Case for support: An investigation of hyper-heuristic methods

Peter Ross, Emma Hart
School of Computing
Napier University
219 Colinton Road
Edinburgh EH14 1DJ

Edmund Burke, Peter Cowling, Graham Kendall, Sanja Petrovic
Automated Scheduling, Optimisation and Planning Research Group
School of Computer Science & Information Technology
University of Nottingham
Jubilee Campus, Nottingham NG8 1BB

Previous Research and Track Record

Note: Peter Ross and Emma Hart have since moved from the University of Edinburgh to Napier University. The AI Applications Institute at the University of Edinburgh was created in 1984 to transfer AI-related technology from the university into practice, working in partnership with a wide range of international businesses and organisations. These include NASA, the US Department of Defense, the RAF and DERA in the UK, Unilever, Hitachi, Pindar and many others. Its has been very successful – for example, the 'Expert Provisioner' system prototyped for and now used by RAF Logistics has, according to its users, saved them over £100 million so far; and the Yellow Pages books are laid out by software developed at AIAI for Pindar. The Evolutionary Computation group at the University of Edinburgh has been carrying out research into many aspects and applications of evolutionary algorithms for the past six years, notably work on timetabling and scheduling where the results have been widely used in practice internationally.

The Automated Scheduling, Optimisation and Planning Research group at the University of Nottingham has been carrying out successful research into the development and application of meta-heuristic techniques and hybridisations to scheduling problems (particularly timetabling) for the last six years, working with several government and commercial organisations. It is internationally recognised within the scheduling community for its work, especially on automated timetabling. The group now has 4 members of academic staff, two research assistants, six PhD students and a secretary. The group has many years of experience, within AI and OR, of meta-heuristic approaches to a variety of scheduling and optimisation problems, and extensive experience of practical and theoretical aspects of the problems that will be tackled in this proposal as well as a wide variety of related problems.

The Research Team

Dr Peter Ross is the Director of AIAI and a senior lecturer within Informatics at the University of Edinburgh, and leads the Evolutionary Computation group there. He was head of the Department of AI until the merger of AI, CS and Cognitive science formed the Division of informatics in 1998. He is an associate editor of the journal Evolutionary Computation (MIT Press) and has been on the programme committee of sixteen international conferences in the area. He has written four books, edited three books of research papers and has published over 60 refereed papers. Recent relevant publications include [TRV99], [HR99], [RHC98], [HRN98a], [HRN98b], [RH98], [HR98], [TWR98].

Edmund Burke has recently been promoted to a Chair in Computer Science at the University of Nottingham. He leads the Automated Scheduling, Optimisation and Planning (ASAP) Research Group at the University of Nottingham and is Editor-in-chief of the Journal of Scheduling (Wiley). He is an Associate Editor on the IEEE Transactions on Evolutionary Computation and is chairman of the steering committee of the international series of conferences on the Practice And Theory of Automated Timetabling (PATAT). He was co-chairman of the programme committee and co-editor of the conference proceedings and Selected Papers volume (published by Springer) for the first two of these conferences (Edinburgh, 1995; Toronto, 1997) and will be for the next one (Konstanz, 2000). Dr Burke is a co-ordinator and chairman of the committee of the EURO Working Group on Automated Timetabling. This group is funded by EURO (The EUropean association of Operational Research societies). He was principal applicant on the proposal that secured this funding. He is a member of the Programme (or refereeing) committees of seventeen international conferences since 1994. He has published over 60 refereed papers during his career and been awarded ten major grants over the last seven years. Recent relevant publications include [BUR99a], [BUR99b], [BUR98a], [BUR98b], [BUR98c], [BUR97], [BUR96a], [BUR96b].

Dr Peter Cowling is a lecturer in Computer Science and a member of the ASAP group at Nottingham. Before joining Nottingham in 1998 he had four years experience in carrying out commercial research into analysis, modelling, optimisation and systems development in scheduling, particularly in relation to the steel industry, at the Université Libre de Bruxelles and AI Systems, Belgium. He has project managed several teams responsible for the mathematical modelling and the design and implementation of heuristics and algorithms for steel industry scheduling, which has resulted in the commercially successful SteelPlanner range of products [BAC95, COW95a], which are currently marketed by IBM consulting and other steel consulting companies, and also has experience in developing commercial solutions for cutting stock problems. He was recently awarded research grants for research into inventory management, timetabling, heuristics and production planning. He has ten refereed papers in print or to appear in international scientific books and journals and contributed to numerous international conferences. He is the book and software review editor of the Journal of Scheduling, was on the programme committee of four recent international conferences, and is a member of the Operational Research Society. Recent relevant publications include [COW99a], [COW99b], [COW99c], [BUR99c], [COW97], [BJO95], [COW95a] and [COW95b].

Dr Sanja Petrovic is a lecturer in Computer Science and a member of the ASAP group at Nottingham. Before this she worked as a Senior Researcher at the Mihajlo Pupin Institute in Belgrade, Yugoslavia. Her main area of research includes the application of artificial intelligence techniques to operational research problems. She is a member of Yugoslav Operational Research Society and Yugoslav Society of Soft Computing and Intelligent Systems, was a member of the Organising Committee of the 9th Mini EURO Conference "Fuzzy Sets in Traffic and Transport Systems" held in Yugoslavia in 1997, and a member of the Organising Committee of the Yugoslav Conference on Operational Research in 1998. She has 7 papers published in distinguished scientific journals and over 20 papers in conference proceedings. She is a permanent referee of the Journal of Multicriteria Decision Analysis. Publications relevant to this proposal include [PET94], [PET95], [PET97], [PET98], [PET99], [RAD97].

Mr Graham Kendall has recently joined the ASAP group as a lecturer in Computer Science. Before this appointment, he worked in industry. He is the author of 5 refereed papers.

Description of the Proposed Research and Its Context

Background

Over the last 40 years or so, academic journals and conference proceedings have contained many novel computer algorithms in areas such as data-mining, cutting, packing, scheduling and other optimisation applications. Most fail to have any significant impact in practice. Typically they are too knowledge-intensive or too complicated for most potential users. Only a very small number of technology-provider businesses can even afford to bring them to market, and only a very few end-user companies have the IT resources and management effort to develop and maintain them for internal use. Often, therefore, users continue to employ simple heuristic methods such as 'largest-first first-fit' in packing tasks, or Brelaz's 'most-difficult-first' algorithm in resource allocation, or the nearest neighbour heuristic in routing problems. Such methods may be relatively straightforward to implement, but individual heuristics do not always perform very well on all problems. Indeed, for a given heuristic it is often possible to engineer a problem or even a type of problem on which its performance is mediocre. There is a serious conflict between using cheap but fragile heuristics on the one hand, and on the other hand using knowledge-intensive methods that can perform very well but are hard to build and maintain.

This proposal concerns *hyper-heuristics*. This is a term we have coined to describe the idea of using a number of different heuristics together, so that the actual heuristic applied may differ at each decision point. In essence, hyper-heuristics are heuristics to *choose* heuristics. This differs from the widely-used term meta-heuristic, which refers to heuristics which control simpler heuristics for a narrow range of problems, rather than the hyper-heuristic approach of choosing between a range of heuristic approaches to robustly solve a wide range of problems. An example of the use of hyper-heuristics is given in [FRC94], in which a genetic algorithm (GA) is applied to open-shop scheduling problems, to evolve the choice of heuristic to apply whenever a task is to be added to the schedule under construction. The process was called 'evolving heuristic choice'. The chromosome is a list of integers; if the i-th gene has value g_i it means to use the g_i-th heuristic at the i-th decision point. The results were very good at that time, and [HR98] improved on them by using different heuristics. Another illustrative example is given in [TRV99], for solving some large real-world exam timetabling problems. In those problems, there are room-capacity constraints as well as the basic constraints that students must not be in two places at once, and there is also a preference that each student's exams be spaced out as far as practicable. The nature of the problems vary; for example, in some problems there are some very large exams and room space is tight, so that it turns out to be best to focus initially on the packing aspect so as to get the awkward exams packed into just a few timetable slots. In other problems space is not tight and such a strategy produces poorer answers. The idea in the cited paper is to construct an algorithm that suits the problem, consisting of two phases. In the first phase one heuristic is used to prioritise the order in which exams are considered and another is used to place each exam. After some threshold is passed, the second phase uses different prioritising and placement heuristics. A GA was used to explore the space of choices of heuristic and threshold. Even for very large problems involving thousands of exams and tens of thousands of students, the system was fast and the results were excellent. Another example is given by very recent work at Edinburgh, to explore the use of hyper-heuristics for tackling vehicle routing problems with time windows, in which customers will only accept a delivery within a specific time-window. The aim is to minimise the number of vehicles necessary as well as the total distance traveled. So far, the results are significantly better than those in [Berger99] but work in this area by the APES group, at Strathclyde and elsewhere (eg [Shaw98]), is likely to provide a better touchstone.

The main hypothesis of this proposal is that hyper-heuristic methods will be appreciably cheaper to implement and easier to use than knowledge-intensive methods and yet deliver better average- and worst-case performance than other heuristic methods. In order to explore this hypothesis, we need a wide-ranging investigation that applies hyper-heuristic ideas to different problem areas, and in a variety of ways, and not only seeks out successes of the approach but also tries to find weaknesses by, for example, trying to construct problems that are awkward for such methods. The examples cited in the previous paragraph all happened to use GAs for heuristic selection, but we do not take it as an article of faith that such evolutionary search methods are the best; often they are not. We do intend to make some use of evolutionary methods as research tools, but we also intend to investigate whether we would recommend non-population-based methods for exploring a space of hyper-heuristics. For example, we suspect that a GA-based approach works well for timetabling and for vehicle routing because, at many decision points, several different heuristics would produce the same decision; and therefore the GA is not searching for a 'needle in a haystack' but for any one of many ways of indirectly representing the same solution. However a GA representation which encodes the choice of heuristic for *every* decision point is clumsy as well as risky because of epistasis; each decision potentially affects the consequences of all following ones. It may be that a better method is to use a modern learning classifier system such as XCS [WIL98], which avoids many of the problems that have bedevilled classifier systems in the past; it has a good credit-assignment scheme, good scaling properties and good control of set size, for example. The idea would be to evolve rules, each of which expresses the problem conditions under which each heuristic should be applied. Conventional classifier systems are rather limited because they cannot express relational conditions; but s-classifiers, in which the conditions are simple s-expressions, can. This could be a much more parsimonious way of expressing hyper-heuristics, arguably akin to how people reason about complex sequential decision tasks, and we are not aware of anyone else having explored this idea as yet. Another important area that we will explore is the extension of the concept of Guided Local Search [VOU99] to consider the choice between a wide range of different heuristics as well as parameter selection for individual heuristics.

The case-based reasoning (CBR) paradigm, where knowledge concerning previously successful solution approaches is reused, has a number of important contributions to classification problems mostly in medical domains [POR90]. Also, we have explored the advantages of using CBR approach in the domain of multicriteria decision making. A CBR system has been developed which assisted in choosing an adequate method for solving a multicriteria problem in hand [PET97], [PET98], [PET99]. Rather than classifying a new problem using necessary and sufficient conditions a CBR system does classification by trying to find the closest matching case in its case base. A large number of existing cases from various domains solved by different meta-heuristic procedures employing different heuristics and parameters can form an initial case base and thus provides a promising starting point. We propose to use case-based reasoning both at a strategic level to provide an initial "rough" matching of appropriate techniques given a problem instance, and also at a tactical/operational level to identify approaches which work well for a region of the search space having a topology similar to an area already explored.

The Variable Neighbourhood Search class of heuristics proposed by Hansen and Mladenovic [HAN97] allow us to use a range of different local search neighbourhood structures, controlled by a very simple mechanism. We have already started to work on a more sophisticated form of heuristic embedding for problems of a routing nature [BUR99d]. Another strand of our research will consider a more sophisticated hyperheuristic to control neighbourhood choice, based upon both on-the-fly testing of the effectiveness of local search neighbourhoods and on an adaptive memory of strategies which have been previously effective. Our primary aim here is to *very* quickly provide adequate solutions to the often large problems arising in practice, where differences between the model and reality make the pursuit of a solution optimal for the model an approach which is inappropriate to generate a solution for the real problem.

Programme and Methodology

The aim of the proposed research is to investigate the possibilities and potential benefits of a hyper-heuristic approach to a range of pseudo-optimisation problems ('pseudo' in the sense that real users rarely care about genuine, provable optimality – 'good enough, soon enough' is the real challenge). To do that:

1. we will find out which methods do and do not make hyper-heuristics work, across a wide range of problems and problem instances, in order to draw general conclusions and advice for the potential user community;
2. specifically, we intend to tackle significant practical problems (to which we already have access) across a range of domains: timetabling, data mining, stock cutting, packing and space allocation, scheduling, in order to allow cross-fertilisation of ideas between these domains.
3. we also intend to explore parameterised families of such problems, in order to gauge the variability in performance, and to use co-evolutionary methods to try to 'break' the hyper-heuristic approach by searching for problems that are especially awkward for it;
4. we will try to demonstrate how quick- and cheap-to-implement knowledge-poor heuristics can be used within a hyper-heuristic framework to provide a methodology suited to fast and cheap development of industrial and commercial systems. This will lead to a prototype hyper-heuristic "toolbox" for the user community.

Work Outline

The proposed research will be carried out at the Universities of Edinburgh and Nottingham. Research will progress in three distinct phases. In the first phase, lasting for sixteen months, the two Universities will develop a wide range of different heuristics for a range of problems. Regular contact between the two sites will have to be maintained to ensure a common architecture for the heuristics developed and that there is no duplication of effort. This first phase will make good use of the complementary expertise of the two sites in different application domains and different heuristic techniques. In this initial phase the work will be divided as follows:

- The Nottingham group will apply hyper-heuristic ideas to problems in cutting, packing and space allocation and in single-machine scheduling with setups (TSP-related) in printed circuit board and steel manufacturing. They will focus on using novel meta-heuristic approaches (such as Variable Neighbourhood Search and Ant systems), hybridising with more established meta-heuristics (such as Tabu search, Simulated Annealing, Genetic Algorithms) and on using case-based reasoning and adaptive memory techniques as techniques for deciding which heuristic to apply at which point of the search process. Using the common, open-ended format for the heuristics developed, the Nottingham group will develop architectures for automatically manufacturing hybrids (and hybrids of hybrids, and so on).

- The Edinburgh group will apply hyper-heuristic ideas to problems in vehicle-routing with time windows, with and without backhauls (benchmark results by other methods are available); and to data mining tasks of classification, in which heuristics can be invoked to (re)process data, determine closeness of points, control complexity and so on. The Edinburgh group will focus on using classifier systems as suggested above, on forms of genetic programming (for example to evolve a ranking function for each heuristic), and on their combination in s-classifiers.
- Both groups will continue to explore a variety of timetabling tasks, using various search methods to try combining heuristic choices. Both groups have considerable experience in this area; timetabling also provides a handy common ground in this phase of the joint project.

This phase offers good scope for producing publications that are targeted on the specific problem areas.

In the second phase, over eight months, the two groups will work closely to consolidate their experience from the first phase, formulate and implement a properly documented common architecture and methodology for hyper-heuristic ideas, intended to be applicable across the variety of different problem areas in this proposal as well as across other domains. Specifically we wish to try to avoid the kind of problem that still hampers the uptake of broad technologies such as genetic algorithms or neural nets, where too many choices are left up to the user and not enough guidance is offered about what to do. The outcome from this phase should be a prototype hyperheuristic system and a well-documented and reasonably general framework for the practical use of hyperheuristics.

In the third phase, lasting twelve months, the groups will refine this common architecture by assessing its performance in their chosen problems areas. This will involve studying how delicate any user choices are, over families of parameterised problems, and trying to 'break' the approach by generating problems that are particularly hard for it, either in terms of user effort or in terms of computational resources. This phase offers a lot of scope for cooperation and friendly competition between the groups. The outcome from this phase should be a well-documented and well-justified framework for hyperheuristics, backed by substantial empirical evidence of its worth across our chosen problem domains. This phase concludes by producing joint reports and publications about the outcome.

Relevance to Beneficiaries

We would hope to cause a fundamental shift in the way that difficult optimisation problems are handled in the real world, not only in the areas that we consider, but also for a much wider range of practically important problems. This is an ambitious aim. The basic idea is particularly motivated by the practical commercial difficulties of using knowledge-rich methods and the relative fragility of cheap knowledge-poor methods, and so is fairly directly concerned with wealth creation. Both groups have continuing links with a good variety of potential end-user businesses, and will seek to keep them involved as the work progresses. AIAI at Edinburgh and ASAP at Nottingham both have good track records of being succesful at this.

Publication of performance data, of best results so far, and of of software will be of direct benefit to researchers and to potential users.

Dissemination and Exploitation

General dissemination of the research outcomes will be by means of academic journal and conference papers, releases of experimental software and instances of appropriate problem test cases via the Internet, and through seminars and talks at other UK universities and at UK and international conferences. We aim to present our research at major heuristics-oriented conferences within the project duration, such as the INFORMS Conference, CEC, GECCO, EURO and the Metaheuristics International Conference. In addition, we will present our work at more problem specific conferences such as the Conference on the Practice And Theory of Automated Timetabling, the UK Planning & Scheduling SIG and so on. We would also submit our research results to relevant international scientific journals.

A web page will be used to disseminate results and software and to present test problems and problem generators for general use as benchmarks. We are mindful that if this project is successful there may well be commercialisable benefits from it, and there are a growing number of ways to handle such technology transfer.

Justification of Resources

The proposed work is of a very specialised and advanced nature. In order to fully justify the potential of this proposal we will require one Research Assistant and one PhD student at each of the two sites for 36 months. At Edinburgh we propose to employ Emma Hart as the RA. She has very considerable experience of timetabling and scheduling through two earlier EPSRC grants (GR/L22232, and GR/J44513 which she joined part-way through) as well as valuable programming experience working in the commercial sector, and valuable prior experience of hyper-heuristics research (see eg [HR99,RHC98,HRN98a,HRN98b,RH98,HR98]). She is on the committee of the timetabling and scheduling group of EVONET, the Europen Network of Excellence in Evolutionary Computing, is chairing a forthcoming wokshop in the area and was one of the two authors of the EVONET report that surveyed the state of the art in that field.

The project group, consisting of the investigators at each institution (Peter Ross, Edmund Burke, Peter Cowling, Graham Kendall and Sanja Petrovic), Emma Hart, the RA appointed at Nottingham and the PhD students, will meet quarterly during the period of the project. We are also requesting cheap net-video equipment to allow the research teams at each institution to contact each other via the web very frequently. External interested parties, such as our collaborators, will be invited to provide feedback at appropriate steering group meetings.

Some technical/CO help will be required at both institutions to assist with the implementation of the proposed methods and techniques. Some part-time secretarial support will also be required at each institution for the day-to-day handling of administrative matters. The proposed work involves the latest developments in heuristic techniques. Indeed, a major aim of this project is to provide a platform for the future of research in this field and it is essential that we develop and maintain links with other researchers at the forefront of this area. We therefore ask for funding to attend various major international conferences and UK conferences (from each institution) – see above. We would also require a small amount of travel funding for visits to related research groups in the UK, to support existing links with teams at City, Leeds, Napier, Southampton and Swansea Universities. Funding is also required for travel from each institution to test locations within the UK and for one portable PC (Nottingham) for use during these visits. The project work requires a high level of processing power and storage capability. Two fast PCs will therefore be required at each institution, each with additional disk space; we expect to use them as dual-boot Win98/Linux systems, and to generate large amounts of data during experiments.

References

[1] H.Terashima-Marín, P.M.Ross and M.Valenzuela-Rendón “Evolution of Constraint Satisfaction Strategies in Examination Timetabling”, Proceedings of the Genetic and Evolutionary Computation Conference 1999 (GECCO-99), pp. 635-642, Morgan Kaufmann, 1999.

[2] E.Hart and P.M.Ross, “An Immune System Approach to Scheduling in Changing Environments”, Proceedings of the Genetic and Evolutionary Computation Conference 1999 (GECCO-99), pp. 1559-1566, Morgan Kaufmann, 1999.

[3] P.M.Ross, E.Hart and D.Corne, “Some observations about GA-based exam timetabling”, Practice and Theory of Automated Timetabling II, eds E.Burke and M.Carter, pp. 115-129, Springer-Verlag Lecture Notes in Computer Science 1408, Springer-Verlag 1998

[4] E.Hart, P.M.Ross and J.Nelson, “Solving a real-world problem using an evolving heuristically driven schedule builder”, Evolutionary Computing, vol 6 no 1, pp. 61-80, MIT Press, 1998

[5] E.Hart, P.M.Ross and J.Nelson, “Producing robust schedules via an artificial immune system”, Proceedings of International Conference on Evolutionary Computing, Anchorage, Alaska, 1998, IEEE Press, 1998

[6] P.M.Ross and E.Hart, “An adaptive mutation scheme for a penalty-based graph-colouring GA”, Parallel Problem Solving from Nature V, eds. A.E.Eiben, T.Back, M.Schoenauer, H-P.Schwefel, Springer-Verlag Lecture Notes in Computer Science 1498, pp. 795-802, Springer-Verlag, 1998

[7] E.Hart and P.M.Ross, “A heuristic combination method for solving jop-shop scheduling problems”, Parallel Problem Solving from Nature V, eds. A.E.Eiben, T.Back, M.Schoenauer, H-P.Schwefel, Springer-Verlag Lecture Notes in Computer Science 1498, pp. 845-854, Springer-Verlag, 1998

[8] A.Tuson, R.Wheeler and P.M.Ross, “An evolutionary/metaheuristic approach to emergency resource distribution in the developing world”, in C.R.Reeves (ed), Proceedings of the International Conference on Artificial Neural Networks and Genetic Algorithms 97, eds G.D.Smith, N.G.Steele, and R.F.Albrecht, pp. 329-332, Springer-Verlag, 1998

[9] H-L Fang, P.M.Ross and D.Corne, “A Promising Hybrid GA/Heuristic Approach for Open-Shop Scheduling Problems”, in *Proceedings of ECAI 94: 11th European Conference on Artificial Intelligence*, A. Cohn (ed), pages 590–594, John Wiley and Sons Ltd, 1994.

[10] J.Berger, M.Sassi and M.Salois, “A Hybrid Genetic Algorithm for the Vehicle Routing Problem with Time Windows and Itinerary Constraints”, Proceedings of the Genetic and Evolutionary Computation Conference 1999 (GECCO-99), pp. 44-51, Morgan Kaufmann, 1999.

[11] P.Shaw, “Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems”, Proceedings of CP-98 (see also http://www.cs.strath.ac.uk/ apes/1998.html, and B. De Backer, V. Furnon, P.J. Kilby, P. Prosser, P. Shaw, “Solving Vehicle Routing Problems using Constraint Programming and Meta Heuristics”, Journal of Heuristics, to appear in 1999)

[12] S.Wilson, “Generalisation in the XCS classifier system”, Proceedings of the Third Genetic Programming Conference (ed J.Koza), pp.665-674, Morgan Kaufmann, 1998.

[13] C.Voudouris and E.P.K.Tsang, “Guided Local Search and its application to the Travelling Salesman Problem”, European Journal of Operational Research, Anbar Publishing, 113(2), pp.469-499, March 1999.

[14] B.Porter, R.Bareiss and R.Holte, ”Concept Learning and Heuristic Classification in Weak-Theory Domains”, Artificial Intelligence 45/1-2, pp 229-263, 1990.

[15] S.Petrovic, Case-Based Reasoning in Multiattribute Analysis, PhD Thesis, University of Belgrade, Mathematical Faculty, 1997

[16] S.Petrovic, ”Case-Based Reasoning in the DSS for Multiattribute Analysis”, Journal of Decision Systems Vol. 7, pp 99-119, Hermes, 1998.

[17] S.Petrovic, ”Classification of Objects Based on Case-Based Reasoning” Yugoslav Journal of Operations Research vol.9, no 1, pp 75-95, 1999

[18] E.K. Burke, J.P. Newall, “A Multi-Stage Evolutionary Algorithm for the Timetable Problem”, IEEE Transactions on Evolutionary Computation, Vol. 3.1, pp 63-74, 1999.

[19] E.K. Burke, A.J. Smith, “Hybrid Evolutionary Techniques for the Maintenance Scheduling Problem”, accepted for publication in the IEEE Power Engineering Society Transactions.

[20] E.K. Burke, M. Carter (eds), “The Practice and Theory of Automated Timetabling II: Selected Papers from the 2nd Int’l Conf. on the Practice and Theory of Automated Timetabling, University of Toronto, August 20th-22nd 1997”, Springer Lecture Notes in Computer Science Series, Vol. 1408, 1998.

[21] E.K. Burke, J.P. Newall, R.F. Weare, “Initialisation Strategies and Diversity in Evolutionary Timetabling”, Evolutionary Computation Journal (special issue on Scheduling), Vol. 6.1, 81-103, 1998.

[22] E.K. Burke, J.H. Kingston and P.A. Pepper, “A Standard Data Format for Timetabling Instances” in [BUR98a], 213-222, 1998.

[23] E.K. Burke, J.H. Kingston, K. Jackson, R. Weare, “Automated University Timetabling: The State of the Art”, The Computer Journal, Vol 40 No. 9, 565-571, 1997.

[24] E.K. Burke, P. Ross, (eds), “The Practice and Theory of Automated Timetabling: Selected Papers from the 1st Int’l Conf. on the Practice and Theory of Automated Timetabling, Napier University, August/September 1995”, Springer Lecture Notes in Computer Science Series, Vol. 1153, 1996.

[25] E.K. Burke, J.P. Newall, R.F. Weare, “A Memetic Algorithm for University Exam Timetabling”, in [BUR96a], 241-250, 1996.

[26] P. Cowling, W. Rezig, Integration of Continuous Caster and Hot Strip Mill Planning for Steel Production, to appear in the Journal of Scheduling.

[27] P. Cowling, Design and Implementation of an Effective Decision Support System: A Case Study in Steel Hot Rolling Mill Scheduling, to appear in: Human Performance in Planning and Scheduling, eds. B.L. MacCarthy and J.R. Wilson, Taylor & Francis.

[28] E.K. Burke, P. Cowling, P. De Causmaecker, G. Vanden Berghe, A Memetic Approach to the Nurse Rostering Problem, to appear in the International Journal of Artificial Intelligence.

[29] P. Cowling, The Total Graph of a Hypergraph, Discrete Mathematics **167/168** pp215-236, 1997.

[30] M. H. Bjørndal, A. Caprara, P. Cowling, et al, Some Thoughts on Combinatorial Optimisation, European Journal of Operations Research **83**, pp253-270, 1995.

[31] P. Cowling, F. Maffioli, A bound for the Symmetric Travelling Salesman Problem through Matroid Formulation, European Journal of Operations Research **83** pp301-309, 1995.

[32] P. Cowling, Optimization in Steel Hot Rolling, Optimization in Industry 3, ed. A. Sciomachen, John Wiley & Sons, Chichester, England, 55-66, 1995.

[33] N. Mladenovic, P. Hansen, Variable neighborhood search Computers & Operations Research, Vol.24, No.11, pp.1097-1100,1997.

[34] E.K. Burke, P. Cowling, R. Keuthen, New Models and Heuristics for Component Placement in Printed Circuit Board Assembly, proceedings of the IEEE International Conference on Information, Intelligence and Systems (ICIIS), Washington, IEEE Computer Society Press, pp133-140, 1999.