SIAG/OPT Views-and-News

A Forum for the SIAM Activity Group on Optimization

Volume 18 Number 2 October 2007

Contents

Optimizers in Industry

My Experiences as an Industrial Research Mathematician

Andrew R. Conn

IBM T.J. Watson Research Center,

Route 134, P.O. Box 218, Yorktown Heights, New York 10598, USA (arconn@us.ibm.com).

1. Setting the scene

What follows is a personal view but an honest one. I trust it will not annoy too many friends and at least they will recognise my voice.

The year was 1989. I had been at the University of Waterloo, except for a postdoctoral fellowship at the Hebrew University, Jerusalem in 1971 and two sabbaticals in France, ever since I joined as a Ph. D. student in 1968. It was fair to say that I was becoming a little restless, in spite of how great the University Waterloo was for me and my research, but it had never occurred to me to work anywhere other than at a university. I had always had reservations about moving to the USA but I had started to feel that my prejudices should be tested and I should find out first-hand their validity or otherwise. In any case I was certainly considering appointments at US universities since it is not easy to leave Waterloo for another Canadian university. Then one day, at Oberwolfach, Ellis Johnson asked me if I would consider going to work at the T. J. Watson Research Center. I said I would, I went down for an interview and accepted a position within a year.

Some points to consider. Working at T. J. Watson is not a typical industrial position. At the time I was recruited, one of my main interests was large-scale nonlinear optimization (the LANCELOT book was nearing completion) and I was anxious to be working on solving problems in which people were extremely interested in the answer. Furthermore, it is really good for ones ego to be asked to apply for a position. The salary I was offered seemed enormous and the working conditions sounded idyllic. Although, as can be expected, most did not turn out exactly as it appeared at the time, I have few regrets concerning the choice I made.

Since the most contentious issues are likely to be associated with my reservations about moving to the US let me just say that many of my prejudices were wrong but many of them were also updated to new ones. Suffice it say that after seventeen years I am still here. Moreover, crazy as it may seem to some, where we live is beautiful and both my wife Barbara and I adore New York City for many reasons. It turned out that my move to IBM almost coincided with the beginning of a significant decline in IBM (albeit, not permanent) in the 90s. So much so that my good friends kept asking me if it was my fault. Consequently, much that was promised me (like I could travel as much as I wanted to — which seemed realistic since in those days I saw Ellis everywhere except at IBM) were never realised. Also, my interpretation of my salary was naive. It turns out to be very expensive to live where we live and then many things that we had taken for granted, like universal health care at moderate cost, good public schools and moderate property taxes and insurance costs transpired to be very different in New York.

Also I had always claimed that teaching was great in theory but often a pain in practise. I still maintain that teaching (at the right level and to good students) is very useful for research but it is a constant demand which may not always fit an ideal agenda. Also, I suspect I enjoyed my teaching much more than my students did. However, it is not unusual for researchers at IBM to teach in a university, and indeed, for three years I taught a course at Yale, had excellent students and was about to build up a group of doctorate students when the department at Yale (Operations Research) was unfortunately closed down, partly at least, as a cost saving measure. Some of my colleagues in Mathematical Sciences are currently teaching at Yale University, Columbia University and New York University whilst maintaining their full-time positions at IBM.

When I left Canada I had the largest NSERC research grant in numerical analysis in the country ($165,374 for three years). Within a year of arriving at IBM I had a DARPA grant (also for three years) of $917,809. Even taking account of the different exchange rates and the fact that NSERC did not allow overheads, the difference was considerable.

2. The environment at IBM research

I strongly believe that an industrial research laboratory like IBM T. J. Watson has much to offer that is not available in a university environment and vice versa. Thus the recent demise of industrial research laboratories as opposed to development laboratories is to be lamented. Very few remain and I do not think this is good for science or the community. Basic research is an important luxury[1] and has to be acknowledged as such. If short term gains are the single most important objective then basic research cannot thrive. One also needs extended periods of uninterrupted thought.

I consider that I am reasonably well acquainted with a cross-section of the western world's leading research universities. I know no place that can offer a more industrious or brilliant group of colleagues in mathematical sciences than those I have at T. J. Watson. That fact is certainly one of the highlights of my move.

The pressure, real or perceived, to complete things in a hurry that are adequate but could be better, while a realistic sign of the times, is one of the largest 'lowlights'.

There is little doubt that I am working for an

[1] One could argue that since it is important it is a necessity and not a luxury, but unfortunately when budgets are tight it is often seen as dispensable.

American company. As a university professor I was essentially self-employed and it was rare that anybody told me what to do. As an IBM employee I may be unmanageable but nevertheless I have a real and rather hierarchical management. Political correctness, intellectual property and legal issues are real and can be a considerable obstacle at times. However, I have IBM's consent to publish articles like this one. In a public education system like Canada's you were subject to the vagaries of the political system but that seemed relatively stable and conservative compared with the effect of the stock market and the economy on the morale of a private company. In the long run both may be equally stable but the short-term variance seems much greater in private industry, which can be both good and bad. There are conflicting agendas that can often appear short-sighted. I often disagreed with education policy but it rarely effected me much locally. Taking care of apparently superficial things can be an annoyance. I unquestionably had better computing support (both hardware and software) when I was at a university and it was significantly easier for me to travel and have exchange visits with colleagues. There are more challenges working in industry and the rewards, both financial and intellectual, are considerable.

On the other hand, there are no shortage of interesting applied projects to work on at IBM and although it may sometimes appear that the customer is too important, when a project is successful the results really matter, the problems are challenging and they can be a significant catalyst for innovations in basic research. It is a clear motivator for learning new things.

3. Projects at IBM research: an application

Early on in my career at IBM I became involved in the design of circuits. The history behind that early involvement is, of itself, interesting. Essentially an electrical engineering colleague at Watson wanted to do some minimax optimization and asked my advise. I gave him some brief discussions and pointers and he took care of the rest himself. At a later date, one of his colleagues was organising an internal meeting across IBM concerning how we handled circuit tuning. On the basis of the earlier encounter it was suggested that I could be considered as a relatively approachable mathematician and I was invited to attend the workshop. I had time to take part in about half of the meeting whereupon the organiser of the workshop asked me what I thought of it. I said it was very interesting but it was too bad that they were using 1960s algorithms in the 1990s and so I was asked to tell them about 1990 algorithms. That was the beginning of an extremely successful collaboration that ended up gaining IBM millions of dollars and resulted in us achieving an IBM award, where I was not only rewarded financially but my wife and I were cared for royally at a retreat with other awardees in the Spring of 2002.

Now you might well think that nothing could be easier than persuading IBM that using modern optimization to design better circuits was a win-win situation. Especially when you consider that the alternative was to purchase outside tools that inevitably meant IBM had to disclose aspects of their circuits (in order to benefit sufficiently from the tools) that then improved the utility of the software that these software vendors were then selling to our competitors. In fact, after the initial success with what one might call a prototype, a considerable part of the next eighteen months was spent fighting for the survival of the project. We did survive because of a number of important non-trivial reasons. Firstly we did receive unfailing support from our *immediate* management. As is often the case, the conflicts were higher up — and in this case relatively easy to explain. Basic to the optimization was the underlying simulation. What we had was a good, fast (about seventy times faster than a more accurate SPICE like simulator) and reliable one, based on piecewise constant approximations and, equally importantly, it also provided reasonably accurate derivative approximations. What others promised, did not exist, would be based upon piecewise linear approximations and would not provide derivatives. The trouble was that 'everyone knows that linear approximations are better than constant ones' and higher management did not appreciate the significance of derivatives to the optimization. One thing I learned early is that, amazingly to me, often what you promise is much more important than what you deliver! I also benefited enormously from having world-class circuit designers and world-class electrical engineers who had the confidence and ability to know that what they were doing was right and eventually to persuade others that this was the case. It is good to realise that, given time, those who understand the details are able to convince those that don't and eventually the correct decision is made. So from a healthy mistrust at the outset on the part of practitioners, they eventually became our biggest supporters. It also impressed upon me the importance of having good software that worked in an environment in which those who had to use it were comfortable. Without which, algorithms twice as good would not have succeeded. I think it is fair to say that the project could not have been as successful without the state-of-the-art optimization, but without the electrical engineers and the familiar interfaces it would not have even got off the ground. Furthermore, this is a project that would not have happened in an academic environment. It needed the investment, level of commitment and breadth of skills that would have been difficult to find in such a cohesive way anywhere other than a place like T. J. Watson. It gave rise to a number of research papers in premier journals, several patents and a best paper award, [2], as well as the IBM recognition.

So, having set the not so atypical scene for a major project let me give you some idea of what it entailed. In high-performance custom-designed chips, the sizes of individual transistors are tuned to maximize the performance of the circuit. The key parameters are transistor widths, which control the amount of current that flows through each transistor. In simple terms, wider transistors generally lead to faster circuits but consume more power and area, and increase the load on previous stages of the circuit. Designers tune their circuits by adjusting the widths of transistors to obtain optimum performance. At the time that we began the project, traditional methods were slow, tedious, manual and error-prone. The computer-aided design tool we eventually came up with, automated circuit tuning while permitting its use on far larger circuits than previous techniques allowed. It used what was then a state-of-the-art optimization technique (LANCELOT, [3]) and it now uses an interior-point filter method (IPOPT, [10]). In addition to making IBM circuit designs better it significantly improved our designers' productivity. Today it is deployed as a standard tool within IBM and assists in the design of all our custom circuits.

4. Projects at IBM research: not primarily an application

So circuit tuning was my first encounter with what could be done in this environment. As an example of another instance, that at least initially is not motivated solely by an application, we were asked to propose research projects that were ideal projects for collaboration between IBM and CMU. Senior management was interested in promoting such a union. The project I had in mind was something I had been thinking about for a while but it was also catalysed by some of my knowledge of circuit tuning since the future technology was certainly going to have to deal with discrete variables in addition to the continuous variables we were currently tuning. This arises, for instance, in dealing with low, regular and high V_t assignment[2]. Consequently I decided that an ideal topic for such a collaborative research project was mixed integer nonlinear programming. The motivation was perfect. IBM research had very strong researchers in linear programming, nonlinear programming and mixed integer linear programming, both from the point of view of theory and software. Furthermore, of course, there was no shortage of challenging applications with discrete and continuous variables and nonlinear functions. Moreover, these practical applications are often sufficiently complex that even sub-optimal solutions can be extremely helpful. Up until recently there had been only a little basic research interaction between the linear and nonlinear sides at IBM. Equally, CMU had expertise on both sides and in addition, they were well-known for their work in chemical engineering that included difficult applications and pioneering work in mixed integer nonlinear programming. But perhaps the most important motivators were that this is an under-researched area (no doubt because of the inherent difficulties to solve such problems) with ever increasing practical importance. In a surprisingly short time we had the enthusiastic support of

[2]For example, speed critical gates are assigned low $V-t$ whereas non-speed critical gates are assigned high $V-t$, thereby reducing leakage.

researchers at CMU and IBM and strong support internally at IBM. The initial objective was to have a reasonable nonlinear mixed integer package available as open source software, along with some test problems, as quickly as possible. The hope being that this would be good for the field and the research momentum within IBM. Along with the open-source software and the promise of applications significant for IBM, we anticipated that some reasonable initial progress would enable the project to become a long term commitment and indeed, this does appear to be the case. Moreover, my actual contributions have been minimal, but under the excellent leadership of my colleague Jon Lee the project is thriving.

Although mixed integer linear programming is already difficult[3] it is not surprising that the nonlinear case is much more difficult, being undecidable even for minimizing a linear function subject to one polynomial equation in ten non-negative integer variables, [8]. Amongst the many issues that make mixed integer nonlinear programming difficult are the fact that the solution typically is not on the boundary, so finding the convex hull of the feasible region is insufficient to determine solutions. In addition, non-convexity means that global optimization is an issue in finding good bounds and finally the subproblems themselves are often very difficult to solve. Moreover, although it is well understood that equivalent formulations mathematically are not equivalent practically, in the case of mixed integer linear programming we have a reasonably sophisticated handle on what makes a good formulation, in the nonlinear case it would appear that the formulation is even more critical but we have far less knowledge of what is appropriate.

Within not much more than a year, with the help of the open source base COIN (`http://www.coin-or.org/index.html`) we were able to make available the first open source mixed integer nonlinear programming software (BONMIN), a hybrid between two classical algorithms for mixed integer nonlinear programming: an outer-approximation-based branch-and-cut-based algorithm and a pure branch-and-bound algorithm. COIN itself began as an IBM initiative but is now a foundation with the partnership of INFORMS. It contains efficient codes for linear programming (CLP), mixed integer linear programming (CBC, CLP, CGL) and nonlinear programming (IPOPT). BONMIN, [1], is also available through NEOS (`http://neos.mcs.anl.gov/neos/solvers/minco:Bonmin/AMPL.html`).

[3]The problem is NP-Complete, but pseudo polynomial-time algorithms (*e.g.*, knapsack) and even polynomial-time for special cases (*e.g.*, network problems with a fixed number of variables) do exist.

5. Projects at IBM research: the present

Not unreasonably IBM prefers that its researchers work on applications of interest to IBM. Clearly circuit tuning fitted that criterion, but once it became an accepted daily tool I yearned to make a splash in a new area (besides, IPOPT was able to do much more than LANCELOT — age is often a negative attribute). In addition to the continuous encouragement of the corporation I feel so strongly that there are so many challenging, economically important and interesting applications that would benefit from the intelligent use of state-of-the-art optimization techniques that it is almost a duty to try to contribute in that applied way. Moreover, one cannot overestimate the pleasure and inspiration one can achieve when working with colleagues who are leading experts in the application area but so outstanding in their breadth that they are easy to interact with on the mathematical side, even when I am a neophyte from the point of view of the application. I may be incorrect in this, but it has certainly been my impression that, from the point of view of real, complex (no pun intended) applications and a marriage with the required applied mathematics, one is not able to match in any university what one can do at a research center like ours. Although it is perhaps self-evident, to be completely successful such a union requires an intimate collaboration between the end-users, the engineers, physicists, geologists, or other experts who understand the application deeply, and the mathematicians. But the reward, from the perspective of research, expanding ones horizons, appreciating one colleagues and ultimately making significant impact on IBM and the applied field, is truly remarkable.

Currently I am endeavouring to do exactly that in what is a new application field for me, namely in petroleum engineering. It certainly is a promising area, in that it is very significant from the perspective of economics, politics and the environment. Moreover, the problems are extremely complex and the industry has plenty of money to invest on better solutions. Equally important from our point of view is that we have world class experts in the application area and in the necessary mathematics. We also have good contacts within IBM and the petroleum industry. Thus we have all the potential to be as successful as we were in circuit tuning, thereby benefitting the corporations, research and the field in general. Currently we are in the early stages and these are facilitated by IBM in that we have what are called 'first-of-a-kind' research projects. These are intended to be collaborations with a customer on difficult problems of importance to them and are expected to lead to significant longer term engagements with a commensurate investment from the client. However the projects are meant to be exploratory, the initial investment of the customer is small and it is well-appreciated that not all of them will lead to further progress, research or engagements, although many of them do. We have successfully completed such a project in assisted seismic matching, inverting for basin modeling, rock physics, and seismic attributes, including seismic amplitude, to match both the observed seismic data and the observed stratigraphy. As the reader can imagine I am not an expert in geophysics — but I am working with somebody who is a leader in the field. We applied advanced numerical optimization techniques to integrate geological and geophysical data and infer the sedimentary parameters that produce a match to seismic data. We wanted to match not just event timing (phase) but also reflection strength (amplitude). This inverse problem of quantitatively matching present-day measurements of structure, stratigraphy, rock physics and/or fluids is inherently ill-posed and computationally difficult. Eventually we achieved a very respectable seismic fit to the synthetic model. This model was designed to be sufficiently realistic to be able to expose many of the challenges in successfully applying state-of-the-art optimization, while simultaneously not being so large that understanding and overcoming inherent issues would mean that progress would slow to a crawl. We were less successful with a model from a shallow, deltaic environment offshore of West Africa, where an erosional feature was filled in with a prograding delta[4], [7]. Nevertheless, accounting for our progress to date, we are reasonably confident that our basic approach and framework are feasible, given modern industrial computing resources, for tackling realistic problems of this nature.

Other work we have done in the general area included a research project with an industrial petroleum research group on automatic history matching. Again this project was exploratory. An important aspect here was the computation of the derivatives in which we used a single algebraic framework to derive forward and adjoint methods. We were also able to benefit enormously from the expertise of the petroleum company's research group leader and their data. Two problems used in a benchmarking study of three automatic history matching tools were provided by the petroleum company. The first was a simple set of synthetic problems designed to test robustness and efficiency in a situation where all tools can be applied exactly the same way. The second was water flooding in a situation close to a real field application. In the first case we obtained a good to excellent match for all wells and in the second case the results were qualitatively similar to commercial software indicating that we can solve problems similar to other tools in the market (amazingly, given that we had a very elementary implementation). Thus once again we have a good basis for further work.

Finally, in discussing this application field, we are currently engaged in a project for a major petroleum company to assist in a maintenance plan for a set of dependent oil platforms where the maintenance activities are dependent in time, corrective maintenance activities arise as the result of inspections and are not always known far in advance and the overall plans need to be defined several years ahead. Any day without production can represent millions of dollars of lost production so there is plenty of incentive to obtain better plans than they have now, even if the results are suboptimal. We are currently working on developing systems for optimizing the maintenance schedule of a single facility but that will ultimately be generalised to multiple de-

[4] A delta that is being built seaward by accumulation or deposition.

pendent facilities. The basic approach is that given a list of maintenance activities, dates when they are performed, their costs and other relevant attributes and treating failures as random events with probability distributions that will be estimated by the petroleum company experts, we can simulate future scenarios of maintenance activities and component failures. We thereby obtain a predictive distribution of the cost of the maintenance schedule including the costs of maintenance (approximately known and fixed) and of shutdowns to repair failed systems (random). We then optimize over maintenance schedules using derivative free techniques (see for example, [5, 9]).

6. A few concluding remarks

As you can see I work in an area rich in problems. It is exciting because of the access we have to extremely knowledgeable people both inside and outside IBM. The mathematical tools include nonlinear optimization, mixed integer nonlinear programming, differential equations, modelling, risk assessment, high performance computing, simulation and statistics. Most of the applications at this sort of level result in archival articles in mathematical journals. Of course, as research professionals we are involved in refereeing, serving on professional boards and committees and, last but not least, promoting the field and our colleagues in universities.

So you see that there is plenty of opportunity to have a rewarding and challenging time working in industrial mathematics, both from the point of view of basic research and applications. Indeed, I am about to complete a third book (with coauthors inside and outside IBM) since joining IBM seventeen years ago, see [3, 4, 6]. I cannot overemphasize the two most wonderful things about my job — unbelievably superb colleagues and a plethora of extremely challenging problems. At the risk of ending on a pessimistic note, as I already mentioned, there used to be many more industrial research places than there are today, which is extremely unfortunate given the need to apply top quality mathematical methods to ever more increasingly complex problems! I encourage those of you so inclined to be a catalyst for strengthening the industrial research centers that exist today and perhaps, the future will enable you to consider joining new ones.

Acknowledgement

The author would like to thank his colleagues, both in Academia and at IBM. There excellence, support, tolerance and good humour are much of what have made my career both rewarding and enormously enjoyable.

REFERENCES

[1] P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Wächter, *An algorithmic framework for convex mixed integer nonlinear programs*, Special issue of Discrete Optimization in memory of George B. Dantzig (1914-2005), editors, E. Balas, A. Hoffman, and T. McCormick, Discrete Optimization, pp. 1–19 (2008), to appear.

[2] A. R. Conn, P. K. Coulman, R. A. Haring, G. L. Morrill, C. Visweswariah, and C. W. Wu, *JiffyTune: circuit optimization using time-domain sensitivities*, IEEE Transactions on Computer-Aided Design of ICs and Systems, 17 (1998), pp. 1292–1309. Also appeared in the book *The Best of ICCAD – 20 Years of Excellence in Computer Aided Design*, Kluwer Academic Publishers, Dordrecht, 2003.

[3] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, LANCELOT*: A Fortran Package for Large-Scale Nonlinear Optimization (Release A)*, Springer Series in Computational Mathematics, Vol. 17, Springer-Verlag, Heidelberg, 1992.

[4] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *Trust-Region Methods*, MPS-SIAM Series on Optimization, SIAM, Philadelphia, 2000.

[5] A. R. Conn, K. Scheinberg, and Ph. L. Toint, *Recent progress in unconstrained nonlinear optimization without derivatives*, Math. Program., 79 (1997), pp. 397–414.

[6] A. R. Conn, K. Scheinberg, and L. N. Vicente, *Introduction to Derivative-Free Optimization*, MPS-SIAM Series on Optimization, SIAM, Philadelphia, 2008, to appear.

[7] M. Imhof and A. K. Sharma, *Quantitative seismostratigraphic inversion of a prograding delta from seismic data*, 75th Annual International Meeting, SEG, Expanded Abstracts, 2005.

[8] J. A. D. Loera, R. Hemmecke, M. Köppe, and R. Weismantel, *Fptas for mixed-integer polynomial optimization with a fixed number of variables*, in SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, New York, USA, ACM Press, pp. 743–748, 2006.

[9] M. J. D. Powell, *Developments of NEWUOA for unconstrained minimization without derivatives*, Tech. Rep. DAMTP 2007/NA05, Department of Applied Mathematics and Theoretical Physics, University of Cambridge, 2007.

[10] A. Wächter and L. T. Biegler, *On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming*, Math. Program., 106 (2006), pp. 25–57.

An Optimizer in the Telecommunications Industry

Mauricio G. C. Resende
Algorithms and Optimization Research Department,
AT&T Labs Research, Florham Park, New Jersey 07932,
USA (mgcr@research.att.com).

1. Introduction

The impact of telecommunications in modern life in the last 100 years is remarkable. Telecommunications has evolved from telegraphy to landline local telephony, to long distance telephony, to mobile telephony, and to the Internet, which now carries voice, video, television, instant messaging, and makes electronic commerce possible. Optimization problems are abundant in the telecommunications industry and the successful solution of these problems has played an important role in telecommunications and its widespread use. Optimization arises in problems as varied as planning and design of optical and wireless networks, routing, restoration, network survivability, e-commerce, and search engine design [31].

In this article, I relate some of the optimization problems I have run into while a Member of Technical Staff at AT&T Bell Labs and AT&T Labs. I have worked in the Algorithms and Optimization Research Department, headed by David S. Johnson, since graduating with a Ph.D. in operations research from the IEOR Department at UC Berkeley. Members of the Algorithms and Optimization Research Department do research on theoretical and experimental algorithmics with a focus on optimization. The department's research spans topics such as computational complexity, approximation algorithms, linear and integer programming, network programming, network design, routing, location, data structures, algorithm engineering, and metaheuristics.

Optimization problems reach us in many ways. Some problems are brought to us by researchers in other fields who often have optimization problems they need to solve. Other problems arise from business-related projects. Sometimes it may be that we have previously done work related to that problem or on a related problem and our tools (solution methods or software) can be directly applied to the problem. In other circumstances, new tools may need to be developed.

The reader will observe that most of the examples described in this article involve the use of metaheuristics [19] to find cost-effective solutions to combinatorial optimization problems. Metaheuristics are high-level procedures that coordinate simple heuristics, such as local search, to find solutions that are of better quality than those found by the simple heuristics alone. The metaheuristic most used in this article is GRASP, or greedy randomized adaptive search procedures [15], and hybridizations of GRASP with other metaheuristics. Other metaheuristics used are path-relinking [18], genetic [28] and memetic [26] algorithms, variable neighborhood search [21], and evolutionary path-relinking [34].

The article is organized as follows. In Section 2. we examine two problems that arise in the management of points of presence of an Internet service provider. This is followed in Section 3. with routing problems, in Section 4. with network design problems, in Section 5. with network migration problems, and in Section 6. with a data mining application. Concluding remarks are made in Section 7..

2. Location problems

An Internet Service Provider (ISP) offering dialup access needs to determine where its modems will be located. Such a location is called a point of presence, or simply a PoP. We describe two problems. In the first, PoP locations need to be determined, while in the second, redundant PoPs need to be identified and shut down.

2.1 PoP placement

In the mid-1990s, when AT&T planned the U.S. rollout of its ISP (AT&T Worldnet), it was faced with the problem of where to locate the PoPs. Given a fixed number of modem pools that could be deployed, a set of about 50,000 potential PoP locations, and the location of each AT&T customer, the task was to determine the location of each modem pool such that the largest number of customers would be able to place a 'free' local call to at least one modem. Free local calls in the U.S. are those made to numbers at most about 15 miles (24 km) away.

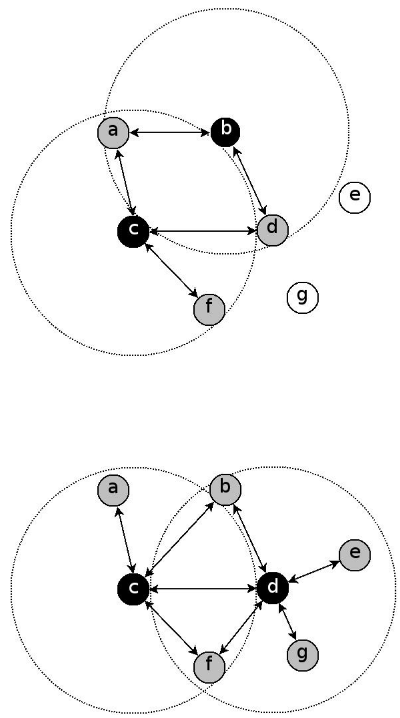

Figure 1: Two PoP placement solutions are compared. The selected PoPs are shown in black, the covered areas are grey, and the uncovered areas are white. The area of coverage is indicated by the dashed circle around the PoP. The bottom solution ($\{c, d\}$) covers all of the areas while the solution on top ($\{b, c\}$) leaves two areas (e and g) uncovered.

For this maximum customer coverage problem we designed an LP-based tool to compute an upper bound on the number of customers that could be covered and a tool [29] that used the metaheuristic GRASP to compute a placement within 1% of the optimum. Besides being very happy with the near-optimum placements produced by the tool, the planners also found the upper bound computation very useful since it enabled them to estimate the minimum number of modem pools needed to achieve a certain level of coverage.

2.2 PoP elimination

Over time, the GRASP-based tool was used to increase the number of PoPs deployed and consequently the coverage. Though a call to a PoP is free to the customer, it is not free to the ISP. Each PoP has an associated network cost, which is the hourly rate paid by the ISP to the network company transporting the access traffic. This hourly rate can vary greatly from PoP to PoP.

Since network costs and coverages of PoPs differ, an opportunity to eliminate PoPs could arise as long as coverage remained unchanged and the cost did not increase. In 2003, we conducted a study to determine if there were any PoPs that could be eliminated. We formulated this as a p-median problem which we solved with a tool based on a GRASP with evolutionary path-relinking [34].

Currently covered customers were grouped into about 70,000 exchanges. Each exchange was a p-median user and each of 1035 PoPs was a p-median facility. The distance, or cost, between a user and a facility was defined to be the network cost, *i.e.*, the product of the PoP rate and the number of hours used by the exchange. Solving the p-median with $p = 1035$ resulted in the network cost with no PoP eliminated. We wanted the smallest value of p that preserved this cost. By solving a series of p-median problems with decreasing values of p, we determined that over 30% of the PoPs could be eliminated while maintaining the same coverage and not increasing the network cost.

3. Traffic routing

Routing of traffic is perhaps the most critical operational problem in telecommunication networks. We consider here two routing problems, the routing of virtual private circuits and Internet traffic routing.

3.1 Routing of private virtual circuits

Telecommunication service providers offer virtual private networks to customers by provisioning a set of permanent (long-term) private virtual circuits (PVCs) between endpoints on a large backbone network. During the provisioning of a PVC, routing decisions are made either automatically by the switch (or router) or by the network designer, through the use of preferred routing assignments and without any knowledge of future requests. Over time, these decisions usually cause inefficiencies in the network and occasional rerouting of the PVCs is needed. The new routing scheme is then implemented on the network through preferred routing assignments. Given a preferred routing assignment, the network will move the PVC from its current route to the new preferred route as soon as this move becomes feasible.

One possible way to create the preferred routing assignments is to appropriately order the set of PVCs currently in the network and apply an algorithm that mimics the routing algorithm used by the switch (or router) to each PVC in that order. However, more elaborate routing algorithms, that take into account factors not considered by the switch, could further improve the efficiency of network resource utilization.

Typically, the routing scheme used to automatically provision PVCs is also used to reroute them in case of network failures, such as trunks (that transport traffic between routers) or cards (located on routers). Therefore, this routing algorithm should be efficient in terms of running time, a requirement that can be traded off for improved network resource utilization when building preferred routing assignments offline.

We solved this problem with variants of a GRASP with path-relinking algorithm for the problem of routing offline a set of PVC demands over a backbone network, such that a combination of the delays due to propagation and congestion was minimized [30, 32]. This problem and its variants are also known in the literature as bandwidth packing problems. The set of PVCs to be routed can include all or a subset of the PVCs currently in the network, and/or a set of forecast PVCs. The explicit handling of propagation delays, as opposed to just handling the number of hops (as in the routing algorithm implemented in some switches) is particularly important in international networks, where distances between backbone nodes vary considerably. The minimization of network congestion is important for providing the maximum flexibility to handle the following situations:

- overbooking, which is typically used by network designers to account for non-coincidence of traffic;
- PVC rerouting, due to link or card failures; and
- bursting above the committed rate, which is not only allowed but sold to customers as one of the attractive features of some services.

3.2 Routing of Internet traffic

Intradomain routing

The Internet is divided into many routing domains, called autonomous systems (ASes). ASes are networks that consist of routers and links connecting the routers. When customer and peer routers are considered, these ASes can have thousands of routers and links. ASes interact to control and deliver Internet Protocol (IP) traffic. They typically fall under the administration of a single institution, such as a company, a university, or a service provider. Neighboring ASes use the Border Gateway Protocol (BGP) to route traffic.

The goal of intradomain traffic engineering consists in improving user performance and making more efficient use of network resources within an AS. Interior Gateway Protocols (IGPs) such as OSPF (Open Shortest Path First) and IS-IS (Intermediate System-Intermediate System) are commonly used to select the paths along which traffic is routed within an AS. These routing protocols direct traffic based on link weights assigned by the network operator. Each router in the AS computes shortest paths and creates destination tables used to direct each IP

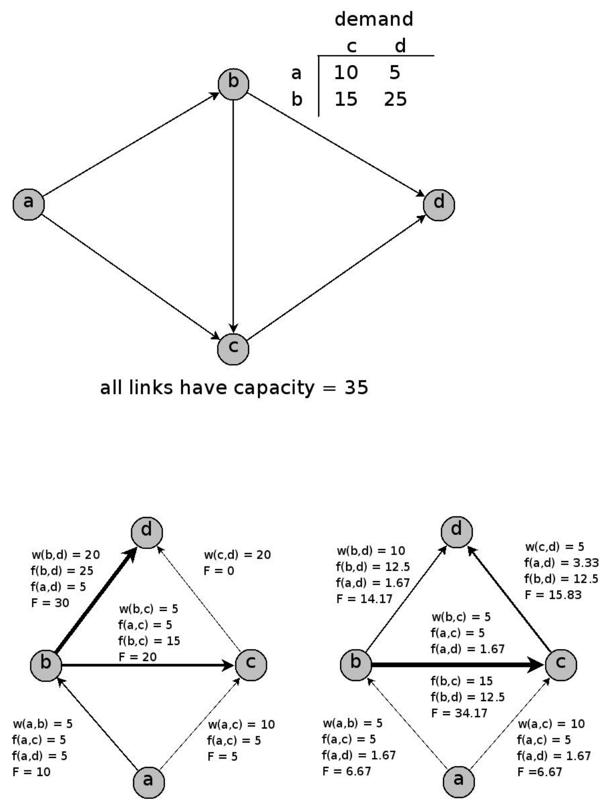

Figure 2: Example of OSPF weight setting. The demands shown next to the network on top of the figure must be routed. All links have capacity 35. Two solutions are shown on the bottom of the figure. $w(u, v)$ is the OSPF weight of link (u, v), the $f(x, y)$ indicated on each link is the amount of traffic originated at router x and going to router y that passes through the link, and F is the total traffic on the link. The solution on the left is better than the one on the right. It has a maximum utilization of 86% (link (b, d)), while the other has a maximum utilization of 98% (link (b, c)).

packet to the next router on the path to its final destination. OSPF calculates routes as follows. Each link is assigned an integer weight ranging from 1 to 65535. The weight of a path is the sum of the link weights on the path. OSPF mandates that each router compute a graph of shortest paths with itself as the root. This graph gives the least weight routes (including multiple routes in case of ties) to all destinations in the AS. In the case of multiple shortest paths originating at a router, OSPF is usually implemented so that it will accomplish load balancing by splitting the traffic flow over all shortest paths leaving from each router. OSPF requires routers to exchange routing information with all the other routers in the AS. Complete network topology knowledge is required for the computation of the shortest paths.

Given a set of traffic demands between origin-destination pairs, the *OSPF weight setting problem* consists in determining weights to be assigned to the links so as to optimize a cost function, typically associated with a network congestion measure.

We proposed two solution methods for this problem: a genetic algorithm [14] and a hybrid genetic, or memetic, algorithm [8] incorporating a local improvement procedure to the crossover operator of the genetic algorithm proposed in [14]. The local improvement procedure makes use of an efficient dynamic shortest path algorithm [10] to recompute shortest paths after the modification of link weights. The memetic algorithm improved upon the pure genetic algorithm, producing better-quality solutions in less time as can be seen in Figure 3. The memetic algorithm was also shown to be more robust than the local search based approach of [16, 17].

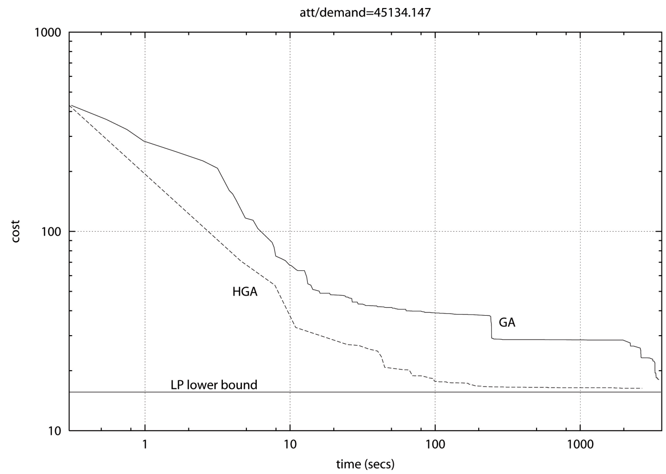

Figure 3: Congestion cost comparison of genetic (GA) and hybrid genetic (HGA), or memetic, algorithms on a 1-hour run on an IP network with 90 routers and 274 links. HGA converges faster to a better solution than the pure GA. LP lower bound is obtained by solving a multi-commodity flow problem where traffic is allowed to take any route.

Interdomain routing

The Internet's two-tiered routing architecture was designed to have a clean separation between the intradomain and interdomain routing protocols. For example, the interdomain routing protocol allows routers at the border of the AS to learn how to reach external destinations, whereas the intradomain protocol determines how to direct traffic from one router in the AS to another. However, the appropriate roles of the two protocols becomes unclear when the AS learns routes to a destination at multiple border routers — a common situation today. Since service providers peer at multiple locations, essentially all of the traffic from customers to the rest of the Internet has multiple egress routers (routers in which traffic leaves the AS). In addition, many customers connect to their provider in multiple locations for fault tolerance and more flexible load balancing, resulting in multiple egress routers for these destinations as well. Selecting among multiple egress points is now a fundamental part of the Internet routing architecture, independent of the current set of routing protocols.

In the Internet today, border routers learn routes to destination prefixes via BGP. When multiple border routers have routes that are 'equally good' in the BGP sense (*e.g.*, local preference, AS path length, *etc.*), each router in the AS directs traffic to its *closest* border router, in terms of the IGP distances. This policy of *early-exit* or *hot-potato* routing is hard-coded in the BGP decision process implemented on each router, offering consistent forwarding of packets.

Although consistent forwarding is clearly an important property for any routing system, we [35] believe that hot-potato routing is disruptive and convoluted. Small changes in IGP distances can sometimes lead to large shifts in traffic, long convergence delays, and BGP updates to neighboring domains.

Network administrators are forced to evaluate the impact of changes in the IGP metrics on BGP routing decisions, rather than viewing the two parts of the routing system separately.

Selecting the egress point and computing the forwarding path to the egress point are two very distinct functions, and we believe that they should be decoupled. Paths *inside* the network should be selected based on some meaningful performance objective, whereas *egress selection* should be flexible to support a broader set of traffic-engineering goals. These objectives vary by network and destination prefix; therefore a mechanism that imposes a single egress selection policy cannot satisfy this diverse set of requirements.

In [35], we propose a new mechanism for each router to select an egress point for a destination, by comparing the candidate egress points based on a weighted sum of the IGP distance and a constant term. The configurable weights provide flexibility in deciding whether (and how much) to base BGP decisions on the IGP metrics. Network-management systems apply linear and integer programming techniques to automatically set these weights to satisfy network-level objectives, such as balancing load and minimizing propagation delays. Our new mechanism is called TIE (Tunable Interdomain Egress) because it controls how routers break ties between multiple equally-good BGP routes. Our solution does not introduce any new protocols or any changes to today's routing protocols, making it possible to deploy our ideas at one AS at a time and with only minimal changes to the BGP decision logic on IP routers.

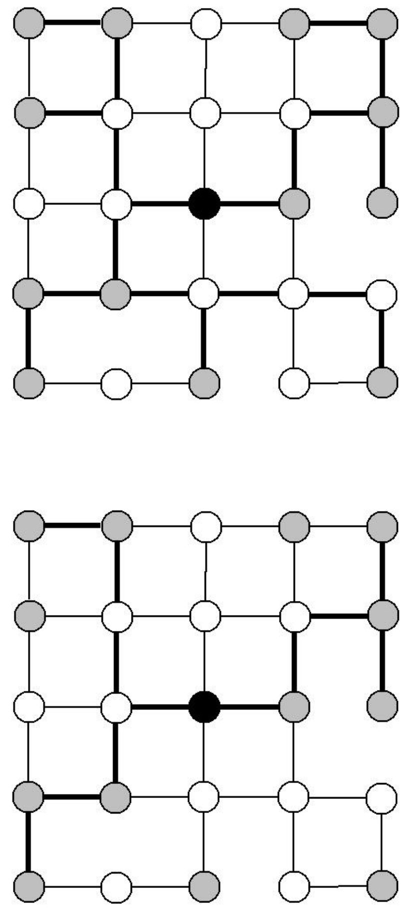

Figure 4: Two fiber networks are compared. Backbone node is shown in black, potential revenue generating premises are shown in grey, street corners and zero revenue producing premises are shown in white. Deployed fiber is shown in bold. The top network provides service to all potential revenue generating premises while bottom network provides service only to profitable premises.

4. Design problems

Network design problems are among the most important applications of optimization in telecommunications. We describe two instances of design problems that we have worked on. In the first, we address the optimization of the tradeoff between revenue generation by a network and the cost of building the network. In the second, we consider survivable IP network design.

4.1 To lay or not to lay fiber

Telecommunication service providers often must decide whether an investment to lay optical fiber cable to provide broadband service to customers is worthwhile. There is a cost associated with laying fiber which depends not only on the total length of the fiber but also on where and how the fiber is laid. For example, underground fiber could be more expensive than overground fiber.

Given a geographic area made up of customer premises and street segments connecting these premises, an estimate of the present value of potential revenue that could be earned from each customer premise, and the present value of the cost of laying fiber in each street segment, a service provider would like to maximize the difference between the revenue earned from customers reached by the fiber and the total cost of the fiber. An objective function value above some specified threshold would indicate the feasibility of the investment.

This type of problem, called the *prize-collecting Steiner tree problem*, can be used to order the attractiveness of different markets when rolling out services such as broadband Internet access, IPTV, and voice over IP. In fact, our motivation for studying this problem was the Telecommunications Act of 1996 which opened up the telecommunications markets in the U.S. and allowed AT&T to compete in local markets. With so many markets to choose to compete in, it was important to determine which were the most attractive.

Work on implementing the approximation algorithm of Goemans and Williamson [20] was done by Johnson *et al.* [24], and we were interested in accessing the quality of the solutions found with the approximation algorithm. In [25], we proposed a cutting planes algorithm to produce strong bounds for the prize collecting Steiner tree problem and showed that most solutions found by the approximation algorithm had a gap with respect to the bounds. In [11], we introduced a new type of GRASP, where randomized construction is done by perturbing the data and applying an approximation algorithm (in this case, the Goemans and Williamson algorithm) on the perturbed problem. In addition to the construction, our heuristic consisted of local search, path-relinking, and a variable neighborhood search post-optimization phase. The GRASP heuristic applies the approximation algorithm at least one time using the original data. Consequently, solutions that it produces are always at least as good as those found by the approximation algorithm. On 84.2% of 114 benchmark test problems, solutions produced by the GRASP heuristic were better than those produced by the approximation algorithm alone. On 91.2% of the instances the solutions found by the GRASP heuristic were provably optimal.

4.2 Survivable IP network design

With the pervasiveness of IP networks in telecommunications, an important question faced by network operators is how to design robust cost-efficient networks on which traffic will be routed with OSPF. Given a network topology (*i.e.*, a set of nodes and a set of arcs where links can be installed), predicted traffic demands, a set of link types to be deployed, each having a different capacity and installation cost per mile, the *survivable IP network design problem* consists in finding a set of OSPF arc weights and the number of each link type deployed on each arc such that network cost is minimized. We further require that in a no-failure or any node/arc failure situation there is enough installed capacity to move all of the predicted traffic.

In [9], we proposed a genetic algorithm to find cost-efficient solutions for this problem for the case in which there is a single link type and number of copies of the link as well the link OSPF weights have to be determined. Since real networks can be built using many different link types (*e.g.*, OC3, OC12, OC48, OC196) having different capacities as well as costs, in [4], we extended the design algorithm described in [9] to handle different link types.

5. Network migration

Network migration arises when traffic is moved from an outdated network to a new network. We consider in this section two applications of network migration: migration of phone traffic from a 4ESS switch based network to a router based IP network and telephone migration from an old PBX switch to a new PBX switch.

5.1 Voice traffic migration

Consider the problem where inter-nodal traffic from an outdated telecommunications network is to be migrated to a new network. Nodes are migrated, one at each time period, from the old to the new network. All traffic originating or terminating at a given node in the old network is moved to a specific node in the new network. Routing is predetermined on both networks and therefore arc capacities are known. Traffic between nodes in the same network is routed in that network.

Nodes are migrated, one at a time, in some predetermined order. When a node is migrated, one or more temporary arcs may need to be set up since the node in the new network to which the traffic is migrated may be adjacent to one or more still active nodes in the old network. A temporary arc remains active until both nodes connected by the arc are migrated to the new network.

In one version of the *network migration scheduling problem*, one seeks an ordering of the migration of the nodes that minimizes the maximum sum of capacities of the temporary arcs. In another version, the objective is to minimize the sum of the total capacities of the temporary arcs over each period in the planning horizon.

We were motivated to look at this problem when AT&T began planning the migration of its switch-based telephone traffic to a new router-based IP network. In [6], we present a GRASP with evolutionary path-relinking for these two variants of the migration problem.

5.2 PBX telephone migration

A PBX, or *private branch exchange*, is a private telephone network such as call forwarding, call recording, call transfer, and voice messaging.

Some PBX features require groups of phone numbers to be defined. These include, for example, multi-line hunt (MLH), call pickup (CPU), intercom (ICOM), series completion (SC), and shared telephone number (STN) groups. An MLH group consists of a cycle of phone numbers. When a call is made to a phone in the cycle and the call is not answered, it is transfered to the next phone in the cycle. This is repeated until someone picks up. A CPU group is a set of phone numbers where any

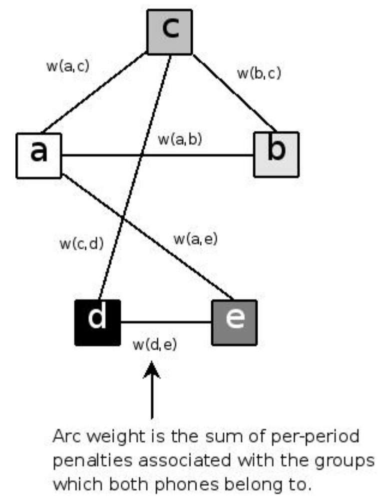

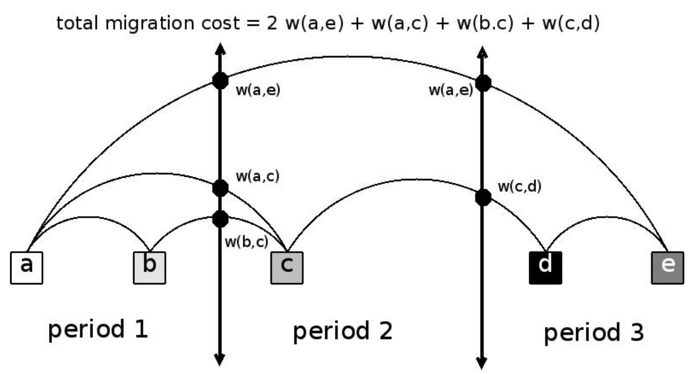

Figure 5: Example of PBX telephone migration. Five telephones are to migrate over a three period horizon. In each period at most two telephones can migrate. In the migration schedule shown, telephones a and b are migrated in period 1, telephone c is migrated in period 2, and telephones d and e are migrated in period 3. Penalties $w(a,b)$ and $w(d,e)$ do not contribute to the cost since both a and b migrate in period 1 and d and e both migrate in period 3. Since a-c, b-c, and d-e migrate in consecutive periods, their penalties are contributed once to the total cost. Since a and e are scheduled two periods apart, their penalty contributes twice to the total cost.

phone in the group can pickup a call made to any other phone in the group. Any phone in an ICOM group can speed dial to any other group member. A SC group is an ordered list of phone numbers. If a call made to the first phone is not answered, it is transfered to the next. This is repeated until someone picks up. If the last phone in the list does not pick up, voice mail answers the call. An STN group is a set of phone numbers for which calls made to them are answered by a single phone (*e.g.*, an assistant). In an enterprise there may exist several MLH, CPU, ICOM, SC, and STN groups and a single phone number may be a member of more than one group.

We consider a problem that arises when an enterprise acquires a new PBX to replace an existing one. Phone numbers need to migrate from the old system to the new system over a given time horizon. Each group has a penalty associated with it. If two phones in a given group migrate in different time periods, then a penalty is incurred. This penalty depends on the set of groups that these phones both belong to as well as the amount of time between the migrations of each phone. We further require that during each time period a specified maximum number of phones are allowed to migrate and assume that there are sufficient periods in the planning horizon to allow for a feasible schedule.

The objective is to schedule the migration plan so that the total migration penalty is minimized. This involves assigning phone numbers to time periods such that no more than the maximum number of phones are assigned to a single period.

We learned of this problem when we were contacted by people implementing such a move for a large investment bank that acquired a new PBX from AT&T. This problem involved migrating over 2500 phones belong to one or more of about 400 groups. Since there was a limit of 375 phones that could be moved per period, the move could not be done in less than eight periods. Since moving phones sharing one or more groups in different time periods could cause a business disruption, we wanted to minimize any possible disruption caused by the migration. To do this we assigned different penalties to phone pairs sharing different groups. Those groups whose disruption would be the most critical had the largest penalties. In [5], we present a GRASP with three local neighborhood structures for this problem.

6. Data mining

The proliferation of massive data sets [2] brings with it a series of special computational challenges to the optimization community. This *data avalanche* arises in a wide range of scientific and commercial applications. With advances in computer and information technologies, many of these challenges are being addressed by diverse inter-disciplinary groups, that include computer scientists, mathematicians, statisticians and engineers, working in close cooperation with application domain experts.

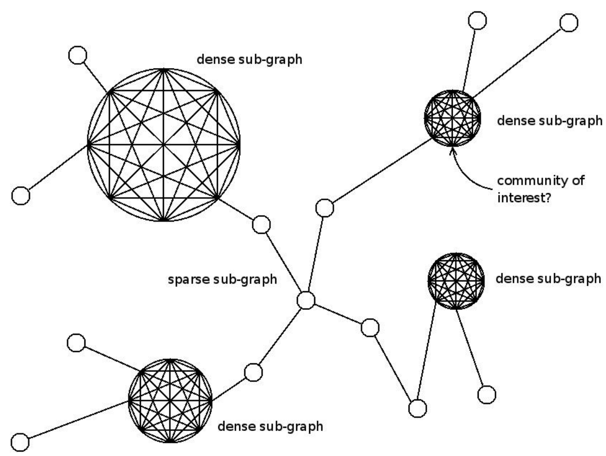

Figure 6: Macro structure of call detail graph. Graph is very sparse with spread-out clusters of very dense sub-graphs. These dense clusters can be interpreted as *communities of interest*.

In 1997, we began exploring the massive graphs associated with telephone calls. When we set out to study these graphs, we had no particular application in mind. We were simply interested in investigating the structure of these graphs. At that time, each call made on the AT&T network generated a record of about 200 bytes with information that included the pair of phone numbers involved in the phone call. The set of all these records is the *call detail database*. About 250 million records were generated on an average day in 1997 with around 320 million on a busy day, corresponding to about 18 terabytes of data per year. Given a time window, a call detail graph can be extracted from the database. In this directed graph, each phone number is a node and for each call placed during the time window an edge exists from the calling number to the called number. Because of privacy concerns, the script that extracted the graph from the database mapped the phone numbers to integers from 1 to n, where n is the number of phones involved in calls.

We describe an experiment involving 12 hours of calls in 1997 [1]. The corresponding graph had 53,767,087 vertices and over 170 million edges. We found 3,667,448 connected components out of which only 302,468 were components of size greater than 3 (there were 255 self-loops, 2,766,206 pairs and 598,519 triplets).

A giant component with 44,989,297 vertices was detected. It is interesting to observe that this is similar to what is predicted by random graph theory even though the call graphs are certainly not random. The giant component has 13,799,430 directed depth first search trees (DFSTs) and one of them is a giant DFST (it has 10,355,749 vertices and 19,072,448 edges). Most of the DFSTs have no more than 5 vertices. The interesting trees have sizes between 5 and 100. Their corresponding induced subgraphs are most of the time very sparse, except for some occasional dense subgraphs with 11 to 32 vertices.

We argued that the largest clique in this component has size not greater than 32. Cliques are either within a subgraph induced by the vertices of a DFST, or distributed among the different DFSTs. We expected the former to occur. There are several large DFSTs, the largest having about 19 million edges. By counting the edges in the trees, one observes that there remain very few edges to go between trees and consequently it is more likely that cliques are within the graphs induced by the nodes of a tree. Since the largest dense subgraph induced by the vertices of a tree had 32 vertices, we did not expect many cliques larger that 32 to be found.

To begin our experimentation, we considered 10% of the edges in the large component from which we recursively removed all vertices of degree one. This resulted in a graph with 2,438,911 vertices and 5,856,224 edges, which fit in memory. In this graph we searched for large cliques with a GRASP for maximum clique. Our first motivation was to identify a lower bound on the size of the maximum clique so that we could delete higher-degree vertices on larger portions of the graph to possibly identify larger cliques. The GRASP was repeated 1000 times, with each iteration producing a locally maximal clique. Though applying local search on every constructed solution may not be efficient from a running time point of view, we applied local search to all constructed solutions to explore its effect in improving clique sizes. Because of the independent nature of the GRASP iterations and since our computer was configured with 20 processors, we created 10 threads, each independently running GRASP starting from a different random number generator seed.

Table 1: Cliques found by `construct` and `local`

size	cliques found by `construct`	`local`	distinct cliques
2	63	62	
3	473	320	
4	95	176	
5	73	103	14
6	116	95	11
7	59	38	25
8	54	63	28
9	22	33	14
10	17	10	9
11	15	38	35
12	10	32	22
13	1	26	18
14	0	3	3
15	0	1	1

Table 1 summarizes the first part of the experimental results. It shows, for each clique size found, the number of GRASP iterations that constructed or improved such solution, and from sizes 5 to 15, the number of distinct cliques that were found by the GRASP iterations. It is interesting to observe that these cliques, even though distinct, share a large number of vertices. Applying a greedy procedure to these cliques to identify a disjoint set of cliques produced one clique of size 15, 12, 9, and 7, four cliques of size 6, and five of size 5.

Next, we considered 25% of the edges in the large component from which we recursively removed all vertices of degree 10 or less. The resulting graph had 291,944 vertices and 2,184,751 edges. 12,188 iterations of GRASP produced cliques of size 26.

Having found cliques of size 26 in a quarter of the graph, we next intensified our search on the entire huge connected component. In this component, we recursively removed all vertices of degree 20 or less. The resulting graph had 27,019 vertices and 757,876 edges.

Over 20,000 GRASP iterations were carried out on the 27,019 vertex – 757,876 edge graph. Cliques of 30 vertices were found. These cliques are very likely to be optimal because we do not expect cliques larger than 32 vertices to be found. The local search can be seen to improve the constructed solution not only for large constructed cliques, but also for small cliques. In fact, in 26 iterations, constructed cliques of size 3 were improved by the local search to size 30.

Finally to increase our confidence that the cliques of size 30 found were maximum, we recursively removed all vertices of degree 30 or less, resulting in a graph with 8724 vertices and about 320 thousand edges. We ran 100,000 GRASP iterations on the graph taking 10 parallel processors about one and a half days to finish. The largest clique found had 30 vertices. Of the 100,000 cliques generated, 14,141 were distinct, although many of them shared one or more vertices.

Quasi-cliques are dense sub-graphs, *i.e.*, they are cliques with a few missing edges. To compute quasi-cliques [3] on this test data, we looked for large quasi-cliques with densities 90%, 80%, 70%, and 50%. Quasi-cliques of sizes 44, 57, 65, and 98, respectively, were found.

It was surprising to us that in only 12 hours of phone calls we found groups of 30 phone numbers where each one either called or was called by all 29 others and groups of 98 where each called or was called by at least half of the other phones. To date, this is the research that we have been involved with that has received the most attention from the media [7, 12, 13, 22, 23].

7. Concluding remarks

In this article, we show a sample of optimization problems that arise in an optimization research department at a telecommunications service provider. Most of the interesting questions we see are NP-hard combinatorial optimization problems. Though we make use of linear and integer programming solvers in many instances, for most cases we use metaheuristics. Metaheuristics, such as GRASP with path-relinking [33], are widely applicable, produce cost-efficient solutions, are relatively easy to implement, and therefore can quickly provide good-quality solutions to problems that arise in practice.

REFERENCES

[1] J. Abello, P. M. Pardalos, and M. G. C. Resende, *On maximum cliques problems in very large graphs*, J. Abello and J. Vitter, editors, *External Memory Algorithms*, DIMACS Series on Discrete Mathematics and Theoretical Computer Science, 50 (1999), pp. 119–130, American Mathematical Society, Providence.

[2] J. Abello, P. M. Pardalos, and M. G. C. Resende, editors, *Handbook of Massive Data Sets*, Kluwer Academic Publishers, Dordrecht, 2002.

[3] J. Abello, M. G. C. Resende, and S. Sudarsky, *Massive quasi-clique detection*, S. Rajsbaum, editor, *LATIN 2002: Theoretical Informatics*, Lecture Notes in Computer Science, 2286 (2002), pp. 598–612, Springer-Verlag, Heidelberg.

[4] D. V. Andrade, L. S. Buriol, M. G. C. Resende, and M. Thorup, *Survivable composite-link IP network design with OSPF routing*, Proceedings of The Eighth INFORMS Telecommunications Conference, Dallas, Texas, 2006.

[5] D. V. Andrade and M. G. C. Resende, *A GRASP for PBX telephone migration scheduling*, Proceedings of The Eighth INFORMS Telecommunications Conference, Dallas, Texas, 2006.

[6] D. V. Andrade and M. G. C. Resende, *GRASP with path-relinking for network migration scheduling*, Proceedings of International Network Optimization Conference (INOC 2007), Spa, Belgium, (2007).

[7] A.-L. Barabási, *Linked: How everything is connected to everything else and what it means for business, science, and everyday life*, Plume, New York, 2003.

[8] L. S. Buriol, M. G. C. Resende, C. C. Ribeiro, and M. Thorup, *A hybrid genetic algorithm for the weight setting problem in OSPF/IS-IS routing*, Networks, 46 (2005), pp. 36–56.

[9] L. S. Buriol, M. G. C. Resende, and M. Thorup, *Survivable IP network design with OSPF routing*, Networks, 49 (2007), pp. 51–64.

[10] L. S. Buriol, M. G. C. Resende, and M. Thorup, *Speeding up dynamic shortest-path algorithms*, INFORMS J. Comput., to appear.

[11] S. A. Canuto, M. G. C. Resende, and C. C. Ribeiro, *Local search with perturbations for the prize-collecting Steiner tree problem in graphs*, Networks, 38 (2001), pp. 50–58.

[12] B. Cipra, *Massive graphs pose big problems*, SIAM NEWS, 32 (1999).

[13] The Economist, *Needles in giant haystacks*, (January 30th – February 5th, 1999), pp. 74,

[14] M. Ericsson, M. G. C. Resende, and P. M. Pardalos, *A genetic algorithm for the weight setting problem in OSPF routing*, J. Comb. Optim., 6 (2002), pp. 299–333.

[15] T. A. Feo and M. G. C. Resende, *Greedy randomized adaptive search procedures*, J. Global Optim., 6 (1995), pp. 109–133.

[16] B. Fortz and M. Thorup, *Internet traffic engineering by optimizing OSPF weights*, Proc. 19th IEEE Conf. on Computer Communications (INFOCOM), IEEE Press, (2000), pp. 519–528.

[17] B. Fortz and M. Thorup, *Increasing internet capacity using local search*, Comput. Optim. and Appl., 29 (2004), pp. 13–48.

[18] F. Glover, *A Template for Scatter Search and Path Relinking*, AE '97: Selected Papers from the Third European Conference on Artificial Evolution, Springer-Verlag, London, (1998) pp. 3–54.

[19] F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, Kluwer Academic Publishers, Dordrecht, 2003.

[20] M. X. Goemans and D. P. Williamson, *The primal dual method for approximation algorithms and its application to network design problems*, D. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*, PWS Publishing Co., Boston, (1996) pp. 144–191.

[21] P. Hansen and N. Mladenović, *Variable neighborhood search*, F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, Kluwer Academic Publishers, Dordrecht, (2003) pp. 145–184.

[22] B. Hayes, *Computing science graph theory in practice: Part I*, American Scientist, 88 (2000), pp. 9–13.

[23] B. Hayes, *Connecting the dots*, American Scientist, 94 (2006), pp. 400–404.

[24] D. S. Johnson, M. Minkoff, and S. Phillips, *The prize collecting steiner tree problem: Theory and practice*, Proc. 11th Ann. ACM-SIAM Symp. on Discrete Algorithms, (2000) pp. 760–769.

[25] A. Lucena and M. G. C. Resende, *Strong lower bounds for the prize collecting Steiner problem in graphs*, Discrete Appl. Math., 141 (2004), pp. 277–294.

[26] P. Moscato and C. Cotta, *A gentle introduction to memetic algorithms*, F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, Kluwer Academic Publishers, Dordrecht, (2003), pp. 105–144.

[27] P. M. Pardalos and M. G. C. Resende, editors, *Handbook of Applied Optimization*, Oxford University Press, New York, 2002.

[28] C. Reeves, *Genetic Algorithms*, F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, Kluwer Academic Publishers, Dordrecht, (2003) pp. 55–82.

[29] M. G. C. Resende, *Computing approximate solutions of the maximum covering problem using GRASP*, J. Heuristics, 4 (1998), pp. 161–171.

[30] L. I. P. Resende and M. G. C. Resende, *A GRASP for frame relay permanent virtual circuit routing*, P. Hansen and C. C. Ribeiro, editors, Proceedings of the III Metaheuristics International Conference (MIC99), Angra dos Reis, Brazil (1999), pp. 397–401.

[31] M. G. C. Resende and P. M. Pardalos, editors, *Handbook of Optimization in Telecommunications*, Springer-Verlag, New York, 2006.

[32] M. G. C. Resende and C. C. Ribeiro, *A GRASP with path-relinking for private virtual circuit routing*, Networks, 41 (2003), pp. 104–114.

[33] M. G. C. Resende and C. C. Ribeiro, *GRASP with path-relinking: Recent advances and applications*, T. Ibaraki, K. Nonobe, and M. Yagiura, editors, *Metaheuristics: Progress as Real Problem Solvers*, Springer-Verlag, Heidelberg, (2005), pp. 29–63.

[34] M. G. C. Resende and R. F. Werneck, *A hybrid heuristic for the p-median problem*, J. Heuristics, 10 (2004), pp. 59–88.

[35] R. Teixeira, T. G. Griffin, M. G. C. Resende, and J. Rexford, *TIE breaking: Tunable interdomain egress selection*, IEEE/ACM Trans. on Networking, 15 (2007), pp. 761–774.

Bulletin

1. Event Announcements

Eighth SIAM Conference on Optimization (SIAM OP08)
May 10-13, 2008
Boston Park Plaza Hotel and Towers
Boston, Massachusetts, USA
http://www.siam.org/meetings/op08/index.php

The Ninth SIAM Conference on Optimization will feature the latest research in theory, algorithms, and applications in optimization problems. In particular, it will emphasize large-scale problems and will feature important applications in networks, manufacturing, medicine, biology, finance, aeronautics, control, operations research, and other areas of science and engineering. The conference brings together mathematicians, operations researchers, computer scientists, engineers, and software developers; thus it provides an excellent opportunity for sharing ideas and problems among specialists and users of optimization in academia, government, and industry.

Funding Agency

SIAM and the Conference Organizing Committee wish to extend their thanks and appreciation to the U.S. National Science Foundation for its support of this conference.

Themes

- Large-scale nonlinear programming
- Mixed integer nonlinear programming
- Conic and convex programming
- Stochastic optimization
- Discrete optimization
- Derivative-free optimization
- PDE constrained optimization
- Applications to biology and life sciences

Organizing Committee

- Kurt M. Anstreicher (Co-chair), University of Iowa
- Michael C. Ferris, University of Wisconsin
- Anders Forsgren, Royal Institute of Technology (KTH), Sweden
- Michel X. Goemans, Massachusetts Institute of Technology
- Eldad Haber, Emory University
- Sven Leyffer (Co-chair), Argonne National Laboratory
- Alexander Martin, Technische Universität Darmstadt, Germany
- Katya Scheinberg, IBM Research
- Claudia Sagastizabal, CEPEL, Brazil
- Jie Sun, National University of Singapore, Singapore

Plenary Speakers

- Etienne de Klerk, Tilburg University, Netherlands: *Exploiting Algebraic Symmetry in Semidefinite Programs: Theory and Applications*
- Matthias Heinkenschloss, Rice University: *PDE Constrained Optimization*
- Jan Modersitzki, University of Lubeck, Germany: *Mathematics Meets Medicine: An Optimal Alignment*
- Annick Sartenaer, Université Notre Dame de la Paix, Belgium: *Multi-Level Optimization*
- Stefan Scholtes, Cambridge University, United Kingdom: *Probability Management: Revisiting an Old Approach to Business Modeling under Uncertainty*
- Pascal Van Hentenryck, Brown University: *Constraint Programming*
- Andreas Wächter, IBM Research: *Nonlinear Integer Optimization*
- Robert Weismantel, University of Magdeburg, Germany: *Algebra, Geometry and Combinatorics of Mixed Integer Optimization Problems*

Invited Mini-Tutorial

Derivative-Free Optimization, by Charles Audet, École Polytechnique Montréal, Canada and Luis N. Vicente, University of Coimbra, Portugal.

Workshop High Performance Optimization Techniques 2008 (HPOPT 2008)
June 11-13, 2008
Tilburg, The Netherlands
`http://lyrawww.uvt.nl/~edeklerk/hpopt2008`

This will already be the tenth HPOPT workshop, and these meetings have in the past proven to be an excellent forum on specialized topics in optimization. The theme of HPOPT 2008 will be "Algebraic Structure in Semidefinite Programming". The meeting will run for three days: one tutorial day on the theme "Using Symmetry in Semidefinite Programming", followed by two days of invited presentations. More details on the workshop are available at web page of the workshop.

All talks at the workshop will be by invitation. The invited speakers who have agreed to participate (as of September 17th, 2007) are:

- M. Kojima (Tokyo Institute of Technology)
- J. Lasserre (CNRS, Toulouse)
- J. de Loera (University of California at Davis)
- M. Muramatsu (The University of Electro-communications, Tokyo)
- Yu. Nesterov (Catholic University Louvain)
- P. Parrilo (MIT)
- F. Rendl (Univ. Klagenfurt)
- M. Schweighofer (Universit de Rennes 1)
- L. Tunçel (University of Waterloo)
- F. Vallentin (CWI, Amsterdam) (Tutorial Lecturer)
- L. Vandenberghe (UCLA)

Instructions on how to register for HPOPT 2008 are available at the workshop web-page. Due to the small scale of the meeting the number of participants will be limited to about sixty, and early registration is therefore recommended.

This announcement has been posted by Etienne de Klerk and Monique Laurent (on behalf of the HPOPT 2008 organizers).

2. Other Announcements

Call for Nomination
SIAG/OPT Prize
`http://www.siam.org/prizes/nominations/nom_siagopt.php`

The SIAM Activity Group on Optimization Prize (SIAG/OPT Prize) will be awarded at the SIAM Conference on Optimization to be held May 10-13, 2008, in Boston, Massachusetts.

The SIAG/OPT Prize, established in 1992, is awarded to the author(s) of the most outstanding paper, as determined by the prize committee, on a topic in optimization published in English in a peer-reviewed journal. The eligibility period is the four calendar years preceding the year of the conference.

Eligibility

Candidate papers must be published in English in a peer-reviewed journal and must bear a publication date in the 2004-2007 calendar years (January 1, 2004 – December 31, 2007). Candidate papers must contain significant research contributions to the field of optimization, as commonly defined in the mathematical literature, with direct or potential applications.

Description of the Award

The award will consist of a plaque and a certificate containing the citation. At least one of the prize recipients is expected to attend the award ceremony and to present the paper at the conference.

Nominations

A letter of nomination, including a bibliographic citation of the paper, should be sent by January 15, 2008, to the address below. Nominations by e-mail are preferred.

SIAM Activity Group on Optimization Prize Committee
Professor Robert Vanderbei, Chair
c/o J. M. Littleton
SIAM
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688
USA

E-mail: littleton@siam.org
Telephone: +1-215-382-9800 ext. 303
Fax: +1-215-386-7999

Selection Committee

Members of the selection committee are: Robert Vanderbei, Princeton University (Chair), Philip Gill, University of California – San Diego, Raphael Hauser, Oxford University, UK, Dick den Hertog, Tilburg University, The Netherlands, and David Karger, Massachusetts Institute of Technology.

Chairman's Column

It is hard to believe that this is the last column that I will write for Views-and-News. The terms of the current SIAG/OPT officers are up at the end of 2007, and SIAM will soon be conducting an election for our replacements. I have enjoyed serving as Chair of our activity group and trust that my successor will have as positive an experience. Executive Director Jim Crowley and the entire SIAM staff have been a pleasure to work with and I would like to thank them for all of their support.

The end of my term as Chair will certainly not, however, be the end of my involvement with SIAG/OPT. Conference co-chair Sven Leyffer and I, along with the other members of the Organizing Committee, have been involved with the planning of the 2008 SIAM Optimization Conference (OP08) for many months. We are delighted with the plenary speakers and invited minisymposia/minitutorial organizers who have already agreed to participate. The deadline for contributed minisymposia, talks and poster presentations is fast approaching and we look forward to planning the conference sessions once all the material is in. The middle of May should be a beautiful time in Boston and the conference schedule was chosen to make it as easy as possible for participants to obtain discounted air tickets without paying for an extra night at a hotel. The conference will be a wonderful opportunity to stay abreast of recent developments while at the same time renewing old acquaintances and making new ones. I am certainly looking forward to OP08 and I hope to see many, many members of our SIAG there!

Kurt M. Anstreicher, SIAG/OPT Chair
Department of Management Sciences
University of Iowa
S210 PBB Iowa City, IA 52242
USA
kurt-anstreicher@uiowa.edu
http://www.biz.uiowa.edu/faculty/anstreicher

Comments from the Editor

It is frequent to hear how pervasive optimization is nowadays in industry. Academic optimizers who care about the applied impact of their research certainly support the promise of this claim, but feel that the potential of optimization is still largely underestimated in practice.

We all know that there are many ways to improve the current situation, but none is more effective than personal contacts between academic researchers and industrial partners. On the other hand, the higher the qualifications of industrial research staff members the easy is to establish and enhance such contacts and collaborations. We do need more qualified optimizers in industry and more of those in a position to promote the field.

It is therefore a great pleasure for me to publish two articles from two reputed optimizers who spent a significant part of their careers in the industry. Their expertise and knowledge of applied optimization is invaluable. I asked them to describe their own experiences and thoughts on the matter. The first article reports the experience of Andrew Conn (IBM Research) as a continuous optimizer and his involvement in circuit design and petroleum engineering. The second article is written by Mauricio Resende (AT&T Labs) and describes his intense portfolio of applications of discrete and network optimization to telecommunications problems.

Some unsuccess in running this issue shows how optimizers are busy in industry. I was unable to obtain all the articles I wanted to, and to get the current ones on time. This is certainly not a complaint. I am extremely grateful to Andy and Mauricio for having spent some of their very busy time writing these excellent contributions.

Luís N. Vicente, Editor
Department of Mathematics
University of Coimbra
3001-454 Coimbra
Portugal
lnv@mat.uc.pt
`http://www.mat.uc.pt/~lnv`