

1 A MULTI-AGENT BASED COOPERATIVE
2 APPROACH TO SCHEDULING AND ROUTING

3 Simon Martin¹, Djamila Ouelhadj², Patrick Beullens³,
4 Ender Ozcan⁴, Angel A. Juan⁵, Edmund.K.Burke¹

5 ¹*Computational Heuristics Operational Research Decision Support (CHORDS) Group,*
6 *University of Stirling, Department of Mathematics and Computer Science, UK*
7 *email: spm@cs.stir.ac.uk*

8 ²*Centre of Operational Research and Logistics,*
9 *University of Portsmouth, Department of Mathematics, UK*
10 *email: djamila.ouelhadj@port.ac.uk*

11 ³*Mathematical Sciences and Southampton Business School and CORMSIS, University of*
12 *Southampton, SO17 1BJ, United Kingdom*
13 *email: P.Beullens@soton.ac.uk*

14 ⁴*Automated Scheduling, Optimisation and Planning Research Group,*
15 *University of Nottingham, Department of Computer Science, UK*
16 *email: Ender.Ozcan@nottingham.ac.uk*

17 ⁵*Department of Computer Science and Telecommunication,*
18 *Open University of Catalonia, Rambla Poblenou, 156 088018 Barcelona, Spain*
19 *email:ajuamp@uoc.edu*

20 **Abstract**

In this study, we propose a general agent-based distributed framework where each agent is implementing a different metaheuristic/local search combination. Moreover, an agent continuously adapts itself during the search process using a direct cooperation protocol based on reinforcement learning and pattern matching. Good patterns that make up improving solutions are identified and shared by the agents. This agent-based system aims to provide a modular flexible framework to deal with a variety of different problem domains. We have evaluated the performance of this approach using the proposed framework which embodies a set of well known metaheuristics with different configurations as agents on two problem domains, Permutation Flow-shop Scheduling and Capacitated Vehicle Routing. The results show the success of the approach yielding three new best known results of the Capacitated Vehicle Routing benchmarks tested, while the results for Permutation Flow-shop Scheduling are commensurate with the best known

values for all the benchmarks tested.

1 *Keywords:* combinatorial optimization, multi-agent systems, scheduling,
2 vehicle routing, metaheuristics, cooperative search, reinforcement learning.

3 1. INTRODUCTION

4 Heuristics are rules of thumb for solving specific computationally hard
5 problems. Researchers and practitioners use heuristics when exact methods
6 fail to produce any solutions with a “reasonable” quality in a “reasonable”
7 amount of time. Heuristics often come with a set of parameters, each requir-
8 ing tuning for an improved performance. Moreover, different heuristics can
9 perform well on different problem instances. Hence, there is a growing num-
10 ber of studies on more general methodologies which are applicable to different
11 problem domains for tuning the parameters (López-Ibáñez et al., 2011; Hut-
12 ter et al., 2007; Ries and Beullens, 2015), generating or mixing/controlling
13 heuristics (Burke et al., 2013; Ross, 2014). In this study, we take an alter-
14 native approach and use cooperating agents, where each agent is enabled to
15 take a different approach with different parameter settings.

16 By cooperative search we mean that (meta)heuristics, executed in parallel
17 as agents, have the ability to share information at various points through-
18 out a search. To this end, we propose a modular agent-based framework
19 where the agents cooperate using a direct peer to peer asynchronous mes-
20 sage passing protocol. An island model is used where each agent has its own
21 representation of the search environment. Each agent is autonomous and
22 can execute different metaheuristic/local search combinations with different
23 parameter settings. Cooperation is based on the general strategies of pattern
24 matching and reinforcement learning where the agents share partial solutions
25 to enhance their overall performance.

26 The framework has the following additional characteristics. By using
27 ontologies (see Section 3.2), we are aiming to provide a framework that is
28 flexible enough to be used on more than one type of combinatorial optimi-
29 sation problem with little or no parameter tuning. This is achieved by using
30 our scheduling and routing ontology to translate target problems into an in-
31 ternal format that the agents can use to solve problems. So far, this approach
32 has been applied successfully to Capacitated Vehicle Routing (CVRP), Per-
33 mutation Flow shop Scheduling (PFSP), reported here and Nurse Rostering
34 reported in Martin et al. (2013).

1 The aim of this study is to develop a modular framework for cooperative
2 search that can be deployed, with little reconfiguration, to more than one type
3 of problem. We also test whether interaction between (meta)heuristics leads
4 to improved performance and if increasing the number of agents improves
5 the overall solution quality.

6 *1.1. Cooperative search in OR: literature*

7 The interest in cooperative search has risen due to successes in finding
8 novel ways to combine search algorithms. Cooperative search can be per-
9 formed by the exchange of states, solutions, sub-problems, models, or search
10 space characteristics. For a general introduction, see e.g. Clearwater et al.
11 (1992); Hogg and Williams (1993); Toulouse et al. (1999); Blum and Roli
12 (2003); Talbi and Bachelet (2006); Crainic and Toulouse (2008). Several
13 frameworks have been proposed recently, incorporating metaheuristics, as in
14 Talbi and Bachelet (2006); Milano and Roli (2004); Meignan et al. (2008,
15 2010), or hyper-heuristics, as in Ouelhadj and Petrovic (2010). Also, El Ha-
16 chemi et al. (2014) explore a general agent-based framework for solution
17 integration where distributed systems use different heuristics to decompose
18 and then solve a problem.

19 In an effort to find ways to combine different metaheuristics in such a
20 way that they cooperate with each other during their execution, a num-
21 ber of design choices have to be made. According to Crainic and Toulouse
22 (2008) an *asynchronous* framework in particular could result in an improved
23 search methodology; communication can then either be many-to-many (di-
24 rect), where each metaheuristic communicates with every other, or it can be
25 memory based (indirect), where information is sent to a pool that (other)
26 metaheuristics can make use of as required.

27 Most cooperative search mechanisms in the OR literature deploy indi-
28 rect communication through some central pool or adaptive memory. This
29 can take the form of passing whole, or possibly, partial solutions, to the
30 pool. See Talbi and Bachelet (2006); Milano and Roli (2004); Meignan et al.
31 (2008, 2010); Malek (2010). Aydin and Fogarty (2004b) applied this ap-
32 proach to job shop scheduling. Recently Barbucha (2014) has proposed an
33 agent-based system for Vehicle Routing Problems where agents instantiate
34 different metaheuristics which communicate through a shared pool.

35 Direct communication, instead, is used only in Vallada and Ruiz (2009);
36 Aydin and Fogarty (2004a), where whole solutions are passed from one pro-
37 cess to another in an island model executing a genetic or an evolutionary

1 simulated annealing algorithm respectively, and in Ouelhadj and Petrovic
2 (2010), where a similar set-up is used for a hyper-heuristic. All three pa-
3 pers addressed the PFSP. Also, this approach is to an extent present in the
4 evolutionary system of Xie and Liu (2009), who investigated the Travelling
5 Salesman Problem. Kouider and Bouzouia (2012) propose a direct com-
6 munication multi agent system for job shop scheduling where each agent is
7 associated with a specific machine in a production facility. Here a problem
8 is decomposed into several sub-problems by a “supervisor agent”. These
9 are passed to “resource agents” for execution and then passed back to the
10 supervisor to build the global solution.

11 Little work has been done on asynchronous direct cooperation where par-
12 tial solutions are rated and their parameters are communicated between au-
13 tonomous agents all working on the total problem. So far, no direct co-
14 operation strategy has been applied to more than one problem domain in
15 combinatorial optimisation. To this end, the agents are truly autonomous
16 and not synchronised. There is a gap in the literature regarding agents co-
17 operating directly and asynchronously where the communication is used for
18 the adaptive selection of moves with parameters.

19 The outline for the rest of the paper is as follows. Section 2 provides
20 formal problem statements for the two case studies. Section 3 describes
21 the proposed modular multi-agent framework for cooperative search, while
22 Section 4 describes how it is implemented. In Section 5 we discuss the exper-
23 imental design. In Section 6 we report the results of the tests where, to the
24 best of our knowledge, for three of the capacitated vehicle routing instances
25 we achieved better results than have been reported in the literature. Finally,
26 Section 7 presents conclusions and suggestions for future work.

27 **2. TEST CASE PROBLEMS**

28 In this section we offer brief problem descriptions of the case studies
29 applied to the agent-based framework proposed in this paper. We chose these
30 instances as they are representative scheduling and routing problems. The
31 algorithms instantiated by the framework are state-of-art implementations
32 (Juan et al., 2014, 2013, 2010a,b). These are all examples of Simheuristics
33 (Juan et al., 2015). This makes them a good fit with the partial solutions
34 identified by the system.

1 *2.1. Permutation flow-shop Scheduling Problem*

2 Given is a set of n jobs, $J = \{1, \dots, n\}$, available at a given time 0, and
 3 each to be processed on each of a set of m machines in the same order,
 4 $M = \{1, \dots, m\}$. A job $j \in J$ requires a fixed but job-specific non-negative
 5 processing time $p_{j,i}$ on each machine $i \in M$. The objective of the PFSP is to
 6 minimise the *makespan*. That is, to minimise the completion time of the last
 7 job on the last machine C_{max} (Pinedo, 2002). A feasible schedule is hence
 8 uniquely represented by a permutation of the jobs. There are $n!$ possible
 9 permutations and the problem is NP-complete (Garey et al., 1976).

10 A solution can hence be represented, uniquely, by a permutation $S =$
 11 $(\sigma_1, \dots, \sigma_j, \dots, \sigma_n)$, where $\sigma_j \in J$ indicates the job in the j^{th} position. The
 12 completion time $C_{\sigma_j,i}$ of job σ_j on machine i can be calculated using the
 13 following formulae:

$$C_{\sigma_1,1} = p_{\sigma_1,1} \tag{1}$$

$$C_{\sigma_1,i} = C_{\sigma_1,i-1} + p_{\sigma_1,i}, \text{ where } i = 2, \dots, m \tag{2}$$

$$C_{\sigma_j,i} = \max(C_{\sigma_j,i-1}, C_{\sigma_{j-1},i}) + p_{\sigma_j,i},$$

$$\text{where } i = 2, \dots, m, \text{ and } j = 2, \dots, n \tag{3}$$

$$C_{max} = C_{\sigma_n,m} \tag{4}$$

14 *2.2. The Capacitated Vehicle Routing Problem*

15 The Capacitated Vehicle Routing Problem (Dantzig and Ramser, 1959)
 16 can be defined in the following graph theoretic notation. Let $G(V, E)$ be an
 17 undirected complete graph where $V = \{v_0, v_1, v_2, \dots, v_n\}$ is the vertex set and
 18 where vertices E is a set of edges.

19 Let the set v_i where $i = \{1, \dots, n\}$ represent the customers who are ex-
 20 pecting to be serviced with deliveries and let v_0 be the service depot. Also
 21 associated with each vertex v_j is a non-negative demand d_j . This value is
 22 given each time a delivery is made. For the depot v_0 there is a zero demand
 23 d_0 .

24 The set E represents the set of roads that connect the customers to each
 25 other and the depot. Thus each edge $e \in E$ is defined as a pair of vertices
 26 (v_i, v_j) . Associated with each edge is a cost $c_{i,j}$ of the route between the two
 27 vertices.

28 Finally there is also a set of unlimited trucks each with same loading
 29 capacity. The aim is to service all the customers visiting them once only and
 30 using as few trucks as possible. In any potential delivery round a customer's

1 demand has to be taken into account. The total demands of customers on
2 the round must not exceed the capacity of the vehicle. This means that it is
3 normally not possible to visit all customers with one truck. As a consequence
4 each delivery round for a truck is called a *route*.

5 The goal of the CVRP problem is to minimise the overall travelling dis-
6 tance to service all customers with varying demand using a given number of
7 trucks, each with the same fixed capacity.

8 This problem was proved to be NP-Hard by Garey and Johnson (1979).

9 *2.3. Benchmark instances*

10 We used the following benchmark instances for testing the experiments
11 described in Section 5. For PFSP, we selected 12 benchmark problems
12 from Taillard (1993). Each Taillard PFSP benchmark instance is labelled
13 as *taiX_j_m*, where X is the instance number and (j, m) , where j indicates
14 the number of jobs, and m the number of machines. In order to facilitate our
15 analysis, we selected 12 of the harder instances two from the (50, 20) pool,
16 two from the (100, 20) pool and then three from the (200, 10) and (200, 20)
17 pools and finally three from the (500, 20) pool of instances for which an op-
18 timal solution is not known. For CVRP, we tested 12 problems from the
19 benchmarks of Augerat et al. (1995). Each instance of this benchmark is
20 denoted as $A - nM - kL$, where M and L indicate the number of delivery
21 points including the depot and the target number of routes, respectively.

22 **3. AGENT-BASED FRAMEWORK**

23 *3.1. Framework architecture and operation*

24 We describe a general agent-based distributed framework where each
25 agent implements a different metaheuristic/local search combination. An
26 agent continuously adapts itself during the search process using a cooper-
27 ation protocol based on the retention partial solutions deemed as possible
28 constituents of future good solutions. These are shared with the other agents.

29 The framework makes use of two types of agent: *launcher* and *meta-*
30 *heuristic* agents.

- 31 • The launcher agent is responsible for queueing the problem instances to
32 be solved for a given domain, configuring the metaheuristic agents, suc-
33 cessively passing a given problem instance to the metaheuristic agents
34 and gathering the solutions from the metaheuristic agents. To achieve

1 this it converts domain specific problem instances into the agent mes-
2 saging protocol using an ontology for scheduling and routing (see Sec-
3 tion 3.2). However the launcher agent plays no actual part in the search,
4 its job is to prepare and schedule problems to be solved by the other
5 agents.

- 6 • A metaheuristic agent executes one of the metaheuristic/local search
7 heuristic combinations that are available. These combinations and their
8 parameter settings are all defined on launching. In this way each agent
9 is able to conduct searches using different combinations and parameter
10 settings from the other agents employed in the search. Each meta-
11 heuristic agent conducts its search using the messaging structure de-
12 fined in the ontology for scheduling and routing and uses no problem
13 specific data and as such is generic.

14 A search proceeds with the launcher reading a number of problem in-
15 stances into memory. It converts them into objects that can be defined by
16 the Ontology for scheduling and routing (section 3.2 below) and then sends
17 each object, one at a time, to the metaheuristic agents to be addressed. For
18 a given problem instance, the metaheuristic agents participate in a commu-
19 nication protocol which is in effect a distributed metaheuristic that enables
20 them to search collectively for good quality solutions. This is a sequence of
21 messages passed between the metaheuristic agents and each message is sent
22 as a consequence of internal processing conducted by each agent. One itera-
23 tion of this protocol is called a *conversation* and is based upon the well-known
24 contract net protocol (FIPA, 2009). In order to arrive at a good solution the
25 agents will conduct 10 such conversations.

26 To understand the pattern matching protocol it is necessary to explain the
27 proposed model for scheduling and routing used throughout the framework.

28 3.2. Scheduling and routing ontology

29 The ontology (Gruber, 1993) plays an important role within our frame-
30 work. It defines a set of general representational primitives that are used
31 to model a number of scheduling and routing problems. The communica-
32 tion protocol and the heuristics are all based on data structures developed
33 from these primitives. This means the framework is modular in that new
34 (meta)heuristics can be easily developed and then deployed on different prob-
35 lems.

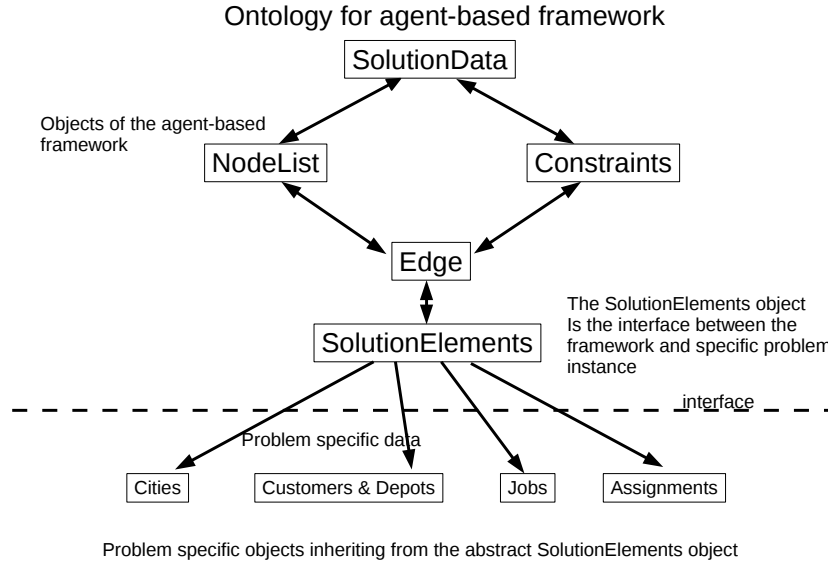


Figure 1: The combinatorial optimisation ontology

1 The ontology used by the framework generalises these notions as abstract
 2 objects.

- 3 • **SolutionElements:** A SolutionElement is an abstract object that can
 4 represent a problem specific object such as a *job* in PFSP or, a *customer*
 5 *or depot* in CVRP.
- 6 • **Edge:** An Edge object contains two SolutionElements objects. These
 7 are used to represent pairs of *jobs* or *customers* in a permutation that
 8 will be in the cooperation protocol to identify good patterns in improv-
 9 ing permutations.
- 10 • **Constraints:** The Constraints interface is between the high level
 11 framework and the concrete constraints used by a specific problem.
 12 These are used to verify a valid permutation.
- 13 • **NodeList:** A NodeList object is a list of SolutionElements objects or
 14 Edges. It represents a *schedule of jobs* in the PFSP. In the case of
 15 CVRP, a NodeList represents a *Route* and is therefore a sub-list of a
 16 full permutation.

1 • **SolutionData:** A SolutionData object is a list of NodeList objects
2 and therefore is the permutation that is optimised by the framework.
3 In this study it represents a *schedule of jobs* in PFSP, or a *collection of*
4 *routes in CVRP*.

5 All message passing in the framework, including the whole ontology, is
6 written in XML. This can be advantageous as many benchmark problems,
7 these days, are also in XML making the interface between problem definition
8 and ontology seamless in practice. Figure 1 shows the structure of the ontol-
9 ogy and how SolutionElements are the interface between the framework and
10 a concrete problem.

11 3.3. Edge selection and short-term memory

12 The framework features a method of Edge selection and short-term mem-
13 ory. A conversation, as has been explained already, is a type of distributed
14 heuristic. Its purpose is to identify constituent features of incumbent solu-
15 tions that are likely to lead to the building of improving solutions.

16 This is achieved by using objects defined in the ontology. SolutionData
17 object in the ontology is built from the sub-objects of NodeLists and Edges
18 and SolutionElements. Thus to represent a permutation of n jobs for PFSP
19 a SolutionData object is built from one NodeList object and which itself is
20 made up $n - 1$ Edges objects which are themselves built from n SolutionEle-
21 ments. Similarly a CVRP representation of n customers is one Solution Data
22 object with x (this number is determined during the search) NodeLists. The
23 NodeLists are built of $n - 1$ Edges and n SolutionElements.

 If we take a permutation of the unique ID numbers of each the Solu-
 tionElements objects we can represent a SolutionData object with 10 ele-
 ments as follows: (3, 4, 6, 7, 5, 8, 9, 0, 1, 2). Furthermore we can break this
 permutation into a collection of Edge objects:

$$(3, 4), (4, 6), (6, 7), (7, 5), (5, 8), (8, 9), (9, 0), (0, 1), (1, 2), (2, 3)$$

24 During a conversation each agent runs its metaheuristic and produces a
25 new incumbent solution. Each agent then breaks this solution into Edge
26 objects and sends them to one of the metaheuristic agents that has been
27 designated as the “initiator” for the duration of that conversation only. All
28 metaheuristic agents are exactly the same and have the potential to take on
29 the role of an initiator in a conversation.

1 The initiator agent collects all the Edge objects from all the other agents
2 into a list and scores them by frequency. Here, frequency is the number of
3 times an Edge appears in the initiators list. The only Edge objects that are
4 retained are the ones that have the same score as the number of agents that
5 are participating in the conversation. The idea here is that if an Edge occurs
6 frequently in all incumbent solutions, it is likely to be an Edge that will be
7 part of an improving solution. These retained good Edges are then shared
8 by the initiator with the other agents.

9 Another feature is the learning mechanism where each agent keeps a short-
10 term memory of good Edges. This is a queue of good Edges that operates
11 somewhat like a Tabu list. An agent's queue is populated during the first
12 conversation with edges from the incumbent solution produced by its meta-
13 heuristic. Thereafter the queue is maintained at a factor, that is 20%, of
14 the size of the candidate solution for the problem instance at hand. In sub-
15 sequent conversations as new edges not already in the list arrive, they are
16 pushed onto the front of the queue while other edges are popped off the back
17 of the queue so that the size of the list does not change.

18 The Edges in the short-term memory are used at the start of each con-
19 versation to modify the performance of agent's metaheuristic to enable it to
20 find better solutions.

21 The basic idea of this learning mechanism is that both the RandNEH and
22 RandCWS heuristics of (Juan et al., 2015) used in this study make use of
23 ordered lists to construct new solutions. These heuristics use biased random
24 functions to choose items from these lists. We use the Edges identified by
25 the learning mechanism to reorder these lists and so influence the way new
26 solutions are constructed.

27 **4. IMPLEMENTATION**

28 The framework is implemented using JADE (Bellifemine et al., 2007). It
29 allows a developer to concentrate on the function and behaviour of agents
30 while it handles inter-agent and inter-platform communication and hus-
31 bandry.

32 The configuration file of a launcher agent lists which problems are to be
33 solved. It also contains how many conversations the metaheuristic agents are
34 going to conduct for a particular problem.

35 At start-up, parameters determine which metaheuristic will be employed
36 as well as any parameter settings associated with it. Once the metaheuristic

1 agents have completed the set number of conversations they each send their
2 best result to the launcher agent. The launcher then prints an output file
3 with the best solution and objective function value.

4 The framework conducts a search where each agent is launched and reg-
5 isters with the JADE platform that hosts the framework. Once this is com-
6 plete, the agents wait for the launcher agent to read in a problem from file.
7 The launcher will then send the problem to each of the metaheuristic agents.
8 Only when the metaheuristic agents receive that problem from the launcher
9 do they embark on a search.

10 4.1. Heuristics used by the agents

11 In this study depending on whether they are solving PFSP or VRP, the
12 agents instantiate the heuristics developed by Juan et al. (2010b,a) respec-
13 tively.

14 In the case of PFSP, the metaheuristic used is the Randomised NEH
15 (RandNEH) algorithm of Juan et al. (2010b). It is a stochastic version of the
16 classic heuristic of Nawaz et al. (1983). Just as the NEH algorithm creates an
17 ordered list of *jobs* sorted from tardiest to quickest, the RandNEH algorithm,
18 instead of choosing jobs in order from the list, chooses them according to a
19 randomised process based on the Triangular probability distribution.

20 While for the CVRP, the metaheuristic used is the Randomised Clarke
21 Wright Savings (RandCWS) algorithm of Juan et al. (2010a). It is a stochas-
22 tic version of the classic savings heuristic of Clarke and Wright (1964). Rather
23 than generating new routes by choosing the greatest relevant saving from the
24 *savings list*, it chooses according to a Geometric distribution where the j^{th}
25 *savings* from the list is chosen by a probabilistic function described in Juan
26 et al. (2010a).

27 Both these algorithms have been integrated into our system according to
28 our framework. This was quite a simple process where the heuristics imple-
29 ment the abstract objects defined in the scheduling and routing ontology.
30 For example, the Edge and Job objects of the RandNEH algorithm are now
31 subclasses of the Edge and SolutionElements abstract classes of the frame-
32 work. Similarly for VRP problems, where the Route, Edge and Customer
33 objects are now subclasses of the NodeElements, Edge and SolutionElements
34 objects of the framework.

35 This means we can use the *good Edges* found as a result of a *conversation*
36 of the framework to modify the Job lists and Saving lists of the RandNEH
37 and RandCWS algorithms respectively.

1 In the case of PFSP, the list of Edges found by the agents is turned
2 into a list of SolutionElements (Jobs) where their order in the Edge list
3 is preserved. The Jobs list generated by the RandNEH algorithm is then
4 reordered with respect to the list of Jobs generated from the Edge list, with
5 the new Jobs being moved to the front of the list. This affects the operation
6 of the RandNEH algorithm where the new Jobs are likely be favoured in the
7 construction of any new improving schedule.

8 It is a similar process for the RandCWS algorithm. However this time
9 the Edges in Edge list are also Super Classes of the Edges in the savings
10 Savings List. Again the Savings List is reordered with respect to the Edge
11 list where these Edges are moved to the head of the Savings List. This again
12 affects the operation of the RandCWS algorithm favouring the *good* Edges
13 found as a result of the Agents' conversations.

14 4.2. Description of a conversation

15 Figure 2 shows the edge selection protocol used by the metaheuristic
16 agents. One complete execution of the algorithm illustrated is a *conversation*.
17 In any conversation, there will be an agent that takes on the role of an
18 initiator and the others are responders. In the very first conversation agent1
19 will always take on the role of initiator. Thereafter, any agent can be the
20 initiator, but it is determined in the previous conversation which agent will
21 be the initiator for the current conversation (see below).

22 In Figure 2 an agent taking on the role of initiator starts a conversation.
23 At the start of a conversation, each agent either takes a list of Edge objects
24 generated from a previous conversation or from one generated by the launch
25 agent (see I_1 and R_1 in Figure 2).

26 The agents then find a new incumbent solutions using their given heuristics
27 in conjunction with the edges provided in the previous step (see I_2 and
28 R_2 in Figure 2).

29 The initiator breaks its incumbent solution into edges and then invites
30 the responder agents to do the same and send them to the initiator, I_3 and
31 R_3 of Figure 2.

32 The receiving agents also send the value of their best-so-far solution. This
33 will be used by the initiator to determine which agent will be the new initiator
34 in the next conversation (see I_4 in Figure 2).

35 In I_4 , the initiator receives the Edge objects from the responding agents
36 and collects them together. Each Edge object is scored and ranked based on
37 frequency. This can be seen in box I_4 of Figure 2 as the function *getScore*.

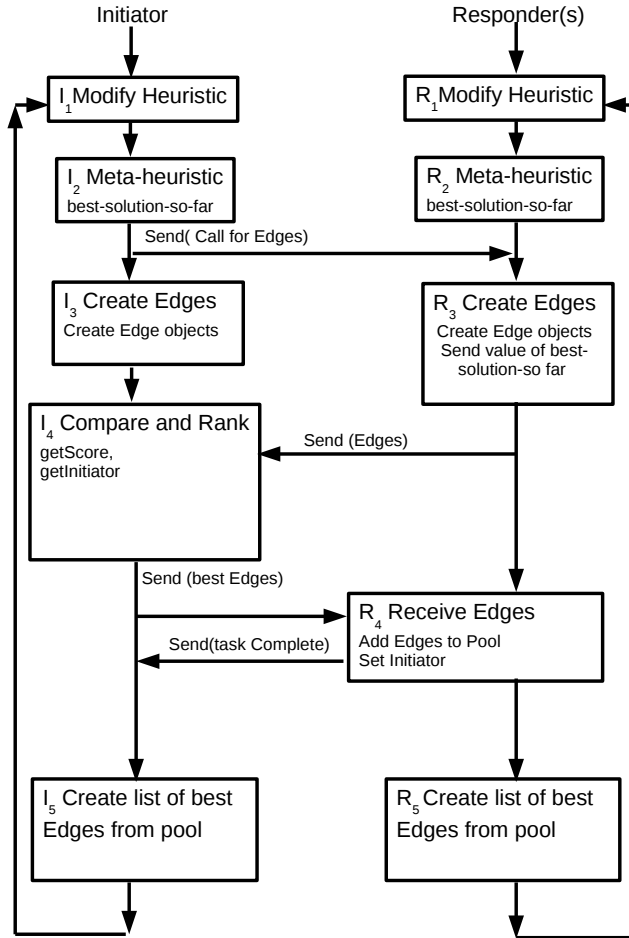


Figure 2: The Cooperation Protocol showing one iteration of a conversation

1 In I_4 of Figure 2, through the function *getInitiator*, the initiator also de-
 2 termines which metaheuristic agent is going to be the initiator in the next
 3 conversation. This is achieved by choosing the agent the best objective func-
 4 tion value to be the initiator.

5 The initiator then sends good *Edge* objects, found during this conversa-
 6 tion, to the receiving metaheuristic agents.

7 Each agent keeps a pool or short-term memory of high scoring *Edge*
 8 objects. The pool acts as a sort of queue and its length is set when the agent

1 is launched. In this study all the agents have a pool size of 20% of length
2 of the instance currently being optimised. During the first conversation each
3 agent populates its pool as good edges are identified. Once the pool is up to
4 size, it is maintained as a queue as described in Section 3.3.

5 The other metaheuristic agents receive the lists of Edge objects from the
6 initiator (see box R_4 in Figure 2). They also update their internal memory's
7 or pools as described above. In box I_5 and R_5 of Figure 2, both initiator
8 and responder metaheuristic agents then each create a new solution by using
9 edges from their updated internal pools. These good edges are passed to the
10 metaheuristic the agent is configured to execute in the current search. The
11 metaheuristic uses these good edges when it is next called at the start of the
12 next conversation (back to I_1 and R_1 of Figure 2). This process repeats and
13 continues until the number of conversations set from the launcher agent are
14 completed.

15 5. EXPERIMENTAL DESIGN

16 In this section we discuss the experimental design.

17 5.1. Launcher agent

18 One launcher agent is invoked in each run. The launcher agent reads
19 from a configuration file the number of agents to be instantiated (see Section
20 5.4) as well as the number of conversations that will be conducted during the
21 test.

22 The launcher agent executes a construction heuristic to build an initial
23 solution for each instance and run: for PFSP a biased-randomised version
24 of the NEH algorithm (Nawaz et al., 1983) with Taillard's speedups imple-
25 mented by (Juan et al., 2010b); and for CVRP, the Randomised CW Savings
26 algorithm (Juan et al., 2014, 2010a). This initial solution is passed on to
27 each of the individual agents.

28 5.2. The number of conversations

29 Juan et al. (Juan et al., 2010b, 2014, 2010a) suggested that, to be effec-
30 tive, the RandNEH and RandCWS heuristics should be run for a maximum
31 time of about 2.5 minutes. We benchmarked their code and observed the
32 same phenomena, hence used the same running time on our machine during
33 our experiments.




Figure 3: Typical Solution trajectories of the RandNEH and RandCWS algorithms

1 This gave us a guide as to how long our system should be run and there-
 2 fore determine the number of conversations that would be needed. The time
 3 taken for the agents to complete a conversation is mainly governed by the
 4 time taken for an agent’s given heuristic to execute. To this end, we con-
 5 ducted tests showing that both heuristics typically have a period of maximum
 6 improvement of about 12 seconds. As an example, Figure 3 plots the solution
 7 trajectories of the PFSP instance tai051 and the CVRP instance A-n45-k9
 8 against time. We can see that these algorithms have their period of greatest
 9 improvement in about the first 12 seconds of operation. Thus we determined
 10 that the system should execute 10 conversations for our system to run for
 11 about the same time as the standalone versions of the RandNEH and Rand-
 12 CWS heuristics. This would also take into account any lag caused by the
 13 asynchronous nature of the system.

14 5.3. Parameter Settings

15 Since the RandCWS and RandNEH methods of Juan et al. were already
 16 written in JAVA, they were integrated with minimum effort as a module
 17 of our agent based system. They utilise the edge selection heuristic of the
 18 agent-based system by taking edges identified during each conversation and
 19 re-ordering the jobs list of the RandNEH algorithms and the savings list of
 20 the RandCWS algorithm as explained in Section 4.1.

21 Both algorithms use a random seed which is a number which introduces
 22 a bias to a random number generator. In the tests for both the PFSP and
 23 CVRP, each agent is configured with exactly the same random seeds (Juan
 24 et al., 2010a,b).

1 However in their article Juan et al. (2011) describe how they combined
2 Monte-Carlo simulation techniques with the Clarke Wright Savings algorithm
3 to develop the probabilistic RandCWS algorithm. It was designed so that it
4 would require little parameter tuning. To this end they describe a parameter
5 α that is used to define different geometric distributions. Such a distribution
6 can then used by the RandCWS heuristic to choose the next edge from the
7 Clarke Wright Savings list as part of its solution building process. The α
8 -parameter is itself chosen at random from a uniform distribution between
9 two values (a, b) where $0 < a \leq b < 1$. In their paper, Juan et al choose
10 α -values from the interval $\alpha \in \{0.05 - 0.25\}$. They show that for any α -value
11 in this interval, the algorithm will give similar and good performance. In
12 correspondence with the authors, it was confirmed that the algorithm will
13 perform less well for α -values of above 2.3, while at the other end of the range
14 α -values close to the 0.05 will perform as any in the cited interval.

15 The intuitive idea for spreading the $\alpha - values$ is to maximise the use of
16 different distributions during a search. While these choices do not effect the
17 solution quality it means the agents will produce slightly different solutions
18 which will produce different edges that will enhance the performance of the
19 distributed edge selection algorithm.

20 In both case studies each metaheuristic is allowed to run for 12 seconds
21 each time it is called.

22 Following Juan et al. (2014, 2013) in what we call our standalone ex-
23 periments, that is the traditional case without cooperative search being
24 used, we compare our cooperating agents with the stand alone by run-
25 ning the experiments for each group for a maximum time of 40 minutes
26 to match the computational effort of the system running 16 agents i.e.
27 $16 \times 150s = 2400s$ (40mins). Thus all agents versus standalone comparisons
28 are made against this worst case scenario.

29 *5.4. Experimental set-up*

30 The main hypothesis to be tested in these experiments is that cooperating
31 agents produce better results than their stand alone equivalents. The results
32 are also compared with state-of-art results for each of these benchmarks. To
33 this end, for each instance of the tests the following scenarios were run:

34 The CVRP tests were conducted as follows with $\alpha - values$ selected on
35 0.01 increments from the set $\{0.03 \text{ to } 0.18\}$

- 36 • Stand alone agent: 1 metaheuristic agent where the $\alpha - value = 0.03$

- 1 • 4 agents: $\alpha \in \{0.03 - 0.06\}$
- 2 • 8 agents: $\alpha \in \{0.03 - 0.1\}$
- 3 • 12 agents: $\alpha \in \{0.03 - 0.14\}$
- 4 • 16 agents: $\alpha \in \{0.03 - 0.18\}$

5 The PFSP tests were conducted similarly but without the need for $\alpha -$
6 *values*.

7 They are tested in this way so that standalone agents running just one
8 metaheuristic at a time can be compared statistically with groups of coop-
9 erating agents in order to test the main hypothesis.

10 Every instance is tested 20 times. The resulting values are then used to
11 evaluate the performance of the test. In particular the average and minimum
12 value of the 20 runs for each problem are taken. These are compared with
13 the known optimal or best values for each problem instance.

14 To test the hypothesis that agents cooperating by edge selection perform
15 better than stand alone agents, Wilcoxon signed rank tests are conducted
16 for each benchmark instance, with a 95% confidence level. We used the
17 Wilcoxon test rather than t-test because we cannot guarantee that the test
18 results will be normally distributed Moore and McCabe (1989). These tests
19 compare the difference between the distributions of 16, 12, 8, and 4 agents
20 cooperating with the stand alone agents. A secondary hypothesis is explored
21 where the performances of groups of 4, 8, 12 and 16 agents are compared using
22 the Wilcoxon signed rank test to ascertain whether increasing the number
23 of agents results in better performance. The following notation is used in
24 tables 2, 3, 6 and 7. Given two algorithms (or different settings for the same
25 algorithm); A versus B, $>$ ($<$) denotes that A (B) is better than B (A) and
26 this performance difference is statistically significant at a 95% confidence
27 level. However, \geq (\leq) denotes that A (B) is better than B (A) although
28 statistical significance could not be supported. Lastly, \approx denotes the case
29 where both approaches consistently achieve the same value.

30 The results for each problem are averaged and the average percentage
31 deviation from the known optimum is calculated. The percentage deviation
32 from a known optimum is calculated in the standard manner:

$$\frac{Method_{solution} - Best_{solution}}{Best_{solution}} \times 100 \tag{5}$$

1 The results are also analysed to find the best result of each group of agents
 2 over the 20 runs of each problem instance.

3 Juan et al. (2014, 2013)

4 5.5. Machines

5 All tests are run on the same Linux cluster using 8 identical machines;
 6 two agents were run per-node of the cluster. The agents are configured to
 7 use 2 GB of memory.

8 6. RESULTS OF EXPERIMENTS

9 6.1. Permutation Flow-shop Scheduling results

10 Table 1 shows the average percentage deviation from the best known or
 11 optimum value for each of the benchmark instances tested, as well as the
 12 percentage deviation for the best value found across the 20 runs. The table
 13 also compares our results with the Hybrid Genetic algorithm of Zobolas et al.
 14 (2009). Here the average value reported by Zobolas et al. (2009) is given as a
 15 percentage deviation from the best known solution (BKS). Despite the fact
 16 that this is a type of hyper-heuristic system where the only parameter tuning
 17 is the number of conversations executed, the PFSP results are competitive
 18 with the state-of-the-art results for these problem instances. It is only in the
 19 larger three instances where our average deviation is not better than that of
 20 Zobolas et al. (2009).

Table 1: The average (avr.) and best percentage deviation from the upper bound over 20 runs for each instance for PFSP. The best values are highlighted in bold

Instance	BKS	Zobolas et al.	1 Agent		4 Agents		8 Agents		12 Agents		16 Agents	
			avr.	best	avr.	best	avr.	best	avr.	best	avr.	best
tai051_50_20	3850	0.77%	0.92%	0.39%	0.84%	0.55%	0.76%	0.47%	0.69%	0.39%	0.63%	0.44%
tai055_50_20	3610	1.03%	0.54%	0.44%	0.67%	0.50%	0.62%	0.28%	0.57%	0.36%	0.50%	0.30%
tai081_100_20	6202	1.63%	1.55%	1.23%	1.52%	1.26%	1.41%	1.06%	1.34%	1.03%	1.30%	1.02%
tai085_100_20	6314	1.57%	1.39%	1.00%	1.34%	1.11%	1.22%	0.97%	1.15%	1.00%	1.11%	0.89%
tai091_200_10	10862	0.24%	0.12%	0.09%	0.09%	0.09%	0.09%	0.09%	0.09%	0.09%	0.09%	0.09%
tai095_200_10	10524	0.03%	0.10%	0.03%	0.09%	0.03%	0.05%	0.03%	0.03%	0.03%	0.03%	0.03%
tai101_200_20	11195	1.34%	1.49%	1.30%	1.38%	1.09%	1.25%	1.01%	1.22%	1.06%	1.19%	0.93%
tai105_200_20	11259	1.04%	1.08%	0.70%	1.02%	0.89%	0.94%	0.78%	0.94%	0.83%	0.88%	0.71%
tai106_200_20	11176	1.11%	1.60%	1.25%	1.55%	1.35%	1.44%	1.27%	1.43%	1.25%	1.42%	1.33%
tai111_500_20	26059	0.73%	0.99%	0.74%	1.01%	0.88%	0.95%	0.86%	0.92%	0.87%	0.88%	0.69%
tai115_500_20	26334	0.82%	0.99%	0.74%	1.01%	0.88%	0.95%	0.86%	0.92%	0.87%	0.88%	0.69%
tai116_500_20	26477	0.49%	0.72%	0.56%	0.69%	0.56%	0.67%	0.60%	0.62%	0.57%	0.61%	0.54%

21 With respect to answering our main hypothesis: “is cooperation by pat-
 22 tern matching better than no cooperation?”, we compared 4 agents cooper-
 23 ating against a stand alone agent (see Section 5.3). In addition we wanted

1 to test if increasing the number of agents produced a statistically significant
 2 improvement in the results. Tables 2 and 3 list these results; in each case we
 3 tested for statistical significance.

4 In table 2, with the exception of the *tai055_50_20* instance, it can be seen
 5 that groups of 8,12 and 16 agents perform better than the stand alone with
 6 statistical significance. However, for the *tai055_50_20* instance, 16 agents
 7 show some improvement, if not statistically, over the stand alone. Further-
 8 more, two instances of 4 agents perform statistically better than the stan-
 9 dalone but the rest all show some improvement but not at the 95% level.

Table 2: Table showing cooperating agents performing better than the standalone equivalent at the 95% in PFSP

Instance	4 vs 1	8 vs 1	12 vs 1	16 vs 1
<i>tai051_50_20</i>	\geq	$>$	$>$	$>$
<i>tai055_50_20</i>	\leq	\leq	\leq	\geq
<i>tai081_100_20</i>	\geq	$>$	$>$	$>$
<i>tai085_100_20</i>	\geq	$>$	$>$	$>$
<i>tai091_200_10</i>	$>$	$>$	$>$	$>$
<i>tai095_200_10</i>	\geq	$>$	$>$	$>$
<i>tai101_200_20</i>	$>$	$>$	$>$	$>$
<i>tai105_200_20</i>	\geq	$>$	$>$	$>$
<i>tai106_200_20</i>	\geq	$>$	$>$	$>$
<i>tai111_500_20</i>	\leq	\geq	$>$	$>$
<i>tai115_500_20</i>	\leq	$>$	$>$	$>$
<i>tai116_500_20</i>	\geq	$>$	$>$	$>$

10 Table 3 explores the possibility that adding more agents leads to better
 11 results. Here we can see that 8 agents perform statistically better than 4,
 12 while 12 agents show some improvement, but not statistically, over 8. The
 13 same is true for 16 over 12 agents. However the instances *tai091_200_10* and
 14 *tai105_200_20* achieve statistical significance as well. By the time we get to
 15 16 versus 4 agents, 16 agents always perform statistically better except for
 16 *tai091_200_10* where statistical significance is not reached. It should also be
 17 noted for *tai091_200_10* while the cooperating agents perform better than
 18 the stand alone, thereafter the all achieve the same value. It is clear that
 19 progressively increasing the number of agents from 4 to 8 to 12 to 16 results
 20 in an increase in performance. However this improvement is not always
 21 statistically significant. If we consider the column of the table where 16

1 agents are compared with 8 we see that the level of improvement gains more
 2 significance. This is suggestive that it is better to increase the number of
 3 agents by a factor of 2.

Table 3: Table showing different groups cooperating agents perform at the 95% confidence level in PFSP

Instance	8 vs 4	12 vs 8	16 vs 12	16 to 8	16 to 4
tai051_50_20	>	≥	≥	>	>
tai055_50_20	>	≥	>	>	>
tai081_100_20	>	≥	≥	>	>
tai085_100_20	>	≥	≥	>	>
tai091_200_10	≈	≈	≈	≈	≥
tai095_200_10	>	≥	≥	≈	>
tai101_200_20	>	≥	≥	≥	>
tai105_200_20	>	≤	>	>	>
tai106_200_20	>	≥	≥	≥	>
tai111_500_20	>	>	≥	>	>
tai115_500_20	>	≥	≥	≥	>
tai116_500_20	>	>	≥	>	>

4 The cooperation mechanism used in this study works by identifying and
 5 sharing of good patterns that form partial solutions to the problem at hand.
 6 These are then passed to a metaheuristic to build a new putative solution
 7 to the problem. Given this, it is interesting to study the patterns (edges)
 8 identified by each agent and compare them to the final solution found by the
 9 system. To this end, the final permutation (Edges which appear in the final
 10 solution and are identified during the search (see table 4)are highlighted in
 11 bold) <**12**, 37, 20, 31, **39,35**, 34, 6, 40, 5, 10, 1, 7, 15, 33, 43, 24, 42, 27,
 12 29, 46, 47, 36, 23, 14, 2, 44, 8, 45, 17, **13**, **22**, 21, 48, 18, 28, 16, 49, 38, 19,
 13 26, 41, 11, 32, 25, 9, 30, 4, **50,3**> of jobs found by the system during one
 14 run of the tai051_50_20 instance is compared with the patterns in table 4.
 15 These are all the unique edges identified during this search. These edges are
 16 identified multiple times but the table only shows them once.

17 Indeed some edges (highlighted in bold) identified by the system do end
 18 up in the final job permutation. Furthermore, we can identify linked edges
 19 such as 50, 3, 12 at the end and beginning of the permutation. However
 20 these are not as many as seen with CVRP results below because of the way
 21 the makespan 4 is calculated as a special cumulative sum of columns of jobs.

Table 4: Patterns found by 4 cooperating agents PFSP for problem tai051_50.20.

Agents	Edges										
agent1	(14,15)	(4,25)	(32,22)	(39,16)	(25,50)	(19,41)	(13,32)	(44,45)	(45,6)	(50,3)	(28,38)
agent2	(35,34)	(9,30)	(5,10)	(2,44)	(12,37)	(1,7)	(4,50)	(10,1)	(24,42)	(50,3)	
agent3	(3,12)	(37,39)	(30,46)	(50,3)	(35,15)	(41,7)	(34,33)	(38,24)	(47,23)	(42,49)	
agent4	(40,21)	(22,13)	(6,42)	(33,40)	(26,2)	(5,14)	(7,18)	(37,28)	(39,35)	(44,11)	

1 *6.2. Capacitated Vehicle Routing results*

2 Table 5 compares the percentage deviation for average and best results for
3 the different groups of agents from the best known solution. The table also
4 compares our percentage deviations for these problem instances with those of
5 Altinel and Öncan (2005) (donated by A) and Juan et al. (2010b) (denoted
6 by B). However, Juan et al. (2010b) only has results for a selection of the
7 instances we tested. They represent the latest work on these benchmark
8 instances so we have included them for comparison. Comparing our results
9 with those of Altinel and Öncan (2005) and Juan et al. (2010b) we can
10 see that agents improve on their results. Furthermore, to the best of our
11 knowledge, in four cases we have found results that are better than the
12 current best known solutions. *A – n39 – k6, A – n45 – k7, A – n55 – k9* and
13 *A – n63 – k9* are highlighted in italics for the best average value and in best
14 for our best overall score.

Table 5: The average (avr.) and best percentage deviation from the optimum/upper bound over 20 runs for each instance for CVRP

Instance	BKS	A and B	Juan et al.	1 Agent		4 Agents		8 Agents		12 Agents		16 Agents	
				avr.	best	avr.	best	avr.	best	avr.	best	avr.	best
A-n38-k5	734.18	3.577%	0.54%	0.07%	0.04%	0.09%	0.04%	0.02%	-0.03%	-0.02%	-0.03%	<i>-0.03%</i>	-0.03%
A-n39-k6	833.14	2.233%	-	0.01%	0.01%	0.01%	0.01%	0.01%	0.01%	0.01%	0.01%	0.01%	0.01%
A-n44-k6	939.33	2.394%	-	0.63%	0.57%	0.70%	0.57%	0.55%	0.39%	0.40%	0.29%	0.32%	-0.12%
A-n45-k6	944.88	1.383%	-	0.92%	0.92%	0.92%	0.92%	0.69%	0.00%	0.20%	0.00%	0.00%	0.00%
A-n45-k7	1147.28	1.842%	0.07%	0.05%	-0.03%	0.07%	-0.02%	0.03%	-0.03%	0.03%	-0.03%	<i>-0.01%</i>	-0.48%
A-n55-k9	1074.46	2.378%	0.14%	0.13%	0.05%	0.26%	0.05%	0.06%	0.05%	0.05%	0.05%	0.05%	0.05%
A-n60-k9	1355.80	1.64%	0.13%	0.50%	0.50%	0.50%	0.50%	0.46%	0.22%	0.40%	0.22%	0.37%	0.22%
A-n61-k9	1039.08	1.654%	0.49%	0.27%	0.26%	0.26%	0.26%	0.25%	0.13%	0.23%	0.12%	0.22%	0.12%
A-n62-k8	1294.28	4.648%	-	0.70%	0.62%	0.76%	0.62%	0.62%	0.62%	0.65%	0.62%	0.62%	0.62%
A-n63-k9	1619.90	2.051%	-	0.75%	0.45%	0.88%	0.69%	0.73%	0.40%	0.53%	0.14%	0.32%	0.14%
A-n65-k9	1181.69	2.392%	0.66%	1.06%	1.05%	1.05%	0.72%	0.92%	0.28%	0.82%	0.64%	0.61%	0.14%
A-n80-k10	1766.50	2.952%	0.2%	1.04%	0.99%	1.04%	0.99%	0.98%	0.77%	0.87%	0.77%	0.85%	0.70%

15 Again we tested for the main hypothesis. We compared groups of 4, 8,
16 12, and 16 agents cooperating against a stand alone agent. As before we
17 tested for statistical significance using the Wilcoxon signed rank test at the
18 95% confidence level. Table 6 lists these results using the same notation as
19 used in table 2 above. As with the PFSP, 4 agents cooperating do not show

1 any improvement from their stand alone equivalent. However, groups of 8, 12
 2 and 16 agents with increasing certainty perform better than the stand alone
 3 agent. Indeed 16 agents all perform better a 95% confidence level except for
 4 the $A - n39 - k6$ instance.

Table 6: Table showing cooperating agents performing better than the standalone equivalent at the 95% in CVRP

Instance	4 vs 1	8 vs 1	12 vs 1	16 vs 1
A-n38-k5	\leq	$>$	$>$	$>$
A-n39-k6	\leq	\geq	\geq	\geq
A-n44-k7	\leq	$>$	$>$	$>$
A-n45-k6	\geq	$>$	$>$	$>$
A-n45-k7	\leq	\geq	\geq	$>$
A-n55-k9	\leq	$>$	$>$	$>$
A-n60-k9	\leq	$>$	$>$	$>$
A-n61-k9	\geq	$>$	$>$	$>$
A-n62-k8	\leq	$>$	\geq	$>$
A-n63-k9	\leq	\geq	$>$	$>$
A-n65-k9	\geq	$>$	$>$	$>$
A-n80-k10	\leq	$>$	$>$	$>$

5 In table 7 we report the results of our tests for the secondary hypothesis.
 6 As with the PFSP results, we can see a gradual improvement as more agents
 7 are added. But again it seems it is necessary to double the number of agents
 8 each time in order to observe improvement in results. The addition of 4
 9 agents each time results in an improvement that is not always statistically
 10 significant. However if the agents are doubled each time in groups of 4, 8
 11 and 16 there is a greater proportion of statistically significant improvement
 12 from the additive case.

13 Finally, we show the patterns generated for a sample on problem instance
 14 $A - n38 - k5$ in table 9 and compare them to the final result of this run in
 15 table 8. We highlight in bold those edges identified by the agents in table
 16 9 that end up in the final solution in table 8. As can be seen there are
 17 many more such edges than for the PFSP. This is because the relationship
 18 between edges and cities is much more direct in the case of CVRP as costs
 19 are calculated as 2D-euclidean distances between cities.

20 From this study we conclude that with no parameter tuning between
 21 case studies our system can produce results which are commensurate with

Table 7: Table showing different groups cooperating agents perform at the 95% confidence level in CVRP

Instance	8 vs 4	12 vs 8	16 vs 12	16 vs 8	16 vs 4
A-n38-k5	>	>	>	>	>
A-n39-k6	≥	≥	≥	≥	≥
A-n44-k7	>	>	>	>	>
A-n45-k6	>	>	>	>	>
A-n45-k7	>	≥	≥	≥	>
A-n55-k9	>	≥	≥	≥	>
A-n60-k9	>	≥	≥	>	>
A-n61-k9	>	≥	≥	≥	>
A-n62-k8	>	≤	≥	≥	>
A-n63-k9	>	>	>	>	>
A-n65-k9	>	≥	>	>	>
A-n80-k10	>	>	≥	>	>

Table 8: Final Solution to CVRP problem A-n38-k5.

Route Name	Routes
Route1	[1, 8, 6, 12, 28, 23, 33, 1]
Route2	[1, 27, 13, 4, 2, 5, 17, 26, 7, 30, 1]
Route3	[1, 9, 34, 36, 24, 31, 11, 22, 1]
Route4	[1, 10, 18, 37, 14, 16, 3, 15, 25, 1]
Route5	[1, 21, 38, 32, 29, 35, 20 , 19, 1]

Table 9: Patterns found by 4 cooperating agents for CVRP problem A-n38-k5.

Agents	<i>Edges</i>
agent1	(35,20) (38,32) (29,35) (20,19) (21,38) (32,29) (1,21) (19,1)
agent2	(30,31) (11,1) (1,19) (31,11) (35,30) (19,35)
agent3	(35,20) (29,19) (1,21) (20,1) (32,29) (38,32) (21,38) (19,35) (1,19)
agent4	(19,1) (32,38) (38,29) (35,20) (21,32) (20,19) (29,35)

1 the state-of-the-art studies in both fields. Furthermore, in four instances with
 2 the CVRP tests we were able to the best of our knowledge beat the current
 3 best results for these instances. We were also able to show for groups of 8, 12
 4 and 16 agents compared with the stand alone equivalent, that cooperation
 5 by pattern finding is better than no cooperation. Finally we are also able to
 6 show that doubling the number agents each time leads to improving results
 7 as shown in Figure 4.




Figure 4: Boxplots of objective values obtained in 10 runs for 16, 12, 8 and 4 agents on a selected instance from the (a) STSP, (b) PFSP, and (c) CVRP problem domains.

8 7. CONCLUSION

9 In this study we propose a general agent-based distributed framework
 10 where each agent implements a different metaheuristic/local search combi-
 11 nation. An agent continuously adapts itself during the search process using
 12 a cooperation protocol based on reinforcement learning and pattern finding.
 13 Good patterns that make up improving solutions are identified by frequency
 14 of occurrence in a conversation and shared with the other agents. The frame-
 15 work has been tested on well known benchmark problems for two tests cases
 16 PFSP and CVRP. In both cases, with no parameter tuning between domains,
 17 the platform performed at least as well as the state-of-art. For CVRP, we
 18 were able, in cases of $A - n38 - k5$, $A - n44 - k6$ and $A - n45 - k7$ to improve
 19 on the best known solutions for these instances¹.

¹<http://neo.lcc.uma.es/vrp/vrp-instances/capacitated-vrp-instances/>

1 We have also shown eight or more agents perform better than a stand
2 alone agent with a 95% confidence level. Furthermore, we have shown with a
3 reasonable level of certainty, if not always with 95% confidence, that an im-
4 provement in performance can be achieved each time you double the number,
5 up to 16, agents used.

6 The distributed computing framework presented can be run on a local
7 network of personal computers each using 2GB memory.

8 The framework also aims to be a generic and modular needing very little
9 parameter tuning across different problem types tested so far. It has been
10 been applied successfully to PFSP and CVRP. It has also been used to model
11 fairness in Nurse Rostering (Martin et al., 2013) using real-world data. This
12 flexibility is achieved by means of an ontology which enables the agents to
13 represent these problems with the same internal structure.

14 This is an interesting and little researched topic that warrants further
15 investigation such as: extending the ontology to apply the framework to
16 new problems; adding more heuristics and metaheuristics and improving the
17 pattern finding protocol.

18 Finally, this framework will be published as an open source project so that
19 other metaheuristics and cooperation protocols can be added and tested by
20 other researchers. The project is called MACS (Multi-agent Cooperative
21 Search) and will be published at the following website: [http://www.cs.
22 stir.ac.uk/~spm/](http://www.cs.stir.ac.uk/~spm/).

23 8. Acknowledgement

24 The study was funded EPSRC Dynamic Adaptive Automated Software
25 Engineering (DAASE) project EP/J017515/1.

26 9. Appendix A. Supplementary material

27 Supplementary data associated with this article can be found, in the
28 online version, at: www.tobeprovided.ac.uk

29 10. References

30 İ. K. Altinel and T. Öncan. A new enhancement of the clarke and wright
31 savings heuristic for the capacitated vehicle routing problem. *Journal of*
32 *the Operational Research Society*, 56(8):954–961, 2005.

- 1 P. Augerat, J.M. Belenguer, E Benavent, A. Corberán, D. Naddef, and G. Ri-
2 naldi. Computational results with a branch and cut code for the capaci-
3 tated vehicle routing problem. *Rapport de recherche- IMAG*, 1995.
- 4 M. Aydin and T.C. Fogarty. A distributed evolutionary simulated annealing
5 algorithm for combinatorial optimisation problems. *Journal of Heuristics*,
6 10(3):269–292, 2004a.
- 7 M. Aydin and T.C. Fogarty. Teams of autonomous agents for job-shop
8 scheduling problems: An experimental study. *Journal of Intelligent Man-*
9 *ufacturing*, 15(4):455–462, 2004b.
- 10 D. Barbucha. A cooperative population learning algorithm for vehicle routing
11 problem with time windows. *Neurocomputing*, 146:210–229, 2014.
- 12 F. L. Bellifemine, G. Caire, and D. Greenwood. *Developing multi-agent sys-*
13 *tems with JADE*. Wiley, 2007. ISBN 0470057475.
- 14 C. Blum and A. Roli. Metaheuristics in combinatorial optimization:
15 Overview and conceptual comparison. *ACM Computing Surveys (CSUR)*,
16 35(3):268–308, 2003.
- 17 E.K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, and
18 R. Qu. Hyper-heuristics: A survey of the state of the art. *Journal of the*
19 *Operational Research Society*, 64(12):1695–1724, 2013.
- 20 G. Clarke and J.W. Wright. Scheduling of vehicles from a central depot to a
21 number of delivery points. *Operations research*, 12(4):568–581, 1964.
- 22 S. H. Clearwater, T. Hogg, and B. A. Huberman. Cooperative problem
23 solving. *Computation: The Micro and the Macro View*, pages 33–70, 1992.
- 24 T. Crainic and M. Toulouse. Explicit and emergent cooperation schemes for
25 search algorithms. *Learning and intelligent optimization*, pages 95–109,
26 2008.
- 27 G.B. Dantzig and J.H. Ramser. The truck dispatching problem. *Management*
28 *Science*, 6:80–91, 1959.
- 29 N. El Hachemi, T. G. Crainic, N. Lahrichi, W. Rei, and T. Vidal. Solution
30 integration in combinatorial optimization with applications to cooperative
31 search and rich vehicle routing. 2014.

- 1 FIPA. Fipa iterated contract net interaction protocol specification, 2009.
2 URL <http://www.fipa.org/specs/fipa00030/index.html>.
- 3 M. R. Garey, D. S. Johnson, and R. Sethi. The complexity of flowshop
4 and jobshop scheduling. *Mathematics of operations research*, 1(2):117–129,
5 1976.
- 6 M.R. Garey and D.S. Johnson. *Computers and Intractability, A Guide to the*
7 *Theory of NP-Completeness*. Bell Telephone Laboratories Inc., 1979.
- 8 T. R. Gruber. A translation approach to portable ontology specifications.
9 *Knowledge acquisition*, 5(2):199–220, 1993.
- 10 T. Hogg and C. P. Williams. Solving the really hard problems with co-
11 operative search. In *Proceedings of the National Conference on Artificial*
12 *Intelligence*, pages 231–231, 1993.
- 13 F. Hutter, D. Babic, H.H. Hoos, and A. J. Hu. Boosting verification by
14 automatic tuning of decision procedures. In *Formal Methods in Computer*
15 *Aided Design, 2007. FMCAD'07*, pages 27–34. IEEE, 2007.
- 16 A. A. Juan, J. Faulin, R. Ruiz, B. Barrios, and S. Caballé. The sr-gcws hybrid
17 algorithm for solving the capacitated vehicle routing problem. *Applied Soft*
18 *Computing*, 10(1):215–224, 2010a.
- 19 A A. Juan, J. Faulín, J. Jorba, D. Riera, D. Masip, and B. Barrios. On the use
20 of monte carlo simulation, cache and splitting techniques to improve the
21 clarke and wright savings heuristics. *Journal of the Operational Research*
22 *Society*, 62(6):1085–1097, 2011.
- 23 A. A. Juan, J. Faulin, J. Jorba, J. Caceres, and J. M. Marquès. Using parallel
24 & distributed computing for real-time solving of vehicle routing problems
25 with stochastic demands. *Annals of Operations Research*, 207(1):43–65,
26 2013.
- 27 A. A. Juan, Helena R. Lourenço, M. Mateo, R. Luo, and Q. Castella. Using
28 iterated local search for solving the flow-shop problem: Parallelization,
29 parametrization, and randomization issues. *International Transactions in*
30 *Operational Research*, 21(1):103–126, 2014.

- 1 A.A. Juan, R. Ruíz, H. R. Lourenço, M. Mateo, and D. Ionescu. A simulation-
2 based approach for solving the flowshop problem. In *Proceedings of the*
3 *Winter Simulation Conference*, pages 3384–3395. Winter Simulation Con-
4 ference, 2010b.
- 5 A.A. Juan, J. Faulin, S. E. Grasman, M. Rabe, and G. Figueira. A review of
6 simheuristics: Extending metaheuristics to deal with stochastic combina-
7 torial optimization problems. *Operations Research Perspectives*, 2:62–72,
8 2015.
- 9 A. Kouider and B. Bouzouia. Multi-agent job shop scheduling system based
10 on co-operative approach of idle time minimisation. *International Journal*
11 *of Production Research*, 50(2):409–424, 2012.
- 12 M. López-Ibáñez, J. Dubois-Lacoste, T. Stützle, and M. Birattari. The irace
13 package, iterated race for automatic algorithm configuration. *IRIDIA,*
14 *Université Libre de Bruxelles, Belgium, Tech. Rep. TR/IRIDIA/2011-004,*
15 2011.
- 16 R. Malek. An agent-based hyper-heuristic approach to combinatorial op-
17 timization problems. In *Intelligent Computing and Intelligent Systems*
18 *(ICIS), 2010 IEEE International Conference on*, volume 3, pages 428–434,
19 2010.
- 20 S.P. Martin, D. Ouelhadj, P. Smet, G. Vanden Berghe, and E. Özcan. Coop-
21 erative search for fair nurse rosters. *Expert Systems with Applications*, 40
22 (16):6674–6683, 2013.
- 23 D. Meignan, J.C. Creput, and A. Koukam. A coalition-based metaheuris-
24 tic for the vehicle routing problem. In *Evolutionary Computation, 2008.*
25 *CEC 2008.(IEEE World Congress on Computational Intelligence). IEEE*
26 *Congress on*, pages 1176–1182. IEEE, 2008.
- 27 D. Meignan, A. Koukam, and J. C. Créput. Coalition-based metaheuristic:
28 a self-adaptive metaheuristic using reinforcement learning and mimetism.
29 *Journal of Heuristics*, pages 1–21, 2010. ISSN 1381-1231.
- 30 M. Milano and A. Roli. Magma: A multiagent architecture for metaheuris-
31 tics. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Trans-*
32 *actions on*, 34(2):925–941, 2004. ISSN 1083-4419.

- 1 David S Moore and George P McCabe. *Introduction to the Practice of Statistics*. WH Freeman/Times Books/Henry Holt & Co, 1989.
- 2
- 3 M. Nawaz, E. E. Enscore Jr, and I. Ham. A heuristic algorithm for the m-
4 machine, n-job flow-shop sequencing problem. *Omega*, 11(1):91–95, 1983.
- 5 D. Ouelhadj and S. Petrovic. A cooperative hyper-heuristic search frame-
6 work. *Journal of Heuristics*, 16(6):835–857, 2010.
- 7 M. Pinedo. *Scheduling: theory, algorithms, and systems*. Prentice-Hall, New
8 Jersey, 2002.
- 9 J. Ries and P. Beullens. A semi-automated design of instance-based fuzzy pa-
10 rameter tuning for metaheuristics based on decision tree induction. *Journal*
11 *of the Operational Research Society*, 66(5):782–793, 2015.
- 12 P Ross. Hyper-heuristics. In *Search Methodologies*, pages 611–638. Springer,
13 2014.
- 14 E. Taillard. Benchmarks for basic scheduling problems. *European Journal of*
15 *Operational Research*, 64(2):278–285, 1993.
- 16 E. G. Talbi and V. Bachelet. Cosearch: A parallel cooperative metaheuristic.
17 *Journal of Mathematical Modelling and Algorithms*, 5(1):5–22, 2006. ISSN
18 1570-1166.
- 19 M. Toulouse, K. Thulasiraman, and F. Glover. Multi-level cooperative search:
20 A new paradigm for combinatorial optimization and an application to
21 graph partitioning. *Euro-Par’99 Parallel Processing*, pages 533–542, 1999.
- 22 E. Vallada and R. Ruiz. Cooperative metaheuristics for the permutation
23 flowshop scheduling problem. *European Journal of Operational Research*,
24 193(2):365–376, 2009.
- 25 X. F. Xie and J. Liu. Multiagent optimization system for solving the trav-
26 eling salesman problem (tsp). *Systems, Man, and Cybernetics, Part B:*
27 *Cybernetics, IEEE Transactions on*, 39(2):489–502, 2009.
- 28 G.I. Zobolas, C. D. Tarantilis, and G. Ioannou. Minimizing makespan in
29 permutation flow shop scheduling problems using a hybrid metaheuristic
30 algorithm. *Computers & Operations Research*, 36(4):1249–1267, 2009.