

Distributed Asynchronous Algorithm for Cross-Entropy-Based Combinatorial Optimization

Poul E. Heegaard*, Otto Wittner**, Victor F. Nicola*** and Bjarne Helvik**

*) Telenor R&D¹ and
Department of Telematics, Norwegian University of Science and Technology, Norway

***) Centre for Quantifiable Quality of Service in Communication Systems²
Norwegian University of Science and Technology, Trondheim, Norway

***) University of Twente, The Netherlands

Abstract

Combinatorial optimization algorithms are used in many and diverse applications; for instance, in the planning, management, and operation of manufacturing and logistic systems and communication networks.

For scalability and dependability reasons, distributed and asynchronous implementations of these optimization algorithms have obvious advantages over centralized implementations. Several such algorithms have been proposed in the literature. Some are inspired by what is known as 'swarm intelligence,' e.g., ant-based optimization algorithms. Others are based on the method of cross-entropy.

In this paper we present a generic distributed and asynchronous cross-entropy-based algorithm for combinatorial optimization. The main advantages over previous algorithms include ease and flexibility of implementation, low overhead, robustness, and speed of convergence. As a result, the algorithm promises a wider applicability and significant efficiency gains compared to its predecessors, particularly for large-sized and real-time problems (such as those arising in distributed network management).

Preliminary comparisons (based on empirical results using standard combinatorial problems) show that the proposed algorithm compares favourably with other existing distributed algorithms with respect to overhead and speed of convergence. Owing to its robustness and convergence properties, the proposed algorithm may be suited for real-time and dynamic applications.

Keywords Combinatorial optimization, Randomization algorithms, Cross-entropy, Swarm intelligence, Ant-based optimization, CE ants, Network management

¹This work was also partially supported by the Future & Emerging Technologies unit of the European Commission through Project BISON (IST-2001-38923).

²*Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence* appointed by The Research Council of Norway, funded by the Research Council, NTNU and UNINETT. <http://www.ntnu.no/Q2S/>

1 Introduction

Combinatorial optimization algorithms are used in many and diverse applications; for instance, in the planning, management, and operation of communication networks. One class of combinatorial problems frequently encountered when managing communication networks is the problem of finding paths. Examples: end to end paths in (virtual) circuit switched networks both for primary paths and backup path in SDH, ATM and MPLS, routes in connectionless networks, shortest (or longest) tours visiting all nodes (STST). (Throughout this paper *path* is used as a collective term encompassing a number of the more specific technical terms path, route, circuit, tour and trajectory.)

Finding the optimal path or combination of paths usually leads to NP-hard combinatorial optimization problems, see for instance [1, 2]. A number of well known methods exist for finding near optimal solutions to these problems, e.g. simulated annealing, [3], tabu search [4] and genetic algorithms [5]. Among recent and promising methods, we have the Ant Colony System [6] proposed by Dorigo et al and the cross-entropy method proposed by Rubinstein [7].

For scalability and dependability reasons, distributed and asynchronous implementations of optimization algorithms have obvious advantages over centralized implementations. By distribution, a single point of failure in a communication network management system can be avoided, and by asynchronous operation the management system itself can be made robust to failures in the network being managed.

One class of potentially distributed and asynchronous algorithms are inspired by social biological systems in nature and is known as *swarm intelligence* [8]. Here to, the truly distributed and asynchronous systems from this class (e.g. [9, 10, 11, 12]) have concentrated on solving the shortest path routing problem. A more general approach is desirable to enable implementation of a wider range of management applications. Constructing systems capable of finding near optimal Hamiltonian cycles in networks, i.e. good solutions to the travelling salesman problem (TSP), can fulfil this generality since TSPs are among the hardest routing problems (NP complete). Hamiltonian cycles are even shown to be directly applicable for network management. In [13, 14] Hamiltonian cycles are argued to be applicable as protection paths (*p-cycle*), enabling fast and simple re-routing of network traffic on link and node failures.

A subclass of swarm intelligence systems are based on Rubinstein's method of cross-entropy [15]. In this paper we present an improved version of the cross-entropy based system presented by Helvik and Wittner in [16], today known as *CE-ants* (Cross-Entropy ants). The new version, denoted *elite CE-ants*, introduces an improved performance function which ensures a better focus on good solutions, while it still keeps the risk of premature convergence to local optima low. The main advantages over the previous version is low overhead, and speed of convergence, two properties of interest in any optimization context. From a network engineering point of view, the speed of convergence and minimum overhead is especially important, e.g. for management systems operating in dynamic networks with limited network resources like in wireless ad-hoc networks. This is even more important than to obtain the global optimal solution as long as the solution is good enough to serve the purpose.

The rest of the paper is organized into four sections. Section 2 presents a summary of the foundation for the CE-ants system, Section 3 presents our new elite approach, Section 4 presents a case study which compares simulation results of the new elite CE-ant system with results for earlier systems, and finally Section 5 summaries by presenting some concluding remarks as well as indications to future work.

2 Cross Entropy Ants (CE ants)

The CE ant system which forms the foundation for the work presented in this paper, is a *swarm intelligence* system [8] and mimics the foraging behavior of ants. It applies a high number of agents with simple behaviors which communicate indirectly and asynchronously. All agents have the mission of searching for cyclic paths and report (by asynchronous messaging) the quality of a path found as measured by a path performance expression.

2.1 Foundations

In [9] Schoonderwoerd & al. introduced the concept of using multiple agents with a behavior inspired by foraging ants to solve problems in telecommunication networks. The concept has been further developed in [10, 12, 11] among others. Schoonderwoerd & al.'s agent system relates to Dorigo & al.'s Ant Colony Optimization (ACO) system [6]. The overall idea is to have a number of simple ant-like mobile agents search for paths between source and destination nodes. While moving from node to node in a network, an agent leaves markings resembling the pheromones left by real ants during ant trail development. This results in nodes holding a distribution of pheromone markings pointing to their different neighbor nodes. An agent visiting a node uses the distribution of pheromone markings to select which node to visit next. A high number of markings pointing towards a node (high pheromone level) implies a high probability for an agent to continue its itinerary toward that node. Using trail marking agents together with a constant evaporation of all pheromone markings, Schoonderwoerd and Dorigo show that after a relatively short period of time the overall process converges toward having the majority of the agents follow a single trail. The trail tends to be a near optimal path from the source to the destination.

2.2 Cost function

In this paper we regard fully meshed networks with n nodes, i.e. every node has a direct link to every other node. The number of Hamiltonian cycles in such a network is $(n - 1)!$. A link connecting two adjacent nodes k, l has a link cost L_{kl} . The link cost may in a realistic communication network topology be in terms of incurred delay by using the path, "fee" paid the operator of the link, a penalty for using a scarce resource like free capacity, etc. or a combination of such measures.

Path i through the network is represented by $\pi_i = \{r_1, r_2 \dots, r_{n_i}\}$ where n_i is the number of nodes traversed. For a Hamiltonian cycle $n_i = n + 1$, $\forall i$ and $r_{1_i} = r_{n_i}$.

The cost function, L , of a path is additive,

$$L(\pi_i) = \sum_{j=1}^{n_i-1} L_{r_j r_{j+1}} \quad (1)$$

2.3 The Cross Entropy Method

In [7] Rubinstein develops a centralized search algorithm with similarities to Ant Colony Optimization [17, 18]. The total collection of pheromone markings in a network at time t is represented by a probability matrix P_t where an element $P_{t,rs}$ (at row r and column s of the matrix) reflects the normalized intensity of pheromones pointing from node r

toward node s . An agent's stochastic search for a sample path resembles a Markov Chain selection process based on P_t .

In a large network with a high number of feasible paths with different qualities, the event of finding an optimal path by doing a random walk (using a uniformly distributed probability matrix) is rare, e.g. the probability of finding the shortest Hamiltonian cycle in a 26 node network is $\frac{1}{25!} \approx 10^{-26}$. Thus Rubinstein develops his algorithm by founding it in rare event theory.

By importance sampling in multiple iterations Rubinstein alters the transition matrix ($P_t \rightarrow P_{t+1}$) and amplifies certain probabilities such that agents eventually find near optimal paths with high probabilities. Cross entropy (CE) is applied to ensure efficient alteration of the matrix. To speed up the process further, a performance function weights the path qualities (two stage CE algorithm [19]) such that high quality paths have greater influence on the alteration of the matrix. Rubinstein's CE algorithm has 4 steps:

1. At the first iteration $t = 0$, select a start transition matrix $P_{t=0}$ (e.g. uniformly distributed).
2. Generate N paths from P_t . Calculate the minimum Boltzmann temperature γ_t to fulfill average path performance constraints, i.e.

$$\min \gamma_t \text{ s.t. } h(P_t, \gamma_t) = \frac{1}{N} \sum_{k=1}^N H(\pi_k, \gamma_t) > \rho \quad (2)$$

where

$$H(\pi_k, \gamma_t) = e^{-\frac{L(\pi_k)}{\gamma_t}}$$

is the performance function returning the quality of path π_k . $L(\pi_k)$ is the cost of using path π_k . $10^{-6} \leq \rho \leq 10^{-2}$ is a search focus parameter. The minimum solution for γ_t will result in a certain amplification (controlled by ρ) of high quality paths and a minimum average $h(P_t, \gamma_t) > \rho$ of all path qualities in the current batch of N paths.

3. Using γ_t from step 2 and $H(\pi_k, \gamma_t)$ for $k = 1, 2, \dots, N$, generate a new transition matrix P_{t+1} which maximizes the "closeness" (i.e. minimizes distance) to the optimal matrix, by solving

$$\max_{P_{t+1}} \frac{1}{N} \sum_{k=1}^N H(\pi_k, \gamma_t) \sum_{ij \in \pi_k} \ln P_{t,ij} \quad (3)$$

where $P_{t,ij}$ is the transition probability from node i to j at iteration t . The solution of (3) is shown in [7] to be

$$P_{t+1,rs} = \frac{\sum_{k=1}^N I(\{r, s\} \in \pi_k) H(\pi_k, \gamma_t)}{\sum_{l=1}^N I(\{r\} \in \pi_l) H(\pi_l, \gamma_t)} \quad (4)$$

which will minimize the cross entropy between P_t and P_{t+1} and ensure an optimal shift in probabilities with respect to γ_t and the performance function.

4. Repeat steps 2-3 until $H(\hat{\pi}, \gamma_t) \approx H(\hat{\pi}, \gamma_{t+1})$ where $\hat{\pi}$ is the best path found.

2.4 Distributed Cross Entropy Method

In [16] a distributed and asynchronous version of Rubinstein’s CE algorithm is developed, today known as *CE ants*. By a few approximations, (4) and (2) may be replaced by autoregressive counterparts based on

$$P_{t+1,rs} = \frac{\sum_{k=1}^t I(\{r, s\} \in \pi_k) \beta^{t-k} H(\pi_k, \gamma_t)}{\sum_{l=1}^t I(\{r\} \in \pi_l) \beta^{t-l} H(\pi_l, \gamma_t)} \quad (5)$$

and

$$\min \gamma_t \text{ s.t. } h'_t(\gamma_t) > \rho \quad (6)$$

where

$$\begin{aligned} h'_t(\gamma_t) &= h'_{t-1}(\gamma_t) \beta + (1 - \beta) H(\pi_t, \gamma_t) \\ &\approx \frac{1 - \beta}{1 - \beta^t} \sum_{k=1}^t \beta^{t-k} H(\pi_k, \gamma_t) \end{aligned}$$

and where $\beta \in \langle 0, 1 \rangle$ controls the history of paths remembered by the system (i.e. replaces N in step 2). See [16] or appendix A in [15] for details on autoregression. Step 2 and 3 in the algorithm can now be performed immediately after a single new path π_t is found, and a new probability matrix P_{t+1} can be generated.

The CE ants algorithm may be viewed as an algorithm where search agents evaluate a path found (and calculate γ_t by (6)) right after they reach their destination node, and then immediately return to their source node backtracking along the path. During backtracking pheromones are placed by updating the relevant probabilities in the transition matrix, i.e applying $H(\pi_t, \gamma_t)$ through (5).

The CE ants algorithm resembles Schoonderwoerd & al.’s original system as well as Dorigo & al.’s AntNet system [10]. However, none of the earlier systems implements a search focus stage (the adjustment of γ_t). The search focus stage ensures fast and accurate convergence even when the problem at hand is of NP hard complexity. It is unclear if Schoonderwoerd & al.’s original system or Dorigo & al.’s AntNet system can (with limited tuning) find good solutions to NP-hard problems, e.g. find near optimal Hamiltonian cycles.

3 Elite selection

In the original CE ants algorithm [16], all samples of paths, $\forall_t \pi_t$, are considered when probabilities are updated by (5) and the control variable, denoted *temperature* γ_t , is updated by (6). This means that even ants following a path with poor cost value with little or no new value with respect to finding the best path, will cause an update of the temperature as well as backtracking and update of the pheromone values. In order to reduce this overhead, i.e. reduce the number of samples in (5), the concept of *elite selection* is introduced. In this section this is presented in details, supported by numerical evidence to demonstrate the effect of different elite selection criteria. The elite selection is applicable in both static and dynamic environments. We denote the new CE ant system *elite CE ants*.

Examples of use of elite CE ants are given in Section 4 where case studies of TSP in fully meshed 26 and 48 node networks are given.

3.1 Static environments

The heuristic idea of elite selection is simply to do updates of temperature and pheromones only when the sampled path has a cost value that is within a certain bound relative to the best path found at iteration t . The challenge is to do this without limiting the search space too much and throwing away partly good samples that might be vital to obtain the optimal, or at least a good, solution.

Without loss of generality, in this paper the “best” value means “minimum” value. Hence, let $L_{best,t}$ denote the minimum value found at time t ,

$$L_{best,t} = \min_{\forall u \leq t} L(\pi_u)$$

Focusing on the *elite*, i.e. the good solutions relative to the best (so far), is formulated as an *elite selection criterion*:

$$L(\pi_t) \leq (1 + \rho_2) \cdot L_{best,t-1} \quad (7)$$

The factor, ρ_2 , can be either a constant or be changed over time. The elite selection will be increasingly restrictive over time as the threshold is decreased when

1. A new best solution is found, $L_{best,t+n} < L_{best,t}$, or
2. The factor, ρ_2 , is decreased when neither of the last $D_{periode}$ ants are accepted. This is equivalent to the increase of search focus introduced in [16]. The $D_{periode}$ is typically set to the product of the number of nodes and average node connectivity.


To get an impression of the performance of elite CE ants, see Figure 1 which presents results from simulations of TSP in a 26 node fully meshed network where:

1. $\rho_2 = 0$: only accept new samples that are equal to or better than $L_{best,t}$, i.e. the best know up till now,
2. $\rho_2 > 0$: accept new solutions that are within a certain bound relative to $L_{best,t}$.


With $\rho_2 = 0$ convergence is fast, but the solutions are inferior compared to both the original approach and elite selection with $\rho_2 > 0$. The reason is most likely that with $\rho_2 = 0$ the selection is too restrictive and throws away partly good solutions that are necessary to obtain the optimal. Using $\rho_2 > 0$ is much more promising. The results in Figure 1 are from simulations with initial $\rho_2 = 0.25$, and where ρ_2 is decreased by 0.05 down to 0 every time $D_{periode} = 26 * 25 = 625$ ants are received without any updates. With these parameters both the convergence rate and the quality of the solution are improved compared to the original CE ants algorithm. Increased convergence rate means reduced overhead which is exactly what we are aiming for.

However, the bad news is that this approach requires tuning of yet another parameter, i.e. ρ_2 . Sensitivity studies have shown that the performance of elite CE ants is sensitive to the initial settings of ρ_2 . The plot in Figure 1 only shows the results for $\rho_2 = 0.25$ which was the best setting for the 26 node case. If this is also valid for other network sizes or other problems is not known.

From (2) it is known that the best solutions in the CE ants method relates to ρ through $e^{-L(\pi)/\gamma} > \rho$ which can be rearranged to $L(\pi) < -\gamma \ln \rho$. As a heuristic, this means that the good solutions should be less than $-\gamma \ln \rho$. To see how this threshold



(a) Average cost values



(b) Best cost values

Figure 1: Plot of cost values, both average and best, using different elite selection criteria. Average over 10 simulation replica.




Figure 2: Comparison of elite selection

relates to the elite selection of (7) the change in $(1 + \rho_2) \cdot L_{best,t}$ is plotted against $-\gamma_t^\# \ln \rho$ in Figure 2. The temperature is here calculated over the samples that are accepted by the elite selection criterion given in (7). This means that the calculation of $\gamma_t^\#$ is according to (6) as before, but now the index t is the number of accepted samples according to (7).

The results in Fig. 2 show a strong correlation between the two thresholds. The $-\gamma_t^\# \ln \rho$ is a smoothed variant of the stepwise threshold $(1 + \rho_2) \cdot L_{best,t}$. Hence, based on (2) and the numerical observations, the elite selection criterion was redefined to only accept agents with a cost

$$L(\pi_t) < -\gamma_{t-1}^* \ln \rho \quad (8)$$

to contribute and backtrack, i.e. update temperature and pheromones. The temperature γ_t^* is determined by (6), now with index t as the number of samples accepted by (8).

The elite selection will be increasingly restrictive over time as the threshold is decreased when

1. Many new good solutions are found that will reduce the temperature, γ_t^* ,
2. The factor, ρ , is decreased when neither of the last $D_{periode}$ ants are accepted. This was introduced in [16] to improve speed of convergence and quality of solution. The $D_{periode}$ is typically set to the product of the number of nodes and average node connectivity.

In Figure 2 it is demonstrated that using the elite selection of (8) gives similar performance as when using (7), eliminating the parameter ρ_2 .

In previous work on centralized, batch-oriented cross entropy [20] an elite set idea applying a finite set of elite samples was tested, without improving performance. With elite CE ants we use a distributed variant of the cross entropy method, where we let

ants with significant contribution to finding good solutions immediately update the pheromones/probabilities. We do not maintain a finite set of elite samples, but have an elite selection bound that is initially broad, and gradually decreased as we approach a good solution. The reduction of the bound depends on the discovery of good solutions. See Figure 2 for an example which illustrates how the elite selection bound is accepting many samples significantly worse than the best in the first phases, for then to slowly decrease towards the best solution when the search is about to converge ($-\gamma_t^* \ln \rho \rightarrow L_{best,t}$ when t is large).

3.2 Dynamic environments

In dynamic environment where link and node failures might occur, the search space and the corresponding optimal solution change to the worse. If the changes is significant, the elite selection as formulated in (8) will reject all ants, even the ones that have found a good solution in the new search space. The reason is because only samples that are better, or slightly worse than the threshold in (8) are accepted. The threshold is calculated based on observations from previous search space. In order to detect these changes, and to accept good samples from the new search space, a simple reformulation of the elite selection criterion is proposed;

$$L(\pi_t) < -\gamma_t \ln \rho \tag{9}$$

where γ_t is the temperature as in (6) calculated over *all* samples, not only the samples selected by the the elite selection criterion. However, the updates of pheromones are as in the static case; only when the sampled path has cost value below the threshold. Observe that the update of pheromones are using the temperature, γ_u^* , which is determined by (6), now with index $u \leq t$ that is the number of accepted samples according to (9).

When the temperature is calculated over all samples, previous studies [21, 22] have shown that the temperature reacts rather quickly to changes in search space, and hence the elite selection threshold will be adjusted shortly after a network failure and start accepting ants with good solutions in the new search space. Current work in progress looks into the use of this approach to manage virtual connections in a packet network exposed to changes in network configurations.

Observe that the temperature over all samples will converge towards the temperature over updating samples when the search is converging. Hence, it is possible to use the elite selection criterion from (9) in static environments as well, but with some reduction in efficiency.

4 Case studies

To test the performance of the CE ants with elite selection, a number of simulation replica of Travelling Salesman Problems (TSP) are conducted. The TSP is chosen because this is known to be an NP complete problem and hence will stress the performance of the proposed method. In current work in progress the focus is on network management problems e.g. like establishment and management of virtual connections in dynamic networks and protection cycles for fault-tolerant purposes [14]. The topologies are taken from TSPLIB [23], one fully meshed network with 26 nodes (fri26) and one with

Table 1: Summarized results from 10 simulations

		Topology	
		<i>fri26</i>	<i>ry48p</i>
No of nodes		26	48
Standard CE-ants	No of tours	122054 (9634)	779046 (53347)
	Best tour	970 (1018)	15201 (15721)
	Converged average	1015 (13)	15663 (141)
Elite CE-ants	No of tours	49942 (10045)	232004 (112564)
	Best tour	941 (1018)	14828 (15752)
	Converged average	1011 (42)	15063 (244)
Rubinstein's CE-method	No of tours	-	345600
	Best tour	-	15509


48 nodes (*ry48p*). The 48 node topology is used to compare our algorithm to Rubinstein's algorithm [20]. However, keep in mind that we require our method to be fully distributed and hence a comparison with centralized methods is not completely fair.

4.1 Results

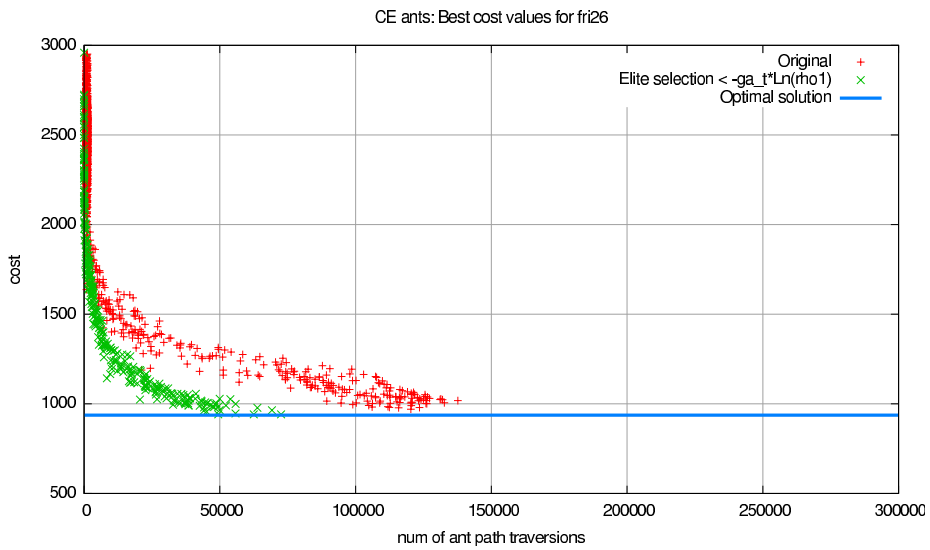
The results are given in Figures 3 and 4. All plots present average results from 10 independent simulation replica. The x-value is the number of tours ants (agents) have completed. We are counting both the tours of the ants that have successfully found a path from source to destination, and the tours of the backtracking ants that are updating the pheromones. This means without elite selection all tours will be counted twice. The reason for counting both searching and backtracking tours is that they both contribute to the overhead, and to be able to compare the original system (where all ants are backtracking) and the elite selection approach. An alternative is to use the CPU time consumption as the comparison criterion. However this would require the source code of the simulator to be revised and optimised, and it would make comparison with results from the literature more complicated. The plots show (a) the cost values at tour x , averaged over 10 replications, and (b) the best cost values found and the number of tours traversed before this. The latter is unsorted, meaning the all best value observations in all simulations are plotted in the same plot.

A summary of result details is given in Table 1. The number of tours, denoted *No of tours*, is the average number of tours before the best sample is found. The standard deviation is given in brackets. The best cost values are reported under *Best tour* being the best of the best tours of the 10 replica and with the worst of the best tours given in brackets. Finally, the average of the average cost value is reported under *Converged average* with standard deviation in brackets. The last rows show results obtained by Rubinstein's original algorithm. (Rubinstein reports better cost values in [20] but with approximately 6 times the number of samples.)

The same parameter settings as was applied in [16] were reused. $\beta = 0.998$, $\rho = 0.01$ and ρ reduction factor = 0.95 were applied in both simulation scenarios.




(a) Average cost values




(b) Best values

Figure 3: TSP in a 26 node example. 10 simulation replica.



(a) Average cost values



(b) Best cost values

Figure 4: TSP in a 48 node example. 10 simulation replica.

4.2 Observations

In the 26 node case, it is observed from the results in Figure 3 that the convergence rate of elite CE ants is approximately 3 times better than the original CE ants, and that it converges to a better solution, both for the average and the best solution found. From Table 1 it can be observed that the quality of the solutions are improved by 3-4%.

By increasing the size of the problem the improvement become even more significant. In the 48 node case, it can be observed from Figure 4 that the convergence rate of elite selection approach is at least 3 times better than the original CE ants, and that it converges to a better solution, both for the average and the best solution found. As for the 26 node case, from Table 1 it can be observed that the quality of the solutions are improved by 3-4%

4.3 Summary

Previous comparisons between CE ants and results reported by Rubinstein [16] concluded that CE ants was comparable with respect to quality of solution, but speed of convergence was not equally good. Up to 5 times more tours had to be traversed before convergence compared to the total number of samples in Rubinstein's algorithm. With the new elite CE ants approach the speed of convergence is significantly improved, and so is the quality of solutions. From a network engineering point of view it is essential to keep the rate of convergence, and the overhead imposed by ants, under control. Hence, the reduced overhead without loss of quality in solution is considered to be a significant new achievement.

5 Concluding Remarks

In this paper we have introduced an improved version of an algorithm known as CE ants, which is well suited for solving routing problems in communication networks as well as other combinatorial optimization problems. The CE ants algorithm is fully distributed and may be implemented by use of simple autonomous mobile agents. Agents act asynchronously and independently and communicate with each other only indirectly using path quality markings (pheromone trails).

In contrast to other "ant-inspired" distributed stochastic routing algorithms, the CE ants algorithm has a mathematical foundation inherited from Reuven Rubinstein's cross-entropy method for combinatorial optimization. Rubinstein proposes an efficient search algorithm using Kullback-Leibler cross-entropy, important sampling, Markov chains and the Boltzmann function.

Our new version of CE ants, denoted *elite CE ants*, has improved performance both in the speed of convergence and the quality of the paths found. By ensuring that insignificant path samples found are not processed, less overhead and a better focus on good paths is achieved. The significance threshold, or the *elite selection criterion*, applied is dynamic and shifts according to the level of convergence in the search process.

The new elite CE ants present good results when tested on a hard (NP-complete) routing problem, the Travelling Salesman Problem. Compared to the original CE ants algorithm, the elite CE ants find near optimal Hamiltonian paths of a quality around 3% better with only 1/3 of the effort. Performance wise the elite CE ants compare equally well to Rubinstein's original centralized algorithm. In addition, as for the original CE

ants, the elite CE ants can perform their search in a fully distributed and asynchronous manner.

The elite focus mechanism introduced performs well in the static network cases studied in the paper. Preliminary studies indicate that by doing some adjustments to how the elite selection criterion is controlled, the elite focus mechanism may perform very well in dynamic network cases as well. In a communication network management context, dynamic networks are considered among the most difficult to control, operate and plan.

Future work include evaluating the performance of the elite CE ants in dynamic network scenarios as well as examine if the elite CE ants can be applied in a network monitoring context. Valuable status information may be derived from studying how the temperature parameters and the elite selection criterion level vary during the search process.

References

- [1] M. O. Ball, *Handbooks in Operation Research and Management Science, Network Models*, vol. 7. North Holland, 1995.
- [2] M. O. Ball, *Handbooks in Operation Research and Management Science, Network Routing*, vol. 8. North Holland, 1995.
- [3] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by Simulated Annealing," *Science* 220, pp. 671–680, 1983.
- [4] F. Glover, *Tabu Search*. Kluwer, 1996.
- [5] D. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, 1998.
- [6] M. Dorigo and L. M. Gambardella, "Ant Colony System: A Cooperative Learning Approach to the Traveling Salesman Problem," *IEEE Transactions on Evolutionary Computing*, vol. 1, April 1997.
- [7] R. Y. Rubinstein, "The Cross-Entropy Method for Combinatorial and Continuous Optimization," *Methodology and Computing in Applied Probability*, pp. 127–190, 1999.
- [8] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.
- [9] R. Schoonderwoerd, O. Holland, J. Bruten, and L. Rothkrantz, "Ant-based Load Balancing in Telecommunications Networks," *Adaptive Behavior*, vol. 5, no. 2, pp. 169–207, 1997.
- [10] G. D. Caro and M. Dorigo, "AntNet: Distributed Stigmergetic Control for Communications Networks," *Journal of Artificial Intelligence Research*, vol. 9, pp. 317–365, Dec 1998.
- [11] J. Schuringa, "Packet Routing with Genetically Programmed Mobile Agents," in *Proceedings of SmartNet 2000*, (Vienna), September 2000.

- [12] T. White, B. Pagurek, and F. Oppacher, "Connection Management using Adaptive Mobile Agents," in *Proceedings of 1998 International Conference on Parallel and Distributed Processing Techniques and Applications (PDAPTA'98)*, 1998.
- [13] D. Stamatelakis and W. Grover, "Theoretical Underpinnings for the Efficiency of Restorable Networks Using Preconfigured Cycles ("p-cycles")," *IEEE Transactions on Communications*, vol. 48, pp. 1262–1265, August 2000.
- [14] O. Wittner, B. E. Helvik, and V. F. Nicola, "Internet failure protection using hamiltonian p-cycles found by ant-like agents," *Journal of Network and System Management, Special issue on Self-Managing Systems and Networks*, vol. Submitted, 2005.
- [15] O. Wittner, *Emergent Behavior Based Implements for Distributed Network Management*. PhD thesis, The Norwegian University of Science and Technology, November 2003.
- [16] B. E. Helvik and O. Wittner, "Using the Cross Entropy Method to Guide/Govern Mobile Agent's Path Finding in Networks," in *Proceedings of 3rd International Workshop on Mobile Agents for Telecommunication Applications*, Springer Verlag, August 14-16 2001.
- [17] M. Dorigo and G. D. Caro, "Ant Algorithms for Discrete Optimization," *Artificial Life*, vol. 5, no. 3, pp. 137–172, 1999.
- [18] M. Zlochin, M. Birattari, N. Meuleau, and M. Dorigo, "Model-based Search for Combinatorial Optimization," IRIDIA IRIDIA/2001-15, Universite Libre de Bruxelles, Belgium, 2000.
- [19] R. Y. Rubinstein, "The Cross-Entropy and Rare Events for Maximum Cut and Bipartition Problems - Section 4.4," *Transactions on Modeling and Computer Simulation*, To appear.
- [20] R. Y. Rubinstein, *Stochastic Optimization: Algorithms and Applications*, ch. Combinatorial Optimization, Cross-Entropy, Ants and Rare Events - Section 7: Noisy Networks. Kluwer Academic Publishers, 2001.
- [21] P. Heegaard, "Performance monitoring of routing stability in dynamic networks." Section 5 in Deliverable 05: "Models for basic services in AHN, P2P and Grid networks" in IST-FET Project BISON (IST-2001-38923), December 2003.
- [22] A. Mykkeltveit, "Realization of a distributed route management system by mobile agents," tech. rep., Dept. of Telematics, NTNU, 2003.
- [23] G. Reinelt, "TSPLIB." Institut für Angewandte Mathematik, Universität Heidelberg, <http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/>, 2001.