Parallelization of the scatter search for the p -median problem [☆]

Félix García-López, Belén Melián-Batista,
José A. Moreno-Pérez, J. Marcos Moreno-Vega ^{*}

Departamento de Estadística, I.O. y Computación, Universidad de La Laguna, La Laguna 38271, Spain

Received 7 June 2002; accepted 19 November 2002

Abstract

This article develops several strategies for the parallelization of the metaheuristic called scatter search, which is a population-based method that constructs solutions by combining others. Three types of parallelization have been proposed to achieve either an increase of efficiency or an increase of exploration. The procedures have been coded in C using OpenMP and compared in a shared memory machine with large instances. The obtained algorithms are tested on the p -median problem.

© 2003 Elsevier Science B.V. All rights reserved.

Keywords: Parallelization; Metaheuristics; Scatter search; p -median problem

1. Introduction

Scatter search (SS) [12,23] is a population-based metaheuristic that constructs solutions by combining others in a set of solutions, named reference set. The basic SS combines solutions of the reference set and runs a local search (LS) procedure to reach a local optimum. A LS consists of selecting a better neighbor solution until

[☆] This work has been partially supported by the Spanish Ministerio de Ciencia y Tecnología under grant TIC2002-04242-C03-01, 70% with FEDER funds.

^{*} Corresponding author.

E-mail addresses: fegarcia@ull.es (F. García-López), mbmelian@ull.es (B. Melián-Batista), jamoreno@ull.es (J.A. Moreno-Pérez), jmmoreno@ull.es (J. Marcos Moreno-Vega).

no such solution exists. The reference set is updated depending on the results of the improvements.

In this paper we propose the parallelization of the SS metaheuristic to achieve either an increase of efficiency or an increase of exploration. In any parallelization some steps of the algorithm are distributed among the available processors. The increase of the efficiency is obtained by performing the same steps in less time than the sequential algorithm, whereas the increase of the exploration is obtained by performing more steps in the same time as the sequential algorithm. Several strategies for parallelizing a SS algorithm are analyzed. We test them with large instances of the p -median problem obtained from the TSPLIB [30]. The corresponding parallel algorithms are coded in C using OpenMP, a model for parallel programming portable across shared memory architectures. OpenMP [29] is based on a combination of compiler directives, library routines and environment variables that are used to specify shared memory parallelism in Fortran and C/C++ programs.

The p -median problem is a discrete location/allocation problem. The objective is to select, from a discrete candidate set of facility points, p locations for the facilities such that the sum of the distances from a set of users to the chosen facility points is minimized. There are not capacities on the facilities. The problem was first formulated by Hakimi [13,14]. It is a prototype of a wide class of hard combinatorial problems, named p -selection problems, where the solutions are generated by selecting p items from a finite universe. The evaluation of the objective function in p -selection problems ranges from a simple calculation of a function value in a small constant time to solve another hard problem or even to perform a simulation process. The standard moves for p -selection problems are the interchange or swap moves that exchange an item in the solution with another one out of the solution.

Next section describes the basics of the SS metaheuristic. In Section 3 we provide a brief overview of logistic problems in general and of the p -median problem, in particular, as a prototype of the class of fixed-size location/allocation problems. In Section 4 we discuss briefly the application of the SS to the p -median problem. Several parallelization strategies are analyzed in Section 5. Computational results and concluding remarks are given in Sections 6 and 7, respectively.

2. The scatter search metaheuristic

SS is a population-based metaheuristic that uses a reference set to combine its solutions and construct others. The method generates a reference set from a population of solutions. Then a subset is selected from this reference set. The selected solutions are combined to get starting solutions to run an improvement procedure. The result of the improvement can motivate the updating of the reference set and even the updating of the population of solutions.

The initial population must be a wide set of disperse solutions. However, it must also include good solutions. Several strategies can be applied to get a population with these properties. The solutions for the population can be created, for instance, by using a random procedure to achieve a certain level of diversity. Then a

simple improvement heuristic procedure must be applied to these solutions in order to get better solutions. The initial population can also be obtained by a procedure that provides at the same time disperse and good solutions like GRASP procedures [6].

A set of good representative solutions of the population is chosen to generate the reference set. The good solutions are not limited to those with the best objective values. By good representative solutions we mean solutions with the best objective values as well as disperse solutions. Disperse solutions should reach different local minima by the local search. Indeed, a solution may be added to the reference set if the diversity of the set improves. So the reference set must consist of a set of disperse and good solutions selected from the populations. The criteria for updating the reference set, when necessary, must be based on comparisons and measures of diversity between the new solutions and the existing solutions.

A subset of solutions from the reference set is selected for applications of a combination method to get good starting solutions for an improvement procedure. In general, the method consists in selecting all the subsets of a fixed size. The combination procedure tries to combine good characteristics of the selected solutions to get new current solutions. The purpose is to get better solutions which are not similar to those already in the reference set.

The possible improvement solution methods applied to the solutions range from simplest local searches to a very specialized search. A very simple procedure is a local search based on basic moves consisting of selecting the best improving move or a first found improving move. The procedure must allow to use tools like recent or intermediate memory, variable neighborhoods, or hashing scanning methods of the neighborhood. Then the method applied could be a Tabu Search [10,11], a Variable Neighborhood Search [16,17] or any sophisticated hybrid heuristic search.

The metaheuristic strategy includes the decision on how to update the reference set taking into account the state of the search. The algorithm must also realize when the reference set does not change and seek to diversify the search by generating a new set of solutions for the population. In addition, the metaheuristic includes the stopping criterion for the whole search procedure. Then the best solution used in the reference set is provided by the method.

Usual stopping conditions are based on allowing a total maximum computational time or a maximum computational effort since the last improvement. The computational effort is measured by the number of iterations, number of local searches or real time.

A high level pseudocode of the sequential scatter search (SSS) is showed in Fig. 1.

Therefore, the SS strategy involves six procedures and three stopping criteria to solve an optimization problem. The procedures are the following:

- (1) *The initial population creation method.* The procedure *CreatePopulation* creates a random initial population *Pop* of good and disperse *PopSize* solutions.
- (2) *The reference set generation method.* The procedure *GenerateReferenceSet* selects *RefSetSize* of the best representative solutions in the population *Pop* to be included in the set *RefSet*.

```

Procedure Sequential Scatter Search
begin
repeat
  CreatePopulation(Pop,PopSize);
repeat
  GenerateReferenceSet(RefSet,RefSetSize);
repeat
  SelectSubset(SubSet,SubSetSize);
  CombineSolutions(SubSet, CurSol);
  ImproveSolution(CurSol, ImpSol);
until (StoppingCriterion1);
UpdateReferenceSet(RefSet);
until (StoppingCriterion2);
until (StoppingCriterion3);

```

Fig. 1. The scatter search metaheuristic.

- (3) *The subset generation method.* The procedure *SelectSubset* generates subsets *SubSet*, which consists of *SubSetSize* good solutions in *RefSet*, to apply the next combination procedure.
- (4) *The solution combination method.* The procedure *CombineSolutions*, provided with parameters used to modulate the intensification and/or diversification, combines the solutions in *SubSet* to get the new current solution *CurSol*.
- (5) *The improvement solution method.* The procedure *ImproveSolution*, provided with parameters to modulate the specialization of the method, improves the current solution *CurSol* to get an improved solution *ImpSol*.
- (6) *The reference set updating method.* The procedure *UpdateReferenceSet* updates the reference set by deciding when and how the obtained improved solutions are included in the reference set replacing some solutions already in it.

In addition to these six procedures, the metaheuristic involves three stopping procedures that implement the criteria to decide when to go to an above step.

- (1) *New reference set criterion.* The first criterion (*StoppingCriterion1*) decides when to generate a new reference set from the population.
- (2) *New population criterion.* The second criterion (*StoppingCriterion2*) decides when to generate a new initial population.
- (3) *Termination criterion.* Finally, the third criterion (*StoppingCriterion3*) decides when to stop the whole search.

3. Selection problems

The selection problems are those problems that can be solved by choosing an optimal set of items from a universe. A generic selection problem consists in choosing the set of items that minimizes a cost function subject to some constraints. The eva-

valuation of some real objectives means solving another hard problem or running a simulation procedure. Also the set of constraints can be so restrictive such that only one selection is feasible or that all the selections of items are feasible. A fixed size, say p , selection problem, named a p -selection problem, consists in choosing the set of p items that minimizes a cost function subject to some constraints; i.e. a selection problem where all the feasible solutions have size p .

Several of the most relevant combinatorial optimization problems can be formulated as p -selection problems for an appropriate p . Among them are the most famous problems in Combinatorial Optimization: the Travelling Salesman Problem, Knapsack Problem, p -Median Problem, Spanning Tree Problem, Shortest Path Problem, Network Flow Problem, Matching Problem and Steiner Problem. Even most of the special versions of these problems can also be formulated as selection problems. Also the Linear Programming Problem consists in the optimal selection of a base, a set of points of which the size is known.

The p -facility location-allocation problems constitute a wide class of logistic problems. Consider a set L of m potential locations for p facilities and a set U of locations of n given users. A p -facility location-allocation problem consists of locating simultaneously p facilities at locations of L and allocate every user u to a facility point in order to minimize a cost function that represents the resource used for serving the demand of every user from the corresponding facility point. One of the most relevant p -facility location problems without capacity limit in the facilities is the p -median problem. The p -median problem consists of locating p facilities in order to minimize the sum of distances from every user to its closest facility.

Given the set $L = \{v_1, v_2, \dots, v_m\}$ of potential locations for the facilities (or location points), and the set $U = \{u_1, u_2, \dots, u_n\}$ of users (or customers, or demand points), the entries of an $n \times m$ matrix $D = (d_{ij})_{n \times m} = (\text{Dist}(u_i, v_j))_{n \times m}$ give the distances travelled (or costs incurred) for satisfying the demand of the user located at u_i from the facility located at v_j , for all $v_j \in L$ and $u_i \in U$. The objective of the p -median problem is to minimize the sum of these distances (or transportation costs), i.e.,

$$\text{minimize} \sum_{u_i \in U} \min_{v_j \in X} \text{Dist}(u_i, v_j)$$

where $X \subseteq L$ and $|X| = p$.

The p -median problem is $\mathcal{NP}$ -hard [9,21]. Many heuristics and exact methods have been proposed to solve it. Exact algorithms are provided by Beasley [1], Cornuejols et al. [3], Galvão [7], Hanjoul and Peeters [19] and Rosing et al. [34], among others. Classical heuristics for this problem often cited in the literature are *Greedy* by Kuehn and Hamburger [22], *Alternate* by Maranzana [26] and *Interchange* by Teitz and Bart [39]. Heuristics based on linear programming have been proposed by du Merle et al. [4] and by Senne and Lorena [37,38]. Another type of heuristics suggested in the literature is based on the relaxed dual of the integer programming formulation of the p -median problem found in the papers by Erlenkotter [5], Galvão [7] and Nauss and Markland [28]. A 1-interchange move is extended into a so-called 1-chain-substitution move, which is applied to p -median problem by Maldenović et al. [27]. Tabu Search heuristics are suggested by Voss [40], where some variants of the

so-called reverse elimination method are discussed and by Rolland et al. [36]. Other heuristic method applied to the p -median problem are the Heuristic Concentration by Rosing and Revelle [35] and the GRASP with path-relinking Resende and Werneck [31]. The Variable Neighborhood Search heuristic and its variants have also been applied to the p -median problems by Hansen and Mladenović [15] and Hansen, et al. [18]. In addition, several hybrids of these heuristics have been suggested. Several parallelization methods of the Variable Neighborhood Search are considered for the p -median problem by García al. [8]. Other heuristic procedures have been proposed by Hodgson [20], Resende and Werneck [32], Rosing [33] and Whitaker [41].

4. Application of SS to the p -median problem

For most instances, the set of potential locations for the facilities and the set of locations of the users coincide, so that $L = U$ and $m = n$. In this case, a solution for the p -median problem consists of selecting a set X of p points from U to locate the facilities. The solution is evaluated by a cost function which is the sum of the distances from the users to the points in the solution. This cost function is given by:

$$\text{Cost}(X) = \sum_{u \in U} \min_{v \in X} \text{Dist}(u, v)$$

Regardless the search technique employed, it is necessary to specify a solution coding, which encodes alternative candidate solutions for manipulation. The choice of the coding that provides an efficient way of implementing the moves and evaluating the solutions is essential for the success of the heuristic. Every solution X is encoded by arranging all the points of U in an array $[v_i : i = 1, \dots, n]$ where v_i is a point in X for $i \leq p$, and it is a point out of X for $i > p$.

4.1. Initial population creation method

The initial population creation method must be designed in order to get a random set of disperse and good solutions. The method we use has several phases. The first phase starts by dividing the set L in several sets. A constructive method to get a good solution consists of starting by selecting an arbitrary initial point u of L and performing $p - 1$ times the following operation: select the farthest facility point to the already selected points. We perform this constructive method for each set of the partition of L . Then we get a solution for each set of the partition. However, since the constructed solution depends on the starting point, we apply it from several starting points to get different solutions. We also apply an improvement procedure to these solutions.

For evaluating the dispersion among solutions, we need to consider a distance between them. This distance is defined using the same objective function. Let $f_Y(X)$ be the objective function for the set of users in Y :

$$f_Y(X) = \sum_{v \in Y} \min_{u \in X} \text{Dist}(u, v)$$

The distance between two solutions X and Y is $\text{Dist}(X, Y) = f_Y(X) + f_X(Y)$.

We use two parameters to modulate the dispersion in the reference set: α is the proportion of solutions of the population that are obtained improving disperse solutions, and β is to modulate, for the remainder solutions, the tradeoff between cost and dispersion. Given the values, previously fixed, of the population size $PopSize$ and the parameters α and β , the population is obtained as follows. We use the above method to get $\lfloor \alpha PopSize \rfloor$ solutions for the population. The remainder solutions up to $PopSize$ are obtained by an scoring procedure. For each solution X we define the score by:

$$h(X) = \text{Cost}(X) - \beta \text{Dist}(X, Pop)$$

where $\text{Dist}(X, P) = \min_{Y \in P} \text{Dist}(X, Y)$. The $\lceil (1 - \alpha)PopSize \rceil$ best solutions according $h(\cdot)$ are included in the population Pop .

4.2. Reference set generation method

The reference set consists of a set of good and disperse solutions selected from the population. The generation of a reference set is done by selecting $RefSetSize_1$ solutions from the best solutions and $RefSetSize_2$ disperse solutions taking into account the above $RefSetSize_1$ solutions ($RefSetSize = RefSetSize_1 + RefSetSize_2$). After including the best $RefSetSize_1$ solutions in $RefSet$, we iteratively include in the $RefSet$ the farthest solution from the solutions already in $RefSet$, repeating this procedure $RefSetSize_2$ times.

4.3. Subset generation method

The selection of a subset for applying the combination usually consists in selecting all the subsets of fixed size r . For our computational experience we use $r = 2$. However, in order to avoid repetitions of combinations when some solutions of the reference set do not vary, we keep information on the combinations performed.

4.4. Combination method

Given the subset of selected solutions from the reference set, the combination method tries to get a solution with the good characteristics of these solutions. First of all, we select the points that are in all these solutions; let X be the set of these points. For every user point u let

$$L(u) = \{v \in L : \text{Dist}(u, v) \leq \beta \text{Dist}_{\max}\}$$

where

$$\text{Dist}_{\max} = \max_{u, v \in L} \text{Dist}(u, v)$$

Choose the point $u^* \in L$ such that $\text{Dist}(X, u^*) = \max_{u \in L} \text{Dist}(X, u)$ and select at random a point $v \in L(u^*)$ that is included in X . This step is iteratively applied until X has size p .

```

Procedure Local Search ( $X$ )
begin
  repeat
    Set  $X' \leftarrow X$ 
    For every  $i$  and  $j$  with  $1 < i \leq p$  and  $p < j \leq n$  do
      Set  $X'_{ij} \leftarrow X' - \{v_i\} + \{v_j\}$ 
      If  $Cost(X'_{ij}) < Cost(X)$  then  $X \leftarrow X'_{ij}$ 
    Until  $Cost(X) = Cost(X')$ 
  end.

```

Fig. 2. Local search.

4.5. Improvement solution method

The improvement solution method is the a local search based on the base moves. A local search method for combinatorial optimization performs a series of moves in solution space, which improve each time the value of the objective function until a local optimum is found. The base moves for local searches for the p -median problem are the swap or interchange moves. For every solution X , given an element v_i in the solution and an element v_j not in the solution, the interchange move consists of dropping v_i from X and adding v_j to X . Let X_{ij} denote the solution obtained from X by interchanging v_i by v_j , where $v_i \in X$ and $v_j \in L - X$. Each solution X has a neighborhood associated $\mathcal{N}(X)$, that consists of the solutions X_{ij} that are reached from X by the base move. At each iteration, the local search procedure obtains an improved solution X' in the neighborhood $\mathcal{N}(X)$ of the current solution X , until no further improvement is found. The local search is implemented by choosing each time the best possible move among all interchange moves.

The pseudocode of the local search is given in Fig. 2.

5. Parallelization methods

The main purpose of parallel processing is to perform computations faster than with a single processor by using a number of processors concurrently. This pursuit has had a tremendous influence on almost all the activities related to computing. The need for faster solutions and for solving larger-size problems arises in a wide variety of applications. In general, the parallel computation can be used to increase the size of the problems that can be solved, to speed up the computations and to attempt a more thorough exploration of the solution space.

In this paper, we analyze three different parallelization strategies for the SS algorithm. The first two strategies reduce the running time of the procedure. The first parallelization reduces the running time of every local search algorithm and the second parallelization performs the local searches in parallel from the results of the combinations. The last parallelization increases the exploration in the solution space by running the SS in parallel for several populations.

5.1. Synchronous parallel scatter search

In the SSS algorithm, in every iteration, the most time consuming part is the local search. Then we propose a synchronous algorithm that enables solving, in parallel, the local searches. We denote the synchronous parallel scatter search algorithm that parallelizes the local search by synchronous parallel scatter search (SPSS). The following template (Fig. 3) shows the pseudocode of the SPSS with n_pr processors. Note that this pseudocode is obtained from the SSS by replacing the local search by the parallel version of the local search. In SPSS, the neighborhood, $\{X_{ij} : 1 \leq i \leq p, p < j \leq n\}$, is divided in n_pr subsets I_r . These subsets are assigned to the processors and each returns an improving neighbor in its subset of the neighborhood. The best solution among the neighbors provided by the processors is chosen as the current solution. This strategy is a low-level parallelism.

5.2. Replicated combination scatter search

We propose another parallel algorithm that also reduces the running time of the procedure. The procedure is parallelized by selecting several subsets from the reference

```

Procedure SPSS
begin
repeat
  CreatePopulation(Pop,PopSize);
  repeat
    GenerateReferenceSet(RefSet,RefSetSize);
    repeat
      SelectSubset(SubSet,SubSetSize);
      CombineSolutions(SubSet, CurSol);
      (* ImproveSolution(CurSol, ImpSol) *)
      Take  $I \leftarrow \{(i, j) : 1 \leq i \leq p, p < j \leq n\}$ .
      Divide  $I$  in  $n\_pr$  subsets  $I_r, r = 1, \dots, n\_pr$ .
      repeat
        Set  $ImpSol \leftarrow CurSol$ 
        For each processor  $r = 1, \dots, n\_pr$ , do in parallel
          Set  $X_r \leftarrow CurSol$ ;
          For every  $(i, j) \in I_r$  do
            Set  $X_{ij} \leftarrow CurSol - \{v_i\} + \{v_j\}$ ;
            If  $Cost(X_{ij}) < Cost(X_r)$ 
              then  $X_r \leftarrow X_{ij}$ ;
          For  $r = 1, \dots, n\_pr$  do
            If  $Cost(X_r) < Cost(CurSol)$ 
              then  $CurSol \leftarrow X_r$ ;
        Until  $Cost(ImpSol) = Cost(CurSol)$ ;
      until (StoppingCriterion1);
    UpdateReferenceSet(RefSet);
  until (StoppingCriterion2);
until (StoppingCriterion3);

```

Fig. 3. Synchronous parallel scatter search.

Procedure RCSS

```

begin
repeat
  CreatePopulation(Pop,PopSize);
  repeat
    GenerateReferenceSet(RefSet,RefSetSize);
    repeat
      For each processor  $r = 1, \dots, n_{pr}$ , do in parallel
        SelectSubset(SubSetr,SubSetSize);
        CombineSolutions(SubSetr, CurSolr);
        ImproveSolution(CurSolr, ImpSolr);
      until (StoppingCriterion1);
      UpdateReferenceSet(RefSet);
    until (StoppingCriterion2);
  until (StoppingCriterion3);

```

Fig. 4. Replicated-combination parallel scatter search.

set that are combined and improved by the processors. These procedures are replicated as many times as the number of available processors. The local optima found by the processors are used to update the reference set. This method is the replicated-combination scatter search (RCSS) that is described in Fig. 4.

5.3. Replicated parallel scatter search

The replicated parallel scatter search (RPSS) consists of a multistart search where the local searches are replaced by SSs using different populations that run on the parallel processors (see Fig. 5). The RPSS parallelization method corresponds to a natural parallelization of the hybrid metaheuristic between a SS and a multistart search.

Procedure RPSS

```

begin
repeat
  Foreach processor  $r = 1, \dots, n_{pr}$ , do in parallel
    CreatePopulation(Pop,PopSize);
  repeat
    GenerateReferenceSet(RefSetr,RefSetSize);
    repeat
      SelectSubset(SubSetr,SubSetSize);
      CombineSolutions(SubSetr, CurSol);
      ImproveSolution(CurSolr, ImpSolr);
    until (StoppingCriterion1);
    UpdateReferenceSet(RefSetr);
  until (StoppingCriterion2);
until (StoppingCriterion3);

```

Fig. 5. The replicated parallel scatter search.

6. Computational results

The algorithms were coded in C and OdinMp [2] (OdinMp is a free and portable implementation of OpenMP for ANSI C), and tested with large instances of the p -median problem. For each instance and each one of the three algorithms we did a single run. The distance matrix was taken from the instance TSPLIB RL1400 that includes 1400 points. The sets of instances are characterized with the number n of points (1400) and the number p of facility points or medians that is reported in first column of Tables 1–3 going from 10 to 100. Some computational results on these instances of the problem have been reported in [8,18], where several heuristics were compared. The algorithms run on the machine *TEIDE* (8 processors ALPHA at 466-MHz, with 2 GB and O.S. DIGITAL UNIX 4.0C) of the *University of La Laguna*.

In Table 1, we report the results for the SPSS algorithm, which runs as many times as the number of available processors (1, 2, 4 and 8, respectively). The objective value, real time in seconds and speed-up are presented for each number of processors. The real time is used to obtain the time reduction provided by the parallel algorithm SPSS. Speed-up is the ratio of sequential time to parallel time to solve a particular problem on a given machine. It is given by:

$$\text{speed-up} = \frac{\text{Time to solve a problem with the sequential code}}{\text{Time to solve the same problem with the parallel code on } n\text{-pr processors}}$$

The results in Table 1 show that, in general, the speed-up increases with the number of processors. For two processors, the time is reduced in all the cases. For 4 processors and $p = 80$, we observed a detrimental anomaly [24,25]; i.e., the real time is greater than the real time for 2 processors. We also observed that, in some cases, using 4 and 8 processing elements, an acceleration anomaly manifests itself in the form of a speed-up greater than 4 and 8, respectively.

This kind of anomalous behaviour is related to the search space to be analysed by the algorithm. Note that for the anomalous cases the objective values change for different number of processors. For the detrimental anomaly with $p = 80$, we get a better objective value with 4 and 8 processors than with 2 processors, which indicates that the search space for 4 and 8 processors is greater than for 2 processors. In the speedup anomaly with $p = 90$ the opposite occurs, the objective value for 4 processors is greater than for two processors, which indicates a smaller search tree. For the case of $p = 40$, the objective values do not change but it does not mean that the search space does not change. This algorithm is strongly affected by the non-determinism intrinsic of the concurrent programming.

Table 2 summarizes the results for the RCSS algorithm, where the headings are the same as in Table 1. The speed-up reached is smaller than the speed-up for the SPSS algorithm, but anyway the speed-up increases with the number of processors for the RCSS algorithm. The objective values obtained with both algorithms are similar. The real times for the SPSS are better than the real times for the RCSS. This fact occurs because during the execution of the RCSS algorithm, the needed barrier synchronizations produce a higher sequentiality than in the SPSS algorithm.

Table 1
Synchronous parallel scatter search

p	1 processor			2 processors			4 processors			8 processors		
	Objective	Time	S-u	Objective	Time	S-u	Objective	Time	S-u	Objective	Time	S-u
10	101 249.47	380	1	101 249.47	221	1.72	101 249.47	137	2.77	101 249.47	88	4.32
20	57 857.55	682	1	57 857.55	399	1.71	57 857.55	260	2.62	57 857.55	161	4.24
30	44 057.55	1245	1	44 057.55	718	1.73	44 065.85	402	3.10	44 084.10	271	4.59
40	35 002.02	1858	1	35 002.02	1047	1.77	35 002.02	461	4.03	35 002.02	283	6.57
50	29 089.71	1443	1	29 089.71	847	1.70	29 089.71	612	2.36	29 089.71	311	4.64
60	25 185.79	1930	1	25 185.79	1119	1.72	25 186.24	634	3.04	25 167.84	628	3.07
70	22 125.46	1864	1	22 125.46	1088	1.71	22 125.46	649	2.87	22 125.46	416	4.48
80	19 884.51	1794	1	19 884.51	1049	1.71	19 870.51	1085	1.65	19 870.51	471	3.81
90	17 987.91	4168	1	17 987.91	2322	1.80	18 006.23	780	5.34	18 002.35	488	8.54
100	16 563.93	3341	1	16 563.93	1912	1.75	16 551.68	1266	2.64	16 554.08	520	6.64

Table 2
Replicated combination scatter search

p	1 processor			2 processors			4 processors			8 processors		
	Objective	Time	S-u	Objective	Time	S-u	Objective	Time	S-u	Objective	Time	S-u
10	101 249.47	380	1	101 249.47	237	1.60	101 249.47	145	2.62	101 249.47	103	3.69
20	57 857.55	682	1	57 857.55	461	1.48	57 857.55	307	2.22	57 857.55	193	3.53
30	44 057.55	1245	1	44 057.55	806	1.54	44 023.23	667	1.87	44 023.23	687	1.81
40	35 002.02	1858	1	35 002.02	1173	1.58	35 002.02	578	3.21	35 002.02	557	3.34
50	29 089.71	1443	1	29 089.71	965	1.50	29 089.71	953	1.51	29 089.71	410	3.52
60	25 185.79	1930	1	25 185.79	1240	1.56	25 175.01	760	2.54	25 167.82	433	4.46
70	22 125.46	1864	1	22 125.46	1263	1.48	22 125.46	828	2.25	22 125.46	486	3.84
80	19 884.51	1794	1	19 884.51	1214	1.48	19 872.28	1015	1.77	19 884.52	519	3.46
90	17 987.91	4168	1	17 987.91	2871	1.45	17 987.91	1389	3.00	17 987.91	537	7.76
100	16 563.93	3341	1	16 563.93	2285	1.46	16 561.58	1113	3.00	16 553.70	827	4.04

Table 3
Replicated parallel scatter search

p	Best known	2 processors		4 processors		8 processors	
		Objective	Time	Objective	Time	Objective	Time
10	101 249.47	101 249.47	428	101 249.47	438	101 249.47	424
20	57 857.55	57 857.55	787	57 857.55	791	57 857.55	779
30	44 013.02	44 013.02	1371	44 013.02	1962	44 013.02	1634
40	35 002.02	35 002.02	2048	35 002.02	2058	35 002.02	1750
50	29 089.71	29 089.71	1667	29 089.71	1672	29 089.71	1605
60	25 160.40	25 160.40	2067	25 160.40	2124	25 160.40	2382
70	22 125.46	22 125.46	2035	22 125.46	2044	22 125.46	3040
80	19 870.29	19 870.29	3557	19 870.29	3591	19 870.29	2796
90	17 987.94	17 987.91	22 359	17 987.91	4730	17 987.91	12 742
100	16 551.20	16 556.58	3477	16 552.48	4076	16 552.48	3838

Table 3 shows the results for the RPSS algorithm. The first column shows the best known objective value. For each number of processors 2, 4 and 8 we report the objective values and real times for the RPSS algorithm. The best objective values among the three algorithms are found by the RPSS. The inexistence of barrier synchronizations allows the fact that the real computation times are similar for different number of processors. The real time reported in this table is the maximum of the computational time of the processors.

7. Conclusions

The combination of SS and parallelism provides a useful tool to solve hard problems. SS is parallelized in several ways. SPSS is obtained by parallelizing the local search. By parallelizing the combinations of solutions we get the RCSS, where the set of possible combinations is divided among the available processors and solved in parallel. The RPSS parallelizes the whole procedure and each processor runs in parallel a SS.

The objective values found with these algorithms are comparable with the best obtained in the literature [8,18]. The SPSS and RCSS algorithms reduced the computational time properly. RPSS increases the diversification. Note that it is possible to increase the intensification by sharing the best solution reached by the processors. Other possible strategy for increasing the diversification is to apply several combination methods in parallel.

References

- [1] J.E. Beasley, A note on solving large p -median problems, *European Journal of Operational Research* 21 (1985) 270–273.
- [2] C. Brunschen, M. Brorsson, OdinMP/CCp-A Portable Implementation of OpenMP for C, in: *Proceedings of First European Workshop on OpenMP EWOMP99*, 1999, pp. 21–26.

- [3] G. Cornuejols, M.L. Fisher, G.L. Nemhauser, Location of bank accounts to optimize float: an analytical study of exact and approximate algorithms, *Management Science* 23 (1977) 789–810.
- [4] O. du Merle, D. Villeneuve, J. Desrosiers, P. Hansen, Stabilized column generation, *Discrete Mathematics* 194 (1999) 229–237.
- [5] D. Erlenkotter, A dual-based procedure for uncapacitated facility location, *Operations Research* 26 (1978) 992–1009.
- [6] T.A. Feo, M.G.C. Resende, Greedy randomized adaptive search procedures, *Journal of Global Optimization* 6 (1995) 109–133.
- [7] R.D. Galvão, A dual-bounded algorithm for the p -median problem, *Operations Research* 28 (1980) 1112–1121.
- [8] F. García-López, B. Melián-Batista, J.A. Moreno-Pérez, J.M. Moreno-Vega, The parallel variable neighborhood search for the p -median problem, *Journal of Heuristics* 8 (2002) 375–388.
- [9] M.R. Garey, D.S. Johnson, *Computers and intractability: a guide to the theory of NP-completeness*, W.H. Freeman and Co, San Francisco, 1979.
- [10] F. Glover, Tabu search-part I, *ORSA Journal on Computing* 1 (1989) 190–206.
- [11] F. Glover, Tabu search-part II, *ORSA Journal on Computing* 2 (1990) 4–32.
- [12] F. Glover, M. Laguna, R. Martí, Fundamentals of scatter search and path relinking, *Control and Cybernetics* 39 (3) (2000) 653–684.
- [13] S.L. Hakimi, Optimum location of switching centers and the absolute centers and the medians of a graph, *Operations Research* 12 (1964) 450–459.
- [14] S.L. Hakimi, Optimum distribution of switching centers in a communication network and some related graph theoretic problems, *Operations Research* 13 (1965) 462–475.
- [15] P. Hansen, N. Mladenović, Variable neighborhood search for the p -Median, *Location Science* 5 (1997) 207–226.
- [16] P. Hansen, N. Mladenović, An introduction to variable neighborhood search, in: S. Voss et al. (Eds.), *Metaheuristics Advances and Trends in Local Search Paradigms for Optimization*, Kluwer, 1999, pp. 433–458.
- [17] P. Hansen, N. Mladenović, Variable neighborhood search: principles and applications, *European Journal of Operational Research* 130 (2001) 449–467.
- [18] P. Hansen, N. Mladenović, D. Pérez-Brito, Variable neighborhood decomposition search, *Journal of Heuristics* 7 (2001) 335–350.
- [19] P. Hanjoul, D. Peeters, A comparison of two dual-based procedures for solving the p -median problem, *European Journal of Operational Research* 20 (1985) 387–396.
- [20] M.J. Hodgson, Toward more realistic allocation in location-allocation models: an interaction approach, *Environment and Planning A* 10 (1978) 1273–1285.
- [21] O. Kariv, S.L. Hakimi, An algorithmic approach to network location problems, part 2. The p -medians, *SIAM Journal on Applied Mathematics* 37 (1969) 539–560.
- [22] A.A. Kuehn, M.J. Hamburger, A heuristic program for locating warehouses, *Management Science* 9 (1963) 643–666.
- [23] M. Laguna, Scatter search, in: P.M. Pardalos, M.G.C. Resende (Eds.), *Handbook of Applied Optimization*, Oxford University Press, 2002, pp. 183–193.
- [24] T. Lai, S. Sahni, Anomalies in parallel branch-and-bound algorithms, *Communications of the ACM* 27 (6) (1984) 594–602.
- [25] T. Lai, A. Sprague, Performance of parallel branch-and-bound algorithms, *IEEE Transactions on computers* C 34 (10) (1985) 962–964.
- [26] F.E. Maranzana, On the location of supply points to minimize transportation costs, *Operations Research Quarterly* 12 (1964) 138–139.
- [27] N. Mladenović, J.A. Moreno-Pérez, J.M. Moreno-Vega, A chain-interchange heuristic method, *Yugoslav Journal of Operational Research* 6 (1996) 41–54.
- [28] R.M. Nauss, R.E. Markland, Theory and application of an optimizing procedure for lock box location analysis, *Management Science* 27 (1981) 855–865.
- [29] OpenMP: A proposed Industry Standard API for Shared Memory Programming, White Paper, October 1997. Available from <http://www.openmp.org/openmp/mp-documents/paper/paper.html>.

- [30] G. Reinelt, TSP-Lib a travelling salesman library, *ORSA Journal of Computing* 3 (1991) 376–384. Available from <<http://www.iwr.uni-heidelberg.de/iwr/comopt/software/tsplib95>>.
- [31] M.G.C. Resende, R.F. Werneck, A GRASP with path-relinking for the p -median problem, *Journal of Heuristics*, in press.
- [32] M.G.C. Resende, R.F. Werneck, On the implementation of a swap-based local search procedure for the p -median problem, Technical Report TD-5E4QKA, AT&T Labs Research, 2002.
- [33] K.E. Rosing, An empirical investigation of the effectiveness of a vertex substitution heuristic, *Environment and Planning B* 24 (1997) 59–67.
- [34] K.E. Rosing, C.S. ReVelle, H. Rosing-Vogelaar, The p -median and its linear programming relaxation: an approach to large problems, *Journal of the Operational Research Society* 30 (9) (1979) 815–823.
- [35] K.E. Rosing, C.S. ReVelle, Heuristic concentration: two stage solution construction, *European Journal of Operational Research* 97 (1997) 75–86.
- [36] E. Rolland, D.A. Schilling, J.R. Current, An efficient tabu search procedure for the p -median problem, *European Journal of Operational Research* 96 (1996) 329–342.
- [37] E.L.F. Senne, L.A.N. Lorena, Lagrangean/surrogate heuristics for p -median problems, in: M. Laguna, J.L. González-Velarde (Eds.), *Computing Tools for Modeling Optimization and Simulation: Interfaces in Computer Science and Operations Research*, Kluwer, 2000, pp. 115–130.
- [38] E.L.F. Senne, L.A.N. Lorena, Stabilizing column generation using Lagrangean/surrogate relaxation: an application to p -median location problems, *European Journal of Operational Research*, in press.
- [39] M.B. Teitz, P. Bart, Heuristic methods for estimating the generalized vertex median of a weighted graph, *Operations Research* 16 (1968) 955–961.
- [40] S. Voss, A reverse elimination approach for the p -median problem, *Studies in Locational Analysis* 8 (1996) 49–58.
- [41] R. Whitaker, A fast algorithm for the greedy interchange of large-scale clustering and median location problems, *INFOR* 21 (1983) 95–108.