Ensuring Fast Adaptation in an Ant-Based Path Management System

Laurent Paquereau and Bjarne E. Helvik

Centre for Quantifiable Quality of Service in Communication Systems*,
Norwegian University of Science and Technology, Trondheim, Norway
{laurent.paquereau,bjarne}@q2s.ntnu.no

Abstract. The Cross-Entropy Ant System (CEAS) is an Ant Colony Optimization (ACO) system for distributed and online path management in telecommunication networks. Previous works on CEAS have focused on reducing the overhead induced by the continuous sampling of paths. In particular, elite selection has been introduced to discard ants that have sampled poor quality paths. This paper focuses on the ability of the system to adapt to changes in dynamic networks. It is shown that not returning ants may cause stagnation as that tends to make stale states persist in the network. To mitigate this undesirable side-effect, a novel pheromone trail evaporation strategy, denoted Selective Evaporation on Forward (SEoF), is presented. By allowing ants to decrease pheromone trail values on their way forward, it enforces a local re-opening of the search process in space upon change when elite selection is applied.

Keywords: CEAS, elite selection, Selective Evaporation on Forward.

1 Introduction

Ant Colony Optimization (ACO) [1] systems are systems inspired by the foraging behaviour of ants and designed to solve discrete combinatorial optimization problems. More generally, ACO systems belong to the class of Swarm Intelligence (SI) systems [2]. SI systems are formed by a population of agents, which behaviour is governed by a small set of simple rules and which, by their collective behaviour, are able to find good solutions to complex problems. ACO systems are characterized by the indirect communication between agents - (artificial) ants - referred to as stigmergy and mediated by (artificial) pheromones. In nature, pheromones are a volatile chemical substance laid by ants while walking that modifies the environment perceived by other ants. ACO systems have been applied to a wide range of problems [1]. The Cross-Entropy Ant System (CEAS) is such a system for path management in dynamic telecommunication networks.

The complexity of the problem arises from the non-stationary stochastic dynamics of telecommunication networks. A path management system should adapt

* "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, UNINETT and Telenor. http://www.q2s.ntnu.no

E. Altman et al. (Eds.): Bionetics 2009, LNICST 39, pp. 24–35, 2010.

to changes including topological changes, e.g. link/node failures and restorations, quality changes, e.g. link capacity changes, and traffic pattern changes. The type, degree and time-granularity of changes depend on the type of network. For instance, the level of variability in link quality is expected to be higher in a wireless access network than in a wired core network.

Generally, the performance of an ACO system is related to the number of iterations required to achieve a given result. Specific to the path management problem in telecommunication networks are the additional requirements put on the system in terms of time and overhead. On changes, the system should adapt, i.e. converge to a new configuration of paths, in short time and with a small overhead. In addition, finding a good enough solution in short time is at least as important as finding the optimal solution, and there is a trade-off between quality of the solution, time and overhead.

Previous works on CEAS have focused on reducing the overhead induced by the continuous sampling of solutions [3]. In particular, *elite selection* has been introduced in [4] to reduce the overhead by allowing only "good" ants to update pheromone trails. The present work addresses the adaptivity of the system. ACO systems are intrinsically adaptive and this characteristic has been used as an argument for applying ACO algorithms to dynamic problems. ACO systems have been shown to be able to adapt to changes and techniques have been developed to prevent from *stagnation*[1] [5]. However, to our knowledge, the adaptivity of ACO systems in itself, for instance in terms of number of iterations needed to converge after a change, has received little attention.

In this paper, a novel extension, denoted *Selective Evaporation on Forward* (SEoF), is introduced to improve the adaptivity of the system. The rest of this paper is organized as follows. Section 2 provides a brief introduction to CEAS. For a comprehensive presentation, the reader is referred to [3]. Section 3 describes elite selection and Section 4 characterizes the stagnation caused by not returning ants. Next, Section 5 presents the SEoF extension. Finally, Section 6 discusses related work and Section 7 concludes.

2 CEAS in a Nutshell

CEAS is an asynchronous and distributed ACO system for path management in telecommunication networks based on the Cross-Entropy (CE) method for stochastic optimization [6]. Ants cooperate to collectively find and maintain minimal cost paths, or sets of paths, between source and destination pairs. Each ant performs a random search directed by the pheromone trails to find a path to a destination. Each ant also deposits pheromones so that the pheromone trails reflect the knowledge acquired by the colony thus enforcing the stigmergic behaviour characterizing ACO systems. Similarly to what happens in nature, good solutions emerge as the result of the iterative indirect interactions between ants.

[1] *Stagnation* refers to a system that fails to adapt, or adapts very slowly, because it has converged too hard to a solution.

Formally, let a network be represented by a bidirectional weighted graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ where $\mathbf{V}$ is the set of vertices (nodes) and $\mathbf{E}$ the set of edges (links). $(v, i) \in \mathbf{E}$ denotes the link connecting node v to node i and $L((v, i))$ is the weight (cost) of link (v, i). Starting from a node s, an ant incrementally builds a path to a destination node d by moving through a sequence of neighbour nodes applying at each node a stochastic decision policy depending on the local pheromone trails and the ant internal state (*biased exploration*). At node v, the probability that an ant decides to move to node i is given by the *random proportional rule*

$$p_{t_v,vi}^{(s,d)} = \frac{\tau_{t_v,vi}^{(s,d)}}{\sum_{j \in \mathcal{N}_v} \tau_{t_v,vj}^{(s,d)}}, \ \forall i \in \mathcal{N}_v \tag{1}$$

where $\tau_{t_v,vi}^{(s,d)}$ is the pheromone trail value at node v for the link (v, i) after t_v pheromone deposits at node v, see below, and $\mathcal{N}_v \subseteq \mathbf{N}_v = \{i \in \mathbf{V} \mid (v, i) \in \mathbf{E}\}$ is the set of neighbours of node v not yet visited by the ant. After it has reached its destination d, the ant backtracks. On its way backward, it triggers pheromone evaporation[2]

$$\tau_{t_v-1,vi}^{(s,d)} \leftarrow \beta \cdot \tau_{t_v-1,vi}^{(s,d)}, \ \forall i \in \mathbf{N}_v, \ \forall v \in \boldsymbol{\pi}_{t,[s,d]} \tag{2}$$

and deposits pheromones (*online delayed* pheromone release)

$$\tau_{t_v,vi}^{(s,d)} \leftarrow \tau_{t_v-1,vi}^{(s,d)} + I\left((v, i) \in \boldsymbol{\omega}_{t,[s,d]}\right) \cdot \Delta\tau_t^{(s,d)}, \ \forall i \in \mathbf{N}_v, \ \forall v \in \boldsymbol{\pi}_{t,[s,d]} \tag{3}$$

where $\beta \in [0, 1)$ denotes the *memory factor*, $\boldsymbol{\pi}_{t,[s,d]} = \langle s, v_1, v_2, \ldots, v_{h-1}, d\rangle$ the sequence of nodes traversed by the ant, $\boldsymbol{\omega}_{t,[s,d]} = \langle (s, v_1), (v_1, v_2), \ldots, (v_{h-1}, d)\rangle$ the sequence of links, $\Delta\tau_t^{(s,d)}$ the amount of pheromones deposited (*pheromone increment*), $I(x) = 1$ if x is true, 0 otherwise, and t is the number of ants that have returned from d[3]. $\Delta\tau_t^{(s,d)}$ is chosen so that (1) minimizes the cross-entropy between two consecutive sets of random proportional rules $\mathbf{p}_t^{(s,d)} = \{p_{t,vi}^{(s,d)}\}_{\forall v,i}$ and $\mathbf{p}_{t-1}^{(s,d)}$ subject to the cost history $\mathbf{L}_{t,[s,d]} = \{L(\boldsymbol{\omega}_{k,[s,d]}) \mid k = 1, \ldots, t\}$, where $L(\boldsymbol{\omega}_{k,[s,d]}) = \sum_{\forall (i,j) \in \boldsymbol{\omega}_{k,[s,d]}} L((i, j))$ is the cost of the path $\boldsymbol{\omega}_{k,[s,d]}$.

$$\Delta\tau_t^{(s,d)} = H\big(L(\boldsymbol{\omega}_t), \gamma_t^{(s)}\big) = e^{-L(\boldsymbol{\omega}_t)/\gamma_t^{(s)}} \tag{4}$$

where $\gamma_t^{(s)}$ is an internal parameter called the *temperature* and determined at d by minimizing it subject to $h_t(\gamma_t) \geqslant \rho$. $h_t(\gamma_t) = \beta \cdot h_{t-1}(\gamma_t) + (1-\beta) \cdot H(L(\boldsymbol{\omega}_t), \gamma_t)$ is the overall auto-regressive performance function and $\rho \in (0, 1)$ a configuration parameter (*search focus*). In the following, the subscripts and superscripts referencing the source and destination nodes are omitted to help readability.

[2] More precisely, $\tau_{t,vi}$ is implemented as an auto-regressive function of γ_t and evaporation is applied on the auto-regressive variables. See for instance [3].

[3] t is incremented when an ant is at its destination, hence the number of pheromone deposits at node v is $t_v = \sum_{k=1}^{t} I((v, \cdot) \in \boldsymbol{\omega}_{k,[s,d]})$.

The temperature γ_t controls the weights given to solutions. For a given temperature, the lower the cost, the larger the pheromone increment. For a given cost value, the lower the temperature, the smaller the pheromone increment, but the larger the relative difference with the increment for a solution of higher cost, see Figure 1. γ_t is self-adjusting. When the network conditions stay unchanged, γ_t asymptotically converges to $\tilde{\gamma}_t$. If network conditions between s and d degrade, e.g. if the quality of the best path is altered so that it is no longer the best or if the best path is no longer available, $\tilde{\gamma}_t$ becomes larger and γ_t gradually increases (*reheating change*). If the network conditions between s and d improve, e.g. when a better path has become available and has been discovered, $\tilde{\gamma}_t$ becomes smaller and γ_t gradually decreases (*cooling change*).

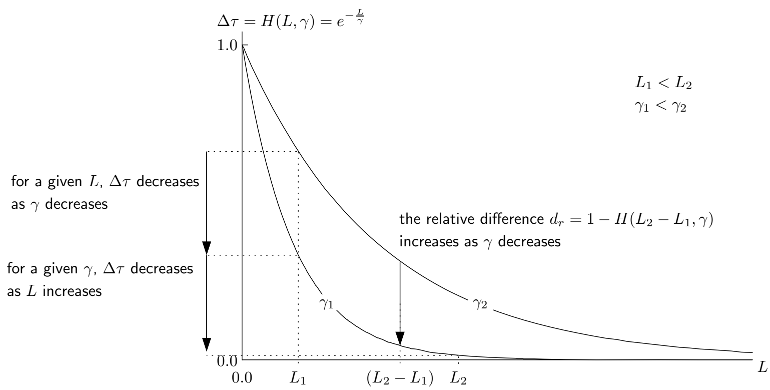

Fig. 1. Changes in pheromone increments as cost and temperature vary

To bootstrap the system, ants do not apply (1) but a *uniformly distributed proportional rule* (*uniform exploration*)

$$p_{t_v,vi} = \frac{1}{|\mathcal{N}_v|}, \ \forall i \in \mathcal{N}_v \tag{5}$$

During normal operation, a given percentage of ants applying (5) is maintained. The purpose of such *explorer ants* is threefold: (i) to ensure that new solutions are discovered, (ii) to prevent the system from converging too hard, and (iii) to maintain sparse pheromone trails on alternative solutions providing roughly up-to-date bootstrapping information in case of a change.

3 Elite Selection

Elite selection consists in only letting ants that have sampled relatively good paths (*elite paths*), i.e. paths which cost is below a certain cut-off level (*elite selection level*), update the temperature and backtrack. The rationale is that relatively poor quality solutions, i.e. which cost is beyond this level, lead to negligible changes in the pheromone trail distributions and, hence, that ants that have sampled those paths can be discarded at the destination without affecting the performance of the system otherwise. Elite selection is therefore primarily an overhead reduction technique. However, it also contributes to improving the convergence speed of the system in terms of number of iterations by focusing the search around the best solutions [4].

Formally, let n be the total number of ants arrived at d from s (*forward ants*), $\boldsymbol{\omega}_k^*$ the path followed by the k^{th} ant, and $\boldsymbol{\Omega}_n^*$ the set of all candidate paths. Elite selection can be formulated as

$$\boldsymbol{\omega}_n^* \in \hat{\boldsymbol{\Omega}}_n \Leftrightarrow L(\boldsymbol{\omega}_n^*) \leqslant \chi_n \tag{6}$$

where $\hat{\boldsymbol{\Omega}}_n \subseteq \boldsymbol{\Omega}_n^*$ is the set of elite paths and χ_n the elite selection level. It is shown in [4] that an appropriate elite selection level is

$$\chi_n = -\gamma_n^* \cdot \ln \rho \tag{7}$$

where γ_n^* is the temperature calculated from the total cost history $\mathbf{L}_t^* = \{L(\boldsymbol{\omega}_k^*) \mid k = 1, \ldots, n\}$. Contrary to the *total temperature* γ_n^*, the *elite temperature* $\gamma_{\hat{t}}$ used in (4) to determine $\Delta\tau_t$ is computed from $\hat{\mathbf{L}}_{\hat{t}} = \{L(\hat{\boldsymbol{\omega}}_k) \mid k = 1, \ldots, \hat{t}\}$ where $\hat{t} \leqslant n$ is the number of elite ants, i.e. ants that have sampled an elite path. χ_n is self-adjusting; γ_n^* adjusts to the network conditions and so does χ_n.

Ants that do not meet the elite selection criterion (6) are discarded, with the exception of explorer ants that are always returned. Hence, the number of ants that backtrack (*backward ants*) is $t \in [\hat{t}, n]$. All the operations performed at the destination are summarized in the flow chart shown in Figure 2.

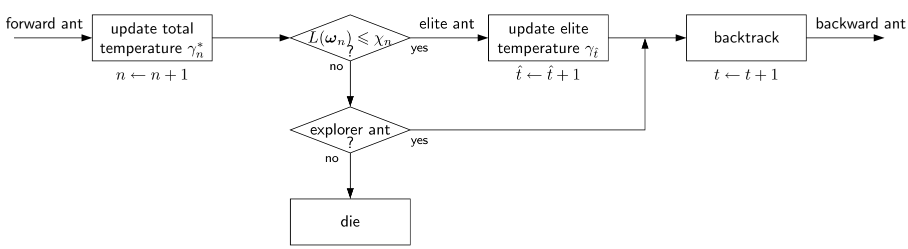

Fig. 2. Elite selection flow chart

4 Stagnation Caused by Non Returning Ants

An ant may not return because of elite selection or accidentally (loss). Looking at the adaptivity of the system, this latter case is similar to the case of an ant discarded by elite selection after the degradation of a path. Hence, it is not further considered in the following.

As long as the highest pheromone values correspond to the best path, the search process is correctly biased and elite selection is an advantageous feature as only good paths get reinforced. When the highest pheromone values do not correspond to the best path anymore, i.e. when the search process has become incorrectly biased, e.g. after a degradation of the path, it turns out to be harmful because it tends to keep stale pheromone trails longer. The crux of the problem is that pheromone trails are only updated by backward ants. Discarding ants at the destination results in no update of the pheromone values.

In particular, stagnation occurs on reheating changes because: (i) the elite set may be temporarily empty after the change, in which case even an ant following the new optimal solution does not meet the elite selection criterion ($\hat{\boldsymbol{\Omega}}_n = \emptyset \Leftrightarrow L(\boldsymbol{\omega}_n^*) > \chi_n,\ \forall \boldsymbol{\omega}_n^* \in \boldsymbol{\Omega}_n^*$), and (ii) the elite temperature is lower than what it will be when the system has converged so pheromone deposits are smaller than what they will be when the system has converged ($\gamma_t < \tilde{\gamma}_t \Leftrightarrow \Delta\tau_t < \tilde{\Delta}\tau_t = H(L(\boldsymbol{\omega}_t), \tilde{\gamma}_t)$). On cooling changes, even though the search process becomes incorrectly biased, elite selection does not cause stagnation because: (i) the elite set is never empty so pheromone trails can always potentially be updated ($\hat{\boldsymbol{\Omega}}_n \neq \emptyset$), (ii) the elite temperature is higher than what it will be when the system has converged so pheromone deposits are larger than what they will be when the system has converged ($\gamma_t > \tilde{\gamma}_t \Leftrightarrow \Delta\tau_t > \tilde{\Delta}\tau_t$), and (iii) elite selection focuses the search on the best paths and thus prevents temperature increase on sampling low quality solutions.

The above cases are demonstrated by simulation of a simple scenario chosen for illustration. CEAS, with and without elite selection, is applied to find the shortest path between nodes s and d in the four node network shown to the left in Figure 3. At $t = 0$, all the links are operational. The link (a, b) is taken down at $t = 3000$ [s] and restored at $t = 11000$ [s]. Node s generates ants at rate $\lambda = 1$ [ant/s] out of which 5% are explorer ants. The path memory factor applied is $\beta = 0.998^4$, and the search focus $\rho = 0.01$. Figure 3 shows the probability $p_{t,sa}$ of forwarding ants to node a at node s. The results are averaged over 30 replications. For $t < 3000$ [s] and $t \geqslant 11000$ [s], the best path is $\langle s, a, b, d \rangle$. Hence,

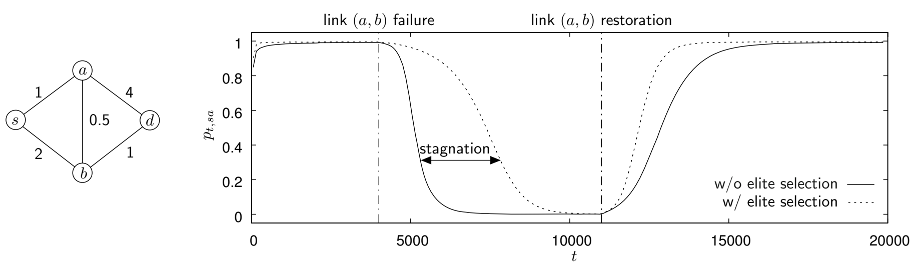

Fig. 3. Characterization of the stagnation caused by elite selection

[4] Such a high β is chosen to emphasize the stagnation phenomenon but is unnecessary to solve such a simple problem.

when the system has converged, $p_{t,sa} \approx 1$. In both intervals, the temperature decreases and using elite selection leads to faster convergence. When (a, b) is down, the best path is $\langle s, b, d \rangle$, hence, when the system has converged, $p_{t,sa} \approx 0$. Immediately after the link break, the search process is incorrectly biased and the temperature should increase, and elite selection causes stagnation[5].

5 Selective Evaporation on Forward

Stagnation occurs when the search process becomes wrongly biased because most of the ants are discarded at the destination and no update is triggered at intermediate nodes. However, the fact that an ant does not backtrack is an information in itself; it means that the sampled path is not, or no longer, an elite path. When the system has converged, the probability that an ant follows an elite path is high. Hence, if the network conditions are stable, the probability that an ant returns is also high. Therefore, if the system has converged, for a node along the sampled path, not receiving a backward ant indicates with a high probability that network conditions have changed. The idea is to exploit this knowledge locally at each node to reflect the change on the pheromone distribution and mitigate stagnation.

The approach followed, denoted *Selective Evaporation on Forward* (SEoF), is a hybrid online pheromone evaporation strategy combining step-by-step evaporation on the selected links and delayed evaporation on the other links. Formally, a forward ant triggers evaporation at node v on the selected link (v, i)

$$\tau_{t_v,vi} \leftarrow \beta \cdot \tau_{t_v,vi}, \tag{8}$$

and, on its way backward, (2) is replaced by

$$\tau_{t_v-1,vi} \leftarrow \beta \cdot \tau_{t_v-1,vi}, \ \forall i \in \mathbf{N}_v \mid (v,i) \notin \boldsymbol{\omega}_t, \ \forall v \in \boldsymbol{\pi}_t. \tag{9}$$

If an ant samples an elite path, the behaviour of the system is unmodified. If an ant samples a path that is not in the elite set, at each node along the path the pheromone trail value on the selected link only is reduced, so the probability that an ant samples the same path again is decreased. When the system has converged, the effect of this extension is marginal since most of the ants follow an elite path. On the other hand, when the pheromone distribution becomes wrongly biased after a change, the probability that an ant follows a path that is not in the elite set is high and gradually decreasing the pheromone trail values along this path does make a difference. Figure 4 shows how this strategy, denoted 'SEoF plain', effectively mitigates stagnation on reheating changes when applied to the scenario used in Section 3. In addition, in this case, the optimal solution is easy to find and reducing the pheromone trail values on paths that are not in the elite set also improves the performance of the system in cooling phases by accentuating the focus on the best solution.

Now, when the system has not yet converged, there is a non-negligible probability that an ant follows a path that is not in the elite set and therefore locally

[5] In this context, stagnation means that the number of iterations needed to converge is greater than the number of iterations that is needed if elite selection is not used.

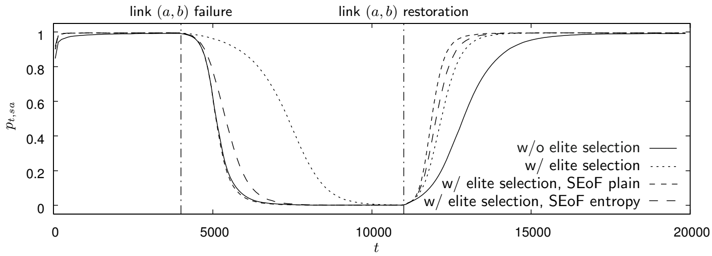

Fig. 4. Mitigating stagnation caused by elite selection using SEoF

reduces the probability of choosing a link although it may be part of a good path. Hence, this simple strategy may divert the system from good solutions and result in a slower convergence and/or convergence to a poorer solution, especially when good solutions are hard to find. See Appendix A for an illustration. To avoid this potential pitfall, (8) is replaced by

$$\tau_{t_v,vi} \leftarrow (\beta + (1-\beta)E_{t_v,v}) \cdot \tau_{t_v,vi} \tag{10}$$

and (9) by

$$\tau_{t_v-1,vi} \leftarrow \begin{cases} \beta \cdot \tau_{t_v-1,vi}, \ \forall i \in \mathbf{N}_v \mid (v,i) \notin \boldsymbol{\omega}_t, \ \forall v \in \boldsymbol{\pi}_t \\ \dfrac{\beta}{\beta + (1-\beta)E_{t_v-1,v}} \cdot \tau_{t_v-1,vi}, \ (v,i) \in \boldsymbol{\omega}_t, \ \forall v \in \boldsymbol{\pi}_t \end{cases} \tag{11}$$

where

$$E_{t_v,v} = \frac{-\sum_{\forall i \in \mathbf{N}_v} p_{t_v,vi} \log p_{t_v,vi}}{\log |\mathbf{N}_v|} \tag{12}$$

denotes the (normalized) entropy at node v.

The entropy reflects how open is the search process at node v. The idea behind this revised scheme is to balance between forward and backward evaporation on selected links depending on how severe the stagnation would be if the search process was wrongly biased. Stagnation may only occur if $E_{t_v,v} < 1$ and is all the more so severe as $E_{t_v,v}$ is low. If the entropy is close to 0 and the pheromone distribution is wrongly biased, it enforces a strong local re-opening of the search process. If the system has not yet converged, the entropy at node v will still be relatively high and evaporation will occur mostly on the way back thus alleviating the diversion effect caused by SEoF. Compared to 'SEoF plain', this revised scheme, denoted 'SEoF entropy', leads to a slightly slower convergence in the case of the scenario used in Section 3 because the pheromone reduction is lessened as the search process is re-opened. Nevertheless, it still significantly mitigates the stagnation caused by elite selection. See Figure 4. Applied to a more complex problem, it effectively reduces the diversion effect caused by SEoF, see Appendix A. More generally, SEoF is shown to improve the performance of CEAS on reheating changes as illustrated in the Appendix B.

SEoF has the following attractive characteristics: it is an online, distributed, self-adjusting, gradual and problem-independent approach, it is based on local information only, and it does not require any extra configuration parameter. Moreover, it retains and makes use of the available information about alternative solutions.

6 Related Work

Pheromone trail evaporation and elitism are not specific to CEAS. Evaporation is a core component of the ACO meta-heuristic allowing a colony to forget about old solutions and integrated in most ACO systems[6]. Elitism has been proposed for static optimization problems to improve the convergence speed by reinforcing the best solution [9,10,11]. However, to the best of our knowledge, selective evaporation is a novel idea and self-adjusting elite selection as it is implemented in CEAS remains original.

Ant Colony System (ACS) [10] also uses a hybrid pheromone update strategy including an online step-by-step reduction of the pheromone trail values on the selected edges. However, the purpose of that reduction is radically different from what is presented in this paper. It has been introduced to counterbalance a strong elitist selection. Historically though, it is interesting to note that delayed pheromone update was early preferred to step-by-step pheromone update (AS [12]). Later, elitism was introduced to improve the performance of the system (Elitist AS [9]), before a hybrid solution combining step-by-step and delayed pheromone update was proposed to cope with drawbacks introduced by elitism.

Strategies for modifying pheromone trail values after a change are found in studies on applying AS to the dynamic Traveling Salesman Problem (TSP). For instance, *pheromone shaking* is a centralized mechanism introduced in [13] to smoothen pheromone distributions when the cost between nodes has changed and equalization strategies to adapt to node insertion/deletion are proposed in [14]. In this context, changes are globally known. Pheromone modifications are applied only once, offline (daemon activity, i.e. not triggered by ants), immediately after a change and affect all pheromone trail values at all nodes.

Other ACO systems for path management in dynamic networks also integrate mechanisms to adapt to changes, but all use problem-specific measures. AntNet [7], for instance, relies on local problem-specific heuristics, AntHocNet [8] on mechanisms borrowed from traditional MANET routing protocols, e.g. explicit route error notifications.

Finally, SEoF enforces a re-opening of the search process in space when elite selection is used. As mentioned in the introduction, another dimension is time. Self-Tuned Ant Rate (STAR) [15] is a complementary extension exploiting elite selection to improve the adaptivity of the system in terms of time to converge after a change by allowing a temporary increase of the rate of forward ants.

[6] Notable exceptions are ACO systems proposed for routing in telecommunication networks including AntNet [7] and AntHocNet [8]. Although the term is used for AntNet in [1], authors refers to normalization, not evaporation.

7 Conclusion

This paper introduces a selective pheromone trail evaporation strategy improving the ability of CEAS to adapt to changes. Combined with elite selection, it results in an explicit pheromone trail reduction on changes which enforces a local re-opening of the search process in space and effectively mitigates the stagnation caused by not returning ants otherwise. To our knowledge, such an approach addressing the adaptivity of an ACO system is original and it is foreseen that the principles are applicable to other ACO systems.

References

1. Dorigo, M., Di Caro, G., Gambardella, L.M.: Ant Algorithms for Discrete Optimization. Artificial Life 5(2), 137–172 (1999)
2. Bonabeau, E., Dorigo, M., Theraulaz, G.: Swarm Intelligence: From Natural to Artificial Systems. Oxford University Press, Oxford (1999)
3. Heegaard, P.E., Wittner, O.J.: Overhead Reduction in Distributed Path Management System. In: Computer Networks. Elsevier, Amsterdam (in press, 2009)
4. Heegaard, P.E., Wittner, O.J., Nicola, V.F., Helvik, B.E.: Distributed Asynchronous Algorithm for Cross-Entropy-Based Combinatorial Optimization. In: Rare Event Simulation and Combinatorial Optimization (RESIM/COP 2004), Budapest, Hungary (2004)
5. Sim, K.M., Sun, W.H.: Ant Colony Optimization for Routing and Load-Balancing: Survey and New Directions. IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans 33(5), 560–572 (2003)
6. Rubinstein, R.Y.: The Cross-Entropy Method for Combinatorial and Continuous Optimization. Methodology and Computing in Applied Probability 1(2), 127–190 (1999)
7. Di Caro, G., Dorigo, M.: AntNet: Distributed Stigmergetic Control for Communications Networks. Journal of Artificial Intelligence Research 9, 317–365 (1998)
8. Ducatelle, F.: Adaptive Routing in Ad Hoc Wireless Multi-hop Networks. PhD thesis, University of Lugano, Switzerland (2007)
9. Dorigo, M., Maniezzo, V., Colorni, A.: Ant System: Optimization by a Colony of Cooperating Agents. IEEE Transactions on Systems, Man, and Cybernetics - Part B 26(1), 29–41 (1996)
10. Dorigo, M., Gambardella, L.M.: Ant Colony System: A Cooperative Learning Approach to the Traveling Salesman Problem. IEEE Transactions on Evolutionary Computation 1(1), 53–66 (1997)
11. Stützle, T., Hoos, H.H.: $\mathcal{MAX} - \mathcal{MIN}$ Ant System. Journal of Future Generation Computer Systems 16, 889–914 (2000)
12. Colorni, A., Dorigo, M., Maniezzo, V.: Distributed Optimization by Ant Colonies. In: 1st European Conference on Artificial Life (ECAL 1991), pp. 134–142 (1991)
13. Eyckelhof, C.J., Snoek, M.: Ant Systems for a Dynamic TSP. In: Dorigo, M., Di Caro, G.A., Sampels, M. (eds.) ANTS 2002. LNCS, vol. 2463, pp. 88–99. Springer, Heidelberg (2002)
14. Guntsch, M., Middendorf, M.: Pheromone Modification Strategies for Ant Algorithms Applied to Dynamic TSP. In: Boers, E.J.W., Gottlieb, J., Lanzi, P.L., Smith, R.E., Cagnoni, S., Hart, E., Raidl, G.R., Tijink, H. (eds.) EvoWorkshops 2001. LNCS, vol. 2037, pp. 213–222. Springer, Heidelberg (2001)

15. Heegaard, P.E., Wittner, O.J.: Self-tuned Refresh Rate in a Swarm Intelligence Path Management System. In: de Meer, H., Sterbenz, J.P.G. (eds.) IWSOS 2006. LNCS, vol. 4124, pp. 148–162. Springer, Heidelberg (2006)

Appendix A: Diversion Effect Caused by SEoF

This appendix illustrates the possible diversion effect caused by SEoF and its mitigation using the entropy of the local pheromone distributions. This effect is more pronounced when good solutions are hard to find. Hence, it is demonstrated by applying CEAS to solve the `fri26` symmetric static TSP taken from TSPLIB[7] and also used in [4]. Note that CEAS has not been specifically designed to solve the TSP. Such a hard (NP-complete) problem is chosen to stress the performance of the system.

Figure 5 shows the mean value of the cost $L(\boldsymbol{\omega})$ of the sampled paths with respect to the number of tours completed by ants, averaged over 26 runs. Error bars indicate 95% confidence intervals. Parameter settings are similar to those used in [4]. In this case, applying 'SEoF plain' significantly reduces the improvement obtained by using elite selection in terms of convergence speed. Applying 'SEoF entropy', the performance of the system is much less impaired.

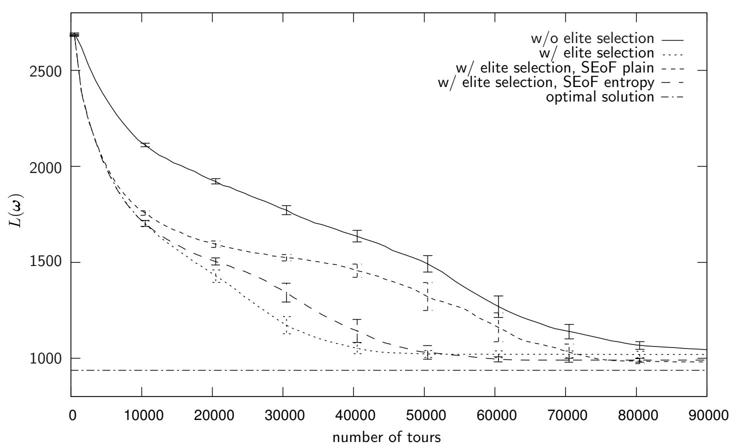

Fig. 5. 26 node TSP example

Appendix B: Fast Adaptation on Reheating Changes

In the simple illustrative example used throughout this paper, the performance of CEAS on reheating changes is worse with elite selection than without, and the performance of the system without elite selection is used as a reference to define stagnation. When the set of candidate solutions is larger, applying elite selection

[7] `http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95`

generally improves the performance of the system also on reheating changes by focusing the search around the best solutions. However, considering stagnation in a broader sense, the causes of stagnation listed in Section 4 still apply and prevent the system from fast adaptation on reheating changes. Now, SEoF is not designed to match the performance of the system without elite selection and it improves the adaptivity of the system in general. This is demonstrated by applying CEAS to find and maintain the minimum cost path between nodes 0 and 9 in the 10 node network depicted in Figure 6.

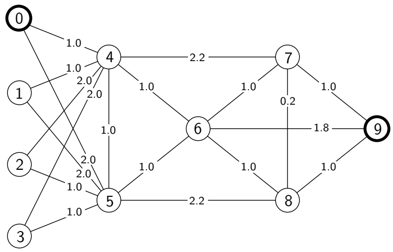

Fig. 6. 10 node network

The cost $L(\boldsymbol{\omega}_t)$ of a path is given by the sum of the delays of each link. Parameters are set as in Section 4. At $t = 0$ [s], all the links have the delay values (in ms) given in Figure 6 and the best path is $\langle 0, 4, 6, 9\rangle$ (3.8 [ms]). At $t = 10000$ [s], the cost of the link between nodes 6 and 9 is increased to 2.2 [ms] (reheating change) and there are then two best paths, $\langle 0, 4, 6, 7, 9\rangle$ and $\langle 0, 4, 6, 8, 9\rangle$ (4.0 [ms]). Figure 7 shows the mean cost of the paths sampled by normal ants averaged over 30 replications. Error bars indicate 95% confidence intervals. Applying SEoF, the time t_{r} to converge after the cost increase is significantly reduced.

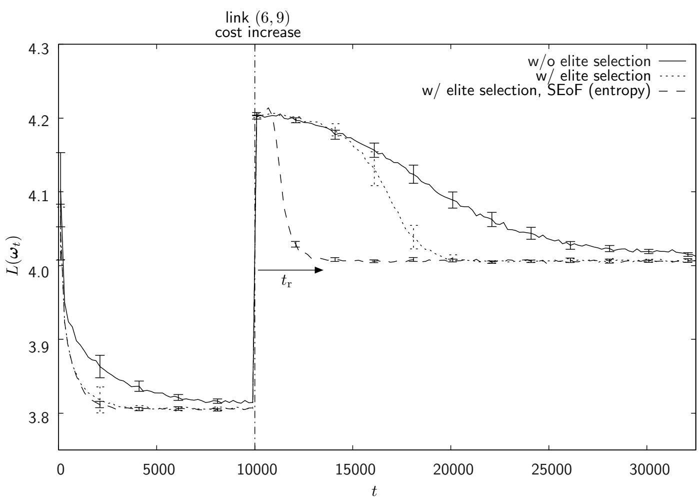

Fig. 7. Average cost of the paths sampled by normal ants