



AT THE FOREFRONT OF ANALYTICS IN AFRICA



INDUSTRIAL SCALE OPTIMISATION IN
OPERATIONS RESEARCH

ORSSA Newsletter June 2014

www.orssa.org.za



FROM THE EDITOR

BY MARK EINHORN (einhorn@sun.ac.za)
 ORSSA NEWSLETTER EDITOR



Mark Einhorn

Greetings to all you ORSSA members out there and welcome to the first edition of the Newsletter for 2014. I trust that you all had a peaceful and enjoyable festive season and that 2014 has been kind to you thus far.

I would like to take this opportunity to wish our new president, Hennie Kruger and his Executive Committee all the best for the coming year, and

I have no doubt that they will take all steps necessary to continue the advancement, and increase the awareness of Operations Research in South Africa.

Right, now that the pleasantries are taken care of, let's delve into the contents of this Newsletter, shall we? I would like to first draw your attention to the very important notice on page 16 regarding the change in banking details of the society. Please do read it.

We get a double-dose of our new president in this issue in the form of his first presidential column as well as a very interesting interview with Hennie in which I was afforded the opportunity to pick his brain on various topics regarding the society and his role as president.

This issue also sees the first installment of a series of articles on the state of OR at various South African universities, the first being Stellenbosch University (no bias I promise!).

The feature article of the issue is written by a previous contributor, Collin Philips of OPSI Systems, and discusses optimisation of the cash in transit process, definitely worth a read and another excellent example of the application of OR techniques in industry.

That's all from me for now, until June, cheers all and enjoy the read!

Features	Page
FROM THE EDITOR	1
FROM THE PRESIDENT'S DESK	2
THE 2014 ORSSA EXECUTIVE COMMITTEE	3
OPERATIONS RESEARCH AT SOUTH AFRICAN UNIVERSITIES PART I: STELLENBOSCH UNIVERSITY	4
CLOUD COLLABORATION FOR THE SUPPLY CHAIN OF TOMORROW — EVENT SYNOPSIS	5
CASH IN TRANSIT OPTIMISATION	6
GETTING TO KNOW OUR PRESIDENT: AN INTERVIEW WITH HENNIE KRUGER	13
BOOK REVIEW: THINKING FAST AND SLOW	15

QUERIES & CONTRIBUTIONS

Any queries or contributions to the Newsletter are most welcome, especially article submissions. For any queries or contributions, please contact the Newsletter editor at einhorn@sun.ac.za.



FROM THE PRESIDENT'S DESK

By HENNIE KRUGER

(Hennie.Kruger@nwu.ac.za)

ORSSA PRESIDENT



Hennie Kruger

Time is flying this year, although it seems like yesterday when the first newsletter of 2014 appeared. With this second edition, we are already approaching the last half of the year. I trust that you all have experienced a productive year up to now and are anticipating an even better second half of the year.

Talking about looking forward to the second half of the year – that is when one of the highlights on the ORSSA calendar is taking place. This year, the annual conference will be taking place at Stonehenge near Parys in the Free State from 14 to 17 September. The 2014 Conference is organised by the Vaal Triangle branch with Lieschen Venter as chairperson, together with a very capable and enthusiastic organising committee. I hope to meet a great number of members at the conference. Registration is already open and although the abstract submission for the peer-reviewed stream has already closed, the abstracts for the normal presentation stream are still open until 11 July 2014. Everybody is encouraged to register and organise their accommodation as soon as possible. More details are available on the ORSSA webpage (www.orssa.org.za), as well as page XX of the newsletter.

This year, we will have three keynote speakers, namely Prof Ted Ralphs from Lehigh University in Pennsylvania, Prof Andreas Bley from the University of Kassel in Germany and Prof Frans Waanders from the North-West University's Potchefstroom Campus. Their presentations will deal with the conference theme, "Accessible Analytics", and will focus on opportunities to expand the application of OR to new areas. Delegates are also invited to attend an optional tutorial on 14 September 2014. The tutorial will be presented by Prof Ted Ralphs and will provide an overview of the COIN-OR optimisation suite.

I want to thank the local organising committee and congratulate them on their arrangement of sponsors. At the same time I want to thank the sponsors for their support of the conference. The sponsors who have already agreed to support the conference include AIMMS, OPSI Systems, Blue Stallion, SASOL and the North-West University (Potchefstroom Campus).

Concerning the different chapters, I am pleased to report that I receive information continually of activities taking

place in their areas; it seems as if the life arteries of the society are alive and well. Thank you very much to everybody who is working so hard to organise the events and who attends them. The digital event that was organised by the Vaal Triangle chapter during March is a new initiative. Participants were able to watch a presentation by an international speaker by means of Google Hangouts and were also able to interact with the speaker. During the AGM in September, all chapter chairs will be giving comprehensive feedback concerning the activities in the various chapters. It is with regret that I have to report the resignation of Louis Dannhauser as chairperson of the Johannesburg branch due to work pressure. Thank you, Louis, for your sacrifices over a long period of time and for the enthusiasm with which you worked in the Johannesburg chapter – we appreciate it and wish you all the best. Hopefully, we may approach you again in the near future to help us in the executive committee. Congratulations to Brahm Bothma who has been elected in Louis's place as new chapter chair. We are all seeing forward to working with you and we know that you will contribute to the success of the Johannesburg chapter. Welcome also as a new member of the executive committee. I would also like to welcome Danie Lotter as the new chapter chair for the Western Cape. Danie is already a member of the executive committee in the Newsletter Business Manager portfolio. Thank you Danie for taking responsibility for both portfolios and we are also looking forward to work with you.

At the moment, the executive committee is busy to determine members' interest regarding the possible certifying as analytical professionals. This is an initiative of INFORMS and as an initial step, we want to see whether there are the need and interest to establish an examination centre in South Africa. An e-mail has been sent to each ORSSA member recently in regard to this. Angela Rademeyer is handling the matter and she and Michelle Fischer (who is already a certified analytical professional) have written an information document which can be read on page 8. I want to encourage all members to consider the advantages of such an international certification and to contact Angela if you are interested in it. Thank you, Angela, for your help and the time that you spend on this matter.

By this time, each of you has received a statement of payable membership fees with the March newsletter. You would have noticed that our banking affairs are now handled by Standard Bank and that there is thus a new account number in which membership fees may be paid in. I want to encourage all members to settle their membership fees, please – especially if your payment is already overdue. The society is run with a rather limited budget and we need everyone's support. If there is any enquiry regarding the statements or any other financial matter,

feel free to contact the treasurer, Tiny du Toit, directly at: Tiny.duToit@nwu.ac.za

Lastly, I want to share with you news of a project which is extremely exciting to me. The idea of the project has already occurred in 2012 and is now well on its way. The project involves the expertise, knowledge and experience of especially the “senior” ORSSA members, but also those of younger members, which will make the end product a real contribution to OR in South Africa. At the moment, I am not going to let the cat out of the bag any further, but may add that Jan van Vuuren and I will be giving more information to ORSSA members during a plenary session

at the conference. Do not, therefore, miss the session if you want to hear more about this exciting project.

I want to conclude by thanking everyone for his or her contribution during the first half of the year and I want to invite each member again to contact me or any other member of the executive committee if you have a suggestion on anything concerning ORSSA. Please let us know if you want to become involved in any activity of the society. Greetings, until we see each other at Stonehenge in September.

With best wishes/Alles van die beste
Hennie Kruger

NOTEWORTHY NEWS

GYS J. WESSELS RECEIVES EDELMAN LAUREATE MEDAL

Gysbert J. Wessels received his medal as an Edelman Laureate at the Honors Reception at this year’s INFORMS Conference on Operations Research and Business Analytics. The Franz Edelman competition recognizes outstanding examples of operations research, management science and advanced analytics in practice. Each year up to six teams compete for the Franz Edelman Award by presenting their work at the conference and publishing a paper in the Interfaces journal. All authors of the finalist papers

earn medals but these medals are only presented in person at the Edelman Reception. Gys was in attendance at this year’s conference and he was finally honoured with a medal for his role on the 1996 Edelman Award winning team from South Africa and as an author of the paper, “Guns or Butter: Decision Support for Determining the Size and Shape of the South African National Defense Force”. Gys is shown seated on the right with the “Class of 2014” Edelman Laureates. Gys is in prestigious company. Included in the photo are Alvin E. Roth winner of the 2012 Nobel Prize in Economics and John D.C. Little best known for his proof of the queuing formula $L = \lambda W$, commonly known as Little’s Law.



The Class of 2014 Edelman Laureates

ORSSA WESTERN CAPE CHAPTER EVENT

On the 10th of April, the Western Cape chapter hosted its second chapter event for the year together with the South African Institute for Industrial Engineering (SAIIE) at the Bergkelder wine cellar, which is famous for its “cellar in the mountain”, in Stellenbosch. The talk was well attended and attracted forty individuals from a number of universities as well as the operations research industry.

Attendees were welcomed with a glass of wine after which they were guided through parts of the Fleur du Cap winery to the cellar located inside the Pa-pegaaiberg where the talks were scheduled to take place. The evening kicked off with a cocktail function which provided the attendees with networking opportunities and a chance to catch up with colleagues over some of the Cape’s finest wine and delightful canapes.



Guests mingle in the Fleur du Cap wine cellars

The theme for the evening was achievements and opportunities in freight transportation modelling and those in attendance were privileged enough to witness talks from two esteemed speakers. The first speaker was Jan Havenga from the Department of Logistics at Stellenbosch Uni-

versity who gave a talk on the role of freight modelling in strategic decision making and the second speaker was Johan Joubert from the Department of Industrial and Systems Engineering at the University of Pretoria who gave a talk on agent-based freight based transportation modelling. The evening, which was well

and truly a success, marked Danie Lotter’s first Chapter event as head of the Western Cape chapter of ORSSA and was thoroughly enjoyed by all those in attendance.

OPERATIONS RESEARCH AT SOUTH AFRICAN UNIVERSITIES PART II: NORTH-WEST UNIVERSITY

by Hennie Kruger (Hennie.Kruger@nwu.ac.za), North-West University

The first official reference to operations research at the Potchefstroom University for Christian Higher Education (PU for CHE)¹ appeared in the yearbook of 1967 as part of Applied Mathematics for a BSc degree – students had to take Elementary Linear Programming as part of Applied Mathematics. The beginning of structured operations research training occurred in 1977 with the first BSc degree in Operations Research which was offered at the Vaal River Branch of the PU for CHE in Vanderbijlpark.

¹ The University at Potchefstroom was known as the Potchefstroom University for Christian Higher Education (PU for CHE). After merging with the University of North-West in Botshatswana in 2004, the name was changed to the North-West University.

At the present-day North-West University (NWU) operations research modules are presented as part of subject disciplines like Statistics, Mathematics, Applied Mathematics and Computer Science. It is especially in Computer Science where specific operations research modules are taught as part of a BSc in IT degree. These modules are found in Decision Support Systems and are presented on second-year, third-year and honours level. It covers a wide range of operations research topics i.e. during the second year linear and non-linear programming, integer programming, goal programming, transportation and assignment models and introductory network modules are offered. The third year modules include topics such as decision making, forecasting, queuing theory, simulation etc. These topics are then further explored on honours level. Students may then further their studies and also do a Masters or



A campus building of North-West University.

PhD degree in Decision Support Systems. A large number of postgraduate students in Decision Support System have completed their studies in recent years and have covered a wide variety of operations research topics.

The Computer Science department has excellent facilities for students in Decision Support Systems which include a dedicated lab for 4th year students. Students pursuing a Masters or PhD degree have their own facilities which comprise of offices which can house between two and six students.

Operations research is also offered at the Vaal Triangle Campus of the NWU where the focus is on data analysis for risk management. Undergraduate modules are also taught in the Computer Science curricula while Operations Research is a curriculum on its own in Computer Science for Masters Students.

A subsequent development at the erstwhile PU for CHE that gave more prominence to the teaching of operational research was the establishment of the Centre for Business Mathematics and Informatics® (BMI) in 1998 (in collaboration with Absa). One of the main objectives of the North-West University's Centre for BMI is to train students for professional careers based on the mathematical and computer sciences. At the moment the flagship programme is the one-year MSc BMI with specialisations in Quantitative Risk Management (QRM, including Actuarial Science), Financial Mathematics (FM) and Business Analytics (BA). The MSc BMI training programme is made up of specialist courses, an industry directed research project and an industry practice module. Dependent on the particular specialisation area, courses may be selected from Practical Risk Management, Enterprise-wide Risk Management, Retail Credit Risk, Advanced Investment Theory, Advanced Pricing of Derivatives, Data Mining Techniques, Business Intelligence, Customer Behavioural Analytics, Predictive Analytics and Multiple Criteria Decision Making. The industry practice module focuses on preparing students for their industry directed research projects and includes training in creative problem solving skills, practical communication skills, and on-the-job training in project management skills, which includes meeting management,

client expectation management and the management of interpersonal relationships. The industry directed research projects are conducted on-site at a specific company in industry. SAS has sponsored the establishment of a SAS Laboratory at the Centre for BMI and recently gave seeding money for a Business Analytics programme with focus on Big Data and Data Science. A large portion of the seeding money was used to procure a powerful server on which high performance analytics on big data sets can be performed.

Since 1998 a total of 297 MSc BMI students have been delivered to industry at a rate of approximately 25 per year. Research programmes support the specialisation areas and include basic and applied risk and business analytics focus areas. The Centre for BMI has negotiated contracts with Absa/Barclays, the Department of Science and Technology and SAS totalling more than R10m per year, which include funds for applied research, bursaries, and capacity building. According to the Director of the Centre for BMI, the BMI training and research programmes are viewed as operational research applied in practice, and especially for the benefit of the financial services industry.

The Centre for BMI's Quantitative Risk Management programmes are accredited by the Professional Risk Managers' International Association (PRMIA) for Level I and II of their Professional Risk Manager (PRM) designation and there is a strong relationship between the goals of the industry directed research projects and that of the new Certified Analytics Professional (CAP) certification introduced by INFORMS.

The establishment of a new department of Industrial Engineering has been approved at the NWU's Potchefstroom campus and the first intake of undergraduate as well as post graduate students will be in 2015. As with any other Industrial Engineering programme, operations research is an integral part of these new curricula and will definitely contribute to operations research education and research activities at the NWU.

INDUSTRIAL SCALE OPTIMISATION IN OR

by Robert Bennetto, freelance consultant and optimisation enthusiast. (robert.bennetto.za@gmail.com)

Mathematical optimisation is a universal concept spanning many disciplines. The OR practitioners' toolbox is constantly being updated in order to deal with industrial optimisation challenges. This article serves as a brief tour of the latest hot topics in the field.

The scene

I recently read an article about how the egg heads at Google were using a massive cluster to train a billion parameter neural network to understand the features of cats based on YouTube videos. There are so many things about that statement that make it curious. More than ten thousand CPU's grinding away in near perfect, symphonic harmony (somewhere in the cloud) is just awesome. I picture Scotty from the Star Trek Enterprise screaming that he's giving it all he's got.

Having grown up with the modern incarnation of the household computer I sometimes find it difficult to picture the early implementations and get a handle on just how rapidly the landscape has changed. The ability to capitalise on the infrastructure provided through cloud-based services is rapidly becoming a necessity in industry; cheaper, faster and easier to use every year. If you aren't using it, you will ultimately be left behind. The follow-on article to the 10,000 CPU YouTube cat malarkey was really an eye opener. The boffins at Stanford¹ used CUDA and 3 machines with graphics cards to train a bigger model with 11 billion parameters. Did they just accomplish the same task (in principle), ten times larger with orders of magnitude less computational resources? The trick is the way neural nets lend themselves to being calibrated in parallel. By sticking to simple arithmetic and using the graphics cards to do the heavy lifting they've shown that just about anyone can play with these massive networks without having to rent² half a server farm for a week.

While this example is about a Machine Learning technique - the principle holds. Constantly asking yourself if you can do it better is a big component of industrial OR. Think of industrial OR as sport; companies competing to show that they can solve problems better than the rest of the pack. They needed that massive cluster of nodes as a stepping-stone to show that it is viable concept in the first place. Once they knew what was to be gained, they could invest

1 Andrew Ng is common to both teams

2 If you would prefer to buy, the CPU cluster costs around \$1m, before the electricity bill.

the effort in speeding it up in order to add more parameters. Now they are leading the pack and will be shaping the Machine Learning research landscape in a very exciting way for the next 10 years.




Figure 1: Visualisation of the LP

Our first love

Bright eyed and full of ambition and a hunger for knowledge, almost every OR practitioner remembers the first time they encountered Linear Programming (LP). Its elegance is in its simplicity. Stated neatly in a single line:

$$\min\{c^T x: Ax \geq b; x \geq 0\}$$

Minimise the cost ($c^T x$ is the objective function) given a set of variables (x is the stuff you can control), coefficients and constraints (A and b define limits you can't break). If you are in the lucky (and often unlikely) position where your decision variables are all continuous and your objective function is linear, you will reap maximum rewards from this relationship with LPs. They are mind-bendingly fast with the performance gains being measured as "millions of times" faster due to hardware and (more so) software improvements over the last 20 years. This means we can go big and solve these problems optimally in polynomial time.

The allure of your first love is ultimately dispelled. If you need to model discrete outcomes, or integer solutions, you can still use the LP framework but with the condition that we can only pick integer solutions, the solid dots in

Figure 1 above. We generally refer to this as a Mixed Integer Program (or MIP) since we have dropped the assumption of linearity. This may sound great but dropping that assumption of linearity leaves us compromised as we can no longer solve the resulting set of equations in polynomial time. The problem of solving MIPs remains NP-complete in the general sense.

For the first 20 years of their commercial use MIPs remained a tool for the academic. You'd be hard pressed to find a paper from the late 80's with an application in OR that doesn't contain a MIP formulation with the authors having solved a very small instance thereof. It served as a framework for expressing problems and demonstrating their complexity but not really solving problems on a scale useful to many practitioners.

The breakthrough in MIP performance has come quite recently, and unbeknown to many. CPLEX 6.5 implemented a host of cutting plane techniques, originally thought to only have theoretical merit (CPLEX is now on version 12.4). The cutting plane implementations have reshaped the performance of commercial MIP solvers. The 2013 Edelman Award winners used CPLEX and CONOPT to solve a Non-linear MIP to perform a cost benefit analysis of flood damage versus dike height upgrades in the Netherlands.

There are many problems which you may have thought were hard but which a modern MIP solver can probably tackle given a good formulation. The advantage of this approach is that we always have sight of the objective function - enabling us to know how far away we are from a theoretical best solution (that fast continuous relaxation we had when we had linear variables).

Knowing how good a solution is, and how much better it could still be is a very powerful idea. It enables us to quantify how good our solution is and to focus on how to quickly prune the search space so that we can tighten the lower bound on the objective. It is the strong mathematical basis that holds MIP's up and keeps them in good favor with many. Unfortunately most problems cannot be cast neatly into the MIP framework and chances are that if you try to force it to fit, the results will be high blood pressure, wasted hours, and an overheated CPU.

So why bother bringing up MIP's in the first place? A couple of reasons come to mind but I think one of the most telling signs is that companies like Gurobi (IBM's primary competitor in this space) have already geared their offering around the cloud. They've made it embarrassingly simple for you to model your problem locally and ship it off to their servers for processing. If you need to run a big distributed MIP search, and don't have the hardware lying around, you

can now order a quick burst of CPU's to throw at it.

Gurobi are leading this distributed MIP space with an incredibly focussed view that this is the future of optimisation. I would argue that there are a few problems. Firstly, you need to be able to model the problem neatly as an MIP - the likelihood of this being achieved drops dramatically the more *features* a client requests. Secondly, computation with commercial optimisers is not cheap. This creates a considerable barrier to entry unless you are an academic or solving problems with big payoffs.

MIP state of the art

There are many noteworthy industrial sized applications where the MIP framework is used to perform very large-scale neighbourhood searches (VLNS for short) attempting to achieve optimality. You can think of this as a hill climber where we use the MIP engine to climb to the top of a hill where we can keep shuffling around which variables participate in the definition of the hill.

If your problem was small enough, all the variables would participate in the definition of the hill and you would end up at the optimum in a single iteration. As the problem grows in size, it takes exponentially longer to climb that hill, so we have to select a subset of the variables to optimise over. Knowing which subset of variables to pick in VLNS is tough, so what is done in practice is to randomly select large subsets and repeatedly optimise over them until we hit some kind of stopping criterion.

Constraint Programming (CP) is like the delinquent child of optimisation. It cares nothing for the mathematics and focusses entirely on the logical inferences we can make about a problem expression (through constraint propagation). Continuous variables are discouraged in CP because they don't propagate well. You can think of the MIP and CP as being from two different schools of thought, one pioneered by a Mathematician, the other by a Computer Scientist.

Combining MIP and CP approaches successfully has often resulted in the new benchmark to beat. The MIP framework takes care of the mathematical bounds of the problem and CP handles all the funny integer bits and weird constraints we were never mad about in the first place. There are a couple of ways to go about this separation of tasks. One way is that the LP can act as the master model with the CP validating nodes for feasibility in the search tree. The other method is to use the CP model as the master and have it guided to optimality by asking the LP to prioritise variables and values. The major downside is that you will need to work on the same problem for a very long time



43rd ORSSA Annual Conference

14–17 September 2014

An advance warm welcome to the 43rd Annual Conference of the Operations Research Society of South Africa (ORS-SA)! The Conference will be hosted by the Vaal Triangle Chapter of ORSSA, and will be held at Stonehenge in Africa, just outside of Parys, from September 14th to 17th, 2014.

The conference will open with an optional tutorial on Sunday afternoon and a welcome reception on Sunday evening September 14th and will close at lunchtime on Wednesday September 17th. Participation over the full spectrum of Operations Research is encouraged, including papers of a more fundamental nature, those on the application of Operations Research techniques in business and industry, about topical issues in Operations Research, and about the philosophy, teaching and marketing of Operations Research.

The keynote speakers at the conference will be announced in due course.

Following the successful introduction of published conference proceedings in 2011, authors will again have the choice of either (a) only presenting papers orally at the conference, or (b) additionally submitting full papers (which will be pre-

sented orally at the conference) for inclusion in the peer-reviewed conference proceedings. Registration, abstract and full paper submissions will open during the first quarter of 2014.

Delegates are responsible for their own travel and accommodation arrangements. Stonehenge in Africa is recommended, as the Society has arranged competitive rates for delegates.



Stonehenge in Africa

Important Dates

17 March	Early bird registration & abstract/paper submission opens
11 April	Abstract submission closes for reviewed papers
18 April	Notification of acceptance of abstracts of reviewed papers and go-ahead to submit full papers for peer-review
16 May	Submission of full papers for inclusion in the conference proceedings closes
25 July	Early bird registration closes
11 July	Abstract submission closes for oral presentation of all papers
18 July	Notification of abstract acceptance for non-reviewed papers
18 July	Notification of acceptance of reviewed papers for proceedings
15 August	Cut-off for qualification of reduced room rates at the hotel
22 August	Registration closes

Please visit the ORSSA website and click on the link *ORSSA 2014* for more information:

www.orssa.org.za

before it's obvious how to get the right interactions and representations between these models working efficiently and effectively.

Admit defeat

You have deadlines not to be tried with and your OR mates waiting for you at the pub. You can't afford to spend ages thinking too deeply about this optimisation task because the client needs results, not awesome³ maths. You lick your wounds and pick the meta-heuristic you'll be most comfortable talking about at the pub with your OR buddies.

Choose your poison

So how bad is this compromise? That depends mostly on how much time you have to pour into your meta-heuristic. It is worth mentioning that not all meta-heuristics are born equal and that each of them usually has roots in a particular problem type. They all have super cool names though which will impress your non-OR friend⁴.

Ant Colony Optimisation (ACO) is very good when solving problems with some kind of routing component, akin to ants having to efficiently collect food. Traditional Genetic Algorithms (GAs) are very strong when you have lots of mostly independent binary variables and when inheriting subsets from parents makes sense. In this sense, ACO is far better suited to solving problems like the Travelling Salesman Problem (TSP) than GAs are. That said, GAs have been philosophically adapted to solve many problems using representations which don't look like DNA strands whilst maintaining the principle of populations and survival of the fittest to control the search.

ACO and Evolutionary Algorithms are very popular population based methods. Simulated Annealing (SA) and Tabu Search (TS) stand out from the crowd when it comes to a more traditional style of search. They both modify a single solution using some sort of alteration mechanism and select the new solution in some probabilistic manner⁵. In TS for example, it's typical to have a graph representation of the problem and define *move* or *swap* operators which correspond to changes in the graph. These operators are tried time and again by the TS where operations which failed to result in improvements are recorded as *taboo* in its short term memory. This direct modelling of the problem is great if you can understand how good move operators should be used. In this regard the other representations are more generic but aren't going to be able to exploit the search space as well as a direct search probably could.

3 often disputably awesome.

4 you only have one, the rest are awesome OR folk.

5 in an attempt to avoid local minima.



Figure 2: Planning 70 000 visits across 12 weeks results in 500 billion binary decision variables. Understanding trade-offs become very important at this scale. Imagery courtesy of OPSI Systems.

I don't believe there is a right answer as to which poison is best. If you were to try answer the question seriously, there are more papers out there comparing performance under different test conditions than you could read in your lifetime.

Sharpen your sword

Very few industrial sized problems can be tackled by black box solvers competitively without some sort of augmentation. Small frequent changes and rigorous testing are the main ingredients for this recipe to be successful. It's common to include some flavour of local search in big implementations. The main benefit is to avoid the "something stupid" effect. No matter how close to optimal a solution is, if the client is ever able to identify something wrong with what your meta-heuristic conjured up, there'll be fighting words. Using local search to polish solutions helps prevent you from having to face the poor decisions your meta-heuristic may make.

Split your eggs

Parallelism is hard. Modern processing speeds flattened out in 2005 and the progress we've enjoyed since has mostly been made with instruction-level parallelism, and more efficient, densely packed transistors. This allows us to throw more cores at the problem consuming only slightly more power. The power efficiency of processors has become a key component in enabling the techno-savvy lifestyles we enjoy (smart phones and notebooks that don't need to be charged every 30 minutes).

Population based algorithms have an edge in the parallelisation arena. It's not to say that there aren't ways to parallelise search techniques like Tabu Search or Simulated Annealing, it's just really hard and it's going to take lots of time and testing to get right (remember your friends are waiting at the pub!). In a typical optimisation task your real CPU crunch is going to be on the objective function evaluations. Population based methods have a convenient separation where all the individuals in the population can normally be evaluated independently. This means we can distribute the computation whenever we want to evaluate an individual (which is normally costly) without having to streamline every piece of code to run in parallel.

The result is that if we want more speed, and our solution design supports being run in parallel, we can just throw more physical CPU's at the problem⁶. It's worth mentioning that the degree of natural, or easy, parallelism in population based optimisers is quite low especially when compared to those deep cat-learning algorithms I mentioned earlier. GPUs support very simple calculations, and by contrast, the modern CPU is a veritable Swiss Army Knife when it comes to handling complex computations effectively. So often what works well on one will not necessarily work well on the other. GPUs don't get out of bed for less than a thousand threads in the morning.

Industrial scale problems are hard and will take time to solve. Taking advantage of the computing architecture available is just one way to get ahead. Dr. Bill Dalley chief scientist, senior vice president of Nvidia research and lead architect of one of the worlds current fastest super computers running LinPack (Titan) said in an interview "*The future of the world is parallel. Serial programs are going to be like COBOL programs, people used to do those but they don't do them anymore.*"

I've spoken mostly about problems which can be well defined. There are many applications which require simulation because there are just too many variables and their interactions aren't known or understood a-priori. Historically this has been a bit of a clumsy approach, mostly because the demands of the technique have been ahead of the hardware. I believe that this has shifted dramatically in the last 10 years because of the programming model available through high degrees of parallelism.

Major motoring manufacturers use simulation to aid in the design of their vehicles (aerodynamics, engine efficiency, road handling, you name it). Given more processing nodes, they will simply increase the granularity of the simulations providing greater accuracy and benefits. Given even more processing power, whole parameter sets can be optimised

over without ever really having to worry about the underlying model. A very exciting time indeed.

The Future - Back to basics

Last year Google and NASA teamed up to buy their own Quantum Computer⁷. It was decided that the best way to test a Quantum Computer is to give it really hard optimisation problems to which it can simultaneously consider all possibilities and produce the optimal. The decision variables are all binary and represented using qubits. These qubits have the additional state of being not only a 0 or 1, but also being simultaneously 1 and 0 at the same time (superposition).

You can think of all the qubits going into a big sieve in their superposition state and as we tap the sieve they drop out the other side as either a 0 or a 1. This is the annealing (or cooling) process and in a single second they can process not only the optimal but also tens of thousands of good alternatives to the problem specified.




Figure 3: The processor at the core of the D-WAVE Quantum Computer.

6 to a point, you will have diminishing returns

7 Built by D-WAVE SYSTEMS INC.

So this is where things get tricky. The current problem specified in the quantum computer is a quadratic assignment problem. That's right, it can only do one thing. This is good because it's a complicated problem which you may or may not need all the features of, however, you may struggle to get your problem to actually fit this format. This conversion of problem representations to this format is an open problem. Quantum computing is limited in this regard in that the very hardware itself mimics the underlying graphical representation which it needs to do in order to solve

The machine installed at NASA has 512 qubits, which doesn't sound like much but it is considerably faster⁸ than CPLEX when solving the Quadratic Unconstrained Binary Optimisation problems (QUBO). It performed less well when problems needed to be translated into its input format which weren't natively expressed as QUBO. Still very impressive when you think that the technology is essentially 50 years younger than its competitor and able to significantly outperform it on certain problems.

Their next generation machines are expected to have 2,048 qubits but there are a lot of questions around how well the performance is going to scale. Did I mention that the processor in this machine is cooled to 0.02 Kelvin? Clearly it is not yet ready for home use just yet.

The playing field

There's an underlying thread in all this. The landscape is still shifting and like scientists before us, we do not yet have the optimal solution. That makes everything exciting. The field of optimisation is fascinating, broad and deep. The last 3 decades have ushered in high level languages to easily express our optimisation problems independently of the hardware we work with. But in having to tackle larger, more complex problems, practitioners are starting to move closer to the hardware to see how best to extract maximum benefit from the horsepower and tech available: distributed cloud architectures, GPU implementations and now qubits.

It's a challenging environment and the need for optimisation to scale up in order to handle the expected growth in data (and our visibility of decisions we can make) over the coming years is pressing. The goal posts of what is considered a "hard problem" are constantly shifting with the landscape and as a result, it is unlikely our work will ever be done!

Explore Further

[Dean et al, 2012] Large Scale Distributed Deep Networks. *NIPS*, 1232–1240.

⁸ 1 second vs 30 minutes. Good to know what you get for \$10m.

[Coates et al, 2013] Deep learning with COTS HPC systems, *Proceedings of The 30th International Conference on Machine Learning*, 1337-1345.

[Bixby, 2010] A brief history of linear and mixed-integer programming computation, *Documenta Mathematica*, 107-121.

ADVERTORIAL: INFORMS CERTIFIED ANALYTICS PROFESSIONAL PROGRAMME

Members may be aware of the new Certified Analytics Professional (CAP) qualification which was launched by INFORMS last year. INFORMS offers the industry's only exam-based certification designed to recognize analytics professionals who want to demonstrate a deep commitment to excellence and continued professional development. The exam is based on a rigorous job task analysis (see page 13) and is not specific to any type of vendor, industry, application, or software. A handbook containing 24 sample exam questions and a list of reference texts used to prepare for the exam, to give you an idea of what is expected, can be obtained from ORSSA. Successful candidates will add CAP after their names.

ORSSA is currently in discussions with INFORMS about the details of establishing an exam centre locally and being able to take the exam in computer based testing centres. In order for us to proceed with investigating the feasibility of being able to write CAP exams in South Africa, INFORMS have requested we establish how many of our members would be interested in obtaining such a qualification. If you are interested or have any comments, queries or suggestions, then please email Angela Rademeyer (angela.rademeyer@opsi.co.za). The CAP benefit and eligibility criteria, as listed on the INFORMS website, include:

Benefits

- Advances your career potential by setting you apart from the competition
- Drives personal satisfaction of accomplishing a key career milestone
- Helps improve your overall job performance by stressing continuing professional development
- Recognizes that you have invested in your analytics career by pursuing this rigorous credential
- Boosts your salary potential by being viewed as experienced analytics professional
- Shows competence in the principles and practice of analytics
- Demonstrates commitment to the field
- Proves to stakeholders that your organization follows industry-standard analytics practice

Eligibility

- Bachelors or Masters degree
- At least five years of analytics work-related experience for bachelors degree holder in a related area, or
- At least three years of analytics work-related experience for Masters (or higher) holder in a related area, or
- At least seven years of analytics work-related experience for Bachelors (or higher) holder in an unrelated area
- Verification of soft skills/provision of business value by employer

We look forward to your response.



ANALYTICS JOB TASK ANALYSIS (JTA)

The CAP™ examination is based on a delineation of common or typical tasks (T) performed and knowledge (K) applied by Analytics Professionals. Domain weights are based on the SMEs assessments of the importance of tasks and the frequency of their performance.

(15%)	Domain I	Business Problem (Question) Framing
T-1	Obtain or receive problem statement and usability requirements	
T-2	Identify stakeholders	
T-3	Determine if the problem is amenable to an analytics solution	
T-4	Refine the problem statement and delineate constraints	
T-5	Define an initial set of business benefits	
T-6	Obtain stakeholder agreement on the problem statement	
(17%)	Domain II	Analytics Problem Framing
T-1	Reformulate the problem statement as an analytics problem	
T-2	Develop a proposed set of drivers and relationships to outputs	
T-3	State the set of assumptions related to the problem	
T-4	Define key metrics of success	
T-5	Obtain stakeholder agreement	
(22%)	Domain III	Data
T-1	Identify and prioritize data needs and sources	
T-2	Acquire data	
T-3	Harmonize, rescale, clean and share data	
T-4	Identify relationships in the data	
T-5	Document and report findings (e.g., insights, results, business performance)	
T-6	Refine the business and analytics problem statements	
(15%)	Domain IV	Methodology (Approach) Selection
T-1	Identify available problem solving approaches (methods)	
T-2	Select software tools	
T-3	Test approaches (methods)*	
T-4	Select approaches (methods)*	
(16%)	Domain V	Model Building
T-1	Identify model structures*	
T-2	Run and evaluate the models	
T-3	Calibrate models and data*	
T-4	Integrate the models*	
T-5	Document and communicate findings (including assumptions, limitations and constraints)	
(9%)	Domain VI	Deployment
T-1	Perform business validation of the model	
T-2	Deliver report with findings; or	
T-3	Create model, usability and system requirements for production	
T-4	Deliver production model/system*	
T-5	Support deployment	
(6%)	Domain VII	Model Lifecycle Management
T-1	Document initial structure	
T-2	Track model quality	
T-3	Recalibrate and maintain the model*	
T-4	Support training activities	
T-5	Evaluate the business benefit of the model over time	

*Tasks performed by analytics professionals beyond certification level

IF THE CAP FITS: BECOMING A CERTIFIED ANALYTICS PROFESSIONAL

by Michele Fisher, CAP®, Sasol (michele.fisher@sasol.com)

I recently attended the INFORMS Conference on Operations Research and Business Analytics in Boston and took the opportunity to sit the Certified Analytics Professional (CAP®) examination. At the same conference, it was announced with fanfare that INFORMS would be rolling out the examination to other countries including South Africa. I was reminded of the quote by George Moore, “A man travels the world over in search of what he needs and returns home to find it.”

I wholeheartedly support the efforts by INFORMS and ORSSA to broaden access to the certification process. The number of professionals able to attend conferences where examinations are being administered is small. I was lucky to have such an opportunity.

I would like to briefly share my experience with the certification process. This may help you decide if certification is for you and inform your feedback to the ORSSA executive. Professional certification has four requirements: education, experience, evidence of soft skills and examination of hard skills.

Education and experience are taken together. The number of years of professional analytics experience that you need depends on your degree and discipline. If you have a Master’s degree (or higher) in a related field then you only need three years of experience. If your degree is from a field not related to analytics then you need seven years of experience. I have a Bachelor of Science in Mathematics and therefore needed five years of experience as an analytics professional. Since I have been working in the field for five times that long, it was easy to show the required analytics experience and provide a reference who would verify that experience. It was only slightly more challenging to dig back into my files to find the official transcript of my marks and degree from university. I scanned this and uploaded the pdf to my application. It did not matter that my transcript was so old that it had been printed on a dot matrix printer.

For the evidence of soft skills, I provided my boss as a contact able to vouch for my skills at communicating with stakeholders, framing problems, working with teams and communicating results. She was contacted by e-mail and found it straightforward to send the required confirmation by return e-mail.

Now all that stood between me and certification was 100 multiple choice questions. While examinations were some-

thing I had always excelled at in school, it had been a long time since I had taken one. I tried to find the time to work my way through the CAP® study guide that was available in draft at that time. Unfortunately it was not a very coherent package with different authors for different chapters and lots of suggested references. The study guide has been polished and is now available for sale. Before you invest in it, I recommend that you have a look at the sample test questions in the CAP® Handbook. If you feel comfortable answering the test questions then you may be ready.

In the end, the handbook contained all the information that I needed. It also has copies of the application forms, but I applied on-line through the INFORMS web-site. When it took me a few days to track down my transcript and confirm my references, I received a friendly e-mail from INFORMS staff offering assistance. They were very helpful throughout the process and things went so smoothly that I barely had time for second thoughts. That was probably a good thing, as I passed the exam and have earned the right to put Certified Analytics Professional after my name.


My journey did not end with the examination. I now must continue to develop my analytics skills and share them in order to maintain certification. Credit is given for volunteer service, work experience, authoring of papers and newsletters, participation in training, etc. As a society, ORSSA provides many opportunities to support recertification and it stands to benefit from the participation of its CAP® members. I believe that this ongoing renewal is the hidden strength of the program that will become more evident as it matures. I know it will make me a stronger professional and I hope that my contributions will make our community stronger.

The views expressed in this article are those of the author and not necessarily of Sasol, INFORMS or ORSSA.

BOOK REVIEW: HANDBOOK OF NEWSVENDOR PROBLEMS: MODELS, EXTENSIONS AND APPLICATIONS.

By Hans Ittman (hittmann01@gmail.com)

Most operations researchers are exposed to the classical newsvendor, or newsboy, problem during their studies. This problem was first stated within a bank context as far back as the 18th century when a bank needed to determine the level of cash reserves it required to cover the demand from its customers. In this problem, the normal distribution was used to estimate the probability that a given amount of cash is sufficient to cover total cash demand. Another, possibly the most well-known, example addresses a newsboy's need to determine the number of newspapers to order to meet future uncertain demand, given that the newspaper becomes obsolete (in other examples, it is perishable) at the end of the day. Given a stationary demand distribution and a known cost-revenue structure, the newsboy can determine the optimal order quantity by either maximizing the expected profit or minimizing the expected cost. Today, a huge amount of research literature is available on this topic and related problem areas; application areas are extensive and include many publications. However, as the editor of this handbook states: "what is absent is a comprehensive reference source of state-of-the-art findings on both theoretical and applied research on the newsvendor problem" (p. v). The goal of this handbook is therefore "to consolidate many latest research findings and applications of the newsvendor problem into one volume" (p. v).



place the traditional risk-neutral performance measure by risk-aversion measures. Although risk-neutral problems can be decomposed, decomposition is not possible in the case this chapter considers. The entire portfolio of products must be considered as a whole; in this way, results are derived that show the impact of risk aversion and demand dependencies on optimal order quantities. In addition, this chapter shows that for general law-invariant coherent measures of risk, fairly accurate approximations of the optimal order quantities for a large number of products with independent demand can be obtained.

When a number of multiple-demand sources are available, decision makers can pool their inventories. Chapter 3 and a few other chapters cover the pooling problem, thus indicating the attractiveness of this mechanism. Inventory can be held at one or more physical locations to serve multiple demand points with fast transshipments to these demand points. This chapter focuses on the inventory pooling problem within the newsvendor framework; its objectives are to study the change in inventory levels when inventories are pooled and develop models to determine optimal inventory levels under such circumstances. It also introduces the copula framework to model various dependency structures among pooled demands.

This book has two parts: Part 1 contains the results of research on basic models and extensions; Part 2 focuses on a variety of practical applications of the extended newsvendor problem. Each part has eight chapters, and each chapter is dedicated to a specific topic in the field. Research related to this area has greatly increased over the past few years, and this handbook presents some of the most current research.

Chapter 1 discusses the multiproduct newsvendor problem (MPNP). In the MPNP, a number of products must be obtained and the demand for each product is assumed to be stochastic. This chapter presents a review of current MPNP research with specific focus on a variety of buyer stocking constraints, product substitution through one- and two-way substitution, and complementary products; it also discusses ideas for future research. Chapter 2 considers a different aspect of the MPNP. The objective is to re-

To improve service offerings, reduce costs, and increase profit margins, many companies are employing collaboration and cooperation strategies. For these companies, the main drivers are the economies of scale they can achieve and risk pooling, such as virtual pooling of inventories in which stocks are kept locally but can be transferred to satisfy demand. Chapter 4 addresses a repeated newsvendor game with transshipments. Retailers act as they do in the classical newsvendor problem; however, after demand has been satisfied, the individual retailers must decide how much leftover stock to share with other retailers. When they do decide to share stock, residual inventories are transshipped to meet residual demands and residual profits are determined based on dual allocations. Chapter 5 presents a review of newsvendor games. It discusses a simple newsvendor game and its extensions, such as large

games and multiple warehouses. The ultimate aim is to determine whether the anticipated benefits can allow a stable collaboration to be formed.

Chapters 6 and 7 consider newsvendor risk preferences. Newsvendor models typically assume risk neutrality. However, we are increasingly seeing attempts to introduce and use other risk preferences to describe newsvendor decision-making behavior. Chapter 7 focuses on risk preferences within the expected utility theory and prospect theory frameworks and directions for further research in the context of these frameworks. Chapter 8 discusses a comparative analysis of the value-of-risk (VaR) measure, the conditional value at risk (CVaR), and the mean-CVaR as objectives and constraints, respectively, in newsvendor models.

Chapters 9 and 10 address a fairly recent phenomenon in which (1) end-of-life products are returned to the original producer, and (2) the producer can, depending on the product's condition, decide to remanufacture the used product. Mobile phones are one of the numerous examples of such products. In one case discussed, a two-period newsvendor-type approach is presented to model the optimal production and remanufacture strategies for a closed-loop supply chain. In such a situation, the link between production and the sales of new products and subsequent supply of used products is important. The second application is the development of a remanufacturing model that assists the newsvendor in making inspection decisions that yield maximum expected profits in a situation with a fixed number of end-of-life items.

Risk pooling, or statistical economies of scale, is a well-known managerial mechanism in different industries. In Chapter 11, using the newsvendor setting, a centralized inventory model is developed whereby inventory is pooled. The inventory allocation to the various retailers in the system is based on each retailer's contribution, which is defined as the number of units the retailer produces and deposits at the central warehouse. A model of a pooled inventory system with priorities shows its advantages over a nonpooled inventory system. The uncertain reliability of machines in metal stamping plants is a huge concern. Chapter 12 discusses an enhancement to the classical newsvendor problem to address this issue; the model incorporates multiple items, machine setup times, and the machine's lack of reliability.

An additional, relevant application is the introduction of carbon emission policies, because high levels of carbon emissions are a major concern in today's world. Chapter 13 addresses this topic by analyzing the classic single-period newsvendor problem under three carbon emission poli-

cies: the mandatory carbon emission capacity, the carbon emission tax, and the cap-and-trade system.

Chapter 14 discusses a practical example in which a fresh product could decay or lose freshness as a result of being transported over long distances. Decisions about the order quantity and selling price are made in sequence; however, the ordering decision is made before the product is ready to be transported, but the selling price depends on the freshness of the product when it arrives at the destination and the timing of the market. The model developed is a variation of the traditional newsvendor problem.

Chapter 15 attempts to address the profit target-setting approach for multiple divisions within a firm. The firm has a number of independent, autonomous divisions, each of which has full responsibility and accountability for its own operations and decisions. Each division is given a profit target and must determine both a production level and a selling price to maximize the probability of achieving the divisional profit targets. It also examines how to set the profit targets.

Finally, Chapter 16 investigates a portfolio approach to the MPNP with budget constraints. The procurement strategy for each newsvendor product is designed as a portfolio contract, which consists of a fixed-price contract and an option contract. With a fixed-price contract, the retailer pays a unit fixed cost for procuring each product; with an option contract, the retailer pays a unit reservation cost up front and receives a commitment from the supplier. For each unit procured under the specified commitment level, the retailer pays a unit execution cost. The objective is to find the optimal quantities for the portfolio contracts to maximize the total expected profit.

The *Handbook of Newsvendor Problems: Models, Extensions and Applications* addresses a wide variety of problems illustrating the breadth and depth of the current research related to the newsvendor problem—research on different models, model enhancements and extensions, and real-world applications. Most of the papers presented clearly have a direct link to real-life situations, thus making this handbook particularly relevant in the supply chain management environment. This publication includes many applications and suggestions for future research, which will be valuable to researchers, practitioners, and postgraduate students who are interested in working on the newsvendor problem.

(This book review appeared in Graham K. Rand (2014) *Book Reviews*. *Interfaces* 44(1):108-117.

Book info: Handbook of News Vendor Problems: Models, Extensions and Applications, Tsan-Ming Choi (Ed.), 2012, Springer, pp 384, \$179.00.