The Optimatics Letter

Issue No. 15: January-March 2002 *Advances in Optimization for Water Distribution System Design & Operations*

Interview GA Optimization Pioneers

Genetic algorithm (GA) optimization of water distribution systems has lately attracted the attention of both large and small water utilities and consultants. Since 1995, Optimatics/Frey has successfully applied its Optimatics OGA software to the unique needs of various water utilities to optimize their distribution systems.

Looking back on six years of GA consulting, we thought it might interest our readers to learn about the origins of GA pipe network optimization. Our interview with the original GA pioneers at Adelaide University (and the founders of Optimatics), Dr. Angus Simpson, Dr. Graeme Dandy and Dr. Laurie Murphy, covers the early days of GA through today.

Frey (Editor): **Angus, how did you first get interested in GA optimization?**

Simpson: I was first exposed to GAs in 1983 and 1984 when I studied with David Goldberg at the University of Michigan. We were both doing our PhDs, both under the supervision of Professor Ben Wylie in Civil Engineering. David was also being jointly supervised by Professor John Holland who created the area of GA optimization back in 1976. David was completing his PhD in the application of GAs to the optimization of the operation of a set of gas compressors in series along a gas transmission pipeline. David was responsible for moving GAs from the Computer Science area into engineering. In 1989, he published a book on genetic algorithms which has sold widely.

Frey: **What prompted you to apply GAs to pressure pipe networks?**

Simpson: I had always had an interest in the optimization of water distribution systems. In 1985, I attended (with you, in fact) an ASCE conference in Buffalo, New York at which a session on the “Battle of the Network Models” was presented. Various researchers around the world were invited to optimize the “Anytown Network” and came to present their results. Attending this conference and in particular the sessions on water distribution network optimization sparked my interest in the field.

After finishing my PhD in water hammer, I took up an academic post at the University of Adelaide in South Australia in early 1987. In 1989, I supervised a fourth year research project that investigated the use of Non-Linear Programming for the optimization of a 300-pipe “real network.” The outcomes were disappointing—the network that was being optimized was too big and we didn’t get any satisfactory results.

In 1990, I supervised another research project with students Laurie Murphy and George Hadji. I gave them a paper by Goldberg and a copy of his book. Laurie became very excited about GAs and their potential. They developed a GA in Pascal based on the pseudo-code in Goldberg’s book. They also developed a Hardy-Cross pipe network solver that they integrated with the GA. The outcome was the first known application of the use of GAs for the optimization of water distribution systems.

Frey: **Laurie, how did that research go and what happened to you next?**

Murphy: The research project was great. Our focus was a 14-pipe, two-reservoir problem that had been previously optimized by pruned enumeration by Gessler. We first enumerated every solution in the search space of over 16,000,000 different combinations. We then used the GA to find the global optimum solution in relatively few evaluations compared with the total size of the search space. It was a real success, but I don’t think George and I received any distinctions—I must talk to Angus about that one day. In 1991, I started on a PhD in the application of GAs to the optimization of the design of water distribution systems with Angus. Graeme Dandy joined us in 1992 as my co-supervisor.

(Continued on page 2)

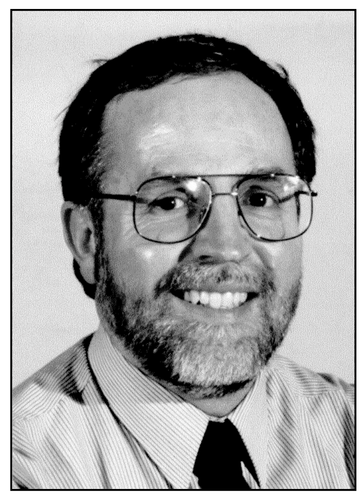

Angus Simpson first learned of GA optimization in 1983 at University of Michigan where Professor John Holland came up with the idea of genetic algorithms to mimic natural selection in 1976.

Published jointly by Frey Water Engineering, Inc., USA and Optimatics Pty Ltd, Australia, consultants specializing in OGA optimization.

Frey: **Graeme, how does your early work in optimization compare with your later experience using GAs?**

Dandy: I completed my PhD in Environmental Engineering Systems at MIT in 1976 studying the optimum design of water quality sampling systems in river networks. I later used conventional optimization methods such as linear and non-linear programming on problems including optimizing pumping from the Murray River to the City of Adelaide.

GAs provide a more flexible approach to optimization. As GAs are guided search procedures, they work in conjunction with simulation models of the systems in question. Essentially, any system that can be simulated can be optimized with GAs. Also, because GAs work with a population of solutions, they can deal with complex search spaces with many non-linearities and logical rules. Conventional optimization procedures for non-linear problems (such as the pipe network design problem) tend to converge to local optima that might be nowhere near the true global optimum solution. GAs also have the advantage that they work well with discrete decisions (e.g., pipe sizes, valve settings and pump choices). Conventional optimization methods have difficulties dealing with these.

Frey: **Laurie, were you able to develop a successful GA strategy to optimize pumping operations as well as pipes, tanks and pump station sizing for the Anytown Network back in 1993?**

Murphy: Yes, we did achieve some pretty good results using GA on that problem. As I recall, I first worked out the optimal tank water level trajectories and then had the GA optimize pumping operations and the locations and sizes for new and duplicate pipes and new tanks. It was a small system, but quite complex. We would use a more straightforward approach to the decision variables today with our Optimatics GA software.

Frey: **For the New York City Tunnels problem, your 1993 research report presents a $38.80 million solution that was found on the 96,750th GA evaluation. Since then others have claimed to have found better solutions, including Savic & Walters ($37.13M in 1997), Lippai, Heaney & Laguna ($38.13M and $37.83M in 1999) and Wu, Boulos, Orr & Ro (same $37.83M and $37.13M solutions in 2001). I was curious about this, so last week I checked all three of these lower-cost solutions in EPANET 2.0. I found that they all fail—they fail to achieve the minimum total head required at Nodes 16, 17 and 19 or Nodes 17 and 19 by up to 0.22 feet. In addition, the Lippai et al and Wu et al $37.83M solution utilizes a 124-inch size to duplicate Pipe 7, but this size is not even an allowable choice in the original 1969 Shaake & Lai problem. Does this bother any of you at all?**

Murphy: Not really. It is great to see the continued interest in pipe network optimization. The New York tunnels problem is a simple benchmark problem to start with and also to refine a proposed optimization technique. I must say, however, that claims that better solutions have been found do not really count for much in that the problem is extremely sensitive to the constants used in the Hazen-Williams head loss equation. Tunnel diameters are large and tunnel flows are high, so differences in head losses are perceptible depending on whether you use a Hazen-Williams head loss formula with diameter to the power 4.87 or 4.871 and flow to power 1.85 or 1.852. This is essentially the key to finding a "better" New York tunnels problem solution than has been previously found.

Simpson: Laurie is right about different coefficients being used in the Hazen-Williams equation. EPANET 2 does use the correct form of the equation. The $38.80 million solution that we found in 1993 remains the lowest cost solution that actually satisfies the minimum allowable pressure constraints.

Dandy: The New York Tunnels problem has been used as benchmark since 1969 to compare various optimization procedures. It would be nice to compare methods using a set of other benchmark problems including more complex ones.

Frey: **Are you happy with how far GA pipe network optimization has come since 1990 when you initiated work in this field?**

Murphy: Couldn't be happier myself. After all, just over 10 years ago we were applying GAs to the New York tunnels problem—a very simple network model of 21 tunnel sections with 2 loops and one pattern of node demands. Today, we are determining the expansions required for the next 20-50 years for city systems such as Albuquerque, San Diego and Detroit to name a few. These

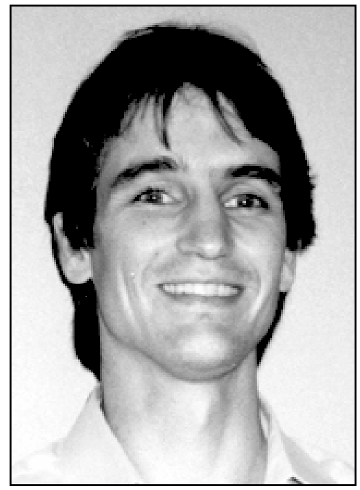

Laurie Murphy was first to solve a water distribution system optimization problem using GA optimization. After solving the Gessler problem in 1990, Laurie went on to complete his PhD applying GA optimization to other benchmark problems including the New York City Tunnels and the Anytown Network.

system models generally have more than 2,000 pipes, multiple sources of water, storage tanks, reservoirs and pump stations, and the expansions that we identify must meet the 24-hour maximum day demands for some future year. You have to be happy with that progress. Having said that, the job is not done—we are improving our software with each new system that we tackle.

Simpson: Since the early 1990s, we have evolved a comprehensive process for optimizing the design and operation of water distribution systems. In addition to choosing pipe diameter sizes for new pipes and duplicate pipes, we can now optimize all aspects of the network including pumps, tanks and pressure reducing valves. Research work is continuing in determining the best way to optimize the operation of pumping in water distribution systems.

Dandy: GA pipe network optimization has certainly come a long way. We are now tackling real systems with up to 20,000 pipes and up to 400 choices including pipes, pumps, tanks and valves. The search space of some of the problems we are tackling is in the order of 10^{200}. When you consider it is estimated that the number of atoms in the universe is around 10^{75}, you can see how complex some of these problems are.

Frey: **What more can be done to convince water utilities and consultants to use GA?**

Dandy: The results obtained speak for themselves. We have consistently saved 10 to 50% of the capital costs of large distribution systems. In all cases we have worked closely with engineers from the utilities and consultants to ensure that the designs and plans that we come up with satisfy their required performance criteria. Anyone who is not making use of this technology today is wasting a significant portion of their capital works budget.

Frey: **Are there limits on problem size or complexity for GA optimization analysis?**

Simpson: Water distribution models are becoming bigger and bigger with some models being as large as 20,000 to 30,000 pipes. Run times for a single simulation using these models can be on the order of minutes. GA optimization involves 100,000 or more simulations. Currently for GA optimization to be practically implemented on such large models, the models need to be skeletonized to the order of 3,000 to 5,000 pipes. The number of decision variables should be maintained below 300 or 400 [unless it is pipes-only]. More research work is required to improve GA techniques for finding optimal solutions for large numbers of decision variables.

Frey: **Is GA optimization difficult to apply?**

Simpson: Yes, the application to real world problems is complex. It is easy to optimize 20 or 30 pipes such as for the New York tunnels problem but real world problems are generally much more complex than this—where pumps, tanks and valves need to be considered. Appropriate formulation and interpretation of the results is 80% of the task while running the software is only 20% of the task. Over the last 11 years we have built up a large base of knowledge on applying GA optimization to the optimal design of water distribution systems.

Frey: **What's ahead for GA optimization?**

Simpson: We have successfully developed the formulation of the optimization of pumping in water distribution systems to minimize operating costs by maximizing the off-peak pumping by combining optimal scheduling and tank trigger levels. This should also be of interest to water utilities—we are continuing research in this area.

Dandy: Utilities are becoming more concerned about water quality in their distribution systems. There is considerable scope to optimize the design of water distribution systems so as to improve water quality. GAs can also be used to optimize system operations. Pumping cost and water quality are important factors to be considered in this case.

Water reuse is becoming increasingly important. Apart from the use of GAs to optimize the distribution systems for reclaimed water, there is also the potential to use them to optimize wastewater treatment plant design. Conjunctive use of surface water and groundwater is another area where GA optimization could be extremely useful.

(Continued on back cover)

Graeme Dandy had applied traditional optimization techniques to water system problems in the 1970s and 1980s. After 10 years working with GA optimization, he finds GAs to be far more powerful and flexible.

For information, contact Frey Water Engineering, Inc., 121 South Chestnut Ave., Suite 200 Arlington Heights, IL 60005-1817 • Phone (847) 670-7970 • Fax (847) 670-7973 • E-mail: info@frey-water.com. You can download past newsletter copies at *www.frey-water.com*.

The Optimatics Letter is published quarterly and is intended to provide information of value to the water industry.

The Optimatics Letter

c/o Frey Water Engineering, Inc.
121 South Chestnut Ave., Suite 200
Arlington Heights, IL 60005-1817
www.frey-water.com

Please pass this newsletter on to key staff involved in distribution system planning and operations.

Also, please call, fax or e-mail us to update names and addresses or to be removed from the mailing list.

The Optimatics Letter *Advances in Optimization for Water Distribution System Design & Operations*

Frey: **What other kinds of optimization research in water looks promising these days?**

Simpson: There are a number of other evolutionary type algorithms that show a lot of promise. Over the last two years we have formulated a new technique based on Ant Colony Optimization (ACO) that is showing significant promise. It seems to be more effective than GAs in finding optimal solutions for standard problems that have been investigated previously in the literature.

Frey: **I'm afraid that's all the space we have in this newsletter. I'd like to thank Angus, Laurie and Graeme for this very interesting look back at the origins of GA pipe network optimization at the University of Adelaide in 1990. If our readers have any questions, you can e-mail the experts at Angus.Simpson@Optimatics.com, etc.**

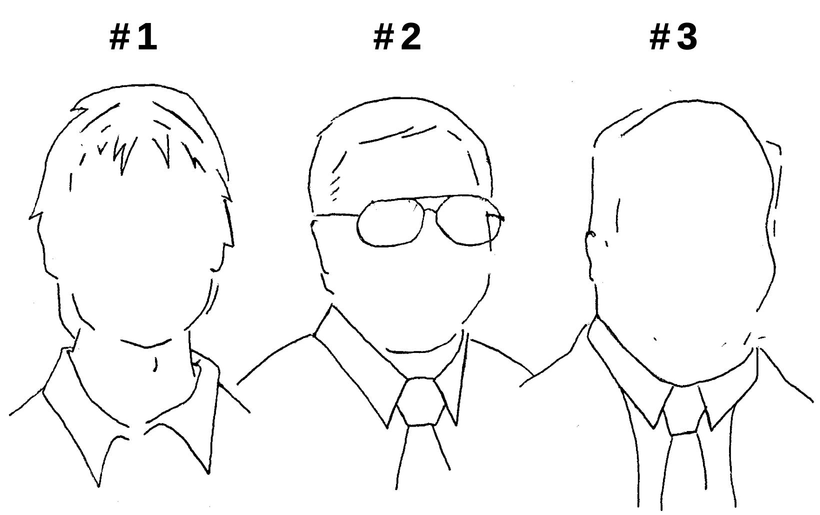

Will the real pioneers of GA pipe network optimization please stand up?