Available online at www.sciencedirect.com

ScienceDirect

Computers & Operations Research 35 (2008) 2049 – 2070

computers & operations research

www.elsevier.com/locate/cor

Genetic local search for multicast routing with pre-processing by logarithmic simulated annealing

M.S. Zahrani[a], M.J. Loomes[b], J.A. Malcolm[a], A.Z.M. Dayem Ullah[c],
K. Steinhöfel[d], A.A. Albrecht[a,*]

[a] *University of Hertfordshire, School of Computer Science, Hatfield, Herts AL10 9AB, UK*
[b] *Middlesex University, School of Computing Science, Hendon, London NW4 4BT, UK*
[c] *Imperial College London, Department of Computing, 180 Queen's Gate, London SW7 2AZ, UK*
[d] *King's College London, Department of Computer Science, Strand, London WC2R 2LS, UK*

Abstract

Over the past few years, several local search algorithms have been proposed for various problems related to multicast routing in the off-line mode. We describe a population-based search algorithm for cost minimisation of multicast routing. The algorithm utilises the partially mixed crossover operation (PMX) under the elitist model: for each element of the current population, the local search is based upon the results of a landscape analysis that is executed only once in a pre-processing step; the best solution found so far is always part of the population. The aim of the landscape analysis is to estimate the depth of the deepest local minima in the landscape generated by the routing tasks and the objective function. The analysis employs simulated annealing with a logarithmic cooling schedule (logarithmic simulated annealing—LSA). The local search then performs alternating sequences of descending and ascending steps for each individual of the population, where the length of a sequence with uniform direction is controlled by the estimated value of the maximum depth of local minima. We present results from computational experiments on three different routing tasks, and we provide experimental evidence that our genetic local search procedure that combines LSA and PMX performs better than algorithms using either LSA or PMX only.

Keywords: Multicast routing; Genetic local search; Simulated annealing; Steiner trees; Quality of service (QoS)

---

1. Introduction

Multicast routing has become an important topic in combinatorial optimisation. A recent overview on multicast routing and associated optimisation algorithms has been presented by Oliveira and Pardalos [1]. The focus of this overview, as in most papers on multicast routing, are on-line algorithms [2,3]. An early summary of problems and technical solutions related to multicast communication was given by Diot et al. [4]. Great effort has been undertaken to incorporate quality of service (QoS) into data communication networks such as ATM and IP networks [5–11]. Many multicast applications, such as video conferencing, distance-learning, and multimedia

---

[*] Corresponding author. Tel.: +44 1707 284 247; fax: +44 1707 284 303.
E-mail address: A.Albrecht@herts.ac.uk (A.A. Albrecht).

0305-0548/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cor.2006.10.001

broadcasting are QoS-sensitive in nature and thus they should benefit from the QoS improvement in the underlying networks.

Designing multicast routing algorithms is a complex and challenging task. Among the various issues involved are the design of optimal routes taking into consideration different cost functions, the minimisation of network load and the avoidance of loops and traffic congestion, and the provision of basic support for reliable transmission. In the present paper, we focus on the design of optimal routes in the off-line mode, as discussed, e.g., in the survey [9] (see Section 2 therein) and [12,13].

The problem of minimising the tree costs of single requests under the constraint that all path capacities are within a user-specified capacity bound, i.e. the requests are executed simultaneously, is referred to as the capacity constrained multicast routing problem (CCMRP) [1,4,14,15].

The CCMRP can be formalised as a constrained Steiner tree problem, which is known to be NP-complete [16]. We note that in applications like video conferencing, multimedia broadcasting, and distance-learning the routing procedure is updated only from time to time, e.g. when new customers register to use one of the services. In such cases, off-line routing algorithms are an appropriate way to solve the routing problem. Since we are dealing with an NP-complete problem, local search methods are a natural choice to tackle the problem; see [9,12,13].

In [17–24], algorithms utilising genetic algorithms (GA) or tabu search are presented. For an overview of search methods, in particular, GA applied to various problem settings in multicast routing, we refer the reader to [25, cf. p. 20–21 therein]. Here, we discuss only a few of the issues raised on this topic. We note that most of the papers are dealing with single trees, but not with routing multiple requests (trees) simultaneously.

The GA proposed in [17,18] assume that several messages all have to be transferred from several sources to multiple destinations, and this has to be executed simultaneously without any order or priority for certain messages. The GA uses a population of chromosomes, where each chromosome is a permutation of the numbers that are assigned to the requests. The algorithms start with a subset of k out of n requests. By using a Steiner tree algorithm, the k requests are routed in the order they appear in the chromosome (partial permutation). Then, to pairs of chromosomes the partially mixed crossover (PMX) operation and the new population (of the same fixed size) is generated by roulette wheel selection, where a sector of a "roulette wheel" is assigned to each offspring whose size is proportional to the fitness measure. The algorithm runs a fixed number of steps, and then k is increased by one in order to check whether $(k+1)$ requests can be scheduled conflict-free. The same procedure is repeated for $(k+1)$ until either $k=n$, or repeated attempts to schedule simultaneously k requests are unsuccessful. The search-based methods from [17,18] are, in part, incorporated into our approach and are discussed in more detail in Section 4.1.

The paper by Ericsson et al. [19] demonstrates a variety of routing problems that can be tackled by GA. The authors apply GA to a routing problem where the link weights are assigned by the network operator, i.e. the problem setting is somewhat different from ours. Then the weight setting problem seeks a set of weights that optimises network performance. Given a set of projected demands, the weight assignment problem, with the objective of minimising network congestion, is NP-hard. The individuals of the population are weight vectors, where the range of components is from 1 to $2^{16}-1$. The crossover operator acts on one elite and one non-elite parent and selects each component of the resulting weight vector according to independently chosen random numbers from (0, 1). The evaluation of the fitness function is rather complicated, since it involves the whole process of routing and the computation of arc loads. The method was successfully tested on the AT&T Worldnet backbone with projected demands, and on several synthetic networks.

Barolli et al. [20] focus on creating a robust path finding solution for mobile ad hoc networks (MANTETs). Since the nodes are mobile, the creation of routing paths is affected by the addition and deletion of nodes, i.e. the topology of the network may change rapidly and unexpectedly. Therefore, QoS is only guaranteed as long as a signal to the node actually exists. The authors propose a genetic algorithm for mobile ad hoc networks (GAMAN) where the network and, respectively, the individuals of the population are represented by trees. The GAMAN algorithm uses the single point crossover and a mutation operation where the "tree junctions" are chosen randomly in the range from zero up to $1/\ell$, for $\ell =$ length of individuals. The algorithm employs the elitist model, where the individual with the highest fitness value in a population is left unchanged in the next generation. The simulation results show that the algorithm is reasonably fast on small to medium size networks.

Yang [21] devised a tabu search algorithm for finding a single, feasible multicast tree efficiently that satisfies a number of QoS constraints. The method is tested on randomly generated networks with 100 nodes (and on 8×8