

Multi Colony Ant Algorithms

Martin Middendorf (martin.middendorf@ku-eichstaett.de)
*Fachgruppe für Informatik, Katholische Universität Eichstätt, Ostenstr. 28,
D-85072 Eichstätt, Germany*

Frank Reischle (frank.reischle@gmx.net) and Hartmut Schmeck
(schmeck@aifb.uni-karlsruhe.de)
*Institut für Angewandte Informatik und Formale Beschreibungsverfahren (AIFB),
Universität Karlsruhe, D-76128 Karlsruhe, Germany*

Abstract. In multi colony ant algorithms several colonies of ants cooperate in finding good solutions for an optimization problem. At certain time steps the colonies exchange information about good solutions. If the amount of exchanged information is not too large multi colony ant algorithms can be easily parallelized in a natural way by placing the colonies on different processors. In this paper we study the behaviour of multi colony ant algorithms with different kinds of information exchange between the colonies. Moreover we compare the behaviour of different numbers of colonies with a multi start single colony ant algorithm. As test problems we use the Traveling Salesperson problem and the Quadratic Assignment problem.

Keywords: Ant algorithm, Quadratic Assignment, Traveling Salesperson Problem

1. Introduction

The Ant Colony Optimization (ACO) metaheuristic has been applied successfully to solve various combinatorial optimization problems as the Traveling Salesperson problem (e.g. Dorigo et al. (1996)), the Quadratic Assignment problem (e.g. Maniezzo et al. (1999) and Gambardella et al. (1999)), the Shortest Common Supersequence problem (Michels et al. (1999)), the Job Shop Scheduling problem (Colorni et al. (1994)), or the Resource-Constraint Project Scheduling problem (Merkle et al. (2000)) (for an overview of ant algorithms see Dorigo et al. (1999)). In ACO principles of communicative behaviour occurring in real ant colonies are used. ACO is an evolutionary approach where several generations of artificial ants search for good solutions. Every ant of a generation builds up a solution step by step going through several decisions. Ants that found a good solution mark their paths through the decision space by putting some amount of pheromone on the edges of the path. The following ants of the next generation are attracted by the pheromone so that they will likely search in the solution space near good solutions.

Ant algorithms are good candidates for parallelization. One approach that is particularly suitable for coarse grained parallelization is a multi colony ant algorithm where every processor holds a colony



© 2001 Kluwer Academic Publishers. Printed in the Netherlands.

of ants. After each generation the colonies exchange information about their solutions. Then, every colony computes new pheromone information which is usually stored in some pheromone matrix. The existing parallel implementations of ant algorithms that follow this approach differ only in granularity and whether the computations for the new pheromone information are done locally in all colonies or centrally by a master processor which distributes a new matrix to the colonies (e.g. Talbi et al. (1999)). A few parallel implementations have been studied so far where the colonies exchange information only after several generations Bullnheimer et al. (1998), Calégari (1999), Michels et al. (1999). The idea is to give the colonies a chance to evolve different pheromone matrices and thus search in different regions of the search space. A similar idea is behind the island model in genetic algorithms (see Kohlmorgen et al. (1999), Whitley et al. (1999)).

In this paper we investigate the influence of different kinds of information exchange between the colonies of a multi colony ant algorithm on the optimization behaviour. We show that it can be advantageous for the colonies to exchange not too much information not too often so that their pheromone matrices can to some extent develop independently. Moreover, the behaviour of multi colony ant algorithms with different numbers of colonies is compared to a multi start approach, where the colonies work independently. It is shown that the choice of the best information exchange method depends on the allowed number of solution evaluations and also on the solution quality that is required at the end. We tested the algorithms on the Traveling Salesperson problem (TSP) and the Quadratic Assignment problem (QAP).

Single colony ant algorithms for the TSP and the QAP are described in Section 2 respectively Section 3. The multi colony approach is explained in Section 4 where we also give an overview of parallel ant algorithms. Experimental results are described in Section 5. A conclusion is given in Section 6.

2. Ant Algorithm for TSP

The Traveling Salesperson problem (TSP) is to find for n given cities a shortest closed tour that contains every city exactly once. Several ant algorithms have been proposed for the TSP (e.g. Bullnheimer et al. (1999), Dorigo et al. (1996), Stützle et al. (1997), Stützle et al. (2000)). In this paper we use a simple ant algorithm as a basis since the aim is to study the general behaviour when multiple colonies are used. We do not intend to find the best ant algorithm for a specific type of problem. Several variants of ant algorithms designed to obtain best results for

the TSP are described by Stützle et al. (2000). A simple ant algorithm with one colony is described in the following.

In every generation each of m ants constructs one solution. The ant starts from a random city and iteratively moves to another city until its tour is complete and the ant is back at the starting point. Let d_{ij} be the distance between the cities i and j . The probability that the ant after arriving at a city i chooses city j from the set S of cities that have not been visited before as the next city is

$$p_{ij} = \frac{(\tau_{ij})^\alpha (\eta_{ij})^\beta}{\sum_{k \in S} (\tau_{ik})^\alpha (\eta_{ik})^\beta}$$

Here τ_{ij} is the amount of pheromone on the edge (ij) , $\eta_{ij} = 1/d_{ij}$ is a heuristic value, and α and β are constants that determine the relative influence of the pheromone and the heuristic on the decision of the ant.

For some constant $m_b \leq m$ the m_b ants having found the best solutions are allowed to update the pheromone matrix. But before that is done some of the old pheromone is evaporated by multiplying it with a factor $\rho < 1$

$$\tau_{ij} = \rho \cdot \tau_{ij}$$

This avoids old pheromone from having a too strong influence on future decisions. Then every ant that is allowed to update adds pheromone to every edge (ij) which is on its tour. The amount of pheromone added to such an edge (ij) is Q/L where L is the length of the tour that was found and Q is a constant:

$$\tau_{ij} = \tau_{ij} + \frac{Q}{L} \quad (1)$$

In addition to the pheromone update by the m_b best solutions of the current ant generation an elitist strategy is used where e additional ants - called elitist ants - add pheromone along the edges of the best solution found so far. The elitist strategy prevents that the amount of pheromone along the edges of the best solution found so far might be reduced too much by evaporation.

The algorithm creates new ant generations until some stopping criterion is met, e.g. a certain number of generations has been done or the best found solution did not change for several generations.

3. Ant Algorithm for QAP

The Quadratic Assignment problem (QAP) is to find an assignment of n facilities to n locations with given distances between the locations and given flows between the facilities. The aim is to minimize the sum of all products of a flow value with its corresponding distance value. Formally, there are n facilities, n locations, and $n \times n$ matrices $A = [a_{ij}]$ and $B = [b_{ij}]$ where a_{ij} is the distance between locations i and j and b_{hk} is the flow between facilities h and k . The problem is to find an assignment of facilities to locations, i.e., a permutation f of $[1, n] = \{1, 2, \dots, n\}$ such that the following measure is minimized

$$\sum_{i=1}^n \sum_{j=1}^n a_{f(i)f(j)} b_{ij} \quad (2)$$

Several authors have applied ACO to the QAP (e.g. Maniezzo et al. (1999), Gambardella et al. (1999), Stützle et al. (2000), see Stützle et al. (1999) for an overview of ant algorithms for the QAP). In this paper a simple ant algorithm is used that works similar to the algorithm for the TSP that is described in Section 2. Here the pheromone value τ_{ij} , $i, j \in [1, n]$ refers to the assignment of facility i to location j . As heuristic value we use $\eta_{ij} = 1/c_{ij}$ where $c_{ij} = \sum_{k=1}^n a_{ik} \cdot \sum_{h=1}^n b_{hj}$. The matrix $C = [c_{ij}]$ is often called coupling matrix. For pheromone update the value of L in formula (1) now refers to the quality of the solutions to the QAP as given by (2).

4. Multi Colony Approaches

The multi colony approaches let several colonies of ants cooperate to find good solutions. Two different approaches have been studied in the literature. The heterogeneous approach is to have colonies of ants with a behaviour that differs between the colonies. This approach has been used e.g. for multi criteria optimization problems where the colonies work with different optimization criteria (see e.g. Gambardella et al. (1999)). In this paper we use the homogeneous approach where all the ants show a similar behaviour (see e.g. Michels et al. (1999)), and where information is exchanged between the colonies every I generations.

Only a few parallel implementations of ant algorithms have been described in the literature. A very fine-grained parallelization where every processor holds only a single ant was implemented by Bolondi et al. (1993). Due to the high overhead for communication this implementation did not scale very well with a growing number of processors.

Better results have been obtained with a more coarse grained variant by Dorigo (1993).

Bullnheimer et al. (1998) propose a parallelization where an information exchange between several colonies of ants is done every k generations for some fixed k . They show by simulations how much the running time of the algorithm decreases with an increasing interval between the information exchange. But it is not discussed how this influences the solution quality.

Stützle (1998) compares the solution quality obtained by several independent short runs of an ant algorithm with the solution quality of one long run whose running time equals the total running times of the short runs. Under some conditions the short runs proved to give better results. Also, they have the advantage that they can run in parallel and also it is possible to use different parameter values for the runs.

Talbi et al. (1999) implemented a parallel ant algorithm for the Quadratic Assignment problem. They used a fine-grained master-worker approach, where every worker holds a single ant that produces one solution. Every worker then sends its solution to the master. The master computes the new pheromone matrix and sends it to the workers.

An island model approach that uses ideas from Genetic Algorithms was proposed in Michels et al. (1999). Here, every processor holds one colony of ants and the locally best solution are exchanged between the colonies after every fixed number of iterations. When a colony receives a solution that is better than the best solution found so far by this colony, the received solution becomes the new best found solution for this colony. This solution influences the colony because because an elitist strategy is applied where during trail update some pheromone is always put on the trail that corresponds to the best found solution.

The results of Krüger et al. (1998) for the Traveling Salesperson problem indicate that it is better to exchange only the best solutions found so far than to exchange whole pheromone matrices and add the received matrices — multiplied by some small factor — to the local pheromone matrix.

Different information exchange strategies between the colonies have been studied in a preliminary study to this paper by Middendorf et al. (2000). One of the strategies consists of establishing a virtual neighborhood between the colonies so that they form a directed ring. During every information exchange step each colony sends its locally best solution to its successor colony in the ring. The local variable storing the best found solution is updated accordingly. A similar approach was done independently in Calégari (1999) where the speed up for different number of colonies and processors was studied.

4.1. STRATEGIES FOR INFORMATION EXCHANGE

We investigate four strategies for information exchange differing in the degree of coupling that is enforced between the colonies by this exchange. Since the results of Krüger et al. (1998) indicate that exchange of complete pheromone matrices is not advantageous all our methods are based on the exchange of single solutions.

- a) Exchange of globally best solution: In every information exchange step the globally best solution is computed and sent to all colonies where it becomes the new locally best solution.
- b) Circular exchange of locally best solutions: A virtual neighbourhood is established between the colonies so that they form a directed ring. In every information exchange step every colony sends its locally best solution to its successor colony in the ring. The variable that stores the best found solution is updated accordingly.
- c) Circular exchange of migrants: As in (b) the processors form a virtual directed ring. In an information exchange step every colony compares its m_b best ants with the m_b best ants of its successor colony in the ring. The m_b best of these $2m_b$ ants are then used to update the pheromone matrix.
- d) Circular exchange of locally best solutions plus migrants: Combination of strategies (b) and (c).

5. Results

We tested the multi colony ant algorithm on the Euclidian TSP instance eil101 from the TSPLIB and the QAP instance tai60b from the QAPLIB. The TSP instance eil101 is symmetric (i.e. $d_{ij} = d_{ji}$ for $i, j \in [1, n]$) and has 101 cities. The smallest tour length for this instance is known to be 629. The QAP instance tai60b is an asymmetric randomly generated instance with 60 facilities/locations. The best known solution for tai60b has a value of about 6.08×10^8 .

The parameter values that were used in our tests are: $\alpha = 1$, $\beta = 5$, and $\rho = 0.95$. The number m_b of ants that are allowed to update the pheromone matrix in a colony, the number e of elitist ants, and the value of the parameter Q were varied with the size of the colony. If not stated otherwise, all given results are averaged over 20 runs each over 500 generations for the TSP, and over 5000 generations for the QAP, respectively.




Figure 1. TSP: Difference σ between pheromone matrices, migration interval 10

5.1. COMPARISON OF INFORMATION EXCHANGE METHODS

Some test runs were performed on the TSP instance to show the influence of the information exchange strategy on the differences between the pheromone matrices of the colonies. Let $T^{(k)}$ be the pheromone matrix of the colony k , $k \in [1, N]$. The average pheromone matrix M is defined by

$$M_{ij} = \frac{1}{N} \sum_{k=1}^N T_{ij}^{(k)}$$

To measure the average difference between the pheromone matrices we took the average σ of the variances between the single elements of the matrices, i.e.,

$$\sigma = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n S_{ij} \quad \text{where} \quad S_{ij} = \sqrt{\frac{1}{N-1} \sum_{k=1}^N (T_{ij}^{(k)} - M_{ij})^2}$$

The test runs were done with $N = 10$ colonies of 10 ants each. Figures 1 and 2 show the results when an information exchange is done after every $I = 10$ and $I = 50$ generations, respectively. The figures show that during the first generations the pheromone matrices in the colonies evolve into different directions which results in a growing average difference σ ($\sigma = 0.09$ (0.23) after 10 generations for $I = 10$ (respectively 50 generations for $I = 50$)). One extreme is method (a) where an exchange of the globally best solution has the effect that σ becomes smaller very fast before it starts to grow slowly. But when the next information exchange takes place σ has not reached the value

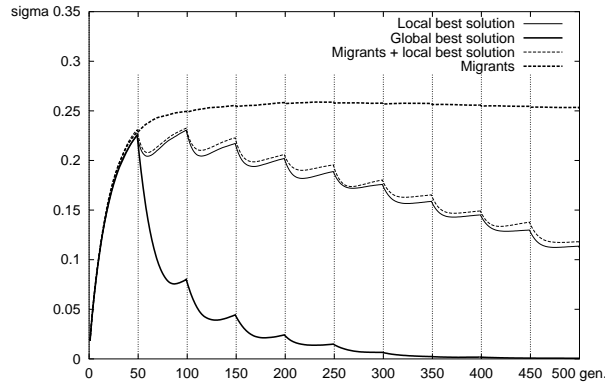


Figure 2. TSP: Difference σ between pheromone matrices, migration interval 50

it had immediately before the last information exchange. Thus, the values of σ every I th generations after the first information exchange are monotonically decreasing. After about 125 (275) generations σ is smaller than 0.01 in the case $I = 10$ (respectively $I = 50$). The other extreme is method (c) where the exchange of migrants has only a small effect on σ and the curves for $I = 10$ and $I = 50$ are quite similar. The value of σ increases over several information exchanges up to about 0.25 and then degrades only very slowly to 0.24 in generation 500 for $I = 10$ whereas it remains at 0.25 for $I = 50$. The curve of method (b) lies between the curves of methods (a) and (c). The colonies can evolve longer into different directions than for method (a) but then the information exchange forces the colonies to develop more and more into the same direction. But there is a big difference between frequent information exchanges ($I = 10$) and rare information exchanges ($I = 50$). For $I = 50$ the matrices show larger differences over more generations (the maximum is in generation 100 with $\sigma = 0.24$) than for $I = 10$ (the maximum is in generation 50 with $\sigma = 0.17$). Moreover, the frequent information exchange leads to a faster decrease of σ . Method (d) behaves similarly to method (b) which means that the circular exchange of the locally best solution is the dominating factor.

In order to show how strict the ant algorithm converges to a single solution the average number of alternatives was measured that an ant has for the choice of the next city. We are interested only in cities that have at least some minimal chance to be chosen next. For the ant k that is placed on city i let $D^{(k)}(i) = |\{j \mid p_{ij} > \lambda, j \in [1, n] \text{ was not visited}\}|$ be the number of possible successors that have a probability $> \lambda$ to be chosen. Then, the average number of alternatives with probability $> \lambda$ during a generation is




Figure 3. TSP: Average number of alternatives D , $\lambda = 0.01$, migration interval 10

$$D = \frac{1}{mn} \sum_{k=1}^m \sum_{i=1}^n D^{(k)}(i) \quad (3)$$

Clearly, $D \geq 1$ always holds for $\lambda < 1/(n-1)$. Note that a similar measure — the λ -branching factor — was used in Gambardella et al. (1995) to measure the dimension of the search space. In contrast to D the λ -branching factor considers all other cities as possible successors and not only those cities that have not yet been visited by the ant. Hence, the measure D introduced here takes into account only the alternatives that the ants really meet, whereas the λ -branching factor is a more abstract measure and problem dependent.

Figures 3 and 4 show the influence of the information exchange strategies on the value of D for $\lambda = 0.01$ when information exchange is done every $I = 10$, respectively $I = 50$ generations. The figures show that the average number of alternatives that have a high probability to be chosen by an ant falls down fast (in all cases below 2 before the 80th generation). But after every information exchange step the value of D grows which means that the exchanged information was exploited by the ants. After generation 150 the D values for method (a) are always lower than for the other methods. Also the increase of D after an information exchange step is less than for the other methods (see Figure 4). As shown above the pheromone matrices of the colonies became very similar quite early when using method (a). Therefore, the exchange of a “strong” global best solution does on the average not lead to the exploration of new regions of the search space by the colonies. It is interesting that during the first 100-150 generations D falls fastest




Figure 4. TSP: Average number of alternatives D , $\lambda = 0.01$, migration interval 50

for method (c) but in generation 500 the D value of about 1.1 in case $I = 10$ is the largest. The D values of methods (b) and (d) with circular information exchange of locally best solutions are quite similar. They are always larger than those of method (a). For frequent information exchange ($I = 10$) the colonies show a convergence behaviour and the value of D falls below the value for method (c). In contrast, the D values for rare information exchange ($I = 50$) fluctuate around 1.1 during the last 100 generations and stay larger than for method (c).

Table 1 shows the lengths of the best found tours after 500 generations with $I = 50$ when using methods (a)-(c) and for the case that no information exchange is done (the results for (d) differ not significantly from those for (b)). For colonies of size 100 three ants update, for colonies of size 10 or 20 two ants update, and for colonies of size 5 ants only one ant updates. In the case of no information exchange it is better to have one large colony than several smaller ones. It seems that in case of the symmetric TSP the greater selection pressure on the ants that are allowed to update in the single large colony is more important than the possibility that several colonies might explore different regions of the search space. This explanation is supported by a study of Boese et al. (1994) who gave evidence that symmetric TSP instances show in general a globally convex structure for the set of local minima. That means there is basically only one interesting region to search for good solutions and better solutions usually lead the ant algorithm nearer to the optimum. Boese et al. (1994) have further shown that adaptive multi-start local search strategies which derive new starting points from the previously-found local minima are better than simple random multi-start methods. This suggests that multi colony ant algo-

Table I. TSP: Different strategies of information exchange: best found solution after 500 generations, $I = 50$

	No information exchange	Exchange of globally best solution	Circular exch. of locally best solutions	Circular exchange of migrants
N=1	640.2	—	—	—
N=5	642.9	640.7	637.1	643.2
N=10	642.9	641.7	637.1	642.8
N=20	648.0	642.9	640.5	645.5

gorithms might profit from the exchange of good solutions. Our results show that this is the case for method (b) were the locally best solutions are exchanged circular and a moderate number (5 or 10) of colonies is used. As shown above when using method (b) the colonies can evolve into different directions for several generations but then the information exchange leads to slowly shrinking differences. On the other hand for methods (a) and (c) there is no advantage to have several colonies over having only one large colony. As shown above the exchange of migrants which are usually worse than the locally best solutions in method (c) is so weak that the colonies can not really profit from it. It should be noted that the picture changes when the information exchange is done more often ($I = 10$) and therefore the chance that good migrants are exchanged is higher. In this case, we found, that a small number of 5 colonies is better than one (the best found solution was 638.7 in this case). For method (c) we have shown above that the colonies can not really explore different search regions since the matrices became very similar after a few generations. But still the information exchange leads to better solutions than with the same number of colonies and no information exchange.

5.2. RESULTS FOR QAP

In this section we discuss whether circular exchange of locally best solutions in a multi colony ant algorithm is also a good strategy for the QAP problem. The QAP problem is interesting because on the one hand side ant algorithms have been successfully applied to it but it shares not the property of symmetric TSP problems that “one big valley” is governing the local minima in the optimization cost surface” (Boese et al. (1994)). Stützle et al. (2000) compared the fitness landscapes of TSP and QAP instances. In particular, they investigated the correlation between the fitness of local minima and their distance

to the nearest global optimum. Formally, this correlation was defined as $\rho(F, D) = Cov(F, D) / (\sqrt{Var(F)}\sqrt{Var(D)})$ where F and D are random variables that probabilistically describe the fitness and the distance of local optima to the nearest global optimum, Var is the variance and Cov is the covariance. Estimations of $\rho(F, D)$ for symmetric TSP and QAP instances based on 3-opt respectively 2-opt local minima have shown that the correlation is generally higher for the TSP instances. In particular, for the QAP instance tai60b that is used also in this paper the estimated correlation was 0.366 whereas it was significantly higher for the six analysed symmetric TSP instances which have correlations in the range [0.450, 0.631]. The lower correlation values for the QAP instances makes it less probable that colonies can profit from information exchanges for this problem.

In the following we present the results for the QAP instance tai60b. In the test runs we used 1 colony with 60 ants, 5 colonies with 12 ants each, 10 colonies with 6 ants each, and 20 colonies with 3 ants each. The number of ants that are allowed to update were $m_b = 3$ (respectively $m_b = 2$, $m_b = 1$, $m_b = 1$) and $e = 3$ (respectively $e = 2$, $e = 1$, $e = 1$).

Figure 5 shows that the differences of the solution quality obtained with different numbers of colonies are very small. In contrast to the results for the TSP there is no clear indication whether several small colonies are better or worse than one large colony. Figure 6 shows the results for the multi colony algorithm when there is no information exchange between the colonies. For this case the results show that the larger the number of colonies is (when the total number of ants per generation is constant) the worse is the quality of the solutions that are found. Similar to the TSP this shows that, when several colonies are used, an information exchange between the colonies is an important factor for a good optimization behaviour.

The influence of information exchange on the differences σ between the pheromone matrices of 10 colonies is depicted in Figure 7. The figure shows that without information exchange there is a stable amount of differences when the matrices of the colonies converge. On the other hand, with information exchange the differences start to decline after about 100 generations. There are only small increases between the information exchange steps. But these small increases can not stop the general decline. If the information exchange is done only every $I = 50$ generations the differences between the matrices became larger. But also in this case after about 200 generations they start to decline.

In order to show how strictly the ant algorithm converges to one solution we measured the average number of alternatives D (defined analogously to equation (3)) that an ant has for the choice of placing the next facility. Figure 8 shows the results for $\lambda = 0.1$. The fastest




Figure 5. QAP: best found solution for circular information exchange every $I = 10$ generations




Figure 6. QAP: best found solution for no information exchange

decrease of D is obtained for 10 colonies without information exchange. During the first 200 generations the decrease for $I = 50$ is nearly as fast as for the case of no information exchange since only 4 information exchanges took place. But then the decrease becomes smaller than for $I = 10$. The reason is that in later generations the exchanged solutions for $I = 50$ differ on the average more from the most probable solutions that are encoded in the pheromone matrices of the receiving colony. This can be concluded from the large D values and the peaks occurring after an information exchange in the curve for $I = 50$ of Figure 8.




Figure 7. QAP: Difference σ between pheromone matrices, $N = 10$

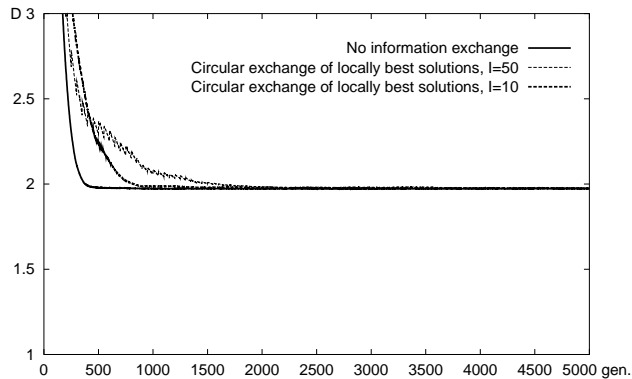


Figure 8. QAP: average number of alternatives D , $\lambda = 0.1$, $N = 10$

We conclude that, as shown for the TSP, information exchange is an advantage for multi colony ant algorithms when applied to the QAP. Also, the optimization behaviour for circular exchange of locally best solutions is quite similar for both problems. Whether a multi colony approach is better than using a single colony is dependent on the problem. It seems that a multi colony approach is less favourable for the QAP. These results complement the results of Stützle et al. (2000) who have shown that the QAP should be relatively more difficult to solve for ant algorithms than the TSP.

exchange is advantageous. Comparing the results for 2 and 3 colonies in figures 9 and 10 shows that 3 colonies are better than 2 colonies for a larger number (10000) of evaluations for every level of solution quality. We can draw the conclusion that the more time is available (i.e. the more evaluations are allowed) the larger the number of colonies (respectively number of independent runs) should be. Further, if a high solution quality is required information exchange between the colonies is important and otherwise no information exchange is better.

6. Conclusion

In this paper we studied the behaviour of simple multi colony ant algorithms with different methods of information exchange for the TSP problem and compared them to an ant algorithm with one colony only. The best information exchange method — circular exchange of locally best solution — was also tested for the QAP problem. We observed that for the TSP problem the multi colony approach with a moderate number of colonies is better than a single colony. Due to different characteristics of the QAP, the multi colony approach is not better but at least not much worse for this problem than having only one large colony. When several colonies are used it is important to have an exchange of good solutions between the colonies. We also compared a multi colony ant algorithm with a multi start ant algorithm (or equivalently a multi colony ant algorithm with independent colonies) when the number of allowed evaluations is bounded. For the TSP it was shown that the more time is available (i.e. the more number of evaluations are allowed) the larger the number of colonies (respectively number of independent runs) should be. Further, if the required solution quality is high information exchange between the colonies is useful but otherwise no information exchange is better.

References

- K. D. Boese, A. B. Kahng, S. Muddu: New Adaptive Multistart Techniques for Combinatorial Global Optimizations; *Operations Research Letters*, 16, 101-113, 1994.
- M. Bolondi, M. Bondaza: Parallelizzazione di un algoritmo per la risoluzione del problema del comesso viaggiatore; Master's thesis, Politecnico di Milano, 1993.
- B. Bullnheimer, R.F. Hartl, C. Strauss: A New Rank Based Version of the Ant System - A Computational Study; *CEJOR*, Vol 7, 25-38, 1999.
- B. Bullnheimer, G. Kotsis, C. Strauss: Parallelization Strategies for the Ant System; in: R. De Leone et al. (Eds.), *High Performance Algorithms and Software in*

- Nonlinear Optimization*; series: Applied Optimization, Vol. 24, Kluwer, 87-100, 1998.
- P.R. Calégarì: Parallelization of population-based evolutionary algorithms for combinatorial optimization problems. PhD thesis, Département D'Informatique, École Polytechnique Fédérale De Lausanne, 1999.
- A. Colnari, M. Dorigo, V. Maniezzo, M. Trubian: Ant system for job shop scheduling; *Belgian Journal Operations Research*, 34, 39-53, 1994.
- M. Dorigo: Optimization, Learning and Natural Algorithms (in Italian). PhD thesis, Dipartimento di Elettronica, Politecnico di Milano, 1992.
- M. Dorigo: Parallel ant system: An experimental study; Unpub. manuscript, 1993.
- M. Dorigo, G. Di Caro: The Ant Colony Optimization meta-heuristic; in: D. Corne, M. Dorigo, F. Glover (Eds.), *New Ideas in Optimization*, McGraw-Hill, 1999, 11-32.
- M. Dorigo, V. Maniezzo, A. Colnari: The Ant System: Optimization by a Colony of Cooperating Agents; *IEEE Trans. Sys., Man, Cybernetics - B*, 26, 29-41, 1996.
- L. M. Gambardella, M. Dorigo: Ant-Q: A Reinforcement Learning approach to the traveling salesman problem; *Proceedings of ML-95, Twelfth Intern. Conf. on Machine Learning*, Morgan Kaufmann, 252-260, 1995.
- L.M. Gambardella, E.D. Taillard, G. Agazzi: MACS-VRPT: a multiple ant colony system with for vehicle routing problems with time windows; in: D. Corne, M. Dorigo, F. Glover (Eds.), *New Ideas in Optimization*, McGraw-Hill, 1999, 63-76.
- L.M. Gambardella, E.D. Taillard, M. Dorigo: Ant colonies for the Quadratic Assignment problem; *Journal of the Operational Research Society*, 50, 167-176, 1999.
- U. Kohlmorgen, H. Schneck, K. Haase: Experiences with fine-grained parallel genetic algorithms; *Ann. Oper. Res.*, 90, 203-219, 1999.
- F. Krüger, M. Middendorf, D. Merkle: Studies on a Parallel Ant System for the BSP Model; Unpub. manuscript, 1998.
- V. Maniezzo, A. Colnari: The Ant system applied to the Quadratic Assignment problem; *IEEE Trans. Knowledge and Data Engineering*, 11, 769-778, 1999.
- D. Merkle, M. Middendorf, H. Schneck: Ant colony optimization for resource-constrained project scheduling; in: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, Morgan Kaufmann, 893-900, 2000.
- R. Michels, M. Middendorf: An Ant System for the Shortest Common Supersequence Problem; in: D. Corne, M. Dorigo, F. Glover (Eds.), *New Ideas in Optimization*, McGraw-Hill, 1999, 51-61.
- M. Middendorf, F. Reischle, H. Schneck: Information Exchange in Multi Colony Ant Algorithms; in: J. Rolim (Ed.) *Parallel and Distributed Computing, Proceedings of the 15 IPDPS 2000 Workshops, Third Workshop on Biologically Inspired Solutions to Parallel Processing Problems (BioSP3)*, Mai 1-5 2000, Cancun, Mexico, Springer-Verlag, LNCS 1800 (2000), 645-652.
- T. Stützle: Parallelization strategies for ant colony optimization; in: A. E. Eiben, T. Bäck, M. Schonauer, H.-P. Schwefel (Eds.), *Parallel Problem Solving from Nature - PPSN V*, Springer-Verlag, LNCS 1498, 722-731, 1998.
- T. Stützle, H.H. Hoos: Improvements on the ant system: Introducing MAX(MIN) ant system; in G. D. Smith et al. (Eds.), *Proc. of the International Conf. on Artificial Neural Networks and Genetic Algorithms*, Springer-Verlag, 245-249, 1997.
- T. Stützle, M. Dorigo: ACO algorithms for the Quadratic Assignment problem. in: D. Corne, M. Dorigo, F. Glover (Eds.), *New Ideas in Optimization*, McGraw-Hill, 1999, 33-50.

- T. Stützle, H.H. Hoos: MAX(MIN) ant system. *Future Generation Computer Systems*, 19, 889–914, 2000.
- E.-G. Talbi, O. Roux, C. Fonlupt, D. Robillard: Parallel ant colonies for combinatorial optimization problems; in J. Rolim et al. (Eds.) *Parallel and Distributed Processing, 11 IPPS/SPDP'99 Workshops*, LNCS 1586, Springer, 239-247, 1999.
- D. Whitley, S. Rana, R.B. Heckendorn: The island model genetic algorithm: On separability, population size and convergence. *Journal of Computing and Information Technology - CIT* 7, 1, 33-47, 1999.
- <http://www.iwr.uni-heidelberg.de/iwr/comopt/soft/TSPLIB/TSPLIB.html>
<http://www.imm.dtu.dk/sk/qaplib/>

Address for Offprints:

Martin Middendorf
Fachgruppe für Informatik
Katholische Universität Eichstätt
Ostenstr. 28
D-85072 Eichstätt
Germany

