# A Guided Cooperative Search for the Vehicle Routing Problem with Time Windows

Alexandre Le Bouthillier<sup> $*\ddagger$ </sup> Teodor G. Crainic<sup> $*\dagger$ </sup> Peter Kropf<sup> $\ddagger$ </sup>

\*Centre de recherche sur les transports, Université de Montréal C.P.6128, succursale Centre-ville, Montréal, QC, CANADA H3C 3J7 {alexleb,theo}@crt.umontreal.ca

<sup>†</sup>Département de management et technologie, Université du Québec à Montréal C.P. 8888, Succursale Centre-Ville, Montral, QC, H3C 3P8 theo@crt.umontreal.ca

<sup>‡</sup>Département d'informatique et de recherche opérationnelle, Université de Montréal C.P.6128, succursale Centre-ville, Montréal, QC, CANADA H3C 3J7 {lebouthi,kropf}@iro.umontreal.ca

### Abstract

We present a framework for a *guided parallel cooperative search* for combinatorial optimization based on the *central memory* multi-thread cooperative search concept. The proposed mechanism endows cooperative search with capabilities to create new information and guide the global search. Based on patterns shared by several of the solutions exchanged among the search threads, information is sent to individual meta-heuristics about promising and unpromising patterns of the solution space. This process results in a better coordination between the individual methods and the possibility to guide the diversification and intensification of the global search. We apply this method to the Vehicle Routing Problem with Time Windows. Experimental results on an extended set of benchmark problem sets illustrate the benefits of the proposed methodology.

**Keywords :** parallel computation, parallel cooperative search, Vehicle Routing Problem with Time Windows

### 1 Introduction

Efficiency and robustness are the primary qualities of good meta-heuristics and parallel computing may significantly enhance these characteristics. Three forms of parallelism can be applied to meta-heuristics: (i) Division of computer-intensive tasks at a low algorithmic level, (ii) Explicit domain decomposition of the solution or search space, (iii) Multi-thread search. An independent multi-thread search produces the best solution among those found by each independent method. Multi-thread cooperative search implements a mechanism that allows information (e.g., solutions) to be exchanged among the search threads. Several such mechanisms have been proposed in the literature and in many cases cooperative multi-search offered superior performance in terms of better solutions and shorter resolution times (Crainic and Toulouse 2003).

While all parallelization strategies may speed up the resolution, cooperative search methods may also increase the robustness of the global search (Crainic and Toulouse 2003). Cooperative search combines the efforts of several independent meta-heuristics by using a so-called solution warehouse (according to the information stored, the names "adaptive" and "central memory" are also used). This device receives "good" solutions from the search threads and, on demand and according to their own internal logic, provides them in return with solutions to, for example, diversify the search. This simple mechanism allows the asynchronous communication and exchange of solutions that influences the search trajectory of each method. Enhanced with simple extraction rules for the returned solutions, simple cooperation was successfully used to address a number of difficult combinatorial problems: network design, multi-commodity location-allocation (Crainic, Toulouse and Gendreau 1995), circuit partitioning and Vehicle Routing Problem with Time Windows (VRPTW) (Le Bouthillier and Crainic 2005).

From a conceptual point of view, the solution warehouse contains useful indirect information for the global search. Thus, for example, the history of solution discovery and the frequency of appearance of certain attributes in solutions of particular quality can be used to create new information about the search space. This process of information creation can transform the solution warehouse in an intelligent data warehouse that holds more complex information and guides individual methods towards promising or unexplored regions.

The goal of this paper is to present a mechanism that endows cooperative search with capabilities to create new information and guide the global search. The proposed identification pattern mechanism sends information to individual meta-heuristics about promising and unpromising patterns of the solution space. By fixing or prohibiting specific solution attribute values in particular search methods, we can focus the search to desired regions. This mechanism may thus be applied to enforce a better coordination between the individual methods and control the diversification and intensification of the global search. We apply this mechanism to the Vehicle Routing Problem with Time Windows (VRPTW). Experimental results on an extended set of benchmark problem sets illustrate the benefits of the proposed methodology.



Figure 1: Abstract Framework of a Cooperative Search

The main contribution of this paper is the introduction of an enhanced cooperative search mechanism that creates new information from exchanged solutions and perform global intensification and diversification phases. The proposed framework does not suppose any specific problem structure and may thus be applied to a wide range of combinatorial problems. The paper also reports very good solutions on benchmark problem sets for the VRPTW: 139 best known results found and a new best known average on the problems of size 200 and 1000.

The paper is organized as follow. Section 2 briefly presents a framework of cooperative search. Section 3 introduces the pattern identification mechanism and the global search phases. Section 4 details the implementation of a guided cooperative search applied to the VRPTW. Section 5 presents the computational results and analyzes them both from the point of view of solving the VRPTW and from that of the performance of the parallel strategy. Conclusions and perspectives are the subject of the last section.

### 2 A Cooperative Search Framework

Figure 1 illustrates the cooperative search framework made up of a number of independent processes (threads) of possibly different types, which communicate through the solution warehouse. A search thread either heuristically constructs new solutions, or executes a neighborhood-based improving meta-heuristic through the search space, or implements a population-based meta-heuristic (e.g., evolutionary algorithms, Scatter Search, Path relinking), or performs post-optimization procedures (e.g., intensive local search) on solutions in the solution warehouse. Improving meta-heuristics, such as Tabu search, aggressively explore the search space, while population-based methods contribute toward increasing the diversity of solutions exchanged among the cooperating methods. When the same meta-heuristic is used by several search threads, the initial solution and particular setting of a number of important search parameters differentiate each search thread from the others.

The cooperation aspect of the parallelization scheme is achieved through asynchro-

nous exchanges of information. Information is shared through a solution warehouse or pool of solutions. In this scheme, whenever a thread desires to send out information (e.g., when a new local optimum is identified), it sends it to the pool. Similarly, when a thread accesses outside information (to diversify the search, for example), it reaches out and takes it from the pool. Communications are initiated exclusively by the individual threads, irrespective of their role as senders or receivers of information. No broadcasting is taking place and there is no need for complex mechanisms to select the threads that will receive or send information and to control the cooperation. The solution warehouse is thus an efficient implementation device that allows for a strict asynchronous mode of exchange, with no predetermined connection pattern, where no process is interrupted by another for communication purposes, but where any thread may access at all times the data previously sent out by any other search thread.

The solution warehouse keeps the information in an order appropriate for the exchange mechanism considered. To fully characterize the cooperation process, one has to specify (i) the information which is to be shared; (ii) the particular methods that makeup the cooperative search; (iii) when and how communications occur; (iv) the utilization each thread makes of the imported information (Crainic and Toulouse 2003). The information exchanged among cooperating procedures has to be meaningful, in the sense that it has to be useful for the decision process of the receiving threads. Information indicative of the current status of the global search or, at least, of some individual meta-heuristic is, in this sense, meaningful.

Two main classes of cooperation mechanisms are found in the literature, based on partial and complete solutions, respectively. Adaptive memory methods (Rochat and Taillard 1995) store partial elements of good solutions and combine them to create new complete solutions that are improved then by the cooperating threads. Central memory approaches exchange complete elite solutions among neighborhood and populationbased meta-heuristics (Crainic, Toulouse, Gendreau 1995, Crainic and Toulouse 2003, Le Bouthillier and Crainic 2005).

In a simple central memory cooperation scheme (e.g., Le Bouthillier and Crainic 2005), threads share information about their respective good solutions identified so far. When a thread improves the imported solution or when it identifies a new best solution, it sends it to the solution warehouse. This scheme is intuitive and simple, and it satisfies the meaningfulness requirement. The selection of the methods involved in cooperative search should be oriented toward obtaining: (i) Good quality solutions; (ii) A broad diversity of solutions to facilitate the discovery of promising regions; (iii) The rapid production of intermediate solutions to feed the information exchange mechanisms; (iv) A mechanism that combine various solutions to create diversity; (v) A mechanism that has the ability to escape local optima.

The solution warehouse is thus the core of the cooperation mechanism. It keeps good solutions and is dynamically updated by the independent search processes. The pool of solutions forms an elite population from which the independent procedures require solutions at various stages of execution. Solutions are ordered according to a

predetermined utility measure that quantifies the solution quality. The solution utility can be its objective value or a combined measure of solution's properties.

Independent methods send their improved solutions to the post-optimization algorithms. Theses solutions are considered in-training until they have been post-optimized and sent as Adult solutions to the solution warehouse. Duplicate solutions received in the solution warehouse are eliminated.

All requests for solutions initiated by the independent processes are sent to the solution warehouse that responds by sending an Adult solution. Solutions are selected randomly according to probabilities biased toward the best based on the same function used to order solutions in the solution warehouse.

The population size in the solution warehouse is set relatively to the problem size and the worst results are eliminated as needed. No direct communications take place between processes thus enforcing their independence and the asynchronous mode of exchange. This scheme makes the cooperation design simpler and, eventually, allows easy modification of the parallel system by adding new methods or dropping inefficient ones. Moreover, it does not assume any specific problem, which makes it equally relevant for problems where solution components may be easily defined (e.g., the routes in vehicle routing problems) and for problems where such structures are much less apparent (e.g., network design). The goal now is to improve upon this simple cooperating scheme by extracting new knowledge from the information exchanged, to yield a more efficient global search.

## 3 Pattern Identification Mechanism and Global Search

In this section, we introduce a mechanism to extract knowledge from the information exchanged and guide each search method towards promising or unexplored regions of the solution space. It uses a pattern identification mechanism on the solutions present in the solution warehouse that is then used to fix or prohibit specific solution attributes (e.g., arcs in network-based problems) for part of the search performed by particular individual meta-heuristics. One may therefore constrain the search space of particular meta-heuristics in cooperation and thus perform global intensification and diversification phases to guide the exploration of the solution space and control the quality and diversity of the solution warehouse population.

To enhance the clarity of the presentation, we focus in the following on problems that may be described in terms on inclusion/exclusion of arcs in given networks. Network design, traveling salesman, and vehicle routing problems are important combinatorial problem classes that belong to this category. We emphasize, however, that the concepts presented in this section may be extended to any problem definition and solution attribute.

### 3.1 Pattern definition

For the class of problems described above, a very simple and general pattern definition may be based on the inclusion of arcs in particular solutions.

Consider the frequency of inclusion of arcs in a given subset of the solution warehouse. In particular, this subset may be the entire population, an elite (e.g., with solution in the 10% best), average (between the 10% and 90% best) or worst (the last 10%) group of solutions. An arc with a high frequency in a given group signals that the meta-heuristics participating to the cooperation have often produced solutions that include that arc. Tagging solutions to identify the last algorithm that sent it, one may induce similar information by subset of participating algorithms as well. When the frequency of inclusion of several arcs is considered, patterns emerge among the solutions of the solution warehouse or the specific group examined.

We define a pattern of length n as a subset of arcs of cardinality n = 1, 2, ...,maximum number of arcs in the problem definition. A frequent (infrequent) pattern relative to a set of arcs is built of arcs with high (low) frequency of appearance in the solutions of the set. High (low) frequency arcs are selected sequentially in decreasing from the highest (increasing from the lowest) frequency value.

We may select patterns from specific subpopulations (e.g., elite, average, and worst) and compare the rate of appearance of a specific pattern between them. These comparisons form the basis of the guidance mechanism proposed in this paper.

### 3.2 Comparing pattern-appearance frequencies among subpopulations

We define an *in-pattern* as a pattern that actually appears in at least one solution of the subset of solutions considered, as opposed to a *statistical pattern* that does not appear in any solution of the subset. A statistical pattern is thus only a consequence of the statistical process of accounting for the frequency of individual arcs

Consider an *in-pattern* of length n common to the three sets of elite, average, and worst solution groups. Two meaningful situations can occur with respect to the frequency of appearance of the pattern in these sets as one move from the elite to the average to the worst subpopulation: The frequency is either increasing or it is stable or decreasing. We call the *in-pattern unpromising* in the first case and *promising* in the latter.

Promising and unpromising patterns may then be used to constrain for a certain time the search space of particular methods in the cooperation and thus induce global intensification and diversification phases, as described in the next subsection.

#### 3.3 Using pattern identification - Global search

Global search intensification and diversification phases may be triggered by fixing (including) or prohibiting (excluding) arcs in the solutions a given meta-heuristic explores during a certain period.

Consider an unpromising pattern of length n. To intensify the search around solutions with "good" attributes, one prohibits the arcs defining the pattern. On the other hand, fixing these arcs for a number of iterations will diversify the search relative to the current set of "good" attribute values. Symmetrically, given a promising pattern of length n, one intensifies the search by fixing the arcs in the pattern, while prohibiting them diversifies it.

We define the global search as the cumulative search effort of the individual methods. To prevent individual methods from converging too rapidly, we favor diversification by prohibiting arcs during the initial phases of the global search. Later on, we encourage intensification by fixing promising arcs to enforce the exploration of theses promising regions of the solution space. We also vary the length of the patterns as a mean of modulating the intensity of global diversification and intensification phases and thus influence the evolution of the diversity and quality of solutions in the solution warehouse.

At the beginning of the search there is not sufficient information gathered as the solution space is insufficiently explored. Therefore, the solution warehouse cannot be considered representative of the search of the individual meta-heuristic, even when the population is diverse. Consequently, initially, we build patterns of increasing length from the average subpopulation only, in order to identify more rapidly promising ones. As the search progresses, patterns from elite and worst subpopulations are built as described above. When a statistical pattern is found, we reduce the length of the pattern until an in-pattern is obtained.

Several modifications must be made to the initial framework presented in the previous section. The solution warehouse now includes a process to identify and manage patterns. This process includes decisions on when to compute particular patterns and to what individual meta-heuristic to send them. It becomes, in fact, the "new" global meta-heuristic corresponding to the global guided cooperative search. As for the solution warehouse, it now contains solutions and pattern information, thus a data warehouse. Communications from the solution warehouse to the individual meta-heuristics are modified as well. The appropriate pattern and instructions on fixing or prohibiting arcs are sent along with the solution (selected according to the original criteria). With respect to the individual meta-heuristics, each needs to be modified to cope with the instructions relative to fixing or prohibiting arcs. Figure 2 illustrates the guided cooperative search framework as applied to the VRPTW.



Figure 2: Guided Cooperative Search applied to VRPTW

## 4 Guided Cooperative Search for the VRPTW

To illustrate the mechanism described in the previous sections, we developed a guided cooperative search for the Vehicle Routing Problem with Time Windows (VRPTW). The application is described in this section. Experimental results are discussed in the next section.

The VRPTW is a well-known combinatorial problem that has been extensively studied and is thus well suited for benchmarking. We address the single depot VRPTW, as illustrated in Figure 3, where one is given a set of customers with known positive demands and specific time intervals when service can be provided. A fleet of homogeneous vehicles of known capacity is available at a given depot to perform this service. The objective is to find a set of closed routes (or tours) that start and end at the depot within its opening hours, such that the total cost of performing the service is minimized, customers are visited and served during their specified time windows, and vehicles are not overloaded.

In the problem version we address, cost is a combination of two factors: the number of vehicles (routes) used and the total distance traveled. A high cost is associated with vehicle utilization to enforce the search towards solutions with reduced number of vehicles. Each customer is visited only once. A vehicle cannot arrive later than the customer's closing time, but is allowed to arrive before the associated opening time, in which case it waits, without explicit penalty, until the customer is ready. Once the service starts, it is carried on until completion, even if the service ending time might be later than the expiration of the time window. Cordeau et al. (2002) review problem variants, formulations, and solution methods for the VRPTW.

Le Bouthillier and Crainic (2005) presented a cooperative parallel method for the VRPTW based on the simple solution warehouse mechanism presented in Section 2. The cooperation involved two Tabu search methods that perform well sequentially, the Unified Tabu (Cordeau, Laporte, and Mercier 2001) and Taburoute (Gendreau, Hertz, and Laporte 1994), two simple evolutionary algorithms with Order and Edge Recombi-



Figure 3: Best known solution for problem C1-6-4

nation Crossover, respectively, as well as a number of post-optimization methods (2-opt, 3-opt, or-opt, and Ejection Chains) that were used to reduce the number of vehicles and the total traveled distance. Four simple construction algorithms were used to provide initial solutions to the population.

We applied the proposed cooperation framework to the parallel method of Le Bouthillier and Crainic (2005). In the VRPTW context, the definition of an in-pattern of length nis straightforward. The problem is defined on a network, where an arc corresponds to a possible movement between two customers or between a customer and a depot. The procedures introduced in the previous section were used to define patterns.

Fixing and prohibiting arcs in the solutions explored by the four meta-heuristics is straightforward as well. Fixing (including) arcs always leads to a non empty solution space that, at the most extreme, may reduce to a single solution that represents a very (too) long pattern. Prohibiting (excluding) patterns may lead to an empty feasible solution space when patterns are too long. To avoid both situations and strike a balance between the number of feasible solutions and the size of the constrained solution space, we limit the pattern length to 25% of the size of the problem.

The computing time allocated to the cooperative method is divided into four phases: Two phases of diversification at the beginning to broaden the search, followed by two intensification phases to focus the search around promising regions. The four phases proceed as follows:

- Phase I. Built unpromising in-patterns of frequent arcs in the average subpopulation and prohibit them in the independent meta-heuristics;
- Phase II. Prohibit arcs from frequent unpromising in-patterns from the worst

subpopulation;

- Phase III. Work with the average subpopulation and fix arcs from frequent promising in-patterns;
- Phase IV. Build frequent promising in-patterns from the elite sub-population and fix the arcs for the meta-heuristic searches.

Pattern lengths are explored in decreasing length in the first two phases and in increasing length in the last two phases, by increments of 1 unit.

### 5 Computational Experiments

The experimentation has a dual objective. On the one hand, we aim to compare the guided cooperative search to the simple version and to the best performing methods proposed in the literature for the VRPTW and, thus, to validate our claim that the proposed method offers competitive performance in terms of both solution quality and computational effort. On the other hand, we also aim to evaluate the impact of guiding the search toward or away from specific patterns and performing diversification and intensification to control the entropy of the population.

A different search method, Taburoute, an Unified Tabu Search algorithm, and two evolutionary algorithms (OX and ER), is run on each of the four processors. The solution warehouse, the post-optimization procedures, the pattern identification method, and the construction methods are run on another processor for a total of five processors in this study.

For Taburoute we use the parameter settings indicated in the original paper for the VRP (Gendreau, Hertz, and Laporte 1994). Tabu tags were set to a length varying between 5 and 10 iterations, 15% of the nodes were evaluated in the p-neighborhood dimensions, the initial solution was selected out of 15 initial generated solutions and a solution from the solution warehouse, a penalty of 1 was used for frequently moved. The parameters for the value function presented in the initial article by Cordeau, Laporte, and Mercier (Cordeau, Laporte, and Mercier 2001) (alpha:=1; beta=1; gamma=1) were used for the Unified Tabu. Finally, an arc mutation probability of 1 percent was used on temporary copies of the parents for the crossovers used by the evolutionary algorithms.

Tests have been carried out on the standard set of test problems proposed by Solomon (1987). The set contains 56 problems of 100 customers each. We also used the extended set produced by Homberger and Gehring (1999) with 300 problem instances that vary from 200 to 1000 customers. The Solomon and extended problems are divided into six categories, named C1, C2, R1, R2, RC1, and RC2. For all problem instances, customers are distributed in a [0,100] square unit. The customers in sets C are clustered together, while those in sets R are distributed randomly. Problems in sets RC combine the two

characteristics. Time windows at the depot are relatively small for problems of type 1, to allow fewer customers to be served by each route; time windows are larger for problems of type 2. The service time is of 10 units by customer for problems of type R and RC, and of 90 units for class C.

Solutions in the solution warehouse are sorted, first by the number of vehicles, second by a weighted sum, C(p), of attributes: the total time required to serve all customers, the associated total distance and total waiting time at customers, and the sum of the slack left in each time window: C(p)=W1\*totalTime + W2\*totalDistance + W3\*totalWait + W4\*totalSlack.

Parameters W1 to W4 were set to 1 in all the reported experiments. This measure combined with the number of vehicles gives us an overall idea of the solution quality (totalTime and totalDistance) and flexibility (totalWait and totalSlack). The last two measures indicate how much slack there exists in the solution and how easily feasible neighboring solutions may be explored.

In previous research (Le Bouthillier and Crainic 2005), cooperative search was found to provide faster results of equivalent or better quality that each of its independent searches. We therefore compare only simple and guided parallel cooperative searches.

Runs of 12 min wall-clock time were performed by the cooperative meta-heuristics for each of the 100 city problems. Longer running times, equal to those reported by Homberger and Gehring (1999) were allowed for the larger problem instances. These times go up to 50 min wall-clock time for the 1000 city problem. To be able to compare to the wall-clock time of the simple cooperative search, we created virtual machines under VMware and forced their CPU clocks to emulate the machines used for the previous study (Le Bouthillier and Crainic 2005). Using faster virtual or physical machines will only decrease the wall-clock time for the same results. A virtual cluster of five Pentium III 850MHz CPU computers with 512MB of RAM under Linux was used. Computations of distances were carried out in double precision. The implementation is machine independent and can be run with MPICH on Unix, Windows, or Linux.

The four global phases that prohibit or fix arcs were each allocated 1/4 of the total wall-clock execution time. In-pattern lengths were at most 25% of the problem size.

Tables 1 and 2 display the average results per class for the cumulative number of vehicles (CNV) and distance (CTD) for the standard Solomon problems and for extended set of problems, respectively. Best results are shown in bold face and, in Table 1, authors are presented in decreasing CNV/CTD value.

Results of the simple Cooperative Search (LC03) and Guided Cooperative Search (LCK05) are compared to those of the best methods (published or not) for the VRPTW on the benchmark site of SINTEF (http://www.sintef.no/static/am/opti/projects/top/) on February 15, 2005: The unified Tabu Search of Cordeau, Laporte and Mercier (2001, denoted CLM), the evolutionary algorithm of Homberger and Gehring (1999, denoted HG), the two stage hybrid local search of Bent and Van Hentenryck (2001, denoted

Author	R1	R2	C1	C2	RC1	RC2	CNV/CTD	Experiment
RT	12.25	2.91	10.00	3.00	11.88	3.38	415	SG. 100MHz
	$1\ 208.50$	961.72	828.38	589.86	$1\ 377.39$	$1\ 117.44$	$57 \ 231$	1 run, 92.2 Min
CLM	12.08	2.73	10.00	3.00	11.50	3.25	407	n/a
	$1\ 210.14$	969.57	828.38	589.86	$1 \ 389.78$	$1 \ 134.52$	57 555	Sun U2 300MHz
LC03	12.08	2.73	10.00	3.00	11.50	3.25	407	5xP850 MHz,
	$1\ 209.19$	963.62	828.38	589.86	$1 \ 389.22$	$1\ 143.70$	57412	$1 \operatorname{run}, 12 \operatorname{Min}$
H99	11.92	2.73	10.00	3.00	11.63	3.25	406	P200 MHz,
	$1\ 228.06$	969.95	828.38	589.86	$1 \ 392.57$	$1\ 144.43$	57 876	10  runs, 13  Min
GH01	12.00	2.73	10.00	3.00	11.50	3.25	406	4 P400 MHz,
	$1\ 217.57$	961.29	828.63	590.33	$1 \ 395.13$	$1 \ 139.37$	57 641	5 runs, 13.5 Min
B01	11.92	2.73	10.00	3.00	11.50	3.25	405	P200 MHz,
	$1\ 222.12$	975.12	828.38	589.86	$1 \ 389.58$	$1\ 128.39$	57 710	1 run, 82.5 Min
LCK05	11.92	2.73	10.00	3.00	11.50	3.25	405	5xP850 MHz,
	$1 \ 214.20$	954.32	828.38	589.86	$1 \ 385.30$	$1\ 129.43$	$57 \ 360$	$1 \operatorname{run}, 12 \operatorname{Min}$
BVH	12.18	2.73	10.00	3.00	11.50	3.25	405	Sun U10 $440 Mhz$
	$1\ 231,\!08$	954.18	828.38	589.86	$1 \ 384.17$	$1\ 124,\!47$	57 272	$5~\mathrm{run},120~\mathrm{Min}$

Table 1: Comparison of average results on 100-customer problems

BVH), the Active Guided Evolution Strategies of Mester and Bräys (2004, denoted MB), Rochat and Taillard (1995, denoted RT), Gehring and Homberger (2001, denoted GH01), and Bräysy (2003, denoted B01). We refer the reader to Le Bouthillier, Crainic, and Kropf (2005) for detailled results on standard and extended Solomon's benchmark sets of problems.

The proposed method, LCK05, yields very good results. It appears in second place in both tables according to the CNV/CTD ratio. Relative to the standard Solomon problems, the method we propose yields a total number of vehicles of 405, which is the lowest number obtained so far and makes the guided cooperative serach one of the best meta-heuristics currently available for the VRPTW.

We observe that, compared to the simple cooperative search, the new method reduced the number of vehicles by 2 and the distance by 52.38 units. We found 24 of the best known results and report the best known CNV. The guided cooperative search reports a new best average for the R1 problem class. For C1 and C2 problem classes of 100 customers, almost all methods found the best known average number of vehicles and the total distance. For the other classes of problems, we found the best known number of vehicles in all instances and at most 0.16% of increase in distance when compared to the other methods.

Problem	MB		LC03		LCK05		HG	
	CNV	CTD	CNV	CTD	CNV	CTD	CNV	CTD
200  total	694	168573	694	$173061,\!631$	<b>694</b>	169958, 49	694	173312,32
400  total	1389	390386	1390	410329,98	1389	396611, 15	1388	409763, 38
600  total	2082	796172	2088	840582,74	2086	809493,46	2076	851680, 33
800  total	2765	1361586	2766	1475435, 14	2761	1443399,01	2755	$1479801,\!56$
1000  total	3446	2078110	3451	2225366, 13	3442	2133644,57	3439	$2236582,\!89$
Total	10376	4794827	10389	$5124775,\!621$	10372	$4953106,\!68$	10352	$5151140,\!48$

Table 2: CNV/CTD for 200-1000 customers

Table 2 displays the results for the cumulative number of vehicles (CNV) and distance (CTD) aggregated by problem size for the extended set of problems. There are not

many authors that addressed the entire problem set, which explains the limited number of entries in the table (the best methods were selected).

We observe that we obtain a new best known CTD value for problems of size 200. For the other problem classes, the difference to the other methods is less than 0.48% in terms of the cumulative number of vehicles. In all cases, we improved upon the simple cooperative search. Since we report a slightly higher total number of vehicles compared to HG, we also report a reduction in the total cumulative distance by 3.94% for all classes.

## 6 Conclusions and Perspectives

We presented an enhanced cooperative search mechanism that creates new information from exchanged solutions and thus performs global intensification and diversification phases. The proposed framework does not suppose any specific problem structure, even though we illustrated the methodology in the context of problems defined by the inclusion/exclusion of arcs in particular networks. We also applied this methodology to the VRPTW and report very good solutions on the extended benchmark problem sets: 139 best known results found and the best known average on the problems of size 200 and 1000.

Experimental results showed that pattern identification method yields good information to guide the global search. Patterns of attributes may be constructed independently of particular solution structures and applied to a wide range of combinatorial problems. Patterns of attributes could then be used in various sequential and parallel methods to orient the search.

Cooperative search is thus quite simple to implement (represents less than 10% of the project in line of code) and the pattern identification method provides an easy mechanism to constrain the search within promising regions of the search space or away from unpromising regions. Very good quality solutions were found in linear speed up by combining good off-the-shelf methods without any particular parameter tuning. The quality of the individual methods influences the global search quality. Improved solutions may be found, however, by generating new information from the frequency of pattern appearance in the best solutions visited and by using the new information to guide the global search.

### 7 Acknowledgements

We are grateful to Jean-François Cordeau for the generous offer of sharing his Unified Tabu search code and to Michel Gendreau for his ideas and assistance with the original version of Taburoute. This research has been partially supported by the Natural Sciences

and Engineering Council of Canada through its Discovery Grants program.

## References

- Bent, R. and Van Hentenryck, P. (2003). A Two-Stage Hybrid Local Search for the Vehicle Routing Problem with Time Windows. to appear in Transportation Science.
- Bräysy, O., Hasle, G., and Dullaert, W. (2003). A multi-Start Local Search Algorithm for the Vehicle Routing Problem with Time Windows. to appear in European Journal of Operational Research.
- Bräysy, O. (2001). A Reactive Variable Neighborhood Search Algorithm for the Vehicle Routing Problem with Time Windows. *To appear in INFORMS Journal on Computing.*
- Cordeau, J.-F., Desaulniers, G., Desrosiers, J., Solomon, M. M., and Soumis, F. (2002). The VRP with Time Windows. In Toth, P. and Vigo, D., editors, *The Vehicle Routing Problem*, SIAM, Monographs on Discrete Mathematics and Applications, chapter 7, pages 157–193. SIAM, Philadelphia, PA.
- Cordeau, J.-F., Laporte, G., and Mercier, A. (2001). A Unified Tabu Search Heuristic for Vehicle Routing Problems with Time Windows. *Journal of the Operational Research Society*, 52:928–936.
- Crainic, T.G. and Toulouse, M. (2003). Parallel Strategies for Meta-heuristics. In F. Glover and G. Kochenberger, editors, *State-of-the-Art Handbook in Metaheuristics*, pages 475–513. Kluwer Academic Publishers, Norwell, MA.
- Crainic, T.G., Toulouse, M., and Gendreau, M. (1995). Parallel Asynchronous Tabu Search for Multicommodity Location-Allocation with Balancing Requirements. Annals of Operations Research, 63:277–299.
- Gehring, H. and Homberger, J (2001). A parallel two-phase metaheuristic for routing problems with time windows. *Asia-Pacific Journal of Operational Research*, 18(1):35–47.
- Gendreau, M., Hertz, A., and Laporte, G. (1994). A Tabu Search Heuristic for the Vehicle Routing Problem. *Management Science*, 40:1276–1290.
- Homberger, J. and Gehring, H. (1999). Two Evolutionary Metaheuristics for the Vehicle Routing Problem with Time Windows. *INFOR*, 37:297–318.
- Le Bouthillier, A. and Crainic, T.G. (2005). A Cooperative Parallel Meta-Heuristic for the Vehicle Routing Problem with Time Windows. *Computers & Operations Research*, 32(7):1685–1708.

- Le Bouthillier, A., Crainic, T.G., and Kropf, P. (2005). Towards a Guided Cooperative Search. Publication CRT-05-09, Centre de recherche sur les transports, Université de Montréal, Montréal, QC, Canada.
- Rochat, Y. and Taillard, É.D. (1995). Probabilistic Diversification and Intensification in Local Search for Vehicle Routing. *Journal of Heuristics*, 1(1):147–167.
- Solomon, M.M. (1987). Time Window Constrained Routing and Scheduling Problems. Operations Research, 35:254–265.