Chapter 6

Discussion & Conclusions

Now all has been heard; here is the conclusion of the matter: Fear God and keep his commandments, for this is the whole duty of man.

Ecclesiastes 12:13, NIV

This chapter reflects on the limitations and, to some extent, the successes of the research presented in the thesis, makes suggestions for further work[1], records the author's conclusions, and lists the contributions of the thesis.

The chapter is divided into six sections. The genetic-algorithm (GA)/heuristic hybrid approach in general is discussed in §6.1, while §6.2 focuses particularly on operator-probability adaptation. The following two sections examine NOMaD's two main application areas: mesh network topology design (§6.3) and routing, fibre choice and wavelength allocation (§6.4). The key conclusions reached in both the earlier chapters and this one, as well as the author's overall conclusions, are recorded in §6.5 and, finally, §6.6 repeats the list of contributions of the thesis first provided in §1.7.

6.1 GA/Heuristic Hybrid Approach

Although the GA/heuristic hybrid approach implemented in NOMaD and presented in this thesis has been reasonably successful for the particular design problems and network sizes considered, there nevertheless remains considerable potential for further work exploring variations to the algorithms and operators examined.

One area is creation of the initial populations. For mesh network topology design, initialisation has been with the random links operator (see p. 80), which inserts a link between each node pair

[1] Additional discussion and suggestions for further work *wrt.* the NOMaD software have been recorded in §E.8

in the network with probability 0.5. Now, with a network of n nodes, the initial expected number of links is $0.5 \times n(n-1)/2$ links, *i.e.* for a 15-node network, 52.5 links. However, Networks 1–5 (§5.1.2 and Fig. 5.3), designed during the initial topology experiments, evolved to have between 36 and 40 links. This reduction in the number of links in a 15-node network during a GA run is further demonstrated in Fig 5.2 on p. 96. Consequently, it may be that initialising the networks with a smaller number of links, similar to the number actually required in the final design, would accelerate the early evolution of the topologies. This could be achieved by adding a link-probability parameter to the random links operator (and for 15-node networks setting it to, say, 0.36). Such a parameter might become even more necessary in larger networks, where otherwise the expected number of links in the initial random population will grow $O(n^2)$, whereas the final designs might be expected to have a number of links determined more by node degree, *i.e.* $O(n)$.

Similar considerations apply to routing, fibre and wavelength allocation, where the initial population is currently created using the random k-shortest paths operator (p. 82). Other possibilities might include seeding the initial population with the results from an heuristic allocation algorithm; or using a random k-lowest hop count routing operator.

Another area meriting examination is the choice and parameter(s) of the selection function. Throughout this thesis, as was explained in §4.2.1, deterministic tournament selection has been used[2], with a tournament size of 4, as there is strong support for tournament selection in the literature [304,305]. However, in future work, initially the influence of tournament size, and subsequently the use of other selection functions, could be explored to identify any improvements over the results reported here.

An alternative to the generational model currently used by NOMaD would be the adoption of *steady-state reproduction* [9,308]. In this approach, rather than the old population being entirely replaced by the new one, instead only a small number of individuals are replaced each iteration of the GA. Typically, only the one or two worst individuals are replaced and, in addition, offspring are often checked before insertion and, if they are identical to existing members of the population, discarded (*steady-state without duplicates*). The main advantage of steady-state reproduction is that superior individuals in the population are retained, and can continue to contribute to evolution, until better individuals replace them. However, whether this potential would be realised in the context of the design problems examined in this thesis has not been established. An additional benefit, given the large memory requirements of the GA/heuristic hybrid apporoach, is that only a single population

[2] In deterministic tournament selection, the best individual in the tournament always wins.

must be held in memory at once, rather than both an old and a new. Consequently, either larger populations or larger networks could be addressed within the same memory constraints.

A novel population dynamic, the *stud GA*, was recently introduced by Khatib & Fleming [309]. They propose allowing only the best individual to mate (*the stud*), applying crossover between the stud and *all* members of the population, and when the population diversity starts to drop, using mutation. While such an approach could not be applied unmodified to a GA/heuristic hybrid, with its many operators, the idea of only mating the best individual may have advantages in a context where powerful local-search/heuristic operators are being utilised. Explicitly recombining the fittest individual with each member of the population in turn before randomly choosing an heuristic operator to apply to the offspring may result in more rapid evolution, and this merits investigation.

Whilst the development of a distributed multi-population approach has already been mentioned, one additional possibility would be to employ different genetic operator combinations in the different sub-populations (*heterogeneous distributed GAs*). This might lead to further improvements beyond those resulting from both the inherent parallelism of the distributed approach and the possible niche/island effect of separate evolution in the different sub-populations. Schnecke & Vornberger [257] successfully used such an approach in their hybrid GA for constrained placement problems. Also, a recent paper by Herrera & Lozano [310] examines several such heterogeneous distributed approaches for real-number representations, establishing their superiority over sequential and homogeneous distributed GAs in this context.

Another area for future work lies in recognising the potential for multi-objective optimisation in the design problems addressed in the thesis. All of the metrics for design assessment in §4.1 combine different elements into a single overall objective, although their relative impact can be adjusted by parameters. Instead, NOMaD could be extended to allow metrics with multiple objectives to be defined and used to guide the evolutionary process [275–277]. In practical network design, it is usually neither possible nor desirable to combine all the objectives and constraints into a single overall metric [10], as the final choice of design may be negotiated and involve factors that cannot easily be included. Consequently, developing a multi-objective approach would enhance the applicability of the work to practical optical network design.

6.2 Operator-probability Adaptation

In §5.3 both the basic operator-probability adaptation algorithm and two of the relevant variants from §4.2.4 were examined in the context of minimum-cost wavelength allocation. Overall, it was concluded that, at least for this design problem and the network sizes examined, the main benefits of operator-probability adaptation are in relieving the experimenter of the burden of setting initial probabilities and in the early performance of the GA/heuristic hybrid, rather than in improvements of the final solution quality obtained. Clearly, however, there are a number of avenues of further work that could be explored.

Both the basic and variant adaptation algorithms could be tested under a wider variety of conditions: with varying network sizes and for different design problems, starting with minimum NWR design and mesh network topology optimisation. Indeed, whilst there is merit in experimenting with the adaptation mechanism in the context of engineering design problems, additional insight might also be gained by using some of the 'standard' evolutionary algorithm test functions [2, 309] or a test-case generator [311] for a systematic study.

A different class of credit-assignment algorithm could also be investigated. Instead of awarding credit equal to the actual improvement in fitness (over the best-so-far, best-in-window, parental or median fitness), a constant amount of credit, say 1 unit, could be awarded each time there was an improvement[3]. This approach would allow operators that were making small improvements in fitness to receive equal reward to those making large improvements. While this may seem counter-productive, it may well be beneficial where fitness improvement events are comparatively rare, and a single large improvement in the credit period (ε) could cause one operator to dominate another one that was nevertheless contributing many small improvements.

Currently, all the credit-assignment algorithms examined reward operators when they produce improved children. However, a future possibility would be to also penalise them when they fail to generate improvements, particularly for those algorithm variants which depend on parental fitness [261].

Further, although the present credit-assignment mechanism rewards parent and grandparent operators, thereby recognising that several operators in a chain may be responsible for an eventual improvement in fitness, these chains of operators are not explicitly tracked by the algorithm. Consider, for example, a GA/heuristic hybrid with three operators, a, b and c. With the present

[3] Suggested to the author by George D. Smith of the University of East Anglia.

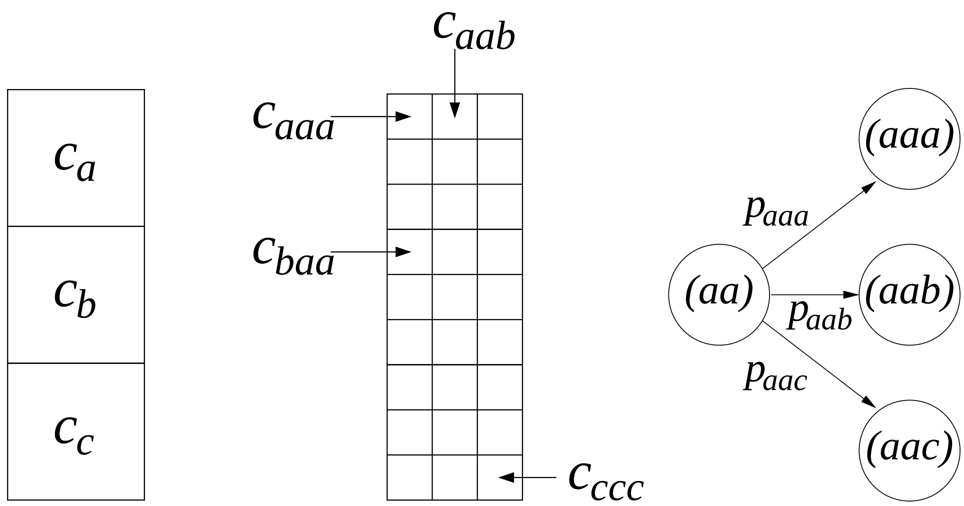

(a) Current algorithm (b) Chain credit (c) Chain probability

Figure 6.1: Operator-chain credit assignment

credit-assignment algorithm, three variables are maintained (Fig. 6.1(a)) to accumulate the credit awarded to the three operators, c_a, c_b and c_c. With the proposed operator-chain credit assignment[4], 27 variables would be used (Fig. 6.1(b)), ranging from c_{aaa} to c_{ccc}. If an operator chain (abc), *i.e.* consisting of grandparent operator a, parent b and then operator c, produced an improvement, the credit would be accumulated in variable c_{abc}. Distinct conditional operator-chain probabilities, derived in the usual way from the accumulated credit, would then be used to choose the operators applied. For example in Fig. 6.1(c), when selecting an operator to apply to an individual whose grandparent and parent operators were (aa), the three conditional probabilities p_{aaa}, p_{aab} and p_{aac} would be employed[5]. While complex, this approach may allow more accurate tracking of successful operator chains, leading to improved performance by a GA/heuristic hybrid utilising this operator-probability adaptation mechanism.

One of the difficulties in designing an effective conventional heuristic is not only must the appropriate operators (steps, perturbations, *etc.*) to incorporate be selected, but also their sequence of application must be determined. One possible approach is to exhaustively try all possible combinations of operators and sequences on some trial design problems. This was attempted by one of the author's M.Sc. students [312, 313] for optical mesh network topology design. Clearly though, such an approach is necessarily limited due to the enormous number of potential operators and se-

[4]This approach was originally inspired by a question from Riccardo Poli of the University of Birmingham.
[5]This treatment assumes a, b and c are unary operators; for binary, the average of the conditional probabilities applying to the two individuals could be substituted.

quences for realistic problems. However, the operator-probability adaptation mechanism described in this thesis may be able to provide insight into both the relative effectiveness and also, if the operator-chain approach was adopted, sequencing of heuristic operators. Consequently, a GA/heuristic hybrid could be developed side-by-side with a conventional heuristic. The operator probability results would allow steps in the heuristic to be evaluated as operators within the GA. This could then lead to the development of a more advanced overall heuristic, using the operators and sequencing identified. Further experimentation with this advanced overall heuristic, which would necessarily be more straightforward to trace and analyse than a GA, might provide inspiration for the development of further operators. These could in turn be assessed in the GA, and so on. It is likely, given the memory and computer time requirements of a population-based approach, that while a conventional heuristic developed in such a context would probably not produce results of the same quality as the GA for small problems, nevertheless it might well be superior to the GA for larger problems, where memory constraints or computation time would disqualify the GA.

6.3 Mesh Network Topology Design

Three studies exploring different approaches to mesh network topology design were reported in §5.1: initial experiments established the superiority of a problem-specific form of crossover; work on local search operators showed the merits of incorporating a flexible local search operator combined with trial-based probability reassignment; finally, the power of iterative heuristic operators was demonstrated, although not without cost in computational terms.

Clearly, future work could examine a wider range of network design problems, perhaps exploring additional heuristic operators for optical mesh network topology design, or exploiting some of the potential improvements in the NOMaD framework itself. Another possibility would be to try and gain additional insight into the design problem itself by characterising the fitness landscape under the different operators [314]. However, a more productive avenue for research may be to apply the GA/heuristic hybrid approach to a different range of network topology design problems.

In those GAs used in §5.1 for mesh network topology design, while the individual links included in the network are selected by the GA, their capacity is only indirectly determined by it. Instead, during assessment by Network Cost Model A (§4.1.1), the traffic requirements are first routed across the network, and then the link capacities set to whatever level is needed to provide reasonable occupancies. In addition, the capacity of the links is costed by steps fitting a continuous function, rather than by arbitrary costs for a fixed set of capacities. However, it would be relatively straightforward to extend NOMaD to include such explicit capacity assignment and arbitrary costs in the evolutionary process (*cf.* Deeter & Smith [148], who used arbitrary costs for a fixed set of link reliabilities).

Although the topology designs in this thesis have been flat mesh networks, nevertheless more constrained topology classes may provide benefits in lower provisioning and maintenance costs, or in simplified network management [19]. One possibility is hierarchical networks, where nodes and links are each associated with one or more discrete levels. NOMaD's network model already provides flexible support for hierarchical networks with up to 32 levels (see §E.4), although this potential has not been used in the author's own work[6]. By developing suitable operators and metrics, the GA/ heuristic hybrid approach could also be applied to the topology design and allocation of hierarchical networks, and the results compared to those here.

A more focused topology class again is that of dual-homing multi-ring hierarchical networks. For example, Fig. 6.2 shows a two-level network of 16 nodes. The nodes have been grouped into four rings at level 1 (R_{11}, R_{12}, R_{13} and R_{14}). From each of the level 1 rings, two nodes have then been selected to form the higher level interconnecting rings (hence *dual-homing*). There are two of these at level 2 (R_{21} and R_{22}), with one of the common nodes from each level 1 ring included in a different higher-level ring, an interconnection strategy referred to as *different common ring* (DCR). One of the author's former Ph.D. students, Katerina Proestaki, has developed several topology design, routing, fibre and wavelength-allocation heuristics for this class of networks, all implemented on NOMaD [20–24]. An interesting area of future work would be to employ the steps in these heuristics as operators within GA/heuristic hybrid algorithms, hopefully providing both improved design results, at least for small networks, as well as additional insight into this class of problems.

6.4 Routing, Fibre & Wavelength Allocation

Routing, fibre choice and wavelength allocation were studied in §5.2. Two metrics for network assessment were used, one based on NWR, and the other on a simplified model of network cost. In both cases, the GAs were compared with recent wavelength-allocation heuristics. For minimum NWR, the GA/heuristic hybrid approach fulfilled its promise, as the blend of a GA with steps from

[6] Although it has been used by Noel Parnis.

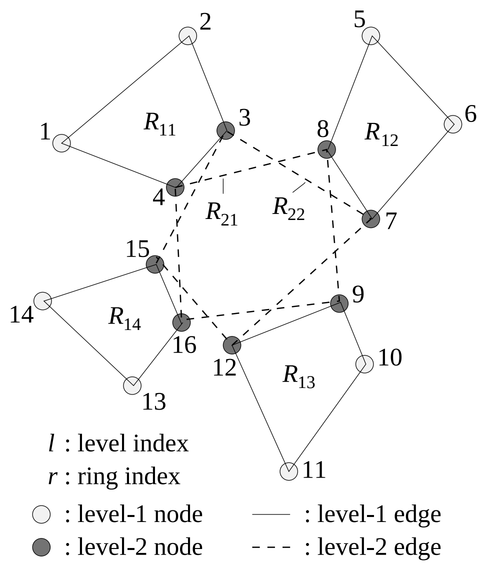

Figure 6.2: Dual-homing multi-ring network with DCR interconnection [20]

one of the heuristics resulted in both improved performance (over the GA alone) and improved results (over the heuristic alone). In addition, the GA results matched those of the best wavelength-allocation heuristic examined. Similar success was achieved with the minimum cost allocation, as in competition with three recent heuristics, the GA/heuristic hybrid approach was able to obtain the lowest network costs for two settings of the cost parameters (including the one judged most realistic), and in a third case, was able to approach the results of the best of the allocation heuristics.

In future, as well as exploring a wider range of network sizes, comparisons could be made with more of the allocation heuristics from the literature including, for example, Nagatsu *et al.*'s wavelength-path (WP) allocation heuristic [296], which is already implemented in NOMaD (see Appendix C), an interesting one by Baroni & Bayvel [300], and a recent one by Alanyali & Ayanoglu [315] that can be adapted to take account of network restoration.

The extension of NOMaD to allow the representation of restoration paths for fibre and wavelength allocation is suggested in §E.8.3. Clearly further experiments could then be performed in this area, and comparisons made with allocation heuristics that also include restoration paths, *e.g.* Nagatsu *et al.* [297, 302]. Likewise, the modelling of restrictions on *e.g.* the number of wavelengths on the fibres could lead in turn to comparisons with similarly designed heuristics, such as others by Nagatsu *et al.* [302, 303].

While the focus in the thesis has been on WP routing, without considering the more straightforward problems of virtual-wavelength-path routing (VWP; see §1.2), a further possibility remains. This is sparse-wavelength-path (SWP) routing, where although wavelength conversion is available in the network, it is only available at some nodes, for example, or where the number of conversions allowed in the network is constrained in some other way [316]. This approach confers some of the flexibility of VWP, particularly during failure conditions, without incurring the high cost and potential transmission limitations of the unrestricted use of wavelength conversion. Appropriate development of NOMaD would allow the GA/heuristic hybrid approach to be applied to this area as well.

The support of NOMaD's network model for both network hierarchy and node clustering is outlined in §E.4, and the potential application of these to topology design of hierarchical networks in general, and dual-homing multi-ring networks in particular, was suggested above. Clearly, it would also be worthwhile to apply NOMaD to the routing, fibre and wavelength allocation of such networks.

In this thesis, network topology design and routing, fibre and wavelength allocation have largely been considered separately, with the network topology being established before the fibre and wavelength allocation is performed. However, it may well be beneficial to combine both into a single integrated design problem (*cf.* for example work by Schnecke & Vornberge [257], who found notable benefits from such an integrated approach in the context of constrained placement problems). This would allow the topology designed to be directly influenced by the allocations, which may well result in a better overall network assessment (*e.g.* lower cost) than tackling the two areas sequentially. For example, in Tables 5.17–5.22 on p. 115, the network costs obtained for Network 6 are always consistently below those for Network 1. However, while both networks have the same node distribution and experience the same traffic demand (see §5.1.2 on p. 98), the topology of Network 1 was obtained with $\alpha = 1$ in Network Cost Model A and that of Network 6 using $\alpha = 0.8$. Consequently, after wavelength allocation, while it might have been expected that Network 6 would have the lower cost with $\alpha = \beta = 0.8$ in Network Cost Model B (Tables 5.20–5.22), Network 1 should have been better for $\alpha = \beta = 1$ (Tables 5.17–5.19). This was not the case, however, and so there is clearly 'room for improvement' that could be exploited by an integrated design approach[7].

6.5 Conclusions

The introduction (Ch. 1) explained the overall direction of the thesis: the application of the genetic-algorithm (GA)/heuristic hybrid approach to two of the main problems in the design of multi-

[7]The author is grateful to Alan Hill of Ilotron Ltd. for a useful discussion on this topic.

wavelength all-optical transport networks: mesh network topology design, and wavelength-path (WP) routing, fibre choice and wavelength allocation. It established the level of background knowledge expected of the reader and briefly discussed the choice of a flat, all-optical, arbitrary-mesh multiple-fibre-link topology with WP routing as the problem domain selected. The background and motivation behind the solution approach of evolutionary algorithms (EAs) in general, and GA/heuristic hybrids in particular, was given, before the overall aims and objectives of the thesis were elaborated.

The chapter on evolutionary algorithms (Ch. 2) included a brief overview and historical perspective on EAs, a tutorial introduction to GAs, as well as three literature surveys. The overview considered early work in the field, GAs, evolutionary programming, evolution strategies, genetic programming, classifier systems and evolutionary hardware, as well as conferences and books on the area. The tutorial provided background on both GAs in general and Goldberg's Simple Genetic Algorithm in particular before introducing in more detail a central theme of the thesis, GA/heuristic hybrids. The first survey gave near-exhaustive coverage of evolutionary telecommunications, *i.e.* the application of EAs to the telecommunications problem domain, up to December 1998. It concluded with some comments on promising directions for future research and encouragement to future authors in regard to both citation of earlier work and the reproducibility of their research. The two remaining sections gave more limited surveys of prior work on GA/heuristic hybrids and operator-probability adaptation, respectively, in areas other than evolutionary telecommunications. Overall, the chapter gave the context for the evolutionary aspects of the author's research, as well as establishing the foundation on which it was built.

The chapter on NOMaD (Ch. 3) provided an overview of the software toolset, including its application areas, design metrics and supported solution approaches. The overall NOMaD architecture, its object-oriented network toolset, interactive scripting capabilities, graphical user interface, testing & verification strategy and an example compiled application were all discussed.

The chapter on GA/heuristic hybrid design (Ch. 4) presented greater detail of NOMaD's design metrics and evolutionary algorithm framework. It is these elements, implemented in the NOMaD software toolset, that encapsulate the author's GA/heuristic hybrid approach to multi-wavelength all-optical transport network design. Details of the design metrics used by the author for network assessment were presented: one for minimum-cost mesh network topology design, another for network wavelength requirement (NWR), and the final one for minimum-cost wavelength allocation. Moreover, the evolutionary algorithm framework was described, including the overall structure of the GA/heuristic hybrid; operators for network topology design; operators for routing, fibre and wavelength allocation; and the operator-probability adaptation mechanism.

The experimental results chapter (Ch. 5) described a number of experiments exploring the suitability of the author's GA/heuristic hybrid approach for multi-wavelength all-optical transport network design. Three studies exploring different approaches to mesh network topology design were reported: initial experiments established the superiority of a problem-specific form of crossover; work on local search operators showed the limited merits of incorporating a flexible local search operator combined with trial-based probability reassignment; finally, the power of iterative heuristic operators was demonstrated, although not without cost in computational terms. In addition, the initial topology approach was used to design the test networks that were to be used in the next group of experiments.

GA/heuristic hybrids for routing, fibre choice and wavelength allocation were then investigated. Two metrics for network assessment were used, one based on NWR, and the other on a simplified model of network cost. In both cases, the GAs were compared with recent wavelength-allocation heuristics published in the literature. For minimum NWR, the GA/heuristic hybrid approach was shown to be as promising as had been hoped, as the blend of a GA with steps from one of the heuristics resulted in both improved performance (over the GA alone) and improved results (over the heuristic alone). In addition, the GA results matched those of the best wavelength-allocation heuristic examined. Similar success was achieved with the minimum cost allocation, as in competition with three recent heuristics, the GA/heuristic hybrid approach was able to obtain the lowest network costs for two settings of the cost parameters (including the one judged most realistic), and in a third case, was able to approach the results of the best of the allocation heuristics.

The final investigation was of NOMaD's operator-probability adaptation mechanism, both in its basic form, where credit is assigned for improvements over the best-so-far individual, as well as two variants that reward improvements over parental and median fitness. At least in the application examined, minimum cost allocation which uses powerful problem-specific operators, the main benefits of operator-probability adaptation were found to be in relieving the experimenter of the burden of setting initial probabilities and in the early performance of the GA/heuristic hybrid, rather than in improvements of the final solution quality obtained.

Earlier sections in this chapter reflected on both the successes and, to a greater extent, the limitations of the work described in this thesis, outlining from these the potential for future research. In §6.1 and §6.2, the GA/heuristic hybrid approach in general was discussed, with operator-probability adaptation considered separately. Future work could enhance initialisation, selection and population dynamics, implement a distributed multi-population GA, or employ multi-objective optimisation. The credit-assignment mechanism could be extended in a variety of ways, including explicit reward of operator chains. In addition, the potential of the GA/heuristic hybrid approach as a testbed for both operators and their sequences within conventional heuristics was also recognised.

In §6.3 and 6.4, NOMaD's two main application areas, mesh network topology design, as well as routing, fibre choice and wavelength allocation, were explored. Whilst recognising the success of the GA/heuristic hybrid approach in the experimental results, a number of directions for future work were nevertheless identified. Addressing hierarchical networks in general, and dual-homing multi-ring networks in particular, was recommended for both topology and wavelength allocation. Moreover, for the latter area, restoration-path allocation, sparse-wavelength-path routing, and the integration of both topology and wavelength allocation into a single design process are all future possibilities.

Overall, the application of the GA/heuristic hybrid approach to optical network design can be considered a success. In particular, problem-specific representation and the incorporation of problem-specific heuristics proved successful for both mesh network topology design, as well as routing, fibre and wavelength allocation in multi-wavelength all-optical transport networks. In many cases, a hybrid GA provided not only better performance than a GA alone, but also better results than conventional heuristics. However, some limitations of the approach were identified, especially the large memory and processing requirements. In addition, the operator-probability adaptation mechanism developed did not provide the expected improvements in result quality, although it still conferred some notable benefits. Finally, several of the questions posed in §1.5 remain open, and in particular whether object-oriented representation is really to be preferred to binary (or other) encoding, but these, and many other questions raised by the research, must remain topics for further work.

6.6 Contributions of the Thesis

The key contributions of the thesis (repeated from §1.7) are:

- A tutorial introduction to GAs and GA/heuristic hybrids (§2.2).
- A near-exhaustive survey of evolutionary telecommunications, *i.e.* the application of EAs to the telecommunications problem domain, up to December 1998 (§2.3) [8].

- The design of the network model and GA/heuristic hybrid framework incorporated in NOMaD (§E.3 and §4.2) [31, 32].
- NOMaD's operator-probability adaptation mechanism (§4.2.4 and §5.3) [33].
- The GA/heuristic hybrid algorithms for optical mesh network topology design, particularly the new link-group crossover operator (§5.1) [34].
- The GA/heuristic hybrid algorithms for minimum network wavelength requirement routing, fibre-choice and wavelength allocation, particularly the new problem-specific operators: wavelength-allocation random k-shortest paths, k-shortest path mutation, path re-route, path shift-out and k-shortest path single-point crossover (§5.2.1) [35].
- The GA/heuristic hybrid algorithms for minimum cost routing, fibre choice and wavelength allocation, using the same new problem-specific operators (§5.2.2) [36, 37].