Chapter 1

AN INTRODUCTION TO EVOLUTIONARY COMPUTATION IN PRACTICE

Tina Yu[1] and Lawrence Davis[2]
[1]*Memorial University of Newfoundland;* [2]*VGO Associates*

Deploying Evolutionary Computation (EC) solutions to real-world problems involves a wide spectrum of activities, ranging from framing the business problems and implementing the solutions to the final deployment of the solutions to the field. However, issues related to these activities are not commonly discussed in a typical EC course curriculum. Meanwhile, although the values of applied research are acknowledged by most EC technologists, the perception seems to be very narrow: success stories boost morale and high profile applications can help to secure funding for future research and can help to attract high caliber students. In this book, we compiled papers from practitioners of EC with the following two purposes:

- Demonstrating applied research is essential in order for EC to remain a viable science. By applying EC techniques to important unsolved/or poorly-solved real-world problems, we can validate a proposed EC method and/or identify its weakness that restricts its applicability.
- Providing information on transferring EC technology to real-world problem solving and successful deployment of the solutions to the field.

1. APPLIED RESEARCH

In Chapter 2, Shingo Takeuchi and Kazuhiro Satitou applied a *multi-objective genetic algorithm* (MOGA) to design products that require disassembly for recycling at the end of their lives, according to recently established environmental regulations. One example of such products is electronic computers, where some of the parts can be reused and others are destined for recycling or a landfill. Product design is normally confronted with multiple objectives and constraints, such as technical specifications and economic returns. Recently, MOGA has become

T. Yu and L. Davis: *An Introduction to Evolutionary Computation in Practice*, Studies in Computational Intelligence (SCI) **88**, 1–8 (2008)
www.springerlink.com

a well-accepted method to obtain a set of non-compromising (Pareto optimal) solutions within a reasonable time frame.

In their case study of the design of a disassemble-able Mac Power G4 computer, the solutions have to satisfy four objectives and four constraints. Among them, two are related to the environmental regulations: maximizing the profit from reusing the disassembled components and minimizing the environmental impacts when the disassembled components are treated. To handle this design task, they expanded their MOGA with a double-loop, where the disassembly sequence planning is embedded within the 3D computer layout design. This modified MOGA delivered 37 alternative designs that have different trade-offs. This indicates that the modified MOGA can be applied to design other products requiring a similar embedded disassembly capability.

Chapter 3 focuses on a different issue which occurs frequently in the structural design of large and complex systems: how to perform component design without carrying out the time-consuming computer simulation of the entire system. In other words, how do we decompose the system such that only component simulations are needed to achieve a good global design? This is a common issue in the structural design of motor vehicles. Based on their extensive experience in delivering EC solutions to automotive industry, Erik Goodman, Ronald Averill and Ranny Sidhu presented two methods to address the decomposition problem.

The first method, called Component Optimization within a System Environment, is for systems with tightly coupled components, such as an automotive chassis. The method decomposes a system into a hierarchy. The boundaries of each component are updated at the end of each component design step. With a small number of iterations, they showed that this method can achieve an improved global design.

The second method is for systems that are composed of quasi-independent components, such as a lower compartment rail. Each component design is carried out independently, yet the components share their design solutions with components at the lower level of the system hierarchy. Since the upper level components are coarse approximations of the lower level components, the information helps reduce the potential design space of the lower level components and allows an improved global design to be obtained in a reasonable time frame. Although the authors only reported their implementation results on automotive component design, the potential of these methods on other large system structural design problems is worth exploring.

The researchers at the Honda Research Institute Europe also face the similar issue of expensive computer simulation required for aircraft structure design. This makes it impractical to use a large number of parameters to represent the design objects. As a result, the flexibility of the structure design is compromised, i.e. not all possible structure shapes are represented. This shortcoming has motivated Stefan Menzel and Bernhard Sendhoff to adopt a new kind of representation, called deformation method, where the design is represented as a lattice of control points of the object. Modification of these control points would produce a new object shape and also generate the grid points of the computational mesh, which is needed for the computational fluid dynamic mesh during simulation. Thus, this representation not only increases the design flexibility but also reduces computational simulation cost.

In Chapter 4, two kinds of deformation methods were investigated: free-form deformations (FFD) and direct manipulation of free form deformations (DMFFD). In FFD, the control points of the object lattice are used to represent the design object. In DMFFD, object points that can be placed on the design shape are used to represent the design object. The object points are then mapped into control points using a least squared algorithm. The authors have used both representations in their evolutionary algorithm to design a turbine stator blade that is part of a gas turbine for a small business jet. Their results show that the evolutionary algorithm has improved the initial design to better meet the turbine specifications. They also reported that the indirect encoding of DMFFD has provided three possible mappings between the object-point genotypes and the final design: many to one, one to one and illegal design. In order to apply DMFFD for industry scale evolutionary design tasks successfully, the authors note, these mapping scenarios need to be analyzed and understood.

The indirect encoding method mentioned in Chapter 4 is a version of embryogeny representations, which is receiving a lot of attention in the EC community recently. Briefly, embryogeny refers to the conversion of the genotypes through a set of instructions to develop or grow the phenotypes. The instructions can be intrinsic to the representation, where they are defined within the genotype, or extrinsic, where an external algorithm is used to develop the phenotype according to the values defined by the genotype. DMFFD is an example of embryogeny representations with an external growth algorithm.

In Chapter 5, Steven Manos and Peter Bentley also used an embryogeny representation with an external (more complicated) growth algorithm in their genetic algorithm to design micro-structured polymer optical fibres (MPOF). Although this is not the first application of evolutionary algorithms using an embryogeny representation, their solutions are innovative and patentable. This work received the Gold award at the 2007 Genetic and Evolutionary Computation Conference Human-Competitive Results competition. This applied research validates the practical value of the embryogeny representation in EC. We can anticipate that more EC applications will adopt this representation in the future.

Micro-structured optical fibres (MOF) are fibres which use air channels that run the length of the fibre to guide light. It is a newly developed technology that is being utilized in various optical communications, such as cable TV and operating theaters. The major task of MOF design is to position the air channels (air holes and their structures) such that they produce the required optical effects for a particular application, such as dispersion flattened fibres. This design task becomes very challenging when the material used is polymer, which allows any arbitrary arrangement of air holes.

To design MPOF, the authors devised a variable length genotype which encodes the positions and sizes of a variable number of air holes and the rotational symmetry of the structures. The conversion of the genotype into a MPOF is by a 4-step decoding algorithm, which not only transforms but also validates the design to satisfy manufacturing specifications and constraints, e.g. each hole must be surrounded by a minimum wall thickness for structural stability. They applied their genetic algorithm with the devised representation to design three different types of MPOF. The evolved designs are interesting and novel. Those that are of particular interest are currently being patented.

While the above work demonstrated the innovative aspect of EC, which can create human competitive designs, there are other areas of design in which human creativity is used to assist in the innovation. One example is consumer product design, such as the design of wallpapers and fabric tiles, where individual tastes play an important role in product acceptance. To make such collaboration possible, Chapter 6 discusses the development of an interactive evolutionary computation tool that can assist professional graphic designers to explore design patterns that go beyond their own imagination.

An interactive EC system opens the evolutionary algorithm's selection-reproduction-evaluation cycle and lets users guide the population toward a particular kind of solution. While the interface between users and an EC tool can have various degrees of flexibility, the one that motivates Carol Anderson, Daphna Buchsbaum, Jeff Potter and Eric Bonabeau is practicality: what feature an EC design tool should have so that graphic designers are more likely to use it in a regular basis to perform their jobs. They have considered many user-friendly features, such as allowing users to seed an initial design population and to freeze parts of the design without being modified by the reproduction operation. They are continuing to enhance the tool so that it can receive a wider acceptance by professional graphic designers.

So far, all the discussed works are in the arena of evolutionary design. This is not surprising since design is one of the most successful applications of EC. However, there are other application areas where EC solutions have been successfully deployed. The next four chapters discuss EC applications in business operations to reduce costs and increase revenues.

In Chapter 7, Cem Baydar applied EC-related technologies to grocery store retail operations, particularly in product pricing. In retail management, a product pricing strategy can impact store profits, sales volume and customer loyalty. A store manager can alter product prices by issuing coupons to achieve target objectives. One simple coupon-issuing strategy is blanket couponing, which offers the same discounts to every customer. In this chapter, Cem proposed an individual-pricing strategy where different discounts are offered to different customers, according to their buying behavior. His simulation results indicated that individual-pricing led to better store performance than the blanket couponing approach.

This result is not surprising, as more focused marketing normally produces results that meet a store's target goals better. However, the task of mapping this untraceable problem (the size of the search space is all possible combinations of coupon values and production items for each customer) into a workable framework using agent-based simulation and population-based simulated annealing is highly non-trivial. The problem-framing skill demonstrated in this project is important for practitioners who deliver real-world EC solutions.

Chapter 8 describes the implementation of a decision support tool for a routine railway operation: railway track intervention planning. Currently, the maintenance and renewal of railway track is mostly carried out by track maintenance engineers assisted by rule-based expert systems. Since the knowledge base of the expert systems is created by human engineers, who are not able to consider all combinations of possible track deterioration mechanisms, the resulting intervention plans are not very reliable and can lead to unnecessary high cost and low levels of safety.

To overcome such shortcomings, Derek Bartram, Michael Burrow and Xin Yao have proposed a data-driven approach to develop the decision support system. Based on historical data of railway track installation, deterioration, maintenance and renewal, they applied machine learning techniques to identify the patterns of failure types, to model the deterioration for each failure type and to determine the most appropriate maintenance for each failure type. The machine learning techniques they used include clustering, genetic algorithms and heuristic learning. Integrating multiple methods to address different needs of a system is becoming the standard approach for solve increasingly complex business problems. Practitioners need to develop the skill of picking the right technology in order to best solve the problem at hand.

Similar to the railway industry, there are areas in petroleum industry operations which rely on expert knowledge. One example is reservoir stratigraphic interpretation based on well log data, which is normally carried out by geologists or geophysicists who are familiar with the field. In Chapter 9, Tina Yu proposed a methodology using computer systems to automate this process.

Mimicking the ways geologists interpret well logs data, the proposed method has two steps: well logs blocking (approximation) and combining multiple logs' information at the same depth level to interpret reservoir properties. In terms of implementation, well logs blocking is carried out by a segmentation algorithm. Each block is assigned a fuzzy symbol to represent its approximate value. To interpret reservoir properties, a fuzzy rule set is generated to examine multiple well logs symbols at the same depth level and determine the property value. The prototype system also integrates various techniques, including a segmentation algorithm, fuzzy logic and a co-evolutionary system, to develop different parts of the interpretation system. Although the initial results based on well log data collected from two West Africa fields are encouraging, the author noted that there is weakness in the co-evolutionary system. More investigation is needed in order to produce quality deployable computer-based stratigraphic interpreters.

Resource scheduling is the task of allocating limited resources to requests during some period of time. In Chapter 10, Darrell Whitley, Andrew Sutton, Adele Howe and Laura Barbulescu discussed permutation based representations and their implementation with a genetic algorithm to solve three different resource scheduling problems. One interesting observation is that the permutations (genotypes) are indirect representations of the schedules (phenotypes). Different mapping algorithms were developed to transform permutations to schedules for different scheduling problems. This representation style is similar to the embryogeny technique presented in Chapter 5, although the work reviewed in this chapter was conducted at an earlier date when the term was not familiar to most EC technologists. Additionally, the authors have discussed various ways that the mapping algorithms, or schedule builder algorithms, can impact the success of their genetic algorithm to find good schedules. Much of the current research in embryogeny techniques are investigating similar issues.

2. TECHNOLOGY TRANSFER

Technology transfer goes beyond applied research and focuses on the soft side of solution deployment, such as social and political issues. In Chapter 11, Arthur Kordon discusses transferring EC technology in corporate environments, such as Dow Chemical.

The chapter starts with a list of values EC contributes to the chemical industry. However, promoting the technology not only requires demonstrating value creation and improved performance but also resolving other non-technical issues. Currently, Dow has one of the most successful teams in deploying EC solutions in the corporate world. Arthur Kordon shared their experiences in how they handle the organization and political challenges to maintain their status in the company. Examples are linking EC to corporate initiatives and addressing skepticism and resistance toward EC technology. These insights are valuable to practitioners who are interested in establishing a sustainable EC team in corporate environments.

Another technology transfer model is the collaboration between academia and industry. With their extensive experience in this type of collaboration, Rajkumar Roy and Jorn Mehnen discussed this model in Chapter 12.

Collaboration between academia and industry is not a new concept. When carried out properly, it produces many rewards: universities receive funding to educate high-caliber students and to develop new technology, while industry receives ready-to-use state-of-art technology that increases company revenues, in addition to receiving highly qualified employees. However, because the two ends of the technology transfer belong to different organizations, trust between the two entities is not always easy to maintain, and lack of trust can cause project failures. In this chapter, the authors listed activities to create and maintain that trust throughout the project cycle. These include documenting expectations, conducting frequent review meetings and providing small deliverables regularly. Finally, the sensitive issue of intellectual property has to be addressed so that both parties can fully commit to the collaborative efforts.

When talking about technology transfer, the most direct way is placing trained EC technologists in jobs that demand EC technology. But how do these two find each other? Chapter 13 is a survey of practitioners of EC to identify potential job sectors and help EC graduates to do job searching.

The survey was conducted from March 2005 to February 2006 by Gregory Hornby and Tina Yu. The major findings from the survey data are: there has been an exponential growth in both EC graduates and practitioners; the main source for finding a job has been networking; while most respondents to the survey are in Europe, the most growth of EC in industry has been in North America; the main application areas of EC techniques are multi-objective optimization, classification, data mining and numerical optimization; and the biggest obstacle for the acceptance of EC techniques in industry is that it is poorly understood. This information, although it cannot be generalized to the entire population of EC practitioners, does provide some direction on where and how to search for an EC job.

One of the most common avenues to transfer emerging technologies, such as EC, to industry and government applications is external consultancy. Since EC was in its infancy in the early 1980s, Lawrence Davis has been working as an external consultant involving many EC projects. With his accumulated 25 years of consultancy experience, he shared twelve lessons learned to improve a project's chances of ultimate success in Chapter 14.

Among the 12 lessons, some are similar to those mentioned in the two previous chapters discussing technology transfer in corporate and in academia-industry collaboration. However, there are others which are unique in consultancy situations. They are mostly related to project perception management, which is a skill most technologists don't learn in academia. Yet, they are critical for a project to obtain funding, gain acceptance by the users and bring to final success. For example, have a project champion and don't speak technically make the project more accessible to non-technical managers. Understand the work process and the system's effect on it prepares the end-users to become familiar with the system and gains their support of final deployment. These non-technical skills are important for practitioners who are interested in consultancy for industrial applications.

With all the effort put forward on applied research and technical transfer, the statistics of successfully deployed EC systems are still very low: 10% according to Lawrence Davis. However, this does not mean that the time and energy spent on the rest of 90% projects were wasted. On the contrary, much has been gained throughout the projects. Davis calls it the Tao of optimization: it can be very valuable to improve company operations and increase revenues without deploying any computer systems. One example is by changing the workflow of a particular operation to reduce expenses.

For our final chapter we have included a success story of deployed EC systems, which have saved Air Liquide millions in operational expenses.

Air Liquide is a multi-national company providing industrial and medical gases and related services. One type of product they supply is liquid gases (oxygen, nitrogen and argon), which are delivered by truck. To coordinate the production and distribution of liquid gases, a system was developed with two parts: a genetic algorithm to schedule production and an ant colony algorithm to schedule distribution. The two systems co-evolve to obtain an integrated schedule that works together to generate the best result. This co-operative co-evolution approach is very effective in solving this complex supply-chain optimization problem. Charles Harper, the Director of National Supply and Pipeline Operations at Air Liquide Large Industries U.S.LP, reported that one of their plants has reported a saving of more than 1.5 million dollars per quarter since Air Liquide began using the system.

The second deployed system optimizes Air Liquide's gas pipeline operation. It uses a combination of techniques, including genetic algorithms, a deterministic heuristic algorithm and brute force search, to address different optimization issues in each sub-system. Using the right technology, as advocated by Lawrence Davis in the previous chapter, enables the team to create a system that generates outstanding performance.

3. CONCLUSION

When we see EC systems working in the field to help people do their jobs better, it brings us a different kind of satisfaction from that gained by doing theorem proving. Regardless of whether you have created EC applications, we hope after reading this book, you will find applied research and technology transfer exciting and we hope that some of the insights in this book will help you to create and field high-impact EC systems.