete pseudo-code for the NEH-Delta is presented in Algorithm 1.

Implementation details and time complexity analysis. The NEH-Delta implementation implicitly constructs the pruning mechanism on the enumeration tree. Every partial permutation schedule δ is represented by a linked list so that, given a partial permutation δ , a new job j_i can be inserted at position p , generating partial permutation δ' , in $O(1)$ amortized time. An important point concerning this time complexity for the construction of δ' is that the elements of δ do not necessarily have to be stored in δ' . In fact, it is possible to store in δ' only a reference (pointer) to δ and the position p in which job j_i should be inserted. However, at the end of an iteration, when the elements of Δ_{next} are copied to Δ , every permutation $\delta' \in \Delta_{next}$ must copy the elements from its originating

Algorithm 1: NEH-Delta

Input : Set J of n jobs,
Set M of m machines,
Integer k ,
Processing times matrix $T \in \mathbb{R}_{M \times J}^+$.

Output: permutation function $\pi : \{1, \dots, n\} \mapsto J$.

begin

- Sort the jobs in set J in decreasing order of its processing time sums resulting in order (j_1, \dots, j_n) ;
- $\Delta \leftarrow \{j_1\}$;
- for each job** $j_i, i = 2 \dots n$ **do**
 - $\Delta_{next} \leftarrow \emptyset$;
 - for each partial permutation** $\delta \in \Delta$ **do**
 - for each position** $p \in \{1, \dots, i\}$ **do**
 - Insert job j_i at the position p of partial permutation δ generating partial permutation δ' ;
 - if** $|\Delta_{next}| < k$ **then**
 - $\Delta_{next} \leftarrow \Delta_{next} \cup \{\delta'\}$;
 - else if** $makespan(\delta') < \max_{\underline{\delta}' \in \Delta_{next}} \{makespan(\underline{\delta}')\}$ **then**
 - $\Delta_{next} \leftarrow \Delta_{next} - \{\underline{\delta}'\}$;
 - $\Delta_{next} \leftarrow \Delta_{next} \cup \{\delta'\}$;
- $\Delta \leftarrow \Delta_{next}$;
- $\pi \leftarrow$ permutation of Δ with minimum makespan;
- return** π ;

end

permutation δ in $O(k \cdot i)$ time. The partial permutation sets Δ and Δ_{next} are implemented as binary heaps using the permutations' makespan as key. Thus, the permutation δ' of maximum makespan in Δ_{next} can be found in $O(\log k)$ time. Similarly, permutations can be inserted or removed from these sets in $O(\log k)$ time.

In the case of makespan ties, the solution that was firstly inserted into the heap always has the priority of being maintained inside the heap. This criterion is used throughout this work as a tie breaking rule. The job sorting initial operation can be accomplished in $O(n \cdot m + n \log n)$ time. Making use of the dynamic programming algorithm of Taillard (1990), the makespan of all partial solutions δ' obtained by inserting job j_i in every position of a previous partial solution δ can be generated and calculated in $O(i \cdot m)$. Consequently, NEH-Delta can be implemented in $O(k \cdot n^2 \cdot (m + \log k))$ time.

4.2 NEH-Alpha

Though the NEH-Delta tends to yield better results increasing k , this is not a rule. In some rare cases the increase of k leads to a worse final makespan², as the better solution is pruned by other partial solutions during the construction phase. This behavior is directly linked to the fact that only the k partial solutions with best makespan are further explored. In order to increase the diversity of the explored solutions, the NEH-Alpha heuristic determines, at each level of the enumeration tree, not an unique set of nodes to be explored but n disjoint sets of nodes. Each set A^p represents the nodes to be explored at level i of the enumeration tree so that the i^{th} job (following the job ordering) is inserted at position p . Furthermore, given a positive integer k defined as a parameter, all sets A^1, A^2, \dots, A^n must have its cardinalities limited to $\lfloor \frac{k}{n} \rfloor$. Making use from the same greedy criterion introduced for the NEH-Delta heuristic, the $\lfloor \frac{k}{n} \rfloor$ nodes selected to compose the set A^p at a level of the enumeration tree are those of minimum cost that are *valid*. The main difference from this strategy to the one employed by NEH-Delta is that, in this approach, nodes only compete with nodes at the same level of the enumeration tree in which the current job is inserted at the same position. The heuristic obtains up to k solutions (in fact, we have a total of $n \cdot \lfloor \frac{k}{n} \rfloor$ solutions) at the end of the heuristic's execution. The one with minimum makespan is the final solution. The complete pseudo-code for the NEH-Alpha is presented in Algorithm 2.

Implementation details and time complexity analysis. Data structures and implementation strategies used for NEH-Alpha are quite similar to the ones of NEH-Delta. The main difference is that partial permutations sets $A^1 \cup \dots \cup A^n$ are implemented as n independent binary heaps so that permutations can be inserted or removed from the permutation sets in $O(\log \lfloor \frac{k}{n} \rfloor)$ time. Consequently, NEH-Alpha can be implemented in $O(k \cdot n^2 \cdot (m + \log \lfloor \frac{k}{n} \rfloor))$ time.

4.3 Computational experiments

All computational experiments in this work were performed using the well known benchmark instances of Taillard (1993). For all instances that already have been solved to optimality, the optimum value is used to compute the gap. For all other instances, the upper bounds shown in Table 2 are used. The algorithms were implemented in C++ and

²As example, we refer to Taillard's benchmark instance tai20_5. The NEH-Delta with $k = 1$ leads to a makespan of 1223, whereas $k = 5$ yields a final solution of makespan 1229.

Algorithm 2: NEH-Alpha

Input : Set J of n jobs,
Set M of m machines,
Integer k ,
Processing times matrix $T \in \mathfrak{R}_{M \times J}^+$.

Output: permutation function $\pi : \{1, \dots, n\} \mapsto J$.

begin

- Sort the jobs in set J in decreasing order of its processing time sums resulting in order (j_1, \dots, j_n) ;
- $A^1 \leftarrow \{j_1\}$;
- for** $p \in \{2, \dots, n\}$ **do**
 - $A^p \leftarrow \emptyset$;
 - for each job** $j_i, i = 2 \dots n$ **do**
 - for** $p \in \{1, \dots, i\}$ **do**
 - $A_{next}^p \leftarrow \emptyset$;
 - for each partial permutation** $\alpha \in A^1 \cup \dots \cup A^i$ **do**
 - for each position** $p \in \{1, \dots, i\}$ **do**
 - Insert job j_i at the position p of partial permutation α generating partial permutation α' ;
 - if** $|A_{next}^p| < k$ **then**
 - $A_{next}^p \leftarrow A_{next}^p \cup \{\alpha'\}$;
 - else if** $makespan(\alpha') < \max_{\alpha \in A_{next}^p} \{makespan(\alpha)\}$ **then**
 - $A_{next}^p \leftarrow A_{next}^p - \{\alpha\}$;
 - $A_{next}^p \leftarrow A_{next}^p \cup \{\alpha'\}$;

- $\pi \leftarrow$ permutation of $A^1 \cup \dots \cup A^n$ with minimum makespan;
- return** π ;

end

compiled under the standard configuration of Visual Studio 2005 Version 8. The experiments were carried out on a Notebook Sony VAIO (VGN-FZ21M), Dual Core 2.0GHz, 3Gb memory, on Windows Vista Home.

Figure 2 shows the average gaps over all instances of Taillard’s benchmark for the NEH-Delta and NEH-Alpha for varying k . The results for the NEH-Alpha are only given for $k \geq 500$, as Taillard’s benchmark contains instances with up to 500 jobs.

An interesting fact is that the NEH-Alpha performs better than the NEH-Delta for low values of k , but is then outperformed by the NEH-Delta for higher k . Figure 3 illustrates the execution times of the algorithms. The linear increase confirms the linear influence of the parameter k in the asymptotic time complexity of the algorithm. Though the execution time grows linearly in k , the gap improvement decreases in both methods.




Figure 2: Average gap over Taillard’s benchmark instances for the NEH-Delta and NEH-Alpha construction phase using varying k

Instance	lb	ub	Instance	lb	ub	Instance	lb	ub
ta051	3771	3847	ta081	6106	6202	ta101	11152	11181
ta052	3668	3704	ta083	6252	6271	ta102	11143	11203
ta053	3591	3640	ta084	6254	6269	ta107	11337	11360
ta054	3635	3719	ta085	6262	6314	ta108	11301	11334
ta055	3553	3610	ta086	6302	6364	ta109	11145	11192
ta057	3672	3704	ta087	6184	6268	ta110	11284	11288
ta058	3627	3691	ta088	6315	6401	ta111	26040	26059
ta059	3645	3741	ta089	6204	6275	ta112	26500	26520
ta060	3696	3756	ta090	6404	6434	ta116	26469	26477

Table 2: Upper and lower bounds for Taillard’s benchmark instances (<http://mistic.heig-vd.ch/taillard/>, May 2008) that have not been solved to optimality yet




Figure 3: Average execution time (in seconds) over Taillard’s benchmark instances for the NEH-Delta and NEH-Alpha construction phase using varying k

5 Improvement Phase

After keeping active k distinct nodes at each level of the enumeration tree, the application of the algorithms NEH-Alpha and NEH-Delta leads to a final set of k distinct solutions. From this set, the permutation schedule of minimum makespan is selected as the final solution. The central point that motivates these new algorithms relies on the expectation of obtaining better final solutions when increasing the value of parameter k . However, increasing the value of k beyond large values (for example 50.000) seems not to pay-off in relation to the computational efforts, as it can be observed in Figure 2. Clearly, it is always possible to achieve optimal solutions by increasing the values of k to $n!$. These results induced the development of new strategies to improve the permutation schedules generated by NEH-Delta and NEH-Alpha, using the knowledge of such threshold values for k to benefit better from the computational effort and focus on new techniques to enhance the quality of the solutions obtained by the construction phase.

5.1 Breadth Search Improvement Phase

As result of the construction phase, after applying either the NEH-Delta or the NEH-Alpha heuristic, a final permutation is selected from a set $S = \{s_1, s_2, \dots, s_k\}$ of candidate solutions. These candidate solutions are constructed in the last iteration of NEH-Delta or NEH-Alpha, after inserting the last job on the permutation schedules. The objective of the Breadth Search Improvement Phase (BSI) is to improve the solutions obtained from the construction phase by working over the whole set S . The strategy adopted by the BSI is composed of a sequence of l improvement iterations. Each of such improvement iterations is constituted of n basic steps. A basic step removes and reinserts the same job from every permutation in S . After removing a job, all n feasible positions in which it can be reinserted are considered. Hence, applying a basic step of a BSI improvement iteration

to a solution $s_i \in S$ creates n possible solutions, including one identical to s_i . However, not all n possible solutions are considered. In fact, from the set of all $n \cdot k$ (possibly new) solutions obtained by removing and reinserting the same job in every permutation of S only k of them with minimum makespan values are selected at the end of a basic step. All other solutions obtained are discarded. The job selected to be removed and reinserted on the i^{th} basic step of a BSI improvement iteration is the i^{th} job in the initial ordering of the classical NEH heuristic. Algorithm 3 outlines the pseudo-code of the method.

Algorithm 3: Breadth Search Improvement Phase

Input : Set $S = \{s_1, s_2, \dots, s_k\}$ of solutions obtained from construction phase,
Integer l ,
Set J of n jobs,
Set M of m machines,
Processing times matrix $T \in \mathbb{R}_{M \times J}^+$.

Output: permutation function $\pi : \{1, \dots, n\} \mapsto J$.

begin

Sort the jobs in set J in decreasing order of its processing time sums resulting in order (j_1, \dots, j_n) ;

$S^1 \leftarrow S$;

for each iteration $t = 1 \dots l$ **do**

$S' \leftarrow \emptyset$;

for each job position $q \in \{1, \dots, n\}$ **do**

for each solution $s \in S^t$ **do**

$s_{rem} \leftarrow$ remove job j_q from solution s ;

for each position $p \in \{1, \dots, n\}$ **do**

$s_{new} \leftarrow$ insert job j_q at position p in s_{rem} ;

if $|S'| < k$ **then**

$S' \leftarrow S' \cup \{s_{new}\}$;

else if $makespan(s_{new}) < \max_{s' \in S'} \{makespan(s')\}$ **then**

$S' \leftarrow S' - \{s'\}$;

$S' \leftarrow S' \cup \{s_{new}\}$;

$S^{t+1} \leftarrow S'$;

$\pi \leftarrow$ permutation of S^{l+1} with minimum makespan;

return π ;

end

Implementation details and time complexity analysis. Each of the l iterations of the BSI is composed of a set of n basic steps. A basic step comprises the deletion of a job from the permutation set, its reinsertion in all n possible positions of every partial permutation and the selection of k solutions to be considered. Hence, the execution time of a basic step is dominated by its n operations of reinsertion and makespan calculations. The deletion of a job can be executed in $O(n)$ time. A job can also be reinserted on a permutation's position p in $O(1)$ amortized time if reinsertions are always taken on subsequent positions. The makespan calculation after reinserting a job can be achieved

in $O(m)$ time using the dynamic programming algorithm of Taillard. The set S' can be maintained as a binary heap, bringing a cost of $O(\log k)$ to delete or insert a permutation in S' . As consequence, a basic step can be executed in $O(m + \log k)$ and the total execution time of the BSI is therefore $O(n^2 \cdot k \cdot l \cdot (m + \log k))$. In fact, this work applies two slightly distinct implementation strategies for the BSI improvement phase. The decision of which strategy should be applied is related to the choice of the algorithm used on the construction phase. In the case that NEH-Delta was applied, the BSI implementation strategy is identical to Algorithm 3. However, when NEH-Alpha is used there is a slight modification on the BSI method. Instead of having an unique set S of solutions for each iteration, k disjoint sets A^p of solutions are considered, following exactly the same idea as applied to the construction phase of NEH-Alpha.

5.2 Depth Search Improvement Phase

This improvement phase aims to improve separately each of the solutions given by the construction phase. As the previously presented improvement phase, this method receives an integer l , indicating the number of times (iterations) that all jobs will be removed and reinserted following the ordering given by the original NEH. In contrast to the BSI phase, the Depth Search Improvement Phase (DSI) inserts each job exclusively at the position that leads to the best makespan and uses the generated solution as starting point for the next job reinsertion. The procedure applied to each of the given starting solutions can be seen as a Local Search with a finite number of steps and a fixed order of the jobs whose neighborhood is examined. Algorithm 4 outlines the pseudo-code of the method.

Algorithm 4: Depth Search Improvement Phase

Input : Set $S = \{s_1, s_2, \dots, s_k\}$ of solutions obtained from construction phase,
Integer l ,
Set J of n jobs,
Set M of m machines,
Processing times matrix $T \in \mathbb{R}_{M \times J}^+$.

Output: permutation function $\pi : \{1, \dots, n\} \mapsto J$.

begin

Sort the jobs in set J in decreasing order of its processing time sums resulting in order (j_1, \dots, j_n) ;

for each solution $s_i \in S$ **do**

for each iteration $t = 1 \dots l$ **do**

for each job $j_i, i = 2 \dots n$ **do**

$s_{rem} \leftarrow$ remove job j_i from solution s_i ;

$s_i \leftarrow$ insert job j_i at its best position in s_{rem} ;

$\pi \leftarrow$ permutation of S with minimum makespan;

return π ;

end

Implementation details and time complexity analysis. For each of the k solutions from the set S given in the input data, the algorithm performs l iterations, i.e. the deletion

and reinsertion of all n jobs will be performed l times. Using Taillard’s dynamic programming algorithm to calculate the makespans for all possible insertion positions, the deletion and reinsertion of a job at the best position costs $O(n \cdot m)$. This leads to a total asymptotic runtime of $O(n^2 \cdot m \cdot k \cdot l)$ for this improvement method. In practice, we can skip further evaluation of a solution if it does not improve during one whole iteration, i.e. after removing and reinserting all n jobs. We also can prune further iterations $t', t' + 1, \dots, l$ on a solution, if the solution’s permutation is the same as one of the already examined ones from other solutions generated after the first $t' - 1$ iterations. An implementation of such redundancy verifications may change the overall time complexity. However, in our experiments, the use of such pruning turned out to be very effective and significantly reduced the execution time on all instances.

5.3 Computational experiments

Figure 4 illustrates the influence of the number of iterations during the improvement phase. Both NEH-Delta and NEH-Alpha (with a fixed $k = 500$) are followed by the Depth and Breadth Search Improvement phases. Two observations can be made. First, the gap improves similarly for each configuration of construction and improvement phase so that a ranking can be established. Second, the average gap seems to converge, i.e. one may limit the number of iterations without a strong influence on the final gap.




Figure 4: Gap improvement during 50 iterations of the two improvement phases combined with the NEH-Delta and NEH-Alpha for $k=500$ (average gaps over Taillard’s benchmark instances)

6 Performance Comparison

In this section, the competitiveness of the algorithms presented in this work is evaluated by comparing them with the most effective heuristics known for the PFSP. In particular, the following algorithms are considered: the NEH-T algorithm of Taillard (1990), the

NEH based algorithm from Kalczynski and Kamburowski (2008) (NEH-KK1), the NEH modification by Dong et al. (2008) (NEH-D), Ruiz and Maroto (2006)’ genetic algorithms GA_RMA and HGA_RMA, the Simulated Annealing algorithm of Osman and Potts (1989) (SA_OP), Widmer and Hertz (1989)’ Tabu Search SPIRIT, Chen et al. (1995)’s genetic algorithm (GA_CHEN), Reeves (1995)’ genetic algorithm (GA_REEV), Murata et al. (1996)’s hybrid genetic algorithm (GA_MIT), Aldowaisan and Allahverdi (2003)’s genetic algorithm (GA_AA) and Chandrasekharan and Ziegler (2004)’s Ant Colony algorithms (M-MMAS and PACO). Furthermore, we compare our results with the ones of Ruiz and Stützle (2007)’s Iterated Local Search (ILS), Iterated Greedy method (IG_RS) and Iterated Greedy with Local Search (IG_RSLs). Our algorithms were configured as follows:

- NEH-Delta and NEH-Alpha with Depth Search Improvement phase, denoted by NEH-Delta/DSI and NEH-Alpha/DSI respectively. The values for k were chosen as follows: $k = 10000$ for the instances with 20 jobs, $k = 5000$ for instances with 50 jobs, $k = 2000$ for instances with 100 jobs and $k = 1000$ for instances with 200 or 500 jobs. The improvement phase was limited to 20 iterations.
- NEH-Delta and NEH-Alpha with Breadth Search Improvement phase, denoted by NEH-Delta/BSI and NEH-Alpha/BSI respectively. The values for k were chosen as follows: $k = 5000$ for the instances with 20 jobs, $k = 2500$ for instances with 50 jobs, $k = 1000$ for instances with 100 jobs and $k = 500$ for instances with 200 or 500 jobs. The improvement phase was limited to 10 iterations.

The execution time for all algorithms was strictly limited to $(n \cdot m \cdot 3/100)$ seconds, the same stopping criterion as used by Ruiz and Stützle (2007) to compare their algorithm with other algorithms for the PFSP. If this time limit was exceeded, the best makespan found so far was considered.

In order to compare the results of this work with the ones of others, one must consider two crucial factors that complicate direct comparison. First, the upper bounds of Tailard’s benchmark instances were constantly updated in the last years. Hence, works based on this set of instances may have used different values of upper bounds to compute the gaps. Second, the performance of the computational resources may directly impact on the results when the experiments are time limited. Clearly, the fairest way to compare results of different works is to use the original implementation of the authors and execute the algorithms under the same conditions. If this is not the case, the above mentioned points must be considered. For algorithms whose results are not impacted by a time limit within their execution (i.e. the algorithms always terminate before the time limit is reached), the impact of computational resources does not have to be considered.

In our comparisons, these crucial factors were considered as best as possible. The NEH-T algorithm was implemented and executed on the same machine as our algorithms to guarantee a fair comparison. Table 3 shows the average gaps per instance group. All NEH-Delta and NEH-Alpha variations considerably improve the results when compared with the original NEH-T. Clearly, the computational effort for these algorithms are much higher.

We now compare with the other algorithms mentioned in the beginning of this section. Table 4 compares their average gaps. The original code of Ruiz and Stützle’s Iterated Greedy Algorithm with Local Search (IG_RSLs) was executed 20 times, yielding an average gap of 0.417%. In order to compare with the results for the other eleven algorithms

Instance	NEH-T	NEH-Delta /DSI	NEH-Alpha /DSI	NEH-Delta /BSI	NEH-Alpha /BSI
20x5	3.30	0.36	0.20	0.39	0.27
20x10	4.60	0.21	0.19	0.22	0.19
20x20	3.73	0.10	0.09	0.03	0.05
50x5	0.73	0.18	0.02	0.18	0.08
50x10	5.07	1.05	0.94	1.53	1.38
50x20	6.68	1.64	1.77	1.89	1.79
100x5	0.53	0.11	0.06	0.21	0.03
100x10	2.21	0.39	0.35	1.03	0.57
100x20	5.34	1.99	1.94	2.28	2.19
200x10	1.26	0.26	0.25	0.47	0.38
200x20	4.42	1.83	1.60	2.30	2.06
500x20	2.07	0.80	0.85	1.22	1.05
Average	3.33	0.74	0.69	0.98	0.84

Table 3: Comparison of the average gaps of the NEH-Delta and NEH-Alpha with the original NEH-T

listed by Ruiz and Stützle (2007), we estimated the difference within the sets of upper bounds and computational resources as explained in the following.

The similarity between sets of upper bound values used in different works is measured by using the average gap of the NEH-T as a reference. Ruiz and Stützle as well as Kalczynski and Kamburowski reported a NEH-T average gap that is very close to the average gap found in the experiments of this work. Thus, it is assumed that the set of upper bounds used by the above authors is similar to the one used in our experiments. Dong et al. seem to have worked with an older set of upper bounds and reported an NEH-T average gap of 2.74%, which is 21.5% less than the average we found. Hence, the gap reported for their algorithm was multiplied by this factor.

Now, possible differences in the computational resources must be considered. Kalczynski and Kamburowski’s NEH-KK1 algorithm does not exceed the given time limit, so it is valid to compare their results directly with the ones of our algorithms. Ruiz and Stützle report their results based on the above explained time limit. The average gap reported for their IG_RSLs algorithm is 0.44%, a difference of exactly 5.227% to the gap we found for this algorithm executing it on our machine. In order to have an estimated average gap for the other algorithms used in the comparisons of Ruiz and Stützle, we discount this 5.227% from the average gaps of these algorithms. Such values are marked as “estimated” in the table.

All algorithms in Table 4, except the NEH-T, NEH-KK1, NEH-D and our algorithms, are metaheuristics, i.e. are non-polynomial and non-deterministic. The metaheuristics of Ruiz and Stützle (IG_RSLs, IG_RS), Ruiz (HGA_RMA) and Chandrasekharan and Ziegler (PACO, M-MMAS) perform extremely well on the set of instances, leading to an average gap lower than one percent. While the NEH-T, NEH-KK1 and NEH-D reach average gaps around three percent, our NEH-Delta and NEH-Alpha approaches lead to gaps lower than one percent, competing well with the other leading metaheuristics.

Algorithm	Avg gap
IG_RSLs	0.42
HGA_RMA	0.54 (estimated value)
NEH-Alpha/DSI	0.69
PACO	0.71 (estimated value)
IG_RS	0.74 (estimated value)
NEH-Delta/DSI	0.74
M-MMAS	0.84 (estimated value)
NEH-Alpha/BSI	0.84
NEH-Delta/BSI	0.98
ILS	1.01 (estimated value)
GA_RMA	1.07 (estimated value)
GA_REEV	1.52 (estimated value)
GA_AA	2.16 (estimated value)
SA_OP	2.24 (estimated value)
GA_MIT	2.30 (estimated value)
NEH-D	2.87 (estimated value)
NEH-KK1	3.15
NEH-T	3.33
GA_CHEN	4.57
SPIRIT	4.83 (estimated value)

Table 4: Comparison of the average gaps of the NEH-Delta and NEH-Alpha with other heuristics and metaheuristics by time limited execution

Farahmand et al. (2009) presented five polynomial deterministic heuristics and Ruiz and Maroto (2005) reviewed further 13 polynomial-time deterministic heuristics for the PFSP. None of those algorithms reported an average gap smaller than the ones listed above for the heuristics introduced in this work, considering Taillard’s benchmark instances.

7 Conclusions

This work introduces new deterministic heuristics for the PFSP based on extensions of the classical NEH algorithm. The development of the new proposed methods was motivated by an analysis of the partial solutions generated during the NEH construction process, exploring the quality of partial solutions close to the ones selected by the NEH heuristic and its relation to the nodes belonging to the enumeration tree of the PFSP. Following such analysis, the algorithm NEH-Delta was proposed with the objective of considering not only the best partial permutation schedules on a level of the enumeration tree, but a set with k promising solutions to be explored. The NEH-Alpha heuristic also preserves k partial solutions active, but, in order to increase the diversity of solutions, keeps it in n disjoint sets representing the position of the last inserted job. Computational experiments demonstrated a significant enhancement on the quality of the solutions generated by NEH-Delta and NEH-Alpha heuristics when compared to the results obtained from the classical NEH. However, it was observed that, for some threshold values of k , the increase of parameter k leads to little improvement on solutions in comparison with the increase of the total execution time. Inspired by this fact, there were proposed two strategies to comprise a phase of improvement of the solutions generated by NEH-Delta and NEH-Alpha, named Breadth Search Improvement (BSI) and Depth Search Improvement (DSI). NEH-Alpha/DSI, NEH-Alpha/BSI, NEH-Delta/DSI and NEH-Delta/BSI are deterministic heuristics, which can be executed in polynomial time taking parameters k and l as a polynomial function in n or m . Following the methodology introduced by Ruiz and Stützle (2007), computational experiments were carried out in order to compare the performance of the new stated methods with currently well performing heuristics for the PFSP.

Regarding the heuristics here proposed, it is interesting to observe that the best results for the construction phase alone is not confirmed when the improvement phase is added. NEH-Delta beats NEH-Alpha for large values of k in the construction phase. One possible explanation for this phenomenon is that the diversity of the solutions can be enforced by augmenting the number k of partial solutions, not requiring an explicit mechanism for this as NEH-Alpha provides. The comparison of the four new NEH heuristics, improvement phase included, says this is not necessarily true. It allows two observations. The first one is that for both search approaches, BSI and DSI, the NEH-Alpha was dominant. This may be due to the larger diversity imposed by the explicit mechanism which now, with the improvement phase, makes a difference. The second one is that the DSI (depth search) seems to be consistently better. Although the purpose of the application of the breadth search was to provide a more diversified search, it seems that selecting a whole new set of k best solutions at each step may, in fact, lead to a less diversified one, when comparing to equivalent sets of the BSI. In other words, assuming that good heuristics are the ones capable of diversifying and intensifying the search, the dominance of NEH-Alpha/BSI was somewhat expected.

Finally, the experimental results attest that NEH-Alpha/DSI, NEH-Alpha/BSI, NEH-

Delta/DSI and NEH-Delta/BSI stand among the most effective heuristics already proposed for the PFSP. In particular, from the best of our knowledge, following the experimental methodology introduced by Ruiz and Stützle, no polynomial time deterministic heuristics proposed so far lead to experimental results close to the ones obtained in this work.

References

- Aldowaisan T and Allahverdi A. New heuristics for no-wait flowshops to minimize makespan. *Computers and Operations Research* 2003. 30 (8):1219–1231.
- Bonney M C and Gundry S W. Solutions to the Constrained Flowshop Sequencing Problem. *Operational Research Quarterly* 1976. 27, No. 4, Part 1:869–883.
- Campbell H G, Dudek R A and Smith M L. A Heuristic Algorithm for the n Job, m Machine Sequencing Problem. *Management Science* 1970. 16:B630–B637.
- Chandrasekharan R and Ziegler H. Ant-colony algorithms for permutation flowshop scheduling to minimize makespan/total flowtime of jobs. *European Journal of Operational Research* 2004. 155 (2):426–438.
- Chen C L, Vempati V S and Aljaber N. An application of genetic algorithms for flow shop problems. *European Journal of Operational Research* 1995. 80:389–396.
- Dannenbring D G. An Evaluation of Flow Shop Sequencing Heuristics. *Management Science* 1977. 23:1174–1182.
- Dong X, Huang H and Chen P. An improved NEH-based heuristic for the permutation flowshop problem. *Computers and Operations Research* 2008. 35:3962 – 3968.
- Farahmand R S, Ruiz R and Boroojerdian N. New High Performing Heuristics for Minimizing Makespan in Permutation Flowshops. *Omega-International Journal of Management Science* 2009. 37:331–345.
- Framinan J M, Gupta J N D and Leisten R. A review and classification of heuristics for permutation flow-shop scheduling with makespan objective. *Journal of the Operational Research Society* 2004. 55:1243–1255.
- Framinan J M, Leisten R and Rajendran C. Different initial sequences for the heuristic of Nawaz, Ensore and Ham to minimize makespan, idletime or flowtime in the static permutation flowshop sequencing problem. *International Journal of Production Research* 2003. 41, Issue 1:121 – 148.
- Graham R L, Lawler E L, Lenstra J K and Kan A H G R. Optimisation and approximation in deterministic sequencing and scheduling: a survey. *Annals of Discrete Mathematics* 1979. 5:287–326.
- Gupta J N D. A Functional Heuristic Algorithm for the Flowshop Scheduling Problem. *Operational Research Quarterly (1970-1977)* 1971. 22, No. 1:39–47.
- Gupta J N D. Heuristic algorithms for multistage flowshop scheduling problem. *AIEE Transactions* 4 1972. 1972:11–18.

- Hejazi S R and Saghafian S. Flowshop-scheduling with makespan criterion: a review. *International Journal of Production Research* 2005. 43:2895–929.
- Hundal T S and Rajgopal J. An extension of Palmer' s heuristic for flowshop scheduling problem. *International Journal of Production Research* 1988. 26:1119–1124.
- Johnson D D, Garey M R and Sethi R. The complexity of flowshop and jobshop scheduling. *Mathematics of Operations Research* 1976. 1:117–29.
- Johnson S M. Optimal two- and three-stage production schedules with set-up times included. *Naval Res. Logist. Quart.* 1954. 1:61–68.
- Kalczynski P J and Kamburowski J. On the NEH heuristic for minimizing the makespan in permutation flowshops. *The International Journal of Management Science* 2005. v35 i1:53–60.
- Kalczynski P J and Kamburowski J. An improved NEH heuristic to minimize makespan in permutation flow shops. *Computers and Operations Research* 2008. 35(9):3001–3008. ISSN 0305-0548. doi:http://dx.doi.org/10.1016/j.cor.2007.01.020.
- Koulamas C. A new constructive heuristic for the flowshop scheduling problem. *European Journal of Operational Research* 1998. 105, Issue 1:66–71.
- Murata T, Ishibuchi H and Tanaka H. Genetic algorithms for flowshop scheduling problems. *Computers and Industrial Engineering* 1996. 30 (4):1061–1071.
- Nawaz M, Jr E E E and Ham I. A heuristic algorithm for the m-machine, n-job flow-shop sequencing problem. *Omega-International Journal of Management Science* 1983. 11:91–95.
- Nowicki E and Smutnicki C. New results in the worst-case analysis for flow-shop scheduling. *Discrete Appl. Math.* 1993. 46(1):21–41. ISSN 0166-218X. doi:http://dx.doi.org/10.1016/0166-218X(93)90156-I.
- Osman I and Potts C. Simulated annealing for permutation flow-shop scheduling. *Omega-International Journal of Management Science* 1989. 17 (6):551–557.
- Page E S. An Approach to the Scheduling of Jobs on Machines. *Journal of the Royal Statistical Society. Series B (Methodological)* 1961. 23, No. 2:484–492.
- Palmer D S. Sequencing Jobs Through a Multi-Stage Process in the Minimum Total Time – A Quick Method of Obtaining a Near Optimum. *OR* 1965. 16:101–107.
- Reeves C R. A genetic algorithm for flowshop sequencing. *Computers and Operations Research* 1995. 22 (1):5–13.
- Ruiz R and Maroto C. A comprehensive review and evaluation of permutation flowshop heuristics. *European Journal of Operational Research* 2005. 165:479–94.
- Ruiz R and Maroto C. A genetic algorithm for hybrid flowshops with sequence dependent setup times and machine eligibility. *European Journal of Operational Research* 2006. 169:781–800.

- Ruiz R and Stützle T. A simple and effective iterated greedy algorithm for the permutation flowshop scheduling problem. *European Journal of Operational Research* 2007. 177:2033–2049.
- Taillard E. Some efficient heuristic methods for the flow shop sequencing problem. *European Journal of Operational Research* 1990. 47 (1):67–74.
- Taillard E. Benchmarks for basic scheduling problems. *European Journal of Operational Research* 1993. 64 (2):278–285.
- Widmer M and Hertz A. A new heuristic method for the flow shop sequencing problem. *European Journal of Operational Research* 1989. 41 (2):186–193.