

# A Cooperative Parallel Metaheuristic for the Capacitated Vehicle Routing Problem

Jianyong Jin<sup>a</sup>, Teodor Gabriel Crainic<sup>b</sup>, Arne Løkketangen<sup>a,\*</sup>

<sup>a</sup>*Molde University College, Specialized University in Logistics, N-6411, Molde, Norway*

<sup>b</sup>*School of Management, UQAM & Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation, Montréal, QC, Canada*

---

## Abstract

This paper introduces a cooperative parallel metaheuristic for solving the capacitated vehicle routing problem. The proposed metaheuristic consists of multiple parallel tabu search threads that cooperate by asynchronously exchanging best found solutions through a common solution pool. The solutions sent to the pool are clustered according to their similarities. The search history information identified from the solution clusters is applied to guide the intensification or diversification of the tabu search threads. Computational experiments on two sets of large scale benchmarks from the literature demonstrate that the suggested metaheuristic is highly competitive, providing new best solutions to ten of those well-studied instances.

*Keywords:* Vehicle routing, Parallel metaheuristic, Cooperative search, Solution clustering.

---

## 1. Introduction

In recent years, cooperative parallel metaheuristics have increasingly been used for solving a variety of difficult combinatorial problems [1]. Such parallel metaheuristics usually use multiple processes (threads) working simultaneously on available processors, with varying degrees of cooperation, to solve a given problem instance. The rationale behind this phenomenon may be twofold. First, it has been demonstrated that such parallel algorithms are

---

\*Corresponding author. Phone: +47 71 21 42 74; Fax: +47 71 21 41 00.

Email address: [arne.lokketangen@himolde.no](mailto:arne.lokketangen@himolde.no) (Arne Løkketangen )

capable of both speeding up the search and improving the robustness (ability of providing equally good solutions to a large and varied set of problem instances) and the quality of the solutions obtained [2]. Second, parallel computing resources have become increasingly available with the advent of computer clusters and multi-core processors. The computer clusters usually consist of a set of identical computers that run standard operating systems and are connected to each other through high speed networks. Many universities nowadays possess such computer clusters. In addition, many commodity laptop and desktop computers today use dual- or quad-core processors. Thus, using parallelism has become an advantageous and practical option. For detailed introduction to parallel metaheuristics, we refer to the book of Alba [3] and the survey paper of Crainic [2].

The capacitated vehicle routing problem (*CVRP*), as the classical version of the vehicle routing problem (*VRP*), aims to determine the minimum total cost routes for a fleet of homogeneous vehicles to serve a set of customers. The CVRP can be defined on a graph  $G = (N, E)$  where  $N = \{0, \dots, n\}$  is a vertex or node set and  $E = \{(i, j) : i, j \in N\}$  is an edge set. Vertex 0 is the depot where the vehicles depart from and return to. The other vertices are the customers which have a certain demand  $d$  to be delivered (or picked up). The travel cost between node  $i$  and  $j$  is defined by  $c_{ij} > 0$ . The vehicles are identical. Each vehicle has a capacity of  $Q$ . The objective is to design a least cost set of routes, all starting and ending at the depot. Each customer is visited exactly once. The total demand of all customers on any route must not exceed the vehicle capacity  $Q$ . Some CVRP instances may have an additional route duration limit constraint, restricting the duration (or length) of any route to a preset bound  $D$ . A detailed introduction of the CVRP and its solution methods can be found in the book of Toth and Vigo [4], and the survey paper of Laporte [5]. Even though a large number of solution methods have been proposed in the literature during last fifty years, it still remains computationally challenging to quickly produce high quality solutions to large scale CVRP instances.

The purpose of this paper is to present a cooperative parallel metaheuristic that takes advantage of modern parallel computing resources for solving large scale CVRP instances. The proposed algorithm incorporates multiple tabu search threads which cooperate by asynchronously exchanging best found solutions through a common solution pool, and includes several novel features. Intensification and diversification of the searches are based on solution clustering. Four variants of reinsertion neighborhood are applied and

46 infeasible solutions may also be sent to the solution pool. These features are  
 47 clearly different from previous work [e.g., 6, 7] and largely contribute to the  
 48 high performance of the proposed metaheuristic. The computational exper-  
 49 iments on two sets of large scale CVRP benchmarks demonstrate that the  
 50 suggested metaheuristic can quickly produce solutions to benchmark prob-  
 51 lems that are highly competitive with the best solutions reported in the  
 52 literature. New best solutions to 10 out of the 32 benchmark instances have  
 53 been identified.

54 The remainder of this paper is organized as follows. In the next section  
 55 the description of the proposed metaheuristic is presented. Then Section 3  
 56 reports the computational results. Finally, concluding remarks are given in  
 57 the last section.

## 58 2. Description of the cooperative parallel metaheuristic

59 In the proposed cooperative parallel metaheuristic (*CPM*), illustrated in  
 60 Figure 1, multiple tabu search (*TS*) threads are run in parallel for solving  
 61 a given CVRP instance. Some of the TS threads are designated to con-  
 62 centrate on intensification while the others are assigned to pursue diver-  
 63 sification. These threads communicate asynchronously through a common  
 64 solution pool.




Figure 1: Framework of CPM

65 The general scheme of CPM is displayed in Algorithm 1. During the  
 66 search process, the solution pool receives solutions sent from the search  
 67 threads. Whenever a solution is received from a search thread, the pool does

68 the clustering, selects a solution and sends it back to the same thread. Each  
69 of the TS threads carries out its search independently and periodically the  
70 search halts and exports its best found solution. It then receives a solution  
71 from the pool and resumes its search from this solution. The detailed de-  
72 scription of the solution pool and the TS threads will be provided in Section  
73 2.1 and 2.2 respectively.

74 The termination of CPM can be controlled in two ways. In the first set-  
75 ting (denoted as TC1), the termination is triggered by the first TS thread.  
76 The metaheuristic terminates after the thread runs for a certain number of  
77 iterations. In the other setting (denoted as TC2), the metaheuristic ter-  
78 minates after the solution pool receives a certain number of non-improving  
79 solutions consecutively. A solution is regarded as non-improving when its  
80 value is not better than that of the current best feasible solution at the pool  
81 or the solution is infeasible.

---

**Algorithm 1: CPM**

---

```

Initialize TS threads and the solution pool
while termination condition not met
    Solution pool do
        Receives solutions.
        Clusters solutions.
        Selects and sends solutions back.
    Each TS thread do asynchronously
        Performs the search.
        Sends best found solution to solution pool.
        Receives new solution to start from.
    end while
Return best feasible solution

```

---

82 In terms of the taxonomy introduced in Crainic and Nourredine [8] for  
83 parallel metaheuristics, CPM fits into the  $pC/KC/MPDS$  classification. The  
84 first dimension  $pC$  indicates the global search is controlled by multiple co-  
85 operative threads. The second dimension  $KC$  stands for knowledge colle-  
86 gial information exchange and refers to the fact that multiple threads share  
87 information asynchronously and knowledge is created from the exchanged  
88 information to guide the cooperating threads. The last dimension  $MPDS$  in-  
89 dicates that multiple search threads start from different points in the solution  
90 space and follow different search strategies.

## 91 2.1. Solution pool

92 To explore a search space effectively and efficiently, a metaheuristic ap-  
93 proach should be able to both intensively investigate areas of the search  
94 space with high quality solutions, and to move to unexplored areas of the  
95 search space when necessary. These goals are usually achieved by intensifi-  
96 cation and diversification mechanisms of the metaheuristic [9]. Glover and  
97 Laguna [10] highlight that intensification is to carefully search the neighbor-  
98 hood of elite solutions while diversification encourages the search process to  
99 generate solutions that differ from those seen before. A solution clustering  
100 approach is used in CPM to implicitly identify common features of solutions  
101 and collect search history information, which then provide a good basis for  
102 selecting promising search areas for intensification and less explored areas for  
103 diversification.

104 During the whole search process, the solutions sent to the solution pool  
105 from the search threads are dynamically clustered into groups according to  
106 their similarities. For the CVRP, the similarity can be measured in terms  
107 of the number of edges solutions have in common. In this way, solutions  
108 kept in the solution pool can be grouped into clusters and in each cluster all  
109 solutions have a certain number of edges in common. Each cluster can thus  
110 approximately represent a region of the search space that CPM has explored.  
111 The features of the solutions in a cluster, such as the number of solutions and  
112 the quality of the solutions, can indicate how thoroughly a search region has  
113 been explored and how promising it may be. Such search history information  
114 is used to guide the starting solution selection for the TS threads so that they  
115 can pursue intensification or diversification effectively.

116 The solution clustering approach has been applied by Voß [11] for the  
117 quadratic assignment problem. In his algorithm, a small number of elite so-  
118 lutions previously found are stored by a clustering approach and are used as  
119 the starting solutions for the intensification phases. In CPM, the solution  
120 clustering approach is applied differently in three aspects. First, all solu-  
121 tions sent to the pool are clustered, regardless of their quality. Second, the  
122 solutions are clustered for both intensification and diversification purposes.  
123 Third, the actual clustering mechanism is different.

### 124 2.1.1. Solution clustering

125 Clustering is often defined as the process of grouping a collection of pat-  
126 terns into dissimilar segments or clusters based on a suitable notion of close-  
127 ness or similarity among these patterns. In CPM, solutions are grouped into

128 clusters based on their similarities. A cluster, in this context, refers to a  
129 collection of solutions that are similar. All solutions sent to the solution pool  
130 are clustered.

131 To support solution exchange and information extraction from the clus-  
132 ters, a set of components are implemented for each cluster.

- 133 • **Feasible solution list:** a list of feasible solutions assigned to a cluster.
- 134 • **Infeasible solution list:** a list of infeasible solutions assigned to a  
135 cluster.
- 136 • **Edge residence counter:** a residence counter for all edges. The resi-  
137 dence counter of an edge is defined as the number of feasible solutions  
138 containing this edge that have been assigned to a cluster. Since the ob-  
139 jective of the search is to identify high quality feasible solutions, only  
140 feasible solutions are used to compute the edge residence counter.
- 141 • **Feasible solution counter:** the number of feasible solutions that have  
142 been assigned to a cluster.
- 143 • **Infeasible solution counter:** the number of infeasible solutions that  
144 have been assigned to a cluster.
- 145 • **Average feasible solution value:** the average value of all feasible  
146 solutions in a cluster.
- 147 • **Average infeasible solution value:** the average value of all infeasible  
148 solutions in a cluster.

149 Whenever a solution enters a cluster, the components of the cluster are  
150 updated. In each cluster, duplicate solutions will be eliminated and the lists  
151 of feasible and infeasible solutions are sorted in ascending order according to  
152 solution value.

153 In addition, the clusters are sorted in ascending order according to the  
154 average feasible solution value. When there are only infeasible solutions in a  
155 cluster, replace the average feasible solution value with the average infeasible  
156 solution value for sorting.

157 To determine whether a solution is similar to the solutions in a cluster,  
158 the similarity between the solution and the cluster is calculated according to  
159 Equation 1.

$$Similarity = \frac{\sum_{(ij) \in E} n_{ij} \times X_{ij}^s}{FSC \times \sum_{(ij) \in E} X_{ij}^s} \quad (1)$$

160 The components of the formula are defined as follows.

- 161 •  $n_{ij}$ : the residence counter of edge  $(i, j)$  of the cluster.
- 162 •  $X_{ij}^s$ : 1 if edge  $(i, j)$  appears in solution  $s$ , 0 otherwise.
- 163 •  $FSC$ : the feasible solution counter of the cluster.
- 164 •  $\sum_{(ij) \in E} X_{ij}^s$ : the number of edges in solution  $s$ .

165 The advantage of computing the similarity in such a way is twofold. First,  
 166 it does indicate how many common edges a solution shares with the solutions  
 167 in a cluster. Moreover, it avoids the heavy computational load of calculating  
 168 the similarities between a solution and every solution in the cluster as other  
 169 clustering approaches do.

170 A solution can be placed into a cluster only if the similarity between the  
 171 solution and the cluster is larger than a minimum value. This value is termed  
 172 the minimal similarity requirement. When a solution is sent to the solution  
 173 pool, three possibilities exist. Initially, there is no existing cluster, a cluster  
 174 will be created and the solution will be directly placed into the cluster. When  
 175 there are existing clusters, the solution will be compared with the first cluster  
 176 in the pool. If the similarity requirement is satisfied, it will be put into the  
 177 cluster. Otherwise, it will be compared with the next cluster following the  
 178 sorted order. The comparison may continue until the solution is placed into  
 179 a cluster or it does not satisfy the similarity requirement with any existing  
 180 cluster. Under such a circumstance, a new cluster will be created and the  
 181 solution will be put into the new cluster.

182 Moreover, a pair of status flags is attached to each solution that enters  
 183 a cluster for signaling whether it has been used. The two flags are used for  
 184 intensification and diversification respectively. After a solution is selected  
 185 and sent to an intensification TS thread, its status flag for intensification  
 186 will be set accordingly. Likewise, the status flag for diversification will be set  
 187 after a solution is sent to a diversification TS thread. By setting these flags,  
 188 usually each solution can be selected and sent to each type of search threads  
 189 only once.

190 Three parameters are used to control the clustering process. The first one  
 191 is the minimal similarity requirement that controls the number of common  
 192 edges the solutions in a cluster share. The second one is the maximal number  
 193 of clusters which are allowed to exist in the solution pool. Too many clusters  
 194 will decelerate the clustering process. Whenever a new cluster is created,  
 195 the clustering procedure will check the number of existing clusters. If there  
 196 are more clusters than allowed, a cluster that does not contain any feasible  
 197 solution or has the largest average feasible solution value will be eliminated.  
 198 The last parameter is the maximal number of solutions in a cluster. When  
 199 there are more solutions than allowed, the worst solution in terms of the  
 200 solution value will be removed. It is essential to restrict the number of the  
 201 clusters and the number of solutions in each cluster for the sake of efficiency,  
 202 especially when a large number of threads are employed.

### 203 *2.1.2. Solution selection for intensification threads*

204 For a cluster that contains mainly high quality solutions, indicated by the  
 205 average feasible solution value of the cluster, the search region represented  
 206 by the cluster is usually worth further intensive investigation. In CPM, the  
 207 cluster having the lowest average feasible solution value is assigned as the  
 208 target of the TS threads that pursue intensification. These search threads  
 209 will only receive starting solutions from this best cluster so that the neighbor-  
 210 hoods of high quality solutions can be thoroughly investigated. During the  
 211 whole search process, this best cluster may dynamically be replaced by the  
 212 newly emerged clusters that have lower average feasible solution value than  
 213 the current one. In this way, the intensification search threads will always  
 214 target the vicinity of current best solutions.

215 Whenever an intensification thread needs a starting solution, the solutions  
 216 in the best cluster will be checked. The intensification flag of each solution is  
 217 examined one at a time, starting from the solution with the lowest solution  
 218 value, following the sorted order. The first unused solution will be selected.  
 219 When there are no unused solutions in the best cluster, a solution is randomly  
 220 selected from the cluster. This selected solution will be sent to the search  
 221 thread.

222 If there are infeasible solutions in the best cluster, they will be examined  
 223 and selected first. The reason for this decision is that preliminary experi-  
 224 ments show there are usually many more feasible solutions than infeasible  
 225 solutions in the clusters. The infeasible solutions will seldom be selected if  
 226 the feasible solutions are checked first.



### 2.1.3. *Solution selection for diversification threads*

The search regions that have been less thoroughly explored can be indicated by the number of feasible solutions that have been put into the clusters. The fewer feasible solutions have been put into a cluster, the less thoroughly the region has been searched. Since the tabu search threads seeking diversification are expected to concentrate mainly on less explored search regions, they will receive starting solutions only from any cluster whose feasible solution counter is below a threshold. We term this threshold the diversification threshold. When the feasible solution counter of a cluster exceeds the diversification threshold, the cluster will be neglected while selecting solutions for the diversification threads.

Whenever a diversification thread requires a starting solution, the solution selection procedure starts with the cluster with the lowest average feasible solution value in the solution pool. If the feasible solution counter of the cluster is below the diversification threshold, the solutions in the cluster are then examined. If an unused solution is found, this solution is sent to the search thread, otherwise, the next cluster following the sorted order will be checked. The examination continues until an unused solution is found or all available clusters have been checked. When an unused solution is not found after examining all clusters, a solution is randomly selected.

When examining the solutions in a cluster, the infeasible solutions will be checked and selected first, as for intensification threads.

## 2.2. *The Tabu search threads*

In this sub-section, the common features and differences between the TS threads included in CPM are provided.

### 2.2.1. *The common features of the tabu search threads*

The TS threads included in CPM are developed on the basis of the granular tabu search that was first introduced by Toth and Vigo [12]. Below, the main features of the TS threads are introduced.

#### **The initial solution**

The initial solution of each TS thread is constructed by using the parameterized Clarke-Wright algorithm described in Yellow [13] with a randomly generated shape parameter. The range for the shape parameter is set to  $(0.5, 2)$  as suggested in Groër et al. [14].

### Objective function and constraint relaxation mechanism

To explore the solution space more thoroughly, infeasible intermediate solutions are allowed. To this end, capacity and route length constraints are relaxed and their violations are penalized in the objective function. This augmented objective function is computed as  $F(s) = C(s) + \alpha Q(s) + \beta D(s)$ , where  $C(s)$  is the total travel distance,  $Q(s)$  and  $D(s)$  denote the total violations of the capacity and route length constraints respectively,  $\alpha$  and  $\beta$  are penalty multipliers. The values of the penalty multipliers are self-adjusted during the course of the search as described by Toth and Vigo [12]. The neighboring solutions, both feasible and infeasible, generated during the search process are evaluated in terms of the augmented objective function.

### Inter-route neighborhood structures

Three neighborhood structures for inter-route operations, which are commonly used in the previously published metaheuristics for the CVRP, are implemented in the TS threads. The basic idea of each neighborhood structure is described as follows.

- Reinsertion [15] refers to relocating a customer node from one route to another route.
- 2-opt\* [16] eliminates two edges from two routes and reconnects the two routes with two new edges.
- CROSS-exchange [17] swaps two route segments between two routes. In CPM the route segments can contain 1-3 nodes.

All the three neighborhood structures are implemented in each TS thread. One of them is randomly selected every TS iteration and each of the neighborhoods has equal probability to be selected.

To speed up the search, the granular neighborhood reduction technique applied in Jin et al. [6] is adopted in CPM. Let  $R(u)$  stand for the route containing node  $u$  in a given solution, and  $(u, x)$  be the partial route from node  $u$  to node  $x$ . Define  $N_u$  as the set of the nearest neighbors of customer  $u$ . Assume node  $v$  is a member of  $N_u$  and  $R(v) \neq R(u)$ . Each neighborhood is generated in the following way.

- Reinsertion: For each customer  $u$ , for each  $v$ , reinsert  $u$  right after  $v$ .

- 293     • 2-opt\*: Let  $x$  be the successor of  $u$  in  $R(u)$  and  $y$  be the successor of  
294          $v$  in  $R(v)$ . For each customer  $u$ , for each  $v$ , replace  $(u, x)$  and  $(v, y)$  by  
295          $(u, v)$  and  $(x, y)$ .
- 296     • CROSS-exchange: The procedure introduced in Taillard et al. [17] is  
297         adopted. To reduce the computational effort, one restriction is im-  
298         posed. For a route  $R_1$ , only a couple of routes are chosen for neighbor-  
299         hood generation. Those routes are selected in the following way. First,  
300         a node  $u$  is randomly selected from the middle part of route  $R_1$ . Then  
301         the routes which contain at least one customer node belonging to  $N_u$   
302         are identified. Only these routes are used to form route pairs with  $R_1$   
303         for neighborhood generation.

304     The size of the nearest neighbors set is randomly chosen within a certain  
305     range at each iteration.

#### 306     **Four types of reinsertion strategies**

307     For local search based metaheuristics for the CVRP, one drawback is that  
308     the edges close to the depot are usually more frequently involved in selected  
309     moves compared to more remote edges. To improve this situation, reinsertion  
310     neighborhood structure for inter-route improvement is implemented in  
311     four distinct ways. The main idea is to partition the customer nodes into  
312     groups according to their distance to the depot, and each group should have  
313     an approximately equal number of nodes. At each iteration, neighborhood  
314     generation and move selection are carried out separately for each group. In  
315     such a way, the frequency of modifying the route structures distant from the  
316     depot can be increased. An example is shown in Figure 2. In the figure,  
317     the small square stands for the depot and the dots represent customer nodes.  
318     The customer nodes are divided into two groups, the dots inside the big circle  
319     belong to the first group while the dots outside the big circle constitute the  
320     other group. Nevertheless, the nearest neighbors (the dots inside the small  
321     circle) of a customer do not need to be in the same group as the customer.

322     The first type of reinsertion strategy is to put all customer nodes in one  
323     group, and at each iteration only one move is performed. For the second  
324     strategy, the customer nodes are partitioned into two groups and two moves  
325     are carried out at each iteration, one from each partition. Likewise, for  
326     the other two strategies, the customer nodes are divided into 3 or 4 groups  
327     and 3 or 4 moves are performed at each iteration respectively. These four  
328     types of strategies are termed Type 1 reinsertion, Type 2 reinsertion, Type 3




Figure 2: Customer nodes partition

reinsertion and Type 4 reinsertion accordingly. For a given TS thread, only one type of reinsertion is applied.

### Solution acceptance and tabu mechanism

Among the neighboring solutions, the best move in terms of the augmented objective function that is non-tabu or satisfies the aspiration criterion is accepted. The aspiration criterion overrides the tabu status of a move if this move leads to a new best solution in the current search.

The tabu list is neighborhood dependent. The tabu tenure  $tt$  of each neighborhood is set to be proportional to the number of nodes in the instance. For reinsertion, if  $u$  is relocated,  $u$  is declared tabu for  $tt$  iterations and any moves relocating  $u$  cannot be performed unless it satisfies the aspiration criterion. For CROSS-exchange move swapping route segments  $(X_1, Y_1)$  and  $(X_2, Y_2)$ , node  $X_1$  and  $X_2$  are declared tabu and any moves involving the two nodes cannot be performed unless it satisfies the aspiration criterion. For 2-opt\* move adding edge  $(u, v)$  and  $(x, y)$ , node  $u$ ,  $v$  and  $y$  are declared tabu, any moves involving any one of these three nodes are forbidden unless it satisfies the aspiration criterion.

### Route refinement

In the TS threads, at each iteration, after an inter-route move, the two modified routes are refined separately by an intra-route improvement procedure. The procedure consists of two simple heuristics developed by implementing 2-opt [18] and reinsertion [15] neighborhood structures in a local search setting. The two heuristics are applied to a route alternately. The heuristic using 2-opt repeatedly eliminates two edges and adds two new edges, improving moves are accepted until no improvement can be found. The heuristic using reinsertion seeks improvement by relocating a node to another position, if a move reduces the route length, then it is accepted. The procedure

356 terminates when no improvement can be found.

### 357 **Solution exchange**

358 The TS threads halt and exchange solutions with the solution pool period-  
359 ically. Each TS thread exports its best found solution and receives a new  
360 solution to resume the search. Each search thread decides when to exchange  
361 solutions with the solution pool according to its own search trajectory, the  
362 communication is asynchronous. No direct communication takes place be-  
363 tween the search threads.

364 The search effort between two solution exchanges is termed a search pe-  
365 riod. During a search period, if the best feasible solution a TS thread has  
366 found is better than its starting solution, this solution will be sent to the  
367 pool, otherwise, the best infeasible solution the thread has found will be sent  
368 to the pool. The rationale behind the decision to exchange infeasible solu-  
369 tions is that infeasible solutions generated by one thread may be improved  
370 by another thread so that better feasible solutions can be found.

371 The TS threads can stop and exchange solutions with the solution pool  
372 after either running for a certain number of iterations or failing to find im-  
373 proving feasible solutions for a certain number of iterations. The first stop-  
374 ping mechanism is denoted as SM1 and the second one is denoted as SM2.  
375 The TS threads for intensification and diversification can adopt the same  
376 (all use SM1 or SM2) or different (some use SM1 and the others use SM2)  
377 stopping mechanisms. Several settings are compared in Section 3.3.

### 378 **Solution representation and transformation**

379 To speed up the computation, the solutions are stored in a data structure of  
380 four arrays, namely next-array, pred-array, start-array and end-array. The  
381 first two arrays keep the successor and predecessor of each node while the  
382 other two record the first customer and the last customer of each route.  
383 Using this structure, changes to a solution can be performed very quickly.  
384 The detailed description of this application can be found in Kytöjoki et al.  
385 [19] and Groër et al. [20]. On the other hand, to simplify the information  
386 exchanged between the search threads and the solution pool, the solutions  
387 are transformed to a giant tour format (without route delimiters) before they  
388 are sent to the solution pool. When a TS thread receives a solution from  
389 the pool, the giant tour is transformed back to the four-array format with a  
390 split algorithm presented by Prins [21]. This split algorithm considers both  
391 vehicle capacity and route length constraints, and all solutions will become  
392 feasible after split.

393 To facilitate the solution clustering and sorting in the pool, the feasibility  
394 and the objective function value of each solution are required. To this end,  
395 these two attributes are attached to the giant tour during the transforma-  
396 tion and are exchanged together. In the pool, each solution is stored as an  
397 augmented giant tour.

### 398 *2.2.2. The differences between the tabu search threads*

399 To explore the search space effectively and efficiently, some of the TS  
400 threads are designated to concentrate on intensification while the others are  
401 assigned to pursue diversification. The intensification threads are imple-  
402 mented with smaller tabu tenures than the threads seeking diversification.  
403 Additionally, each TS thread uses only one of the 4 types of reinsertion  
404 neighborhoods.

### 405 *2.2.3. Overall search process*

406 Each tabu search thread begins its search with an initial solution  $s_i$  cre-  
407 ated by using the parameterized Clarke-Wright algorithm. At each iteration,  
408 first a neighborhood structure and the size of the nearest neighbors set (de-  
409 noted as  $LS$ ) are randomly selected. When Type 1 reinsertion, 2-opt\* or  
410 CROSS-exchange neighborhood is selected, one set of neighbors is generated  
411 and the least cost non-tabu solution  $\bar{s}$  is selected to replace the current solu-  
412 tion  $s$ . Then the attributes of the reverse moves are declared tabu and the  
413 routes just modified are refined by the intra-route improvement procedure.  
414 Furthermore, if the current solution  $s$  is feasible and better than the best  
415 feasible solution  $s_f^*$ , replace  $s_f^*$  with  $s$ . When the current solution  $s$  is infea-  
416 sible and better than the best infeasible solution  $s_{inf}^*$ , replace  $s_{inf}^*$  with  $s$ .  
417 For the other three types of reinsertion (Type 2, 3 and 4), the neighborhood  
418 generation and move selection (Step 5-10 in Algorithm 2) will be performed  
419 several times in accordance with the number of groups that the customer  
420 nodes have been divided into.

421 In addition, periodically each tabu search thread stops to exchange so-  
422 lutions with the solution pool and uses the received solution to replace its  
423 current solution and the best solutions, and the tabu lists are re-initialized.  
424 This process is summarized in Algorithm 2.

## 425 **3. Computational results**

426 In this section we describe the experimental platform, the test data sets,  
427 the algorithm configurations and compare the experimental results against

---

**Algorithm 2:** Tabu search

---

- 1: Construct  $s_i$ , set  $s_f^* = s_i$ ,  $s_{inf}^* = s_i$ ,  $s = s_i$ ,
  - 2: Initialize tabu lists and penalty multipliers
  - 3: **while** termination condition not met **do**
  - 4:     Select a neighborhood and  $LS$
  - 5:     Generate and evaluate neighboring solutions
  - 6:     Select a neighboring solution  $\bar{s}$  that minimize  $F(\bar{s})$  and is non-tabu or satisfies the aspiration criterion, set  $s = \bar{s}$
  - 7:     Declare the attributes of the reverse moves tabu for  $tt$  iterations
  - 8:     Refine the routes modified
  - 9:     Update  $s_f^*$  and  $s_{inf}^*$
  - 10:    Update penalty multipliers
  - 11:    If reaching the iteration limit, halt and exchange solution with the solution pool, reset  $s_f^*$ ,  $s_{inf}^*$ ,  $s$ , and tabu lists
  - 12: **end while**
- 

428 the results of the state-of-the-art methods and the best known solutions  
429 (*BKS*) reported in the literature. The current best known results have been  
430 updated to include the new best solutions identified by Groër et al. [14], Jin  
431 et al. [7] and Vidal et al. [22].

432 The analysis of the impact on the performance of several algorithmic  
433 components and the evaluation of the parallel speedup are provided in this  
434 section as well.

### 435 3.1. Experimental platform and implementation issues

436 The proposed metaheuristic is implemented in C++ and uses the message  
437 passing interface (*MPI*) for the inter-processor information exchange. The  
438 results were obtained by running the algorithm on a compute cluster in which  
439 each node consists of two AMD 6172 processors with 12 cores at 2.1 GHz.

440 The basic configuration of CPM has 8 threads. Among them, one thread  
441 is used for the solution pool, 4 TS threads using Type 1, 2, 3 and 4 reinsertion  
442 respectively are for diversification and the remaining 3 TS threads using  
443 Type 1, 2, and 3 reinsertion respectively are for intensification. These 7 TS  
444 threads are regarded as the basic search threads. When more processors are  
445 employed by CPM, the 7 basic search threads can be duplicated and run  
446 on the available processors. The standard configuration of CPM (denoted  
447 as CPM standard) utilizes 24 threads with one for the solution pool, 14 for  
448 diversification and 9 for intensification.

### 3.2. The test data sets

The computational tests were carried out on the CVRP benchmarks of Golden et al. [23] and Li et al. [24]. The 20 benchmark instances of Golden et al. [23] have 200 to 483 customers. The first eight instances also have route length restrictions. Each instance is based on a simple geometric structure: eight instances have customers located in concentric circles around the depot, four instances have customers located in concentric squares with the depot located in one corner, four instances have customers located in concentric squares around the depot, and four instances have customers located in a six-pointed star around the depot. The benchmark instances of Li et al. [24] have 560 to 1200 customers and route length restrictions, and their geometric structure is based on concentric circles around the depot. For each instance under each experimental scenario, CPM was executed 10 times with different random seeds. The average result and best result of these 10 runs are reported.

### 3.3. Evaluating search stopping mechanisms

As mentioned in Section 2, the overall search of CPM can be terminated according to two conditions, TC1 and TC2. For the TS threads, there are two stopping mechanisms (SM1 and SM2) for deciding when to exchange solutions with the solution pool. In addition, some of the TS threads are designated to concentrate on intensification while the others are assigned to pursue diversification. Considering these three aspects, six ways of controlling the search of CPM are compared. The settings of the six variants are shown in Table 1.

Table 1: Search stopping mechanisms

Variant	1	2	3	4	5	6
Overall search	TC1	TC1	TC1	TC2	TC2	TC2
Diversification TS threads	SM1	SM2	SM1	SM1	SM2	SM1
Intensification TS threads	SM1	SM2	SM2	SM1	SM2	SM2

In general, TC1 and SM1 can explicitly control the search effort while TC2 and SM2 may stop the search dynamically according to the progress of the TS threads. Preliminary testing gave no significant difference among the six variants in terms of the solution quality and the search time. Therefore we choose variant 1 to perform the remaining computational experiments since it allows us to explicitly control the search effort.



479 *3.4. Algorithm calibration*

480 The parameters of CPM were selected according to the computational  
 481 results of preliminary experiments on the benchmarks of Golden et al. [23].  
 482 A number of different alternative values were tested and the ones selected are  
 483 those that gave the best computational results concerning both the quality  
 484 of solutions and the computational times needed to achieve these solutions.  
 485 The selected parameter values are given in Table 2.

Table 2: Parameter values for CPM

Parameter	Value
Tabu tenure of of reinsertion, 2-opt* and CROSS-exchange	0.03 N  for diversification threads 0.01 N  for intensification threads
Nearest neighbors set size	(10 + random [0, 10])
Solution exchange	$200 \times \sqrt{ N }$ iterations.
Minimal similarity requirement	0.7
Maximal cluster number	100
Maximal number of solutions in a cluster	300
Diversification threshold	100
Termination condition	$150000 \times \sqrt{ N }$ for $ N  < 500$ $30000 \times \sqrt{ N }$ for $ N  > 500$

|N| represents the instance size.

486 *3.5. Results for the benchmarks of Golden et al. [23]*

487 In Table 3 we compare the results for the 20 benchmark instances of  
 488 Golden et al. [23]. In the table, the first column describes the instances  
 489 (instance number and number of nodes). The second column lists the best  
 490 known solutions previously reported in the literature. The third and fourth  
 491 columns provide the best results presented by Groër et al. [14] and Vidal et al.  
 492 [22]. The remaining columns give the average results, average wall-clock time,  
 493 standard deviations and best results of CPM standard. The third last row  
 494 presents the average deviation of all instances from the best known solutions.  
 495 The second last row provides the average wall-clock time per instances for

a single run. The last row shows the number of runs performed for each instance in each algorithm.

From the table, we see that CPM standard has found new best solutions to 7 instances (numbers in bold font) while the average deviation of the best results from the best known solutions is 0.00%. In terms of this metric, the best results generated by CPM standard are better than those of Groër et al. [14] and Vidal et al. [22]. The wall-clock time required by CPM standard appears shorter than what was used in Vidal et al. [22] and longer than for Groër et al. [14].

Table 3: Comparison of results for benchmarks of Golden et al. [23]

Instances	Previous best known	Groër et al. (2011) 129p	Vidal et al. (2012)	CPM standard			
				Aver.	Time (min)	SD.	Best
1(240)	5623.47	5623.47	5623.47	5623.65	22.05	0.38	5623.47
2(320)	8404.61	8435.00	8404.61	8434.78	34.22	14.45	8405.81
3(400)	11036.22	11036.22	11036.22	11036.22	44.64	0.00	11036.22
4(480)	13592.88	13624.52	13624.52	13620.30	60.87	10.90	<b>13590.00</b>
5(200)	6460.98	6460.98	6460.98	6460.98	15.83	0.00	6460.98
6(280)	8400.33	8412.90	8412.9	8404.06	26.76	4.97	8400.33
7(360)	10102.70	10195.59	10102.70	10134.93	39.01	11.48	10107.49
8(440)	11635.30	11649.89	11635.30	11635.34	54.61	0.00	11635.34
9(255)	579.71	579.71	579.71	580.04	19.43	0.29	579.71
10(323)	736.26	737.28	736.26	737.16	28.82	0.46	<b>735.66</b>
11(399)	912.84	913.35	912.84	912.72	41.33	0.28	<b>912.03</b>
12(483)	1102.69	1102.76	1102.69	1103.20	58.29	1.28	<b>1101.50</b>
13(252)	857.19	857.19	857.19	858.57	18.06	1.20	857.19
14(320)	1080.55	1080.55	1080.55	1080.55	25.08	0.00	1080.55
15(396)	1337.92	1338.00	1337.92	1340.13	36.33	1.38	<b>1337.87</b>
16(480)	1612.50	1613.66	1612.50	1614.73	48.14	1.80	<b>1611.56</b>
17(240)	707.76	707.76	707.76	707.80	16.39	0.07	707.76
18(300)	995.13	995.13	995.13	998.90	25.01	0.96	997.58
19(360)	1365.60	1365.60	1365.60	1366.12	32.60	0.37	1365.60
20(420)	1818.25	1818.25	1818.32	1819.76	41.93	1.10	<b>1817.89</b>
Aver. deviation (%)		0.10	0.02	0.11			0.00
Time (min)		5.00	58.56		34.47		
Runs per instances		5	10	10			10

### 3.6. Results for the benchmarks of Li et al. [24]

The results for the 12 benchmark instances of Li et al. [24] are presented in Table 4. The format of this table is identical to the one of Table 3. For this set of instances, CPM standard has found new best solutions to three instances (numbers in bold font). In terms of the average deviation from the best known solutions, both the average and best results of CPM standard excel the best results of Mester and Bräysy [25] and Groër et al. [14]. The wall-clock time required by CPM standard turns out shorter than what was used in Mester and Bräysy [25] and longer than for Groër et al. [14].

Table 4: Comparison of results for benchmarks of Li et al. [24]

Instances	Previous best known	Mester and Bräysy (2007)	Groër et al. (2011) 129p	CPM standard			
				Aver.	Time (min)	SD.	Best
21(560)	16212.74	16212.74	16212.83	16214.12	14.25	1.05	16212.83
22(600)	14575.19	14597.18	14584.42	14562.10	19.35	11.98	<b>14539.79</b>
23(640)	18801.12	18801.12	18801.13	18853.80	18.18	106.69	18801.13
24(720)	21389.33	21389.33	21389.43	21390.96	22.41	1.39	21389.43
25(760)	16739.84	17095.27	16763.72	16733.07	33.05	16.91	<b>16709.44</b>
26(800)	23971.74	23971.74	23977.73	23981.30	28.73	1.26	23980.12
27(840)	17408.66	17488.74	17433.69	17380.24	38.05	24.26	<b>17343.38</b>
28(880)	26565.92	26565.92	26566.03	26569.96	33.14	1.88	26567.23
29(960)	29154.34	29160.33	29154.34	29157.42	40.85	1.98	29154.33
30(1040)	31742.51	31742.51	31742.64	31746.20	51.17	1.60	31742.64
31(1120)	34330.84	34330.84	34330.94	34333.66	62.63	2.12	34330.94
32(1200)	36919.24	36928.70	37185.85	37188.36	72.36	16.29	37162.54
Aver. deviation (%)		0.23	0.09	0.07			-0.01
Time (min)		104.30	5.00		36.18		
Runs per instances		1	5	10			10

### 3.7. Impact of algorithmic components

To examine the impact on the performance of CPM of the main algorithmic components, a set of experiments was conducted. In each experiment, the CPM standard was altered to deactivate or remove some components respectively. The experiments are described below.

- Only use Type 1 reinsertion (R1): All TS threads employ Type 1 reinsertion strategy. In this experiment, the other three reinsertion strategies are removed from CPM.

- 522 • Only use Type 2 reinsertion (R2): All TS threads employ Type 2 rein-  
523 sersion strategy.
- 524 • Only use Type 3 reinsertion (R3): All TS threads employ Type 3 rein-  
525 sersion strategy.
- 526 • Only use Type 4 reinsertion (R4): All TS threads employ Type 4 rein-  
527 sersion strategy.
- 528 • Only exchange feasible solutions (OF): All TS threads only exchange  
529 feasible solutions with the pool. When a TS thread does not improve  
530 its starting solution, the solution can still be sent to the solution pool  
531 so as to keep the frequency of solution exchange identical.
- 532 • No guidance (NG): In this experiment, both intensification and diver-  
533 sification mechanism are deactivated. Set cluster size threshold for  
534 diversification to a large number (e.g. 5000) so that the threads for  
535 diversification can obtain solutions from any clusters all the time no  
536 matter how many solutions a cluster contains. Set the tabu tenures  
537 and solution selection rule of the intensification threads identical to  
538 those for the diversification threads.

539 These modified versions were tested on benchmarks of Golden et al. [23]  
540 and the average results are compared against those of CPM standard. The  
541 comparison is shown in Table 5. At first glance the impact on the perfor-  
542 mance of the algorithmic components may seem small, but it is in fact crucial  
543 because for these well-studied instances, even minute improvements are diffi-  
544 cult to obtain. The results thus show that all these algorithmic components  
545 contribute to the high performance of CPM.

Table 5: Impact of the algorithmic components

Experiment	CPM st.	R1	R2	R3	R4	OF	NG
Aver. deviation from BKS (%)	0.11	0.67	0.17	0.19	0.23	0.15	0.15
Average wall-clock time per instance (min)	34.47	33.00	33.87	34.45	35.12	34.63	34.53

### 3.8. Effect of search effort on performance

To examine the performance of the proposed metaheuristic when dissimilar search effort is employed, the CPM standard was executed with several settings on benchmarks of Golden et al. [23]. In each setting, the total number of iterations are altered. The results are showed in Table 6.

Table 6: Comparison of the effect of search effort

Iterations ( $1000 \times \sqrt{ N }$ )	30	60	90	150	180
Aver. deviation from BKS of best results (%)	0.08	0.05	0.01	0.00	0.00
Aver. deviation from BKS of average results (%)	0.18	0.14	0.12	0.11	0.11
Average wall-clock time per instance (min)	6.84	13.87	20.51	34.47	40.95

In the table, the first row provides the number of iterations for each experiential setting. The second row presents the average deviations from the best known solutions of the best results for each setting. Likewise, the average deviations of the average results for each setting are shown in the third row. The last row provides the average wall-clock time per instance for each setting. From the results, it is noticeable that the quality of the solutions obtained gradually improves until a peak level is reached as the search effort increases. Additionally, we can see that CPM standard can identify solutions at similar quality to those of Groër et al. [14] and Vidal et al. [22] even when less search effort is utilized.

### 3.9. Measuring the parallel speedup

The parallel speedup is one of the most widely used measures of a parallel algorithm’s effectiveness. This metric is defined as the ratio between the sequential and parallel times. The sequential time is the amount of time required for running the algorithm on a single computer, and the parallel time refers to the amount of time required for the parallel computation when using multiple processors. However, in this experiment, we do not compare the parallel time against the sequential time since it may impair the effectiveness of CPM to run it on a single computer. We instead compare the quality of average results and the wall-clock time when using a different number of processors and a fixed amount of search effort (i.e., the total iterations per thread times the number of processors). This experiment is carried out on benchmarks of Golden et al. [23] and the fixed search effort is set to  $((150000 \times \sqrt{|N|}) \text{ iterations}) \times (24 \text{ processors})$ . The results are shown in

575 Table 7. From the results, it is observable that, for up to 240 processors, in-  
576 creasing the number of processors generally allows CPM to discover solutions  
577 of approximately equivalent quality in roughly linearly reduced time.

Table 7: Analyzing the parallel speedup

Number of processors	8	16	24	72	120	240
Iterations ( $1000 \times \sqrt{ N }$ )	450	225	150	50	30	15
Aver. deviation from BKS (%)	0.113	0.109	0.110	0.111	0.116	0.117
Average wall-clock time per instance (min)	101.44	50.88	34.47	11.72	7.00	3.63

## 578 4. Conclusions and Perspectives

579 In this paper, we have presented a cooperative parallel metaheuristic for  
580 the capacitated vehicle routing problem. The computational experiments on  
581 the two sets of large scale CVRP benchmarks show that the suggested meta-  
582 heuristic is quite effective and competitive in comparison to state-of-the-art  
583 methods from the literature. Though the benchmarks used have been well-  
584 studied, the proposed parallel metaheuristic is still able to identify new best  
585 solutions to 10 of the 32 instances within reasonable computational time.  
586 From a parallel computation point of view, the proposed parallel metaheuris-  
587 tic is efficient and flexible as it can employ at least up to 240 processors and  
588 achieve a roughly linear speedup.

589 In addition, several new features are introduced in this paper. First, us-  
590 ing the structural information (edge residence counters) of solutions enables  
591 the solution clustering to be carried out rapidly, even when thousands of so-  
592 lutions are involved. Search history information can thus be extracted from  
593 the solution clusters, helping the search to better achieve intensification and  
594 diversification. Second, the novel way of implementing reinsertion neighbor-  
595 hoods based on the distance to the depot of customers makes it easier to  
596 seek improvement for different parts of the solutions. Moreover, for coopera-  
597 tive search, it appears beneficial to also exchange infeasible solutions among  
598 search threads. The combination of these features largely contributes to the  
599 high performance of the proposed metaheuristic.

600 In this paper, we focus on cooperation and information exchange and use  
601 the cooperation concept where all exchanges proceed through the common

602 solution pool. The solution clustering based intensification and diversifica-  
 603 tion introduced is very general and problem domain independent. It can be  
 604 easily applied in different contexts for solving other combinatorial problems.  
 605 The modification required is to select the solution elements on which com-  
 606 mon features are built and use them in a similarity measure for the problem  
 607 in question. Future work will focus on pursuing other approaches of ap-  
 608 plying the information extracted from the solution clustering and adopt the  
 609 proposed parallel metaheuristic to other classes of problems.

## 610 **Acknowledgements**

611 The authors thank Compute Canada and the Norwegian Metasenter for Com-  
 612 putational Science (NOTUR) for providing computing resources to conduct the  
 613 experiments of this research.

## **References**

- [1] Le Bouthillier A, Crainic TG. A cooperative parallel metaheuristic for the vehicle routing problem with time windows. *Computers & Operations Research* 2005;32:1685–708.
- [2] Crainic TG. Parallel solution methods for vehicle routing problems. In: Golden B, Raghavan S, Wasil E, editors. *The Vehicle Routing Problem: Latest Advances and New Challenges*. New York: Springer; 2008, p. 171–98.
- [3] Alba E, editor. *Parallel Metaheuristics: A New Class of Algorithms*. Hoboken, NJ: John Willey & Sons; 2005.
- [4] Toth P, Vigo D. *The vehicle routing problem*. SIAM Monographs on Discrete Mathematics and Applications; Philadelphia: PA; 2002,.
- [5] Laporte G. Fifty years of vehicle routing. *Transportation Science* 2009;43(4):408–16.
- [6] Jin J, Crainic TG, Løkketangen A. A parallel multi-neighborhood cooperative tabu search for capacitated vehicle routing problems. *European Journal of Operational Research* 2012;222(3):441–51.

- [7] Jin J, Crainic TG, Løkketangen A. A guided cooperative parallel tabu search for the capacitated vehicle routing problem. In: Norsk Informatikkonferanse NIK 2011. Tapir Akademisk Forlag: ISBN 978-82-519-2843-4; 2011, p. 49–60.
- [8] Crainic TG, Nourredine H. Parallel metaheuristics applications. In: Alba E, editor. *Parallel Metaheuristics*. Hoboken, NJ: John Willey & Sons; 2005, p. 447–94.
- [9] Blum C, Roli A. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Comput Surveys* 2003;35(3):268–308.
- [10] Glover F, Laguna M. *Tabu Search*. Kluwer; 1997.
- [11] Voß S. Solving quadratic assignment problems using the reverse elimination method. In: Nash S, Sofer A, editors. *The Impact of Emerging Technologies on Computer Science and Operations Research*. Boston: Kluwer; 1995, p. 281–96.
- [12] Toth P, Vigo D. The granular tabu search and its application to the vehicle routing problem. *INFORMS Journal on Computing* 2003;15:333–46.
- [13] Yellow PC. A computational modification to the savings method of vehicle scheduling. *Operational Research Quarterly* 1970;21(2):281–93.
- [14] Groër C, Golden B, Wasil E. A parallel algorithm for the vehicle routing problems. *INFORMS Journal on Computing* 2011;23:315–30.
- [15] Savelsbergh MWP. The vehicle routing problem with time windows: minimizing route duration. *INFORMS Journal on Computing* 1992;4:146–54.
- [16] Potvin JY, Rousseau JM. An exchange heuristic for routing problems with time windows. *Journal of the Operational Research Society* 1995;46:1433–46.
- [17] Taillard E, Badeau P, Gendreau M, Geurtin F, Potvin JY. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science* 1997;31:170–86.



- [18] Flood MM. The traveling-salesman problem. *Operations Research* 1956;4:61–75.
- [19] Kytöjoki J, Nuortio T, Bräysy O, Gendreau M. An efficient variable neighborhood search heuristic for very large scale vehicle routing problems. *Computers & Operations Research* 2007;34:2743–57.
- [20] Groër C, Golden B, Wasil E. A library of local search heuristics for the vehicle routing problem. *Mathematical Programming Computation* 2010;2:79–101.
- [21] Prins C. A simple and effective evolutionary algorithm for the vehicle routing problem. *Computers & Operations Research* 2004;31(12):1985–2002.
- [22] Vidal T, Crainic TG, Gendreau M, Lahrichi N, Rei W. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research* 2012;60(3):611–24.
- [23] Golden BL, Wasil EA, Kelly JP, I. M. Chao IM. The impact of meta-heuristics on solving the vehicle routing problem: algorithms, problem sets, and computational results. In: Crainic T, Laporte G, editors. *Fleet management and logistics*. Boston: Kluwer; 1998, p. 33–56.
- [24] Li F, Golden BL, Wasil EA. Very large-scale vehicle routing: New test problems, algorithms, and results. *Computers & Operations Research* 2005;32:1165–79.
- [25] Mester D, Bräysy O. Active-guided evolution strategies for large-scale capacitated vehicle routing problems. *Computers & Operations Research* 2007;34:2964–75.