NOMaD: AN OPTICAL NETWORK OPTIMISATION, MODELLING AND DESIGN TOOLSET

M. C. Sinclair*

Introduction

In recent years there has been a growing interest in the potential of multiwavelength all-optical transport networks to provide the huge bandwidth necessary if broadband services are widely adopted. As part of this effort, this paper provides an overview of a toolset for optical network optimisation, modelling and design (NOMaD) [1, 2]. The toolset is used as part of the author's research into the application of genetic-algorithm (GA)/heuristic hybrid optimisation techniques to optical network design, as well as in several research projects. These include two under the European Commission funded research programme in Advanced Communications Technologies and Services (ACTS): WOTAN (Wavelength-agile Optical Transport and Access Network) and OPEN (Optical Pan-European Network); and the Fujitsu Telecommunications Europe Ltd. "Future Broadband Networks" project. Recent results using GA/heuristic hybrids for the design of multi-wavelength all-optical transport networks are described, including both topology [3], and routing, fibre and wavelength allocation [4, 5].

Overall architecture

NOMaD has been developed as an overall architecture to be extended by several people in a variety of directions to serve the needs of their individual projects [1]. Consequently, its users have not just been given a single unchangeable piece of software, but rather a framework which they can actively add to for themselves. As a result, NOMaD has had to be easy to understand, flexible and extensible. The overall layered architecture is shown in Fig. 1.

To aid both understanding and extensibility, the toolset has been developed using object technology. The analysis and design documented using Booch Notation, and the implementation done in the C++ programming language. This process has been aided by the use of a CASE tool, Rational Rose/C++, which generates skeletal C++ code directly from the diagrams and specifications created by the author.

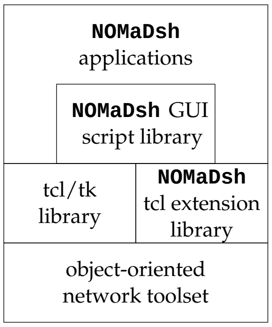

Figure 1: Overall NOMaD architecture

For flexibility, rather than only providing a C++ interface, a NOMaD-specific tcl extension has also been developed. Tcl is a scripting language designed to be extended by the incorporation of application-specific commands, coded in C or C++. Consequently, `NOMaDsh` can be used either interactively to access the underlying network toolset—creating networks, modifying them, assessing their cost/performance, *etc.*—or via scripts written using a combination of tcl and the NOMaD-specific commands. In addition, a library of graphical user interface (GUI) scripts developed using the tk toolkit (tcl/tk) allow `NOMaDsh` scripts to edit and display networks. Further details on NOMaD's tcl scripting layer and GUI can be found in [1].

*M. C. Sinclair is with the Dept. of Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, Essex CO4 3SQ. Tel: 01206-872477; Fax: 01206-872900; Email: `mcs@essex.ac.uk`; URL: `http://esewww.essex.ac.uk/~mcs`

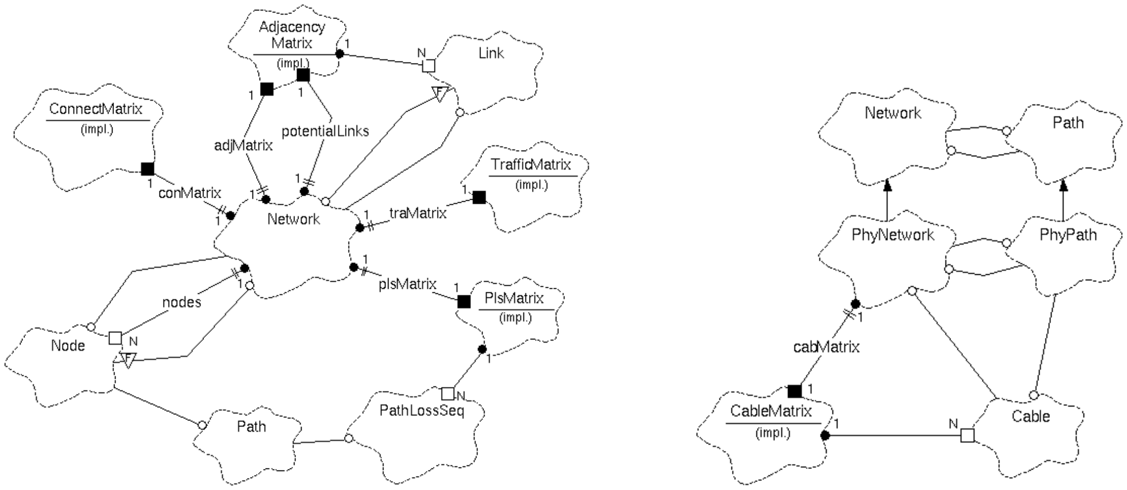

(a) `Network` (b) `PhyNetwork`

Figure 2: Class diagrams

Genetic-algorithm/heuristic hybrids

The simple or 'traditional' GA is based on a simplification of the mechanisms of natural selection and genetics [6]. It uses a population of binary strings, each one of which encodes a potential solution in the problem space. Every generation, the fitness of the strings is determined by decoding them and assessing their worth as solutions; the strings then reproduce in accordance with their fitness, with some strings exchanging part of their genetic material (crossover), and all strings subject to small perturbations at low probability (mutation); the new population of child strings then replaces the old. This process continues until a stopping condition, such as a maximum number of generations, is reached; the solution corresponding to the best string found in the course of a run is taken as the result. However, according to Davis [7]: "Although genetic algorithms using binary representation and single-point crossover and binary mutation are robust algorithms, they are almost never the **best** algorithms to use for any [specific] problem". Another possibility is to adopt a problem-specific encoding, and then devise problem-specific operators that together combine the best existing heuristics for the problem with an overall GA framework. It is almost always possible to develop a problem-specific GA/heuristic hybrid that will outperform either the heuristic(s) or a traditional GA alone, and it is this approach which is used in NOMaD.

Object-oriented network toolset

Clearly, NOMaD's object-oriented network toolset is too large to present in detail here. Consequently, the focus is on those elements that, on the one hand, represent the structures undergoing adaptation, and on the other, implement the overall GA/heuristic hybrid framework itself.

The main `Network` class diagram is given in Fig. 2(a) above, which shows the `Network` class, composed not only of `Node`s, `Link`s, and `PathLossSeq`s (ordered sequences of `Path`s), but also several implementation classes used to represent the `Network`'s adjacency matrix, connection matrix (which records `Link` levels in hierarchical networks) and the traffic requirements. Fig. 2(b) shows the `PhyNetwork` class, derived from `Network`, which uses `PhyPath`s (derived from `Path`) and a matrix of `Cable`s to represent the cable, fibre and wavelength-routed structure of a multi-wavelength optical network. Within NOMaD, rather than using a binary representation, `Network` and `PhyNetwork` objects (for topology and wavelength allocation, respectively) are themselves the structures undergoing adaptation—a highly-structured and very problem-specific representation that requires no decoding or interpretation.

Fig. 3 is the main GA class diagram, illustrating the `GeneticAlgorithm` class, with its `Population` of `Individual`s and its `OperatorPool` of `Operator`s. An `Individual` simply consists of a single `Network` (or `PhyNetwork`) object and some additional accounting information. Operators themselves are created by deriving classes from either `UnaryOp` or `BinaryOp`. In general, those derived from `UnaryOp` are analogous to mutation in a traditional GA, but can be powerful heuristics that utilise arbitrary amounts of problem-specific information. Likewise, `BinaryOp` subclasses are analogous to traditional crossover, although once again, there is considerable potential for problem-specific heuristics.

As in Davis [7], the probabilities of individual operators being applied to an `Individual` adapt through the course of a run using a credit-assignment algorithm, rather than being chosen with a fixed probability throughout [2]. Whenever an `Individual` is created whose fitness exceeds that of the best `Individual` up to the end of the previous generation (`bestIdv`), the improvement in fitness is credited to the `Operator` responsible. In addition, a decreasing fraction of the credit (say, 0.5 and 0.25) is awarded to the parent `Operator` (*i.e.* the `Operator` that created this `Individual`'s parent, `parentOp`) and grandparent `Operator` (*i.e.* `grandParOp`), to ensure that two- and three-`Operator` sequences are rewarded appropriately. After a few generations (say, 4), a small proportion (say, 0.15) of the `Operator` probabilities is reassigned according to the average credit earned per child each `Operator` has generated. In addition, `Operator`s are given a minimum probability (say, 0.05) to ensure they do not lose the opportunity to gain some credit if they, having decayed to the minimum level, are later found to be useful to further evolve the `Population`. An investigation into the effectiveness of NOMaD's operator-probability adaptation algorithm was recently presented in [2].

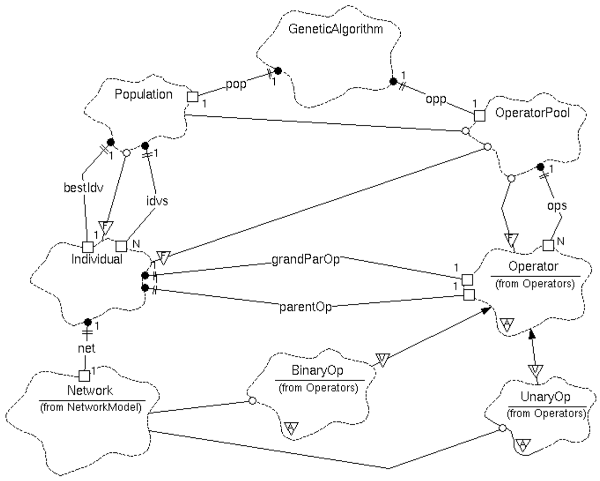

Figure 3: `GA` class diagram

Network topology design

Mesh network topology design starts with a set of node locations and a traffic matrix, and determines which of the node pairs are to be directly connected by a link. The design is guided by an objective function, which allows the 'fitness' of the candidate networks to be evaluated. In [3], NOMaD was applied to the more specific problem of topology design for multi-wavelength all-optical transport networks. The nodes were optical cross-connects, static traffic was used (as is appropriate in a transport network), and two-shortest-node-disjoint-path routing. The restoration routes were fully resourced, and their presence enforced by a penalty incorporated into the objective, thereby ensuring that the network designed would survive the failure of any single component (node or link). The objective itself was a simplified model of network cost [8], based on a summation of node and link costs, but no attempt was made to include fibre or wavelength allocation.

The GA/heuristic hybrids used five operators [3]: single-link mutation, single-point link crossover, link-group crossover, degree-two operator and degree-decrement operator. Link-group crossover selects a group of nodes in one network (those whose distance from a randomly-selected node is less than or equal to the distance to a second randomly-selected node) and then exchanges the links used within that group with those used within the same group of nodes in the other network. It was believed that this would lead to better building blocks than a single-point crossover which assumes a linear encoding of the network topology, as was employed earlier [8]. The degree-two operator alters a network by adding links to nearest nodes, where necessary, making all nodes at least of degree two (*i.e.* connected to two other nodes), to ensure adequate reliability and avoid a cost penalty. The degree-decrement operator simply deletes a random link from the highest-degree node (if there is one of degree greater

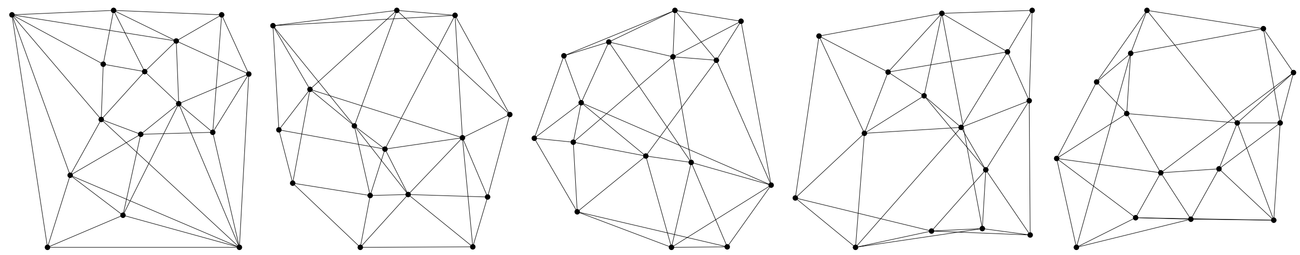

(a) Network 1 (b) Network 2 (c) Network 3 (d) Network 4 (e) Network 5

Figure 4: Network topologies

than two), hopefully thereby reducing the network cost.

To assess the operators [3], five test design problems (node locations and traffic requirements) were generated using the approach described in [9], although further modified to ensure reasonable node separations. Each of the problems had fifteen nodes, used a 1,000 km × 1,000 km area and carried an overall traffic of 1,500 Gbit/s. Three GAs were employed: one used the mutation and single-point crossover operators; the second, mutation and link-group crossover; the third, mutation, link-group and the two problem-specific mutation operators, degree-two and degree-decrement. Overall, link-group crossover was found to be uniformly better than single-point crossover. Surprisingly, however, combining the degree-two and degree-decrement operators with link-group only produced an improvement for one network; in all the others, it actually seemed to reduce performance slightly. The best networks produced [3], or better results obtained since, are illustrated in Fig. 4.

It is hoped in future work to develop more operators for topology design, including some based on branch exchange heuristics (which add, remove and exchange links to improve the objective) plus more powerful local search operators. If earlier, purely heuristic work on topology design for optical mesh networks can be located, then this will also be compared with the GA/heuristic hybrids developed.

Minimum network wavelength requirement design

Another area to which NOMaD has been applied is determining the route, fibre and wavelength to be used by each of the individual traffic requirements in the optical-path layer of a multiwavelength all-optical transport network [4]. For this problem, it was assumed that node (optical cross connect) locations had been selected, and a suitable topology determined, including the number of fibres available in the links (cables) between node pairs. Further, that static traffic requirements needed to be allocated to as many 10 Gbit/s (duplex) wavelength channels as required, each of which would be routed independently on a single wavelength, end-to-end (*i.e.* wavelength-path routing), although restoration paths would not be allocated. The objective function adopted [4] was one of minimum network wavelength requirement (NWR, *i.e.* the number of distinct wavelengths in the network). A low NWR is desirable as a design objective due to the technological limitations, and resulting cost implications, of carrying more than a small number of wavelengths over national or even international distances. However, the actual metric used also included a penalty added for each unsatisfied traffic requirement, to avoid the false minimum of carrying no traffic at all, which would otherwise have an NWR of zero.

The GA/heuristic hybrids used four operators [4]: wavelength-allocation path mutation, wavelength-allocation path single-point crossover, wavelength-allocation path reroute operator and wavelength-allocation path shift-out operator. For a random requirement, path mutation simply reroutes on a randomly chosen k-shortest path (where k is an operator parameter), using the lowest available wavelength on that path (irrespective of fibre choice). Path single-point crossover operates on two networks. The first is modified such that up to a randomly selected individual requirement (considered in order of their end node ids), all the paths remain unchanged, but from then on the paths used are those in the second network, except that the lowest wavelength actually available is used in each case. (However, if any requirement was unsatisfied in the second network, then a random k-shortest path is used instead.) The second network undergoes the complementary recombination. The path reroute operator locates the first path in the network which has the highest wavelength number (found by searching the cables in order of their end node ids), and attempts to reroute it on a lower (and the lowest available) wavelength number using one of the k-shortest paths (where k is an operator parameter, and the shortest path is selected if there is more than one at the lowest available wavelength number). It is based on Step 2a in Wauters and Demeester's HRWA algorithm [10]. The path shift-out operator locates the first path in the network which has the highest wavelength number, and attempts to shift it to a lower wavelength number by shifting out (rerouting) all the paths blocking that particular wavelength number. The wavelength chosen is the one that results in the smallest increase in overall path length, given the constraint that all rerouting must be to lower wavelength numbers than the original wavelength of the path being shifted down. It is based on Step 2b in HRWA [10].

To provide comparative benchmarks [4], four recent wavelength allocation algorithms were selected: a largely traditional GA, although with an embedded allocation heuristic [11], and three recent heuristics [12, 10]. One of these, Wauters and Demeester's HRWA algorithm [10], provided the inspiration for the reroute and shift-out operators. To assess the different approaches, it was decided to use both the network in [11] and five test networks. The latter were those obtained from the network topology design work [3] and illustrated in Fig. 4, with each link (cable) consisting of two fibres. Two GAs were employed: the first used path mutation and path single-point crossover, whereas the second also included the two powerful problem-specific mutation operators, reroute and shift-out. Clearly, the first is close to the traditional GA of [11], whereas the second is a true GA/heuristic hybrid. Overall, the blend of a GA with steps from Wauters & Demeester's HRWA heuristic [10] resulted in both improved performance (over the GA alone) and improved results (over the HRWA heuristic alone). In addition, the hybrid results matched those of the best heuristic as yet implemented by the author.

Further work in this area will include comparing the hybrid with other wavelength allocation heuristics, as well as developing more operators (perhaps inspired by those same heuristics!).

Minimum cost wavelength allocation

Although in [4], low network wavelength requirement was considered to be desirable as a design objective, in [5] it was argued that an objective that consists of a simplified assessment of the cost of the whole network including, at an appropriate level, the cost of exceeding the minimum network wavelength requirement, is superior. The cost model adopted was an extension of that used in [8], with elements to allow for multiple fibres, non-linear dependence on link capacity, and the relative importance of minimising link wavelength requirements [5]. While a penalty was included for unsatisfied traffic requirements, allocation of restoration paths was not. The same three wavelength allocation heuristics as in [4] were adopted as benchmarks, and NOMaD made use of the same set of four operators.

For the relative assessment [5], the same five test networks were adopted (see Fig. 4). Both the GA/heuristic hybrid and the three heuristics were applied with three different levels of importance given to link wavelength requirements. Overall, the hybrid was able to obtain the lowest network costs both with the cost model set for a realistic dependency on link wavelength requirements, and also when only taking account of required capacity; even with the dependency on link wavelength requirements at their highest, the hybrid was able to match the results of the best heuristic yet implemented by the author.

It is hoped in the future, as well as developing more wavelength allocation operators and comparing with a wider set af earlier heuristic approaches, to also investigate the impact of non-linear cost models on the results obtained.

Conclusions & further work

An optical network optimisation, modelling and design toolset has been described, as well as its successful application to three optical network problems: network topology design, minimum network wavelength requirement design, and minimum cost wavelength allocation. Overall, the genetic-algorithm/heuristic hybrid approach appears to perform well in this domain, at least for the problem sizes attempted. Further work, as well as exploring the issues outlined above for the individual problems, should also address the issue of scalability. Additional possibilities include combining both topology and wavelength allocation into a single design process, and a re-implementation of NOMaD in Java to obtain the benefits of multi-threading, an integrated graphical user interface, and portability.

Acknowledgements

The author is grateful to Prof. Rob Massara for his guidance of my Ph.D. project, and to Katerina Proestaki and Noel Parnis (all three from E.S.E., University of Essex) for many interesting discussions during the continuing development of NOMaD.

References

1. Sinclair, M.C.: "NOMaD: Initial architecture of an optical network optimisation, modelling and design tool", Proc. 12th UK Performance Engineering Workshop, Edinburgh, UK, September 1996, pp. 157–167
2. Sinclair, M.C.: "Operator-probability adaptation in a genetic-algorithm/heuristic hybrid for optical network wavelength allocation", Proc. IEEE Intl. Conf. on Evolutionary Computation (ICEC'95), Anchorage, Alaska, USA, May 1998, pp. 840-845
3. Sinclair, M.C.: "NOMaD: Applying a genetic-algorithm/heuristic hybrid approach to optical network topology design", Proc. Intl. Conf. on Artificial Neural Networks and Genetic Algorithms (ICANNGA'97), Norwich, UK, April 1997, pp. 301–305
4. Sinclair, M.C., "Minimum network wavelength requirement design using a genetic-algorithm/heuristic hybrid", *Electronics Letters* v34 n4 February 1988 pp. 388–389
5. Sinclair, M.C.: "Minimum cost routing and wavelength allocation using a genetic-algorithm/heuristic hybrid approach", Proc. Sixth IEE Conf. on Telecommunications, Edinburgh, UK, March/April 1998, pp. 67–71
6. Goldberg, D.E.: "Genetic algorithms in search, optimization and machine learning" (Addison-Wesley, 1989)
7. Davis, L.: "Handbook of genetic algorithms" (Van Nostrand Reinhold, 1991)
8. Sinclair, M.C.: "Minimum cost topology optimisation of the COST 239 European optical network", Proc. Intl. Conf. on Artificial Neural Networks and Genetic Algorithms (ICANNGA'95), Alès, France, April 1995, pp. 26–29
9. Griffith, P.S., Proestaki, A. & Sinclair, M.C.: "Heuristic topological design of low-cost optical telecommunication networks", Proc. 12th UK Performance Engineering Workshop, Edinburgh, UK, September 1996, pp. 129–140
10. Wauters, N. & Demeester, P.: "Design of the optical-path layer in multiwavelength cross-connected networks", *IEEE Journal on Selected Areas in Communications* v14 n5 1996 pp. 881–892
11. Tan, L.G. & Sinclair, M.C.: "Wavelength assignment between the central nodes of the COST 239 European optical network", Proc. 11th UK Performance Engineering Workshop, Liverpool, UK, September 1995, pp. 235–247
12. Wuttisittikulkij, L. & O'Mahony, M.J.: "A simple algorithm for wavelength assignment in optical networks", Proc. 21st Eur. Conf. on Opt. Comm. (ECOC'95), Brussels, 1995, pp. 859–862