



CIRRELT

Centre interuniversitaire de recherche
sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

Parallel Meta-Heuristics and Cooperative Search

Teodor Gabriel Crainic

August 2017

CIRRELT-2017-58

Bureaux de Montréal :
Université de Montréal
Pavillon André-Aisenstadt
C.P. 6128, succursale Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Parallel Meta-Heuristic and Cooperative Search

Teodor Gabriel Crainic*

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT) and Department of Management and Technology, Université du Québec à Montréal, P.O. Box 8888, Station Centre-Ville, Montréal, Canada H3C 3P8

Abstract. The chapter presents a general view of parallel meta-heuristics for optimization. It recalls the main concepts and strategies in designing parallel meta-heuristics and identifies trends and promising research directions. The focus is on cooperation-based strategies, which display remarkable performances, in particular strategies based on asynchronous exchanges and the creation of new information out of exchanged data to enhance the global guidance of the search.

Keywords: Parallel computation, parallel meta-heuristics, functional decomposition, search space separation, independent multi-search, synchronous and asynchronous cooperative search, integrative cooperative search.

Acknowledgments. Partial funding for this project has been provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), through its Discovery Grant and the Discovery Accelerator Supplements programs, and the Strategic Clusters program of the Fonds de Recherche Québécois Nature et Technologies (FRQNT) The author thanks the two institutions for supporting this research.

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

* Corresponding author: TeodorGabriel.Crainic@cirrelt.ca

1 Introduction

The development of meta-heuristics that take advantage of parallel computing aims for two major goals. The first is common to all parallel computing development efforts: solve larger problem instances, faster. That is, address larger problem instances than what is achievable by sequential methods, and do this in reasonable computing times. The second is proper to approximate solution methods, e.g., simple heuristics, meta-heuristics, and matheuristics, and it concerns the method's so-called robustness, that is, its capability to offer a consistently high level of performance over a wide variety of problem settings and instance characteristics. In appropriate settings, e.g., the cooperative multi-search strategies (Section 6), parallel meta-heuristics proved to be much more robust than sequential versions. Moreover, they also generally require less extensive, and expensive, parameter-calibration efforts.

The objective of this chapter is to present an overview of parallel meta-heuristics field in a unified manner. It thus recalls the main concepts and general strategies for the design of parallel meta-heuristics, including the main approaches to instantiate them for neighborhood- and population-based meta-heuristics. Note that the chapter focuses on the *design* of the new class of algorithms parallel meta-heuristics create, and, thus, not on their implementation on particular computing architectures. We do, however, identify new trends, challenges, and opportunities that some of the new computing-platform developments bring to the field. We complete the chapter with a number of major open questions and research challenges.

As the chapter follows and updates Crainic and Toulouse (2010), it focuses on more recent developments (typically, from 2005 to 2017) and, in particular, on cooperation-based strategies, which display remarkable performances for a broad range of optimization problems. In addition to the references provided in the following sections, the reader may consult a number of surveys, taxonomies, and syntheses, e.g., Alba (2005); Alba et al (2013); Crainic and Hail (2005); Crainic (2008); Crainic and Toulouse (2010); Crainic et al (2014); Crainic (2017); Melab et al (2006); Pedemonte et al (2011); Talbi (2009).

The chapter is organized as follows. Section 2 is dedicated to a general discussion of the potential for parallel computing in meta-heuristics, a brief description of performance indicators for parallel meta-heuristics, and the taxonomy used to structure the presentation. Section 3 addresses strategies focusing on accelerating computing-intensive tasks without modifying the basic algorithmic design. Methods based on the explicit separation of the search space are treated in Section 4, while strategies based on the simultaneous exploration of the search space by several independent meta-heuristics constitutes the topic of Section 5. Cooperation principles and strategies are discussed in Section 6 and are detailed in Sections 6.1, 6.2, and 7. We conclude in Section 8.

2 Meta-heuristics and Parallelism

This section is dedicated to a brief overview of the main potential sources for parallel computing in meta-heuristics, followed by a discussion of performance indicators for parallel meta-heuristics. The section concludes with the criteria used in this paper to describe and classify parallelization strategies for meta-heuristics.

2.1 Sources of parallelism

Parallel/distributed/concurrent computing means that several processes work simultaneously on several processors addressing a given problem instance and aiming to identify the best (or a) solution for that instance. Parallelism thus follows from a decomposition of the total computational load and the distribution of the resulting tasks to available processors. According to how “small” or “large” the tasks are in terms of algorithm work or search space, the parallelization is called *fine-* or *coarse-grained*, respectively.

The decomposition may concern the algorithm, the search space, or the problem structure. *Functional parallelism* (Section 3) corresponds to the first case, according to which computing-intensive parts of the algorithm are decomposed into a number of tasks (processes), working on the same data or on a dedicated part, are allocated to different processors and run in parallel, possibly exchanging information. The concurrent execution of the innermost loop iterations, e.g., evaluating neighbors, computing the fitness of individuals, or having ants forage concurrently, provides the main source of functional parallelism for meta-heuristics. This is often also the only source of readily available parallelism in meta-heuristics, the execution of most other steps in the algorithm depending on the status of the search, e.g., what has been performed up to the respective point and the values of the decision variables, which requires either the computation of the previous steps to be completed, or the synchronization of computations; and synchronization generally yields significant delays, which may make such parallel computation non relevant. Traditionally, functional parallelism was therefore interesting as a low-level component of hierarchical parallelization strategies, or when addressing problem settings requiring a significant part of the computing effort to be spent in inner-loop algorithmic components. The rapid development of the utilization of the *graphical processing units* (*GPU*), ubiquitous within most computers, is changing this statement as very impressive reductions in computing times may be obtained (Section 3).

Search *space separation*, constitutes a second major class of parallel strategies. We find under this umbrella the two other cases mentioned above. The general idea is to decompose the problem domain, or the associated search space (for brevity reasons and without loss of generality, the latter term is used in this chapter), and to address the problem on each of the resulting components using a particular solution methodology.

Indeed, there are no data dependencies between the evaluation functions of different solutions and, thus, these may be computed in parallel. Moreover, theoretically, the parallelism in the solution or search space is as large as the space itself when a processor is assigned to each solution. Obviously, the latter strategy is not practical and the search space is separated into subspaces assigned to different processors. Such a separation still leaves a search space for each processor too large for explicit enumeration, however, and, thus, an exact or heuristic search method is required to implicitly explore it.

Space separation is exploited in many of the strategies described in this chapter, but raises a number of issues with respect to an overall meta-heuristic search strategy, e.g., how to separate; how to control an overall search conducted separately on several components of the original space; how to create a complete solution out of the ones obtained on each component; how to allocate resources for an efficient exploration avoiding, for example, regions with poor-quality solutions. The answers to these questions yield several classes of algorithms described in the following sections. These may be grouped, however, into two main approaches: *domain decomposition* and *multi search*. The former explicitly separates the space yielding a number of subproblems to be addressed simultaneously, their solutions being then combined into solutions to the original problem, while the latter performs the separation implicitly, through concurrent explorations by several methods, named *solvers* in the following, which may exchange information or not.

The meta-heuristic or exact solvers involved in a multi-search meta-heuristic may address either the complete problem at hand, or explore partial problems defined by decomposing the initial problem through mathematical programming or attribute-based heuristic approaches. In the former case, the decomposition method implicitly defines how a complete solution is built out of partial ones. In the latter case, some processors work on the partial problems corresponding to the particular sets of attributes defined in the decomposition, while others combine the resulting partial solutions into complete solutions to the original problem. Multi-search strategies, particularly those based on cooperation principles, are at the core of most successful developments in parallel meta-heuristics and are the object of the later sections of this chapter.

2.2 Performance measures

The traditional goal when designing parallel solution methods is to reduce the time required to “solve”, exactly or heuristically, given problem instances or to address larger instances without increasing the computational effort. For exact solution methods that run until the optimal solution is obtained, this translates into the well-known *speedup* performance measure, computed as the ratio between the wall-clock time required to solve the problem instance in parallel with p processors and the corresponding solution time of the best-known sequential algorithm. A somewhat less restrictive measure replaces the latter with the time of the parallel algorithm run on a single processor. See Barr and

Hickman (1993) for a detailed discussion of this issue, including additional performance measures.

Speedup measures are more difficult to define when the optimal solution is not guaranteed or the exact method is stopped before optimality is reached, which is obviously also the case for meta-heuristics. Moreover, most strategies to build parallel meta-heuristics yield solutions that are different in value, composition, or both from those of the sequential versions (when they exist). Hence, an equally important objective for parallel meta-heuristics is to to what extent they outperform their sequential counterparts in terms of solution quality and, ideally, computational efficiency. In other words, the parallel method should not require a higher overall computation effort than the sequential method or should justify the effort by higher quality solutions.

Search robustness is another characteristic increasingly expected of parallel heuristics, robustness with respect to a problem setting being meant in the sense of providing “equally” good solutions to a large and varied set of problem instances, without excessive calibration, neither during initial development, nor when addressing new instances. Multi-search methods, particularly those based on cooperation, generally display a behavior different from those of the sequential methods involved, offering enhanced performances compared to sequential methods and other parallelization strategies in terms of solution quality and method robustness (see Crainic and Toulouse, 1998, 2003, for a discussion of these issues). They are thus generally acknowledged as proper meta-heuristics Alba (2005).

2.3 Parallel meta-heuristics strategies

We adopt the classification of Crainic and Hail (2005), generalizing that of Crainic et al (1997), to describe the different parallel strategies for meta-heuristics. This classification is sufficiently general to encompass the principal parallel meta-heuristic classes, while avoiding a level of detail incompatible with the scope and dimension limits of the chapter.

The three dimensions of the classification define how the global problem-solving process is controlled, how information is exchanged among processes and how, eventually, new information is created, and the diversity of searches involved, respectively. Table 1 synthesizes the dimensions and categories of the classification, which are now detailed.

The first dimension, *Search Control Cardinality*, specifies whether the global search is controlled by a single process or by several processes that may collaborate or not. The two categories are identified as *1-control (1C)* and *p-control (pC)*, respectively.

The second dimension, relative to the type of *Search Control and Communications*, addresses the issue of information exchanges and the utilization of the exchanged in-

Table 1: The parallel meta-heuristics taxonomy

Dimension	Categories			
<i>Control Cardinality</i>	<i>1C</i>		<i>pC</i>	
<i>Control & Commu- nications</i>	<i>RS</i>	<i>KS</i>	<i>C</i>	<i>KC</i>
<i>Differentiation</i>	<i>SPSS</i>	<i>SPDS</i>	<i>MPSS</i>	<i>MPDS</i>

formation to control or guide the search. In parallel computing, one generally refers to *synchronous* and *asynchronous* communications. In the former case, all concerned processes stop and engage in some form of communication and information exchange at moments (number of iterations, time intervals, specified algorithmic stages, etc.) exogenously determined, either hard-coded or imposed by a control (master) process. In the latter case, each process is in charge of its own search, as well as of establishing communications with other processes, and the global search terminates once all individual searches stop. Four categories are defined to reflect the quantity and quality of the information exchanged and shared, as well as the additional knowledge derived from these exchanges (if any); two for synchronous settings, *Rigid (RS)* and *Knowledge Synchronization (KS)*, and, symmetrically, two for asynchronous strategies, *Collegial (C)* and *Knowledge Collegial (KC)*.

More than one solution method or variant (e.g., with different parameter settings) may be involved in a parallel meta-heuristic, and such solvers may be (meta-)heuristics or exact solution methods. The third dimension thus indicates the *Search Differentiation* or diversity: do solvers start from the same or different solutions, and are they the same or not? Note that one characterizes two solvers as “different” even when based on the same methodology (e.g., two tabu searches or genetic algorithms) if they use different search strategies in terms of components (e.g., neighborhoods or selection mechanism) or parameter values. The four classes are: *SPSS*, *Same initial Point/Population, Same search Strategy*; *SPDS*, *Same initial Point/Population, Different search Strategies*; *MPSS*, *Multiple initial Points/Populations, Same search Strategies*; *MPDS*, *Multiple initial Points/Populations, Different search Strategies*, where “point” relates to neighborhood-based, single-solution methods, while “population” is used for population-based ones.

3 Low-Level Parallelization Strategies

Functional-parallelism-based strategies, exploiting the potential for task decomposition within the inner-loop computations of meta-heuristics, aim to accelerate the search, without modifying the algorithmic logic, the search space and behavior of the sequential meta-

heuristic. Hence the label “low level” often associated with such strategies. Typically, the exploration is initialized from a single solution or population, and proceeds according to the sequential meta-heuristic logic, while a number of intensive-computation steps are decomposed and simultaneously performed by several processors.

Most low-level parallel strategies belong to the 1C/RS/SPSS class and are usually implemented according to the classical *master-slave* parallel programming model. A “master” program executes the (1-control) sequential meta-heuristic, separating and dispatching computation-intensive tasks to be executed in parallel by “slave” programs. Slaves perform the tasks and return the results to the master which, once all the results are in, resumes the normal logic of the sequential meta-heuristic. The master thus has complete control on the algorithm execution; it decides the work allocation for all other processors and initiates communications. No communications take place among slave programs.

The neighborhood-evaluation procedure of the local search heuristics, used alone or as component of neighborhood- or population-based meta-heuristics (implementing advanced “schooling” for offspring in the latter case) is generally targeted in 1C/RS/SPSS designs. The master groups the neighbors into tasks and sends them to slaves. Each slave then executes the exploration/evaluation procedure on its respective part of the neighborhood, and sends back the best, or first improving, neighbor found. The master waits for all slaves to terminate their computations, selects the best move and proceeds with the search. The appropriate granularity of the decomposition, that is, the size of the tasks, depends upon the particular problem and computer architecture, but is generally computationally sensitive to inter-processor communication times and work-load balancing. Thus, for example, Davidović and Crainic (2015) discusses several decomposition policies for the permutation-based local search neighborhood applied to the scheduling of dependent tasks on homogeneous processors, and shows that the uniform partition usually called upon in the literature is not appropriate in this context characterized by neighborhoods of different sizes. The authors also show that a fixed coarse-grained non-uniform decomposition, while offering superior results, requires calibration each time the problem size or the number of processors varies. The best performing strategy, called *dynamic fine-grained* by the authors, defines each neighbor evaluation as a single task, the master dynamically dispatching these on a first-available, first-served basis to slave processors as they complete their tasks. The strategy partitions the neighborhood into a number of components equal to the number of available processors, but of unequal size with a content dynamically determined at each iteration. The dynamic fine-grained strategy provides maximum flexibility and good load balancing, particularly when the evaluation of neighbors is of uneven length. The uniform distribution appears more appropriate when the neighbor evaluations are sensibly the same, or when the overhead cost of the dynamic strategy for creating and exchanging tasks appears too high.

Similar observations may be made regarding population-based meta-heuristics. In

theory, all genetic-algorithm operators may be addressed through a 1C/RS/SPSS design, and the degree of possible parallelism is equal to the population size. In practice, the computations associated to most operators are not sufficiently heavy to warrant parallelizing, while overhead costs may significantly reduce the degree of parallelism and increase the granularity of the tasks. Consequently, the fitness evaluation is often the target of 1C/RS/SPSS parallelism for genetic-evolutionary methods, usually implemented using the master-slave model.

The 1C/RS/SPSS parallelism for ant-colony and, generally, swarm-based methods lies at the level of the individual ants. Ants share information indirectly through the pheromone matrix, which is updated once all solutions have been constructed. There are no modifications of the pheromone matrix during a construction cycle and, thus, each individual ant performs its solution-construction procedure without data dependencies on the progress of the other ants. Many parallel ant-colony methods proposed in the literature implement some form of 1C/RS/SPSS strategy according to the master-slave model (e.g., Doerner et al, 2006, and references herein). The master builds tasks consisting of small colonies of one or a few ants, and distributes them to the available processors. Slaves perform their construction heuristic and return their solution(s) to the master, which updates the pheromone matrix, returns it to the slaves, and so on. To further speed up computation, the pheromone update can be partially computed at the slave level, each slave computing the update associated to its solutions. This fine-grained version with central matrix update outperformed the sequential version of the algorithm in most cases. It is acknowledged, however, that it does not scale when implemented on “traditional” (i.e., exploiting the central processing units - CPUs) processors, and that, similarly to other meta-heuristics, it is outperformed by more advanced multi-search methods.

Scatter search and path relinking implement different evolution strategies, where a restricted number of elite solutions are combined, the result being enhanced through a local search or a full-fledged meta-heuristic, usually neighborhood-based. Consequently, the 1C/RS/SPSS strategies discussed above apply straightforwardly as in García-López et al (2005, 2006, 2003) for the p -median and the feature-selection problems. A different 1C/RS/SPSS strategy for scatter search may be obtained by running concurrently the combination and improvement operators on several subsets of the reference set. Here, the master generates tasks by extracting a number of solution subsets, which are sent to slaves. Each slave then combines and improves its solutions, returning its results to the master for the global update of the reference set. Each subset sent to a slave may contain the exact number of solutions required by the combination operator or a higher number. In the former case, the slave performs an “iteration” of the scatter search algorithm García-López et al (2005, 2006, 2003). In the latter case, several combination-improvement sequences could be executed and solutions could be returned to the master as they are found or all together at the end of all sequences. Load-balancing capabilities should be added to the master to avoid differences in work quantity and computing times

between slaves.

To conclude, low level, 1-control parallel strategies are particularly attractive when neighborhoods or populations are large, or the neighbor or individual evaluation is costly. Computing time gains may then be obtained, as illustrated by many early contributions discussed in the surveys indicated in the Introduction. Even more impressive gains may be obtained by taking advantage of the current computing platforms integrating multi-core central processing units (CPUs - the “traditional” processor) and graphical processing units (GPUs) enhanced with data streaming, i.e., hardware data parallelism providing the means for each processor to perform the same task on different parts of the distributed data (e.g., Brodtkorb et al, 2013a,b). This hardware technology offers the possibility of extensive very low-level parallelization reminiscent of the work performed for the massively parallel computers of the late eighties. The neighborhood evaluation in local search heuristics, the fitness evaluation of evolutionary methods, and the evolution of individuals in swarms may clearly benefit from such a hardware-oriented parallelization, spectacular speedups having been observed (e.g., Melab et al, 2011; Brodtkorb et al, 2013a,b; Delévacq et al, 2013; Cecilia et al, 2013; Tan and Ding, 2016; Van Luong et al, 2013). A number of remarks are in order, however. First, the utilization of this technology is not straightforward, and work must be dedicated to its conceptual, technical and experimental aspects. Second, there is also the need to examine the sequential and parallel meta-heuristic designs to identify and value where this technology would bring the most benefits, besides those already identified. The work of Rios et al (2017) is a step on this research path. Finally, as discussed in the following sections, more advanced multi-search strategies outperform low-level strategies in most cases, in particular with respect to solution quality. Consequently, hierarchical settings combining multi-search strategies and 1C/RS/SPSS evaluation procedures, all on CPU-based architectures, are generally used currently. More research is needed in this area to account for the massively parallel possibilities of GPUs.

4 Domain Decomposition

We group under this title the strategies that separate the search space explicitly. The basic idea is intuitively simple and appealing: separate the search space into smaller subspaces, address the resulting subproblems by applying the sequential meta-heuristic on each subspace, collect the respective partial solutions, and reconstruct an entire solution out of the partial ones. This apparently simple idea may take several forms, however, according to the type of separation performed, the permitted links among the resulting subproblems, the possible iterative modification of the separation and the type of control of the parallel meta-heuristic.

Regarding the separation type, the resulting subspaces may constitute a *partition*

of the complete space (disjoint subspaces, their union being the full space), or a *cover* allowing a certain amount of overlap among the subspaces. Note that covers may be defined implicitly by allowing the search within a given subspace to reach out to some part of one or several other subspaces through, e.g., neighborhood moves or individual crossovers.

The separation may be obtained by identifying a subset of variables, and corresponding constraints, eventually, and discarding or fixing the other variables and constraints, the goal being to obtain smaller, easier to address subproblems. Note that it is not always possible, even desirable, to discard. Thus, if one may easily discard the customers in a Vehicle Routing Problem (VRP) that do not belong to a given subspace (the depot must be included in each subspace) and solve the resulting partial VRPs separately, doing the same is much more difficult to implement when considering the commodities and arcs of a Multicommodity Capacitated Network Design problem (MCND). Separation by variable fixing (and projection of the corresponding constraints) appears more flexible as one still works on smaller subproblems, but considering the complete vector of decision variables, some of which are fixed. It is also a more general approach, as we find it in advanced cooperative search methods (e.g., Lahrichi et al, 2015).

Strict partitioning restricts the solvers to their subspaces, resulting in part of the search space being unreachable and the loss of exploration quality for the parallel meta-heuristic. Covers, through explicit or implicit overlapping, partially address this issue; indeed, to guarantee that all potential solutions are reachable, one must make overlapping cover the entire search space, which would negate the benefits of decomposition. To avoid these drawbacks, one can change the separation and start again. This idea translates into a strategy encountered quite frequently in strict partitioning, where the separation is modified periodically, and the search is restarted using the new decomposition. A complete-solution reconstruction feature is almost always part of the procedure. Note that this approach provides also the opportunity to define non-exhaustive separations, i.e., where the union of the subspaces is smaller than the complete search space.

This strategy is naturally implemented using master-slave 1C/RS schemes, with MPSS or MPDS search differentiation. The master process determines the separation and sends partial subsets (or information to define them out of the initial space - this reduces the communication overhead) to slaves, synchronizes them and collects their solutions, reconstructs complete solutions, modifies the separation, and determines when stopping conditions are met. Slaves concurrently and independently perform the search on their assigned subsets. Most implementations addressed problem settings for which a large number of iterations can be performed in a relatively short time and restarting the method with a new decomposition does not require an unreasonable computational effort (e.g., Gendreau et al, 2001, for real-time ambulance fleet management), a full-fledged meta-heuristic being generally used on each subspace.

Explicit space separation may also be performed in a pC, collegial decision-making, framework with MPSS or MPDS search-differentiation. The separation in a pC/KS strategy is collegially decided and modified through information-exchange phases (e.g., round-robin or many-to-many exchanges) activated at given synchronization points. The KS label comes from exchanging not only the best solutions in each subspace (e.g., routes in a VRP), but also from the so-called context information (e.g., un-serviced customers and empty vehicles in a VRP Taillard, 1993) that is used to modify the separation. In the initial proposition by Taillard (1993) for the VRP (simulated on a sequential machine), the customer set was partitioned, vehicles were allocated to the resulting regions, each subproblem was solved by an independent tabu search, synchronization occurred after a number of iterations that varied according to the total number of iterations already performed, and exchanges took place between adjacent processors (corresponding to neighboring regions). The method allow at the time to address successfully a number of problem instances, but the synchronization inherent in the design of the strategy hindered its performance. A parallel ant-colony approach for the VRP based on this idea was presented in Doerner et al (2006) with good speedup results when the number of customer increased.

Domain decomposition methods induce different search behavior and solution quality compared to those of the sequential meta-heuristic. Such methods appear increasingly needed as the dimensions of the contemplated problem instances continue to grow. Given the increased complexity of the problem settings, work is also required on how to best combine search-space separation and the other parallelization strategies, cooperation in particular. The Integrative Cooperative Search (Lahrichi et al, 2015) is a step in this direction (see Section 7).

5 Independent Multi-search

We dedicate a section to the *Independent multi-search* as it was among the first parallelization strategies proposed in the literature, and is also the most simple and straightforward p-control parallelization strategy, generally offering an interesting performance.

Independent multi-search seeks to accelerate the exploration of the search space toward a better solution (compared to sequential search) by initiating simultaneous solvers from different initial points (with or without different search strategies). It thus parallelizes the classical multi-start strategy by performing several searches simultaneously on the entire search space, starting from the same or from different initial solutions, and selecting at the end the best among the best solutions obtained by all searches. Independent multi-search methods thus belong to the pC/RS class of the taxonomy. No attempt is made to take advantage of the multiple solvers running in parallel other than to identify the best overall solution at the final synchronization step.

The efficiency of independent multi-search follows from the sheer quantity of computing power it allows one to apply to a given problem (Battiti and Tecchioli, 1992; Taillard, 1994; Stutzle, 1998; ten Eikelder et al, 1999). The surveys identified in the Introduction describe numerous contributions of applying the pC/RS independent multi-search strategy to a variety of combinatorial optimization problems.

Independent multi-search offers an easy access to parallel meta-heuristic computation, offering a tool when looking for a “good” solution without investment in methodological development or coding. Such methods are generally outperformed by cooperative strategies, however, through mechanisms enabling the independent solvers to share, during the search, the information their exploration generates. As explained in the following sections, this sharing and the eventual creation of new information out of the shared one, yields in most cases a collective output of superior solutions compared to independent and sequential search.

6 Cooperative Search

Cooperative multi-search has emerged as one of the most successful meta-heuristic methodologies to address hard optimization problems (e.g. Talukdar et al, 2003; Alba, 2005; Crainic, 2005; Crainic and Hail, 2005; Crainic, 2008; Crainic and Toulouse, 2008; Talbi, 2006; Crainic and Toulouse, 2010). While independent multi-search strategies seek to accelerate, compared to sequential search, the exploration toward a better solution by initiating simultaneous searches from different initial points, cooperative search strategies go further and integrate *cooperation mechanisms* to share, *while the search is in progress*, the information obtained from this diversified exploration of the same problem instance. The sharing and, eventually, creation of new information out of the exchanged data (Section 7), yields in many cases a collective output with better solutions than a parallel independent search.

Cooperative-search strategies are thus defined by the solver components engaged in cooperation, their interaction mechanism, and the nature of the information shared. The solvers define trajectories in the search space from possibly different initial points or populations, by using possibly different search strategies (including, possibly exact methods). The information-sharing cooperation mechanism specifies how these independent solvers interact, how the exchanged information is used globally (if at all), and how each solver acts on the received information, using it within its own search and, thus, transforming it before passing it to other solvers.

The information-sharing cooperation mechanism specifies how these independent solvers interact, the global search behavior of the cooperative parallel meta-heuristic emerging from the local interactions among them, which makes it a “new” meta-heuristic in its own

right (Crainic and Toulouse, 2008). The similarity between this behavior and that of systems where decisions emerge from interactions among autonomous and equal “colleagues” has inspired the name *collegial control* associated to cooperative-search strategies in the taxonomy used in this chapter. The various cooperation mechanisms proposed in the literature are described in the next sections.

Exchanged information must be *meaningful* and exchanges must be *timely*. The goals are 1) to improve the performance of the receiving solvers, and 2) to create as much as possible a global, “complete” image of the status of the cooperative search to enable guiding it, through its participating solvers, toward a better performance in terms of solution quality and computational efficiency than the simple concatenation of results obtained by non-cooperating solvers. Toulouse et al (1996) proposed a list of questions related to addressing this challenge. The list is still relevant when designing cooperative parallel strategies: What information is exchanged? Between what processes is it exchanged? When is information exchanged? How is it exchanged? How is the imported data used? Implicit in their taxonomy and explicitly stated in later papers, the issue of whether the information is modified during exchanges or whether new information is created completes this list.

“Good” solutions are the most often exchanged type of information, usually taking the form of the overall best solution or the current-best solution of a solver being sent to the others. It was observed, however, that sending out all current-best solutions a solver identifies is often counter productive, particularly when the solver performs a series of improving moves or generations, as solutions are generally “similar” (particularly for neighborhood-based procedures), and the receiving solvers have no chance to act on the in-coming information (unless special receiving mechanisms are embedded in all solvers) before receiving a new solution, or may embark on explorations similar to that of the sending solver. It was also observed that always sending the overall best solution to all cooperating solvers is generally bad as it rapidly decreases the diversity of the search, increasing the amount of worthless computational work (many solvers will search in the same region) and bringing an early “convergence” to a not-so-good solution. Sending out the local optima after a series of improving moves, exchanging groups of solutions, and implementing random selection procedures for the solutions to send out, the latter generally biased toward good or good-and-different solutions, are among the strategies aimed at addressing these issues.

Context information may also be shared profitably when embedded in the mechanisms used to guide the search. Context information refers to data collected by a solver during its own exploration, such as the statistical information relative to the presence of particular solution elements in improving solutions (e.g., the medium and long-term memories of tabu search), the impact of particular moves on the search trajectory (e.g., the scores of the moves of large adaptive neighborhood search), population diversity measures, individual resilience across generations, etc. A limited number of studies indicate

the interest of context-information exchanges (see Section 7), but more research is needed on this topic.

Cooperating solvers may exchange information directly or indirectly. *Direct* exchanges of information occur either when the concerned solvers agree on a meeting point in time to share information, or when a solver broadcasts its information to one or several other solvers without prior mutual agreement. The latter case is to be avoided as it requires solvers to include capabilities to store received information without disturbing their own search trajectories until they are ready to consider it. Failure to implement such mechanisms generally results in bad performances, as observed for strategies combining uncontrolled broadcasting of information and immediate acceptance of received data.

Indirect exchanges of information are performed through independent data structures that become shared resources of data solvers may access asynchronously and according to their own internal logic to post and retrieve information. Such data structures are called *blackboard* in the computer-science and artificial-intelligence vocabulary, while *memory*, *pool*, and *data warehouse* (*reference* and *elite set* are also sometimes used) are equivalent terms found in the parallel meta-heuristic literature. The term *memory* is used in this chapter.

Centralized-memory mechanisms have been used in most parallel meta-heuristic contributions. They receive, eventually process, and post information received from all cooperating solvers, which, in turn, may retrieve this information independently. Distributed memory mechanisms may be contemplated, where a number of memories are inter-connected, each servicing a number of solvers. Such hierarchical structures, with several layers of solvers and memories, appear interesting when a large number of processors is involved, when computations are to take place on grids or loosely coupled distributed systems, and for integrative cooperation strategies. Issues related to data availability, redundancy, and integrity must then be addressed, as well as questions relative to the balancing of workloads and the volume of information exchanged. More research is needed on this topic.

Communications proceed according to an interaction topology represented by a *communication graph* specifying the processes that may engage in direct exchanges and, thus, directing the flow of information within the cooperative search. Each node of the graph represents a solver or a memory. Edges define pairs of solvers or a solver-memory pair. The projection of this graph on the physical interconnection topology of the parallel computer executing the parallel program is generally part of the implementation design.

When and how information is shared specifies the frequency of cooperation activities, who initiates them and when, and whether the concerned solvers must synchronize, i.e., each stopping its activities and waiting for all others to be ready, or not. These two cases are identified as *synchronous* and *asynchronous* communications, respectively, and

are discussed in the following sections. A general observation for both cases, however, is that exchanges should not be too frequent to avoid excessive communication overheads as well as premature “convergence” to local optima (Toulouse et al, 1999a, 2004, 1998, 2000).

Two observations to conclude this general discussion about cooperation. First, it is worth noticing that cooperation is somewhat biased toward intensifying the search in regions of the space that have already been explored and where interesting solutions have been identified. This is particularly true for “simple” cooperation mechanisms based on synchronization or that exchange current-best solutions only. It is thus important to equip the cooperation mechanisms with diversification capabilities, e.g., probabilistic or diversity-driven selection of exchanged solutions (as proposed by Vidal et al, 2012, for example) or creation of new solutions and guidance information (Lahrichi et al, 2015).

Second, the main principles of cooperative parallelization are the same for neighborhood- and population-based meta-heuristics, even though denominations and implementation approaches may differ. We thus structure the presentation that follows based on these principles and general strategies, rather than by meta-heuristic class. The next two subsections discuss the classic synchronous and asynchronous strategies, while the advanced methods based on creation of new information out of the shared one are the topic of Section 7.

6.1 pC/KS Synchronous Cooperative Strategies

Synchronous cooperation follows a p-control (pC), knowledge synchronous (pC/K) strategy, with any of the SPDS, MPSS or MPDS search differentiation approaches. All participating solvers stop their activities at particular moments and engage in an information-exchange phase, which must be completed before any solver can restart its exploration from that synchronization point. Synchronization moments may be determined by conditions imposed exogeneously to all solvers (e.g., number of iterations from the last synchronization point), or detected by an a priori designated solver.

The goal of synchronous cooperative strategies is to re-create a state of complete knowledge at particular points in the global search and, thus, to hopefully guide it into a coordinated evolution toward the problem solution. This goal is generally only partially attained, however, even though these strategies have generally outperformed the sequential versions as well as simpler parallelization strategies. Moreover, synchronization results in significant time inefficiencies as communications are initiated only when the slowest search thread is ready to start. Asynchronous information sharing thus intuitively appears more promising and, indeed, cooperation based on asynchronous exchanges, described in the following sections, generally outperformed synchronous methods. Consequently, few contributions relying on synchronous cooperation were proposed in recent

years.

We therefore restrict this section to recalling the main concepts used in synchronous cooperation, some of which found their way into more advanced strategies, encouraging interested readers to consult the surveys indicated in the Introduction for details and references.

Synchronization may use a complete communication graph or a more restricted, less densely connected communication topology (e.g., ring, torus, and grid graph). *Global* exchanges of information among all solvers take place in the former case, while information follows a *diffusion* process through direct, *local*, exchanges among neighboring processes in the latter.

In a restricted view of the concept, a number of proposed pC/KS cooperative search meta-heuristics based on global exchanges use a designated *communication master* process, which may or not include one of the participating solvers. The communication master manages the synchronization mechanism in a master-slave implementation. It initiates the global search starting the solvers, stops all solvers at synchronization points, gathers the information, updates the global data, verifies the termination criteria of the search and, either effectively terminates it or distributes the shared information (a good solution, generally, the overall best solution in many cases) and sends a signal to the solvers to continue the search (e.g., García-López et al, 2002; Rego, 2001). For coarse-grained island implementations of cooperating genetic methods, synchronization means the communication master initiates the *migration* operator to exchange among the independent populations the best or a small group of some of the best individuals in each (Czech, 2000; Solar et al, 2002). For ant-colony systems, this strategy divides the colony into several sub-colonies individually assigned to solvers, the master updating the pheromone matrix, and starting a new search phase, based on the received solver results (Drias and Ibri, 2003). A more sophisticated approach was proposed in Niar and Fréville (1997), where the master dynamically adjusted the search parameters of cooperating tabu searches according to the results each had obtained so far. The method performed well on the 0-1 Multi-dimensional Knapsack Problem, which is encouraging, as the idea of dynamic adjustment of the search parameters may be generalized to more sophisticated cooperation mechanisms.

A truer global pC/KS cooperative scheme empowers solvers to initiate synchronization. Once it reaches a pre-determined status, a solver thus sends the stopping signal, broadcasts its data (current best solution or group of solutions, in most cases), followed by similar broadcasts performed by the other solvers. Once all information is shared, each solver performs its own import procedures on the received data and proceeds with its exploration of the search space until the next synchronization event. Most synchronous coarse-grained island parallelizations of genetic-based evolutionary methods fall under this category, where migration operators are applied at regular intervals (e.g., Flores

et al, 2003; Hidalgo et al, 2003, the later implementing a hierarchical method with the fitness computation performed at the second level through a master-slave implementation; the overhead due to the parallelization of the fitness became significant for larger numbers of processors). For ant-colony application, where each colony evolves its own pheromone matrix, global synchronization mean that, after a fixed number of iterations, colonies exchange elite solutions that are used to update the pheromone matrix of the receiving colony (Michels and Middendorf, 1999; Middendorf et al, 2002).

Synchronization based on global exchanges of information assumes that making available to all solvers the entire information shared will result in superior performances. Other than the often excessive communication overhead, the main drawback is that solvers relying heavily on the same information end up by exploring the same regions of the search space, resulting in loss of diversity and efficiency. Two approaches have been proposed to overcome this drawback.

First, do not share and use uniquely the local best solutions, as in the pC/RS/MPDS iterated tabu search proposed for the VRP by Cordeau and Maischberger (2012). In this work, solvers synchronized after a number of consecutive iterations without improvement within the individual improvement phases. Synchronization involved the exchange of the good solutions obtained by the solvers and, then, each individual solver built a new starting solution by selecting routes probabilistically among those received and its own. Computational results showed this method to be flexible and efficient for several classes of routing settings with several depots, periodicity of demands, and time windows.

The second approach is based on diffusion. In such strategies, direct communications at synchronization points are possible only with neighboring solvers, i.e., with nodes adjacent in the sparse communication graph. The quantity of information each solver processes and relies upon is thus significantly reduced. Information is still shared between non-adjacent solvers but at the reduced diffusion speed of chains of local exchanges and data modification by the intervening solvers. This idea was less explored compared to the global-exchange strategy, even though synchronous cooperative mechanisms based on local exchanges and diffusion have a less negative impact on the diversity of the search-space exploration, and have yielded good results (e.g., Tongcheng and Chundi, 2002; Middendorf et al, 2002).

6.2 pC/C Asynchronous Cooperative Strategies

Historically, independent and synchronous cooperative methods were the first multi-search approaches to be developed. The focus has shifted, however, to asynchronous cooperation strategies, which may be considered as defining the “state-of-the-art” in parallel multi-search meta-heuristics.

A cooperation strategy is asynchronous when programs initiate cooperation activities according to their own internal logic, without coordination with other solvers or memories. Thus, e.g., a solver may make available its current best solution by posting it on a memory, or may ask for an external solution when it failed to improve the quality of its best solution for a certain number of iterations.

Asynchronous communications provide the means to build cooperation and information sharing among search threads without incurring the overheads associated with synchronization. They also bring adaptability to cooperation strategies, to the extent that the parallel cooperative meta-heuristic may more easily react and dynamically adapt to the exploration of the search space than independent or synchronous search. These benefits come with potential issues one must care for. For example, the information related to the global search that is available when a solver must take an action may be less “complete” than in a synchronous environment. On the other hand, too frequent data exchanges, combined with simple acceptance rules for incoming information, may induce an erratic solver behavior, the corresponding search trajectories becoming similar to random walks. Hence the interest for applying information-sharing based on quality, meaningfulness, and parsimony principles (Crainic et al, 1996, 1997; Toulouse et al, 1996).

Asynchronous cooperative strategies follow either pC/C or pC/KC collegial principles, the main difference between the two being that in the latter “new” knowledge is inferred on the basis of the information exchanged between solvers; pC/KC strategies are addressed in the next section.

In most pC/C asynchronous strategies in the literature, the shared information corresponds to a locally improving solution or individual(s), the most successful contributions sharing local optima only. The principles mentioned above also resulted in mechanisms to diversify the shared information (Crainic et al, 1996). Thus, always selecting the best available solution out of an elite set of good solutions, sent by potentially different solvers, proved less efficient in terms of quality of the final solution than a strategy that selected randomly, but biased by quality, among the same elite set.

When to initiate and perform cooperation activities, as well as how to use the incoming information is particular to each type of meta-heuristic. Most strategies proposed in the literature follow the same idea, however, to send and request information jointly. There is no need to do this, however, even though it can decrease the amount of communication. It may thus be interesting for neighborhood-based methods to make available right away their newly found local optima or improved overall solutions, and not wait for the algorithmic step where examining external information is appropriate. Similarly, population-based methods could migrate a number of individuals when a significant improvement is observed in the quality and diversity of their elite group of individuals. Regarding the request of external information, it may be based on a pre-fixed number

of iterations, but this approach should be restricted to meta-heuristics without search-diversification steps, e.g., tabu search based on continuous diversification. In most other cases, the principle of parsimonious communications implies selecting moments when the status of the search changes significantly, e.g., when the best solution or the elite subpopulation did not improve for a number of iterations. At such moments, solvers generally engage into some form of *search-diversification* phase, e.g., diversification in tabu search, change of neighborhood in variable neighborhood search, and complete or partial re-generation of population in population-based meta-heuristics, which involves the choice or modification of the current solution to initiate a new phase. External information, which generally includes at least one good solution, may prove particularly interesting at that moment. How it is to be used depends on the particular logic of the receiving solver; it may be used to initiate a diversification phase, to modify the search trajectory through a combination with a “local” solution, or to modify the solver behavior in the long run through an insertion in an elite set or population. As already mentioned, however, one tries to avoid frequent imports followed by a replacement of the current solution or population, which will result in a random search.

Direct and indirect exchange pC/C strategies may be used with any meta-heuristic. Historically, however, most genetic-based evolutionary asynchronous cooperative meta-heuristics relied on direct communications over complete communication graphs (Cantú-Paz, 2005). These methods generally implement a coarse-grained island model, migration being triggered by conditions within individual populations, selected migrant individuals being directed toward either all other populations or a dynamically-selected subset. The work of Vallada and Ruiz (2009) illustrates this approach, where migration is initiated by an island that identified a new best solution, which it sends to all other islands. The migrant individual is accepted by the solver of another island only when different from the local population and better than the worst individual in that population. We also mention the work of Izzo et al (2009) who introduced genetic solvers with different strategies, which was a novelty in the GA-island field (previously, all island populations were evolved by the same algorithm), and observed significant improvements compared to more traditional island-based pC/C models. The parallelization of ant-colony methods may use the same approach, where partitions of the initial colony play the role of islands. The contribution of Ling et al (2012) is interesting in this context for novel way of selecting the receiving subcolony (island). Here, a solver initiates an exchange when the evolution of its colony becomes stagnant (no longer improving) by selecting an exchange partner probabilistically based on the relative distance (the most different best solution) and fitness (of the best solution); it then requests the current best solution from the selected partner, and, upon reception, updates its pheromone matrix and continues the search.

Notice that complete communication graphs are not compulsory. Indeed, one could use particular graphs and information-diffusion processes tailored to the problem at hand. Yet, despite encouraging results (e.g., Sevkli and Aydin, 2007, proposing VNS pC/C

strategies over uni and bidirectional ring topologies), too few experiments have been reported yet.

Historically, the sharing of information in most asynchronous cooperative search strategies outside the genetic-evolutionary community is based on some form of indirect communications through a centralized device - data repository / processor -, often called *central memory* (Crainic, 2005; Crainic et al, 1996, 1997). A solver involved in such a cooperation deposits good solutions, local optima generally, into the central memory, from where, when needed, it also retrieves information sent by the other cooperating solvers. Classical retrieval mechanisms are based on random selection, which may be uniform or biased to favor solutions with high rankings based on solution value and diversity. The central memory accepts incoming solutions for as long as it is not full, acceptance becoming conditional to the relative interest of the incoming solution compared to the “worst” solution in the memory, otherwise. Diversity criteria are increasingly considered, a slightly worse solution being preferred if it increases the diversity of solutions in the central memory. Population culling may also be performed (deleting, e.g., the worst half of the solutions in memory).

Central-memory-based cooperative search strategies are described in the literature for most meta-heuristic classes. To the best of our knowledge, Crainic et al (1996) were the first to propose a central-memory approach for asynchronous tabu search in their comparative study for a multi-commodity location problem with balancing requirements. Their method, where individual tabu searches sent to the memory their local-best solutions when improved, and imported a solution selected probabilistically biased by rank before engaging in a diversification phase, outperformed in terms of solution quality the sequential version as well as several synchronous and broadcast-based asynchronous cooperative strategies. The same approach was applied to the fixed cost, capacitated, multicommodity network design problem with similar results (Crainic and Gendreau, 2002).

pC/C with some form of central memory were proposed for a variety of problems, including cutting (Blazewicz et al, 2004), container loading (Bortfeldt et al, 2003), labor-constrained scheduling (Cavalcante et al, 2002), VRP with time windows (VRPTW) (Le Bouthillier and Crainic, 2005). On the other hand, several studies focused on pC/C strategies with some form of central memory for particular classes of meta-heuristics, e.g., simulated annealing (e.g. Lee and Lee, 1996; Sanvicente-Sánchez and Frausto-Solís, 2002; Banos et al, 2013, the latter for multi-objective problem setting), VNS (e.g., Crainic et al, 2004; Polacek et al, 2008, the latter proposing a self-adapting mechanism for the main search parameters based on recent performance, and solution selection out of the ten-best present in memory), GRASP (e.g., Ribeiro and Rosseti, 2007, with cooperation based on applying path relinking to solutions from memory) and tabu search (e.g., James et al, 2009, with memory hosting a reference set and long-term global memories built on short-term local memories sent by solvers).

Notice that cooperating solvers need not belong to the same meta-heuristic class. The next section will show several examples where different meta-heuristics collaborate within pC/KC strategies. We find, in the classical pC/C case, contributions following the same broad strategy described above when calling sequentially on meta-heuristics belonging to different types. The two-phase approach of Gehring and Homberger (2002) for the VRPTW is a typical example of such a method, where each solver first applies an evolution strategy to reduce the number of vehicles, followed by a tabu search to minimize the total distance traveled. A somewhat different two-phase pC/C parallel strategy was proposed in Bastos and Ribeiro (1999) for the Steiner problem, where each phase, using reactive tabu search and path relinking, respectively, implemented the pC/C asynchronous central memory strategy, all processes switching from the first to the second phase simultaneously.

Multi-level cooperative search proposes a different pC/C asynchronous cooperative strategy based on controlled diffusion of information Toulouse et al (1999b). Solvers are arrayed in a linear, conceptually vertical, communication graph and a local memory is associated with each. Each solver works on the original problem but at a different level of aggregation or “coarsening”, the first-level solver working on the complete original problem. It communicates exclusively with the two solvers directly above and below, that is, at higher and lower aggregation levels, respectively. The local memory is used to receive the information coming from the immediate neighbors and to access it at moments dynamically determined according to the internal logic of the solver. In the original implementation, solvers were exchanging improved solutions, incoming solutions not being transmitted further until modified locally for a number of iterations to enforce the controlled diffusion of information. Excellent results have been obtained for various problem settings including graph and hypergraph partitioning (Ouyang et al, 2000, 2002), network design (Crainic et al, 2006b), feature selection in biomedical data (Oduntan et al, 2008), and covering design (Dai et al, 2009). It is noteworthy that one can implement multi-level cooperative search using a central memory by adequately defining the communication protocols. Although not yet fully defined and tested, this idea is interesting as it opens the possibility of richer exchange mechanisms combining controlled diffusion and general availability of global information.

The central-memory pC/C asynchronous cooperation strategy is generally offering very good results, yielding high-quality solutions. It is also computationally efficient as no overhead is incurred for synchronization. No broadcasting is taking place and there is no need for complex mechanisms to select the solvers that will receive or send information and to control the cooperation. It has also proved efficient in handling the issue of premature “convergence” in cooperative search, by diversifying the information received by the solvers through probabilistic selection from the memory and by a somewhat large and diverse population of solutions in the central memory; solvers may thus import different solutions even when their cooperation activities are taking place within a short time span. The central memory is thus an efficient algorithmic device that allows for a

strict asynchronous mode of exchange, with no predetermined connection pattern, where no solver is interrupted by another for communication purposes, but where any solver may access at all times the data previously sent out by the other solvers.

The performance of central-memory cooperation and the availability of exchanged information (kept in the memory) has brought the question of whether one could design more advanced cooperation mechanisms taking advantage of the information exchanged among cooperating solvers. The pC/KC strategies described in the next section are the result of this area of research.

7 pC/KC Cooperation Strategies - Creating New Knowledge

Cooperation, particularly in the central-memory asynchronous form, offers many possibilities for algorithm development. Particularly noteworthy are the flexibility in terms of the different meta-heuristic and exact methods that can be combined, and the population of elite solutions being hosted in the central memory and continuously enhanced by the cooperating solvers. One can thus select cooperating methods that complement each other, some of which heuristically construct new solutions, execute neighborhood-based improving meta-heuristics, evolve populations of solutions, or perform post-optimization procedures on solutions in memory.

The study of Crainic and Gendreau (1999) illustrates the interest of these ideas. The authors combined a genetic solver and several solvers executing the pC/C tabu search for multicommodity location-allocation with balancing requirements of Crainic et al (1996). The tabu searches were aggressively exploring the search space, building the elite solution set in the central memory, while the genetic method contributed toward increasing the diversity, and hopefully the quality, of the solutions in the central memory, which the cooperating tabu search methods would then import. The genetic method was launched once a certain number of elite solutions identified by the tabu searches were recorded in the central memory, using this memory as initial population. Asynchronous migration subsequently transferred the best solution of the genetic pool to the central memory, as well as solutions of the central memory toward the genetic population. This strategy did perform well, especially on larger instances. It also yielded an interesting observation: while the best overall solution was never found by the genetic solver, its inclusion allowed the tabu search solvers to find better solutions, more diversity among solutions in memory translating into a more effective diversification of the global search.

Crainic and Gendreau (1999) and several other studies showed that it is beneficial not only to include solvers of different types in the cooperation, but also to use the

elite population these solvers built in memory to construct an approximate image of the status of the global search, e.g., to learn about the parts of the search space already explored, the relations between the values of certain decision variables (e.g., arcs in a VRP or design solution) and the value of the corresponding solution, the performance of the cooperating solvers on the particular instance given the information they receive from the central memory, etc. This information may then be used to create new knowledge, new and diverse solutions, solution components, “ideal” target solutions, etc., and guide the search. Population-based meta-heuristics are particularly appropriate to generate solutions that add quality and diversity to an elite set.

Cooperative strategies including mechanisms to create new information and solutions based on the solutions exchanged belong to the p-control knowledge collegial (pC/KC) class. Most contributions to this field have solvers work on the complete problem and make the bulk of the section. We conclude the pC/KC section with a discussion on recent developments targeting multi-attribute problem settings where the problem at hand is decomposed and solvers work on particular parts of the problem or on integrating the resulting partial solutions into complete ones.

Historically, two main classes of pC/KC cooperative mechanisms are found in the literature, both based on the idea of exploiting a set of elite solutions, and their attributes, exchanged by cooperating solvers working on the complete problem, but differing in the information kept in memory. *Adaptive-memory* methods (Rochat and Taillard, 1995) store and score partial elements of good solutions and combine them to create new complete solutions that are then improved by the cooperating solvers. *Central-memory* methods exchange complete elite solutions among neighborhood and population-based meta-heuristics and use them to create new solutions and knowledge to guide the cooperating solvers (Crainic, 2005; Crainic and Toulouse, 2003; Crainic et al, 1996). The latter method generalizes the former and, the vocabulary used in the various papers notwithstanding, the two approaches are becoming increasingly unified.

The adaptive-memory terminology was coined by Rochat and Taillard in a paper (Rochat and Taillard, 1995) proposing tabu search-based heuristics for the VRP and the VRPTW that are still among the most effective ones for both problems (see Badeau et al, 1997; Glover, 1996; Taillard et al, 1997, for more on adaptive-memory concepts) The main idea is to keep in memory the individual components (vehicle routes in VRP) making up the elite solutions found by the cooperating solvers, together with memories counting for each component its frequency of inclusion in the best solutions encountered so far, as well as its score, and rank among the population in memory, computed from the attribute values, in particular the objective value of its respective solutions. Solvers construct solutions out of probabilistically selected (biased by rank) solution components in memory, enhance it (tabu search in the initial contribution), and deposit their best solutions in the adaptive memory The probabilistic selection yields, in almost all cases, a new solution made up of components (routes) from different elite solutions, thus inducing

a diversification effect. A number of early developments provided insights into algorithmic design. Worth mentioning are Schulze and Fahle (1999) for the VRPTW, who proposed a set-covering heuristic to select the solution components in memory used to generate the new initial solution of a cooperating solver and Gendreau et al (1999), for real-time vehicle routing and dispatching, actually implementing a hierarchical, two-level parallel scheme: a pC/KC/MPSS cooperating adaptive memory meta-heuristic at the first level, while each individual tabu-search solver implemented the route decomposition of Taillard (1993) with the help of several slave processors on the second level.

Generalizing the pC/C and adaptive-memory strategies, pC/KC central-memory mechanisms keep full solutions, as well as attributes and context information sent by the solvers involved in cooperation. Solvers, which indirectly exchange complete elite solutions and context information through the central memory, may perform constructive, improving and post-optimization heuristics (e.g., Le Bouthillier and Crainic, 2005; Le Bouthillier et al, 2005; Lahrichi et al, 2015), neighborhood- (e.g., tabu search Di Chiara, 2006; Jin et al, 2012, 2014; Lahrichi et al, 2015) and population-based methods like genetic algorithms (e.g., Le Bouthillier and Crainic, 2005; Le Bouthillier et al, 2005; Di Chiara, 2006; Lahrichi et al, 2015) and path relinking (e.g., Crainic et al, 2006a), as well as exact solution methods (Groër and Golden, 2011) on possibly restricted versions of the problem.

The particular solvers to include in cooperation depend on the application. They should be efficient for the problem at hand, of course. Additionally, they should also aim to cover different regions of the search space in such a way that they contribute not only to the quality but also to the diversity of the elite population being built in the central memory.

Other than the information received from the cooperating solvers, the central memory keeps newly created information out of these exchanged data. Statistics-building, information-extraction and learning, and new solution-creation mechanisms provide this new “knowledge”. Memories recording the performance of individual solutions, solution components, and solvers may be added to the central memory, and guidance mechanisms based on this knowledge may be gradually built.

Central-memory mechanisms thus perform two main tasks: data-warehousing and communications with solvers, on the one hand, and information-creation and search-guiding, on the other hand. To distinguish between the two, we single out the latter as the *Search Coordinator (SC)*. The simplest SC mechanism was used in the pC/C strategies of the previous section, where solutions in memory were ordered and rank-biased randomly extracted to answer solver requests. The functions of the SC in pC/KC methods include creating new solutions, extracting appropriate solution elements, building statistics on the presence and performance of solutions, solution elements, and solvers (these belong to the family of memories, well-known in the meta-heuristic community), creating the

information to return when answering solver requests, the latter being part of the so-called *guidance mechanisms*.

The cooperative meta-heuristic proposed by Le Bouthillier and Crainic (2005) for the VRPTW used a simple pC/KC mechanism, involving four solvers, two simple genetic algorithms with order and edge recombination crossovers, respectively, and two tabu search methods that perform well sequentially, Unified Tabu Search (Cordeau et al, 2001) and TABUROUTE (Gendreau et al, 1994). The cooperating solvers shared their respective best solutions identified so far. The SC in central memory performed post-optimization (2-opt, 3-opt, Or-opt, and ejection-chain procedures to reduce the number of vehicles and the total traveled distance) on the received solutions before making them available for sharing. Solvers requested solutions from the central memory when needed, i.e., the genetic algorithms for crossover operations, the Unified Tabu at regular intervals, and TABUROUTE at diversification time. This algorithm, without any calibration or tailoring, proved to be competitive with the best meta-heuristics of its day in linear speedups.

Le Bouthillier et al (2005) enhanced this SC with an innovative learning and guidance mechanism. The authors aimed for a mechanism that, not only returned meaningful information to solvers, but was also independent of particular problem characteristics, e.g., routes in their VRPTW application, and could be broadly applied to network-based problem settings. The SC mechanism is thus based on an atomic element in network optimization, the arc. Starting from the classical memory concepts pioneered for tabu search (Glover, 1989, 1990; Glover and Laguna, 1997), Le Bouthillier et al (2005) combined two ideas: first, that an arc appearing often in good solutions and less frequently in bad solutions may be worthy of consideration for inclusion in a tentative solution, and vice versa, and, second, that this worthiness increases when the behavior appear stable in time. The authors thus considered the evolution of the frequency of inclusion of arcs in solutions of different quality, that is, in the elite (e.g., the 10% best), average (between the 10% and 90% best), and worst (the last 10%) groups of solutions in the central memory. *Patterns* of arcs were then defined representing subsets of arcs (not necessarily adjacent) with similar frequencies of inclusion in particular population groups. Guidance was obtained by transmitting arc patterns to the individual solvers indicating whether the arcs in the pattern should be “fixed” or “prohibited” to intensify or diversify the search, respectively. The solvers accounted for these instructions by using the patterns to bias the selection of arcs for move or reproduction operations. A four-phase fixed schedule (two phases of diversification at the beginning to broaden the search, followed by two intensification phases to focus the search around promising regions) was used (see Le Bouthillier, 2007, for a dynamic version of this mechanism), with excellent results in terms of solution quality and computing efficiency compared to the best-performing methods of the day.

The pC/KC/MPDS method proposed in Groër and Golden (2011) for the VRP illustrates how specialized solvers may address different issues in a cooperative meta-heuristic,

including the generation of new knowledge. Two types of solvers were defined. The so-called heuristic solvers improved solutions received from the SC associated with the central memory (called master in Groër and Golden, 2011), through a record-to-record meta-heuristic (Chao et al, 1995; Golden et al, 1998; Li et al, 2005). On completing the task, the solvers returned the 50 best solutions found and the corresponding routes (a post-optimization procedure was first run on each route). Simultaneously, exact solvers aimed to identify new solutions by solving series of set covering problems starting from a limited set of routes. Each time a set covering problem was solved, the solution was returned to the central memory and the set of the current 10 best solutions was retrieved for the next run. Set-covering solvers had also access to the ordered list of best routes in memory and they selected within to complete their problems. The number of routes selected to set up a set covering problem was dynamically modified during the search to control the corresponding computational effort. The method performed very well, both in terms of solution quality and computational effort (an almost-linear speedup was observed).

Jin et al (2014) proposed a different SC mechanism for a pC/KC meta-heuristic with tabu search solvers for the VRP. Data sharing was relatively simple; solvers periodically (after a number of iterations or when the solution did not improve for a number of iterations) sent best solutions to the central memory, and received a solution back from it, the search being resumed from the received solution. The SC mechanism aimed to identify and extract information from the solutions in memory to guide solvers toward intensification and diversification phases. This was obtained by dynamically (on reception) clustering solutions according to the number of edges in common. Thus, solutions in a given cluster share a certain number of edges, this cluster of edges and solutions being assumed to represent a region of the search space. Search history indicators were associated with clusters giving the number of solutions in the cluster and the quality of the solutions. This information was used to infer how thoroughly the corresponding region had been explored and how promising it appeared. Clusters were sorted according to the average solution value of their feasible solutions, and the cluster with the lowest value, that is, with the largest number of very good solutions, was selected for intensification, while the solution with the lowest number of good solutions was selected for diversification. A solution was then selected in the corresponding cluster and it was sent to the requesting solver. Excellent results were obtained in terms of solution quality and computation effort (an almost linear speedup was observed with up to 240 processors) compared to the state-of-the-art methods of the day.

We complete this section by addressing recent developments targeting multi-attribute, “rich”, problem settings where the problems at hand display a large number of attributes characterizing their feasibility and optimality structures. Traditionally, such problems were simplified, or sequentially solved through a series of particular cases, where part of the overall problem was fixed or ignored, or both. The general idea of the new generation of pC/KC meta-heuristics is to decompose the problem formulation along sets of decision

variables, which is called *decision-set attribute decomposition* in Lahrichi et al (2015). The goal of this decomposition is to obtain simpler but meaningful problem settings, in the sense that efficient solvers, can be “easily” obtained for the partial problems either by opportunistically using existing high-performing methods or by developing new ones. The central-memory asynchronous cooperative search framework then brings together these partial problems and their associated solvers, together with integration mechanisms, reconstructing complete solutions, and search-guidance mechanisms.

According to our best knowledge, Crainic et al (2006a) (see also Di Chiara, 2006) were the first to propose such a methodology in the context of designing wireless networks, where seven attributes were considered simultaneously. The proposed pC/KC/MPDS meta-heuristic had tabu search solvers working on limited subsets of attributes, the others being fixed, and a genetic method combining the partial solutions generated by the tabu search procedures into complete solutions to the initial problem.

The general method, called *Integrative Cooperative Search (ICS)*, was introduced in Lahrichi et al (2015) (for earlier versions, see Crainic et al, 2009a,b) and illustrated through an application to the multi-depot periodic vehicle routing problem (MDPVRP) (Mingozzi, 2005; Vidal et al, 2012). The main components of ICS, to be instantiated for each application, are 1) the decomposition rule; 2) the *Partial Solver Groups (PSGs)* addressing the partial problems resulting from the decomposition; 3) the *Integrators* selecting partial solutions from PSGs, combining them, and sending the resulting complete solutions to the *Complete Solver Group (CSG)*; and 4) the CSG, providing the central memory functionalities of ICS. Notice that, in order to facilitate the cooperation, a unique solution representation, obtained by fixing rather than eliminating variables when defining partial problems, is used throughout ICS.

The selection of the decision-sets for decomposition is specific to each application case, decision variables being clustered to yield known or identifiable optimization problem settings. Thus, an opportunistic rule decomposed the MDPVRP along the depot and period decision sets to create two partial problems, a periodic VRP (PVRP) and a multi-depot VRP (MDVRP), high-quality solvers being available in the literature for both problems.

The PSG may contain one or several solvers targeting particular subsets of attributes. Thus, two PSGs were defined in Lahrichi et al (2015), one for the PVRP and the other for the MDVRP. Each PSG was organized according to the pC/KC paradigm and was thus composed of a set of *Partial Solvers*, a central memory where elite solutions were kept, and a *Local Search Coordinator (LSC)* managing the local central memory and interfacing with the *Global Search Coordinator*. Two algorithms were used as partial solvers, the hybrid genetic algorithm HGSADC of Vidal et al (2012) and GUTS, a generalized version of the Unified Tabu Search Cordeau et al (2001).

Integrators build complete solutions by mixing partial solutions with promising features obtained within the PSGs. Integrators aim for solution quality, the transmission of critical features extracted from the partial solutions, and computational efficiency. The simplest Integrator consists of selecting high-quality partial solutions (with respect to solution value or the inclusion of particular decision combinations) and passing them directly to the Complete Solver Group. Population-based meta-heuristics make natural integrators, as well as solvers of optimization formulations combining solutions or solution elements (e.g., set covering for VRP) to yield complete solutions to the problem at hand. The work of El Hachemi et al (2015) belongs to the latter category, proposing particular optimization models for rich VRP settings, which preserve desired critical variables (desired attributes), present in partial solutions, when selecting and combining routes.

Several Integrators can be involved in an ICS meta-heuristic, increasing the diversity of the population of complete solutions. Four Integrators were thus included in the MDPVRP application, the simple one passing good solutions to the CSG, the crossover and individual education (enhancement) operators of HGSADC, and two of the methods proposed by El Hachemi et al (2015), the first transmitting the attributes for which there was “consensus” in the input solutions, while the second “promoted” them only through penalties added to the objective function. The last three integrators started from pairs of partial solutions randomly selected among the best 25% of the solutions in the central memories of the two PSGs.

The Complete Solver Group includes the central memory, where the complete solutions are stored, together with the context information and the guiding solutions built by the Global Search Coordinator (GSC). Complete solutions are received from Integrators and, when solvers are present in the CSG, these solutions are further enhanced. The GSC 1) builds the contextual information (e.g., the frequency of appearance of each (customer, depot, pattern) triplet in the complete solution set for the MDPVRP, together with the cost of the best solution containing it), 2) generates new guiding solutions to orient the search toward promising features, and 3) monitors the status of the solver groups, sending guiding instructions (solutions) when necessary.

Monitoring is performed by following the evolution of the PSGs (e.g., the number of improving solutions generated during a certain time period) to detect undesired situations, such as loss of diversity in the partial or complete populations, stagnation in improving the quality of the current best solution, awareness that some zones of the solution space - defined by particular values for particular decision sets - have been scarcely explored, etc. Whenever one of these situations is detected, the GSC sends guidance “instructions” to the particular PSG. The particular type of guidance is application specific, but one may inject new solutions or elements, modify the values of the fixed attributes for the PSG to orient its search toward a different area, change the attribute subset under investigation (i.e., change the decomposition of the decision-set attributes),

or modify/replace the solution method in a Partial Solver or Integrator. The last two should not occur too frequently. In Lahrichi et al (2015), guidance took the form of three solutions, which were either randomly selected from the complete solution set, or were built by the GSC out of promising solution elements with respect to the search history.

Lahrichi et al (2015) reported very good results even when compared to the state-of-the-art meta-heuristic. The experimental results also indicated that 1) one should use solvers with similar time performances in order to have them contributing reasonably equally to the cooperation; 2) when using genetic solvers in a PSG it is preferable for long runs to define a local population for each such solver, while using the central memory as population for all cooperating genetic solvers appears better for short runs; and 3) embedding good solvers in the CSG enhances slightly the already excellent performance of the ICS parallel meta-heuristic.

8 Conclusions

This chapter presented an overview and state-of-the-art survey of the main parallel meta-heuristic ideas, discussing general concepts and algorithm design principles and strategies. Four main classes of parallel meta-heuristics strategies were described: low-level decomposition of computing-intensive tasks with no modification to the original algorithm, decomposition of the search space, independent multi-search, and cooperative multi-search, the latter encompassing synchronous, asynchronous collegial and knowledge-creating asynchronous collegial strategies. It is noteworthy that this series also reflects the historical sequence of the development of parallel meta-heuristics, which are now acknowledged, cooperative search strategies in particular, as making up their own class of meta-heuristics.

It must be emphasized that each of these strategy classes fulfills a particular type of task and all are needed at some time. Thus, the idea that everything seems to be known regarding low-level parallelization strategies is not true. Most studies on accelerating computing-intensive tasks targeted the evaluation of a population or neighborhood in classic meta-heuristic frameworks but, as a number of more recent studies show, the best strategy to accelerate a local-search procedure may prove less effective when the local search is embedded into a full meta-heuristics or hierarchical solution method. On the other hand, the evolution of computing infrastructure, in particular, the integration of graphical processing units within computing platforms, opens up interesting but challenging perspectives. In both cases, more research is needed to understand their behavior and identify the most appropriate combination of strategies, particularly low-level and cooperative search, for various meta-heuristics, problem settings, and computing platforms.

Search-space decomposition also seems to have been thoroughly studied, and has been overlooked in the last years, maybe due to the rapid and phenomenal increase in the memory available and the speed of access. Let us not forget, however, that most optimization problems of interest are complex and that the dimensions of the instances one faces in practice keep increasing. Research challenges exist in dynamic search-space decomposition and the combination of cooperative search and search-space decomposition. The Integrative Cooperative Search is a first answer in this direction, but more research is needed.

Asynchronous cooperation, particularly when relaying on memories as communication mechanisms, provides a powerful, flexible and adaptable framework for parallel meta-heuristics that consistently achieved good results in terms of computing efficiency and solution quality for many meta-heuristic and problem classes. A number of challenging research issues are worth investigating.

A first issue concerns the exchange and utilization of context data locally generated by the cooperating solvers, to infer an image of the status of the global search and generate appropriate guiding instructions. Thus, contrasting the various local context data may be used to identify regions of the search space that were neglected or over explored. The information could also be used to evaluate the relative performance of the solvers conducting, eventually, to adjust the search parameters of particular solvers or even change the search strategy. So-called “strategic” decision variables or parameters could thus be more easily identified, which could prove very profitable in terms of search guidance.

A related issue concerns the learning processes and the creation of new information out of the shared data. Important questions concern the identification of information that may be derived from the exchanged solutions and its usefulness in inferring the status of the global search, and determining the appropriate guiding information to be sent to solvers. Research in this direction is still at the very beginning but has already proved its worth, in particular in the context of the integrative cooperative methods.

A third broad issue concerns the cooperation of different types of meta-heuristics, as well as the cooperation of meta-heuristics with exact solution methods. The so-called hybrid and matheuristic methods, representing the former and latter types of method combination, respectively, are trendy in the sequential optimization field. Very few studies explicitly target parallel methods, however. How different methods behave when involved in cooperative search and how the latter behaves given various combinations of methods is an important issue that should yield valuable insights into the design of parallel meta-heuristic algorithms, cooperative ones in particular. A particularly challenging but fascinating direction for cooperative search and ICS is represented by the multi-scenario representation of stochastic optimization formulations, for which almost nothing beyond low-level scenario-decomposition has been proposed yet. Transversal

studies comparing the behavior and performance of particular parallel meta-heuristic strategies over different problem classes, and of different parallel strategies and implementations for the same problem class, would be very valuable in this context, as in the broader field of parallel meta-heuristics.

Acknowledgements

The author wishes to acknowledge the contributions of colleagues and students, in particular Professors Michel Gendreau, Université de Montréal, Canada, and Michel Toulouse, the Vietnamese-German University, Vietnam, who collaborated over the years to the work on parallel meta-heuristics for combinatorial optimization. All errors are, however, solely and entirely due to the author.

Partial funding for this project has been provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), through its Discovery Grant and the Discovery Accelerator Supplements programs, and the Strategic Clusters program of the Fonds de Recherche Québécois Nature et Technologies (FRQNT) The author thanks the two institutions for supporting this research.

References

- Alba E (ed) (2005) *Parallel Metaheuristics: A New Class of Algorithms*. John Wiley & Sons, Hoboken, NJ
- Alba E, Luque G, Nasmachnow S (2013) Parallel Metaheuristics: Recent Advances and New Trends. *International Transactions in Operational Research* 20(1):1–48
- Badeau P, Gendreau M, Guertin F, Potvin JY, Taillard E (1997) A Parallel Tabu Search Heuristic for the Vehicle Routing Problem with Time Windows. *Transportation Research Part C: Emerging Technologies* 5(2):109–122
- Banos R, Ortega J, Gil C, Fernandez A, de Toro F (2013) A Simulated Annealing-Based Parallel Multi-Objective Approach to Vehicle Routing Problems with Time Windows. *Expert Systems with Applications* 40(5):1696–1707
- Barr RS, Hickman BL (1993) Reporting Computational Experiments with Parallel Algorithms: Issues, Measures, and Experts Opinions. *ORSA Journal on Computing* 5(1):2–18
- Bastos MP, Ribeiro CC (1999) Reactive Tabu Search with Path-Relinking for the Steiner Problem in Graphs. In: Voß S, Martello S, Roucairol C, Osman, IH (eds) *Meta-*

- Heuristics 98: Theory & Applications, Kluwer Academic Publishers, Norwell, MA, pp 31–36
- Battiti R, Tecchiolli G (1992) Parallel Based Search for Combinatorial Optimization: Genetic Algorithms and TABU. *Microprocessors and Microsystems* 16(7):351–367
- Blazewicz J, Moret-Salvador A, Walkowiak R (2004) Parallel Tabu Search Approaches for Two-Dimensional Cutting. *Parallel Processing Letters* 14(1):23–32
- Bortfeldt A, Gehring H, Mack D (2003) A Parallel Tabu Search Algorithm for Solving the Container Loading Problem. *Parallel Computing* 29(5):641–662
- Brodtkorb AR, Hagen TR, Schulz C, Hasle G (2013a) GPU Computing in Discrete Optimization. Part I: Introduction to the GPU. *EURO Journal on Transportation and Logistics* 2(1-2):129–157
- Brodtkorb AR, Hagen TR, Schulz C, Hasle G (2013b) GPU Computing in Discrete Optimization. Part II: Survey Focused on Routing Problems. *EURO Journal on Transportation and Logistics* 2(1-2):159–186
- Cantú-Paz E (2005) Theory of Parallel Genetic Algorithms. In: Alba, E (ed) *Parallel Metaheuristics: A New Class of Algorithms*, John Wiley & Sons, Hoboken, pp 425–445
- Cavalcante CBC, Cavalcante VF, Ribeiro CC, Souza MC (2002) Parallel Cooperative Approaches for the Labor Constrained Scheduling Problem. In: Ribeiro C, Hansen P (eds) *Essays and Surveys in Metaheuristics*, Kluwer Academic Publishers, Norwell, MA, pp 201–225
- Cecilia JM, García JM, Nisbet A, Amos M, Ujaldón M (2013) Enhancing Data Parallelism for Ant Colony Optimization on GPUs. *Journal of Parallel and Distributed Computing* 73(1):42–51
- Chao IM, Golden B L, Wasil EA (1995) An Improved Heuristic for the Period Vehicle Routing Problem. *Networks* 26(1):25–44
- Cordeau JF, Maischberger M (2012) A Parallel Iterated Tabu Search Heuristic for Vehicle Routing Problems. *Computers & Operations Research* 39(9):2033–2050
- Cordeau JF, Laporte G, Mercier A (2001) A Unified Tabu Search Heuristic for Vehicle Routing Problems with Time Windows. *Journal of the Operational Research Society* 52(8):928–936
- Crainic TG (2005) Parallel Computation, Co-operation, Tabu Search. In: Rego C, Alidaee B (eds) *Metaheuristic Optimization Via Memory and Evolution: Tabu Search and Scatter Search*, Kluwer Academic Publishers, Norwell, MA, pp 283–302

- Crainic TG (2008) Parallel Solution Methods for Vehicle Routing Problems. In: Golden, BL, Raghavan, S, Wasil, EA (eds) *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer, New York, pp 171–198
- Crainic TG (2017) Parallel Meta-Heuristic Search. In: Marti, R, Pardalos, PM, Resende, MGC (eds) *Handbook of Heuristics*, Springer, New York, in print
- Crainic TG, Gendreau M (1999) Towards an Evolutionary Method - Cooperating Multi-Thread Parallel Tabu Search Hybrid. In: Voß S, Martello S, Roucairol C, Osman, IH (eds) *Meta-Heuristics 98: Theory & Applications*, Kluwer Academic Publishers, Norwell, MA, pp 331–344
- Crainic TG, Gendreau M (2002) Cooperative Parallel Tabu Search for Capacitated Network Design. *Journal of Heuristics* 8(6):601–627
- Crainic TG, Hail N (2005) Parallel Meta-Heuristics Applications. In: Alba, E (ed) *Parallel Metaheuristics: A New Class of Algorithms*, John Wiley & Sons, Hoboken, NJ, pp 447–494
- Crainic TG, Toulouse M (1998) Parallel Metaheuristics. In: TG Crainic, G Laporte (eds) *Fleet Management and Logistics*, Kluwer Academic Publishers, Norwell, MA, pp 205–251
- Crainic TG, Toulouse M (2003) Parallel Strategies for Meta-heuristics. In: F Glover, G Kochenberger (eds) *Handbook in Metaheuristics*, Kluwer Academic Publishers, Norwell, MA, pp 475–513
- Crainic TG, Toulouse M (2008) Explicit and Emergent Cooperation Schemes for Search Algorithms. In: Maniezzo, V, Battiti, R, Watson, J-P (eds) *Learning and Intelligent Optimization*, Springer-Verlag, Berlin, *Lecture Notes in Computer Science*, vol 5315, pp 95–109
- Crainic TG, Toulouse M (2010) Parallel Meta-Heuristics. In: Gendreau, M, Potvin, J-Y (eds) *Handbook of Metaheuristics (2nd Edition)*, Springer, pp 497–541
- Crainic TG, Toulouse M, Gendreau M (1996) Parallel Asynchronous Tabu Search for Multicommodity Location-Allocation with Balancing Requirements. *Annals of Operations Research* 63:277–299
- Crainic TG, Toulouse M, Gendreau M (1997) Towards a Taxonomy of Parallel Tabu Search Algorithms. *INFORMS Journal on Computing* 9(1):61–72
- Crainic TG, Gendreau M, Hansen P, Mladenović N (2004) Cooperative Parallel Variable Neighborhood Search for the p -Median. *Journal of Heuristics* 10(3):293–314
- Crainic TG, Di Chiara B, Nonato M, Tarricone L (2006a) Tackling Electrosmog in Completely Configured 3G Networks by Parallel Cooperative Meta-Heuristics. *IEEE Wireless Communications* 13(6):34–41

- Crainic TG, Li Y, Toulouse M (2006b) A First Multilevel Cooperative Algorithm for the Capacitated Multicommodity Network Design. *Computers & Operations Research* 33(9):2602–2622
- Crainic TG, Crisan GC, Gendreau M, Lahrichi N, Rei W (2009a) A Concurrent Evolutionary Approach for Cooperative Rich Combinatorial Optimization. In: *Genetic and Evolutionary Computation Conference - GECCO 2009*, July 8-12, Montréal, Canada, ACM, CD-ROM
- Crainic TG, Crisan GC, Gendreau M, Lahrichi N, Rei W (2009b) Multi-thread Integrative Cooperative Optimization for Rich Combinatorial Problems. In: *The 12th International Workshop on Nature Inspired Distributed Computing - NIDISC'09*, 25-29 May, Rome, CD-ROM
- Crainic TG, Davidović T, Ramljak D (2014) Designing Parallel Meta-Heuristic Methods. In: *Despotovic-Zrasic, M, Milutinovic, V, Belic, A (eds) High Performance and Cloud Computing in Scientific Research and Education*, IGI Global, Hershey, PA, U.S.A., pp 260–280
- Czech ZJ (2000) A Parallel Genetic Algorithm for the Set Partitioning Problem. In: *8th Euromicro Workshop on Parallel and Distributed Processing*, pp 343–350
- Dai C, Li B, Toulouse M (2009) A Multilevel Cooperative Tabu Search Algorithm for the Covering Design Problem. *Journal of Combinatorial Mathematics and Combinatorial Computing* 68:35–65
- Davidović T, Crainic TG (2015) Parallel Local Search to Schedule Communicating Tasks on Identical Processors. *Parallel Computing* 48:1–14
- Delévacq A, Delisle P, Gravel M, Krajecki M (2013) Parallel Ant Colony Optimization on Graphics Processing Units. *Journal of Parallel and Distributed Computing* 73(1):52–61
- Di Chiara B (2006) Optimum Planning of 3G Cellular Systems: Radio Propagation Models and Cooperative Parallel Meta-heuristics. PhD thesis, Dipartimento di ingegneria dell'innovazione, Università degli Studi di Lecce, Lecce, Italy
- Doerner KF, Hartl RF, Benkner S, Lucka M (2006) Cooperative Savings Based Ant Colony Optimization - Multiple Search and Decomposition Approaches. *Parallel Processing Letters* 16(3):351–369
- Drias H, Ibri A (2003) Parallel ACS for Weighted MAX-SAT. In: *Mira, J, Álvarez, J (eds) Artificial Neural Nets Problem Solving Methods - Proceedings of the 7th International Work-Conference on Artificial and Natural Neural Networks*, Lecture Notes in Computer Science, vol 2686, Springer-Verlag, Heidelberg, pp 414–421
- El Hachemi N, Crainic TG, Lahrichi N, Rei W, Vidal T (2015) Solution Integration in Combinatorial Optimization with Applications to Cooperative Search and Rich Vehicle Routing. *Journal of Heuristics* 21(5):663–685

- Flores CD, Cegla BB, Caceres DB (2003) Telecommunication Network Design with Parallel Multi-objective Evolutionary Algorithms. In: Proceedings of the 2003 IFIP/ACM Latin America Networking Conference - Towards a Latin American Agenda for Network Research, ACM, New York, NY, pp 1–11
- García-López F, Melián-Batista B, Moreno-Pérez JA, Moreno-Vega JM (2002) The Parallel Variable Neighborhood Search for the p -Median Problem. *Journal of Heuristics* 8(3):375–388
- García-López F, Melián-Batista B, Moreno-Pérez JA, Moreno-Vega JM (2003) Parallelization of the Scatter Search for the p -Median Problem. *Parallel Computing* 29:575–589
- García-López F, García Torres M, Melián-Batista B, Moreno-Pérez JA, Moreno-Vega JM (2005) Parallel Scatter Search. In: Alba, E (ed) *Parallel Metaheuristics: A New Class of Metaheuristics*, John Wiley & Sons, Hoboken, NJ, pp 223–246
- García-López F, García Torres M, Melián-Batista B, Moreno-Pérez JA, Moreno-Vega JM (2006) Solving Feature Subset Selection Problem by a Parallel Scatter Search. *European Journal of Operational Research* 169(2):477–489
- Gehring H, Homberger J (2002) Parallelization of a Two-Phase Metaheuristic for Routing Problems with Time Windows. *Journal of Heuristics* 8(3):251–276
- Gendreau M, Hertz A, Laporte G (1994) A Tabu Search Heuristic for the Vehicle Routing Problem. *Management Science* 40(10):1276–1290
- Gendreau M, Guertin F, Potvin JY, Taillard ÉD (1999) Tabu Search for Real-Time Vehicle Routing and Dispatching. *Transportation Science* 33(4):381–390
- Gendreau M, Laporte G, Semet F (2001) A Dynamic Model and Parallel Tabu Search Heuristic for Real-time Ambulance Relocation. *Parallel Computing* 27(12):1641–1653
- Glover F (1989) Tabu Search – Part I. *ORSA Journal on Computing* 1(3):190–206
- Glover F (1990) Tabu Search – Part II. *ORSA Journal on Computing* 2(1):4–32
- Glover F (1996) Tabu Search and Adaptive Memory Programming – Advances, Applications and Challenges. In: Barr R, Helgason R, Kennington J (eds) *Interfaces in Computer Science and Operations Research*, Kluwer Academic Publishers, Norwell, MA, pp 1–75
- Glover F, Laguna M (1997) *Tabu Search*. Kluwer Academic Publishers, Norwell, MA
- Golden B L, Wasil EA, Kelly JP, Chao IM (1998) Metaheuristics in Vehicle Routing. In: Crainic TG, Laporte G (eds) *Fleet Management and Logistics*, Kluwer Academic Publishers, Norwell, MA, pp 33–56

- Groër C, Golden B (2011) A Parallel Algorithm for the Vehicle Routing Problem. *INFORMS Journal on Computing* 23(2):315–330
- Hidalgo JI, Prieto M, Lanchares J, Baraglia R, Tirado F, Garnica O (2003) Hybrid Parallelization of a Compact Genetic Algorithm. In: *Proceedings of the 11th uromicro Conference on Parallel, Distributed and Network-Based Processing*, pp 449–455
- Izzo D, Rucinski M, Ampatzis C (2009) Parallel Global Optimisation Meta-heuristics Using an Asynchronous Island-model. In: *CEC'09 - IEEE Congress on Evolutionary Computation*, IEEE, pp 2301–2308
- James T, Rego C, Glover F (2009) A Cooperative Parallel Tabu Search Algorithm for the Quadratic Assignment Problem. *European Journal of Operational Research* 195(3):810–826
- Jin J, Crainic TG, Løkketangen A (2012) A Parallel Multi-Neighborhood Cooperative Tabu Search for Capacitated Vehicle Routing Problems. *European Journal of Operational Research* 222(3):441–451
- Jin J, Crainic TG, Løkketangen A (2014) A Cooperative Parallel Metaheuristic for the Capacitated Vehicle Routing Problem. *Computers & Operations Research* 44:33–41
- Lahrichi N, Crainic TG, Gendreau M, Rei W, Crisan CC, Vidal T (2015) An Integrative Cooperative Search Framework for Multi-Decision-Attribute Combinatorial Optimization. *European Journal of Operational Research* 246(2):400–412
- Le Bouthillier A (2007) *Recherches coopératives pour la résolution de problèmes d'optimisation combinatoire*. PhD thesis, Département d'informatique et de recherche opérationnelle, Université de Montréal, Montréal, QC, Canada
- Le Bouthillier A, Crainic TG (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 TG, Kropf P (2005) A Guided Cooperative Search for the Vehicle Routing Problem with Time Windows. *IEEE Intelligent Systems* 20(4):36–42
- Lee SY, Lee KG (1996) Synchronous and Asynchronous Parallel Simulated Annealing with Multiple Markov Chains. *IEEE Transactions on Parallel and Distributed Systems* 7(10):993–1007
- Li F, Golden B L, Wasil EA (2005) Very Large-Scale Vehicle Routing: New Test Problems, Algorithms, and Results. *Computers & Operations Research* 32(5):1165–1179
- Ling C, Hai-Ying S, Shu W (2012) A Parallel Ant Colony Algorithm on Massively Parallel Processors and its Convergence Analysis for the Travelling Salesman Problem. *Information Sciences* 199:31–42

- Melab N, Talbi EG, Cahon S, Alba E, Luque G (2006) Parallel Metaheuristics: Models and Frameworks. In: EL-Ghazali Talbi (ed) *Parallel Combinatorial Optimization*, John Wiley & Sons, New York, pp 149–162
- Melab N, Luong TV, Boufaras K, Talbi EG (2011) Towards ParadisEO-MO-GPU: A Framework for GPU-based Local Search Metaheuristics. In: Cabestany J, Rojas I, Joya G (eds) *Advances in Computational Intelligence - Lecture Notes in Computer Science 6691*, Springer, pp 401–408
- Michels R, Middendorf M (1999) An Ant System for the Shortest Common Supersequence Problem. In: Corne D, Dorigo M, Glover F (eds) *New Ideas in Optimization*, McGraw-Hill, pp 51–61
- Middendorf M, Reischle F, Schmeck H (2002) Multi Colony Ant Algorithms. *Journal of Heuristics* 8(3):305–320
- Mingozi A (2005) The Multi-depot Periodic Vehicle Routing Problem. In: Zucker JD, Saitta L (eds) *Abstraction, Reformulation and Approximation*, no. 3607 in *Lecture Notes in Computer Science*, Springer, Berlin / Heidelberg, pp 347–350
- Niar S, Fréville A (1997) A Parallel Tabu Search Algorithm For The 0-1 Multidimensional Knapsack Problem. In: *11th International Parallel Processing Symposium (IPPS '97)*, Geneva, Switzerland, IEEE, pp 512–516
- Oduntan IO, Toulouse M, Baumgartner R, Bowman C, Somorjai R, Crainic TG (2008) A Multilevel Tabu Search Algorithm for the Feature Selection Problem in Biomedical Data Sets. *Computers & Mathematics with Applications* 55(5):1019–1033
- Ouyang M, Toulouse M, Thulasiraman K, Glover F, Deogun JS (2000) Multi-Level Cooperative Search: Application to the Netlist/Hypergraph Partitioning Problem. In: *Proceedings of International Symposium on Physical Design*, ACM Press, pp 192–198
- Ouyang M, Toulouse M, Thulasiraman K, Glover F, Deogun JS (2002) Multilevel Cooperative Search for the Circuit/Hypergraph Partitioning Problem. *IEEE Transactions on Computer-Aided Design* 21(6):685–693
- Pedemonte M, Nesmachnow S, Cancela H (2011) A Survey of Parallel Ant Colony Optimization. *Applied Soft Computing* 11(8):5181–5197
- Polacek M, Benkner S, Doerner KF, Hartl RF (2008) A Cooperative and Adaptive Variable Neighborhood Search for the Multi Depot Vehicle Routing Problem with Time Windows. *Business Research* 1(2):207–218
- Rego C (2001) Node Ejection Chains for the Vehicle Routing Problem: Sequential and Parallel Algorithms. *Parallel Computing* 27(3):201–222
- Ribeiro CC, Rosseti I (2007) Efficient Parallel Cooperative Implementations of GRASP Heuristics. *Parallel Computing* 33(1):21–35

- Rios E, Ochi LS, Boeres C, Coelho VN, Coelho IM, Faria R (2017) Exploring parallel multi-GPU local search strategies in a metaheuristic framework. *Journal of Parallel and Distributed Computing* DOI: <https://doi.org/10.1016/j.jpdc.2017.06.011>
- Rochat Y, Taillard ED (1995) Probabilistic Diversification and Intensification in Local Search for Vehicle Routing. *Journal of Heuristics* 1(1):147–167
- Sanvicente-Sánchez H, Frausto-Solís J (2002) MPSA: A Methodology to Parallelize Simulated Annealing and Its Application to the Traveling Salesman Problem. In: Coello Coello C, de Albornoz A, Sucar L, Battistutti O (eds) *MICAI 2002: Advances in Artificial Intelligence, Lecture Notes in Computer Science*, vol 2313, Springer-Verlag Heidelberg, pp 89–97
- Schulze J, Fahle T (1999) A Parallel Algorithm for the Vehicle Routing Problem with Time Window Constraints. *Annals of Operations Research* 86:585–607
- Sevcli M, Aydin ME (2007) Parallel Variable Neighbourhood Search Algorithms for Job Shop Scheduling Problems. *IMA Journal of Management Mathematics* 18(2):117–133
- Solar M, Parada V, Urrutia R (2002) A Parallel Genetic Algorithm to Solve the Set-Covering Problem. *Computers & Operations Research* 29(9):1221–1235
- Stutzle T (1998) Parallelization Strategies for Ant Colony Optimization. In: Eiben, AE, Back, T, Schoenauer, M, Schwefel, H-P (eds) *Proceedings of Parallel Problem Solving from Nature V, Lecture Notes in Computer Science*, vol 1498, Springer-Verlag, Heidelberg, pp 722–731
- Taillard ED (1993) Parallel Iterative Search Methods for Vehicle Routing Problems. *Networks* 23(8):661–673
- Taillard ED (1994) Parallel Taboo Search Techniques for the Job Shop Scheduling Problem. *ORSA Journal on Computing* 6(2):108–117
- Taillard ED, Gambardella LM, Gendreau M, Potvin JY (1997) Adaptive Memory Programming: A Unified View of Metaheuristics. *European Journal of Operational Research* 135(1):1–10
- Talbi EG (ed) (2006) *Parallel Combinatorial Optimization*. Wiley-Interscience, Wiley & Sons, Hoboken, NJ
- Talbi EG (ed) (2009) *Metaheuristics: From Design to Implementation*. John Wiley & Sons, Hoboken, NJ
- Talukdar S, Murthy S, Akkiraju R (2003) Asynchronous Teams. In: F Glover, G Kochenberger (eds) *Handbook in Metaheuristics*, Kluwer Academic Publishers, Norwell, MA, pp 537–556

- Tan Y, Ding K (2016) A Survey on GPU-Based Implementation of Swarm Intelligence Algorithms. *IEEE Transactions on Cybernetics* 46(9):2168–2267
- ten Eikelder HMM, Aarts BJJ, Verhoeven MGA, Aarts EHL (1999) Sequential and Parallel Local Search for Job Shop Scheduling. In: Voß S, Martello S, Roucairol C, Osman IH (eds) *Meta-Heuristics 98: Theory & Applications*, Kluwer Academic Publishers, Norwell, MA, pp 359–371
- Tongcheng G, Chundi M (2002) Radio Network Using Coarse-Grained Parallel Genetic Algorithms with Different Neighbor Topology. In: *Proceedings of the 4th World Congress on Intelligent Control and Automation*, vol 3, pp 1840–1843
- Toulouse M, Crainic TG, Gendreau M (1996) Communication Issues in Designing Cooperative Multi Thread Parallel Searches. In: Osman IH, Kelly JP (eds) *Meta-Heuristics: Theory & Applications*, Kluwer Academic Publishers, Norwell, MA, pp 501–522
- Toulouse M, Crainic TG, Sansó B, Thulasiraman K (1998) Self-Organization in Cooperative Search Algorithms. In: *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*, Omnipress, Madison, WI, pp 2379–2385
- Toulouse M, Crainic TG, Sansó B (1999a) An Experimental Study of Systemic Behavior of Cooperative Search Algorithms. In: Voß S, Martello S, Roucairol C, Osman IH (eds) *Meta-Heuristics 98: Theory & Applications*, Kluwer Academic Publishers, Norwell, MA, pp 373–392
- Toulouse M, Thulasiraman K, Glover F (1999b) Multi-Level Cooperative Search: A New Paradigm for Combinatorial Optimization and an Application to Graph Partitioning. In: Amestoy P, Berger P, Daydé M, Duff I, Frayssé V, Giraud L, Ruiz D (eds) *5th International Euro-Par Parallel Processing Conference*, Lecture Notes in Computer Science, vol 1685, Springer-Verlag, Heidelberg, pp 533–542
- Toulouse M, Crainic TG, Thulasiraman K (2000) Global Optimization Properties of Parallel Cooperative Search Algorithms: A Simulation Study. *Parallel Computing* 26(1):91–112
- Toulouse M, Crainic TG, Sansó B (2004) Systemic Behavior of Cooperative Search Algorithms. *Parallel Computing* 30(1):57–79
- Vallada E, Ruiz R (2009) A Cooperative Metaheuristics for the Permutation Flowshop Scheduling Problem. *European Journal of Operational Research* 193(2):365–376
- Van Luong T, Melab N, Talbi EG (2013) GPU Computing for Parallel Local Search Metaheuristic Algorithms. *IEEE Transactions on Computers* 62(1):173–185
- Vidal T, Crainic TG, Gendreau M, Lahrichi N, Rei W (2012) A Hybrid Genetic Algorithm for Multi-Depot and Periodic Vehicle Routing Problems. *Operations Research* 60(3):611–624