Newsletter

Operations Research Society of South Africa
Operasionele Navorsingsvereniging van Suid-Afrika

March 2009

www.orssa.org.za

VUMA is a world leading windows-based software package for mine ventilation, cooling and environment control, and as such is applicable to a full variety of underground mining methods.

VUMA Products are based on current R & D, using state-of-the-art procedures and technology. All VUMA algorithms have been verified in recent work and over a development period of some two decades.

Position Available: Scientific Software Engineer/Developer

A position for a scientific software engineer/developer currently exists at VUMA Software Adco in Bryanston, Johannesburg.

Description

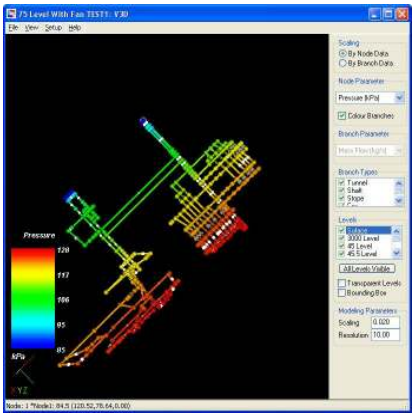

The ideal candidate should hold a degree in the Mathematical Sciences and have an interest in developing software of an Engineering/Scientific nature. A minimum of three years work experience in a software development environment is required, and knowledge or experience in Delphi and C++ Builder would be advantageous. Remuneration is market related according to experience and qualification.

The position offers a dynamic and diverse working environment in the fields of software development, modeling and simulation.

More information available on the web at www.vuma.co.za
Candidates should send their CV accompanied by a letter of motivation to info@vuma.co.za.

FROM THE EDITOR

Contactable at: zane@sun.ac.za

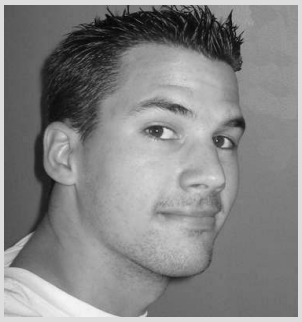

Zane Simpson

This is my first newsletter as Editor, taking over from Basie Kok, and I really hope you enjoy it reading it. I want this newsletter to remain an enjoyable read as it has been. One small change I have made starting with this newsletter, is to include a comic in every newsletter, which has a mathematical "theme".

The member profile is Jacques du Toit, a PHD student in Operations Research at Stellenbosch University. The main article in this newsletter is by Jonas Stray, from the University of Boras in Sweden, on sugarcane harvesting schedules, and there is also yet another great book review by Hans Ittman.

Features

COMIC

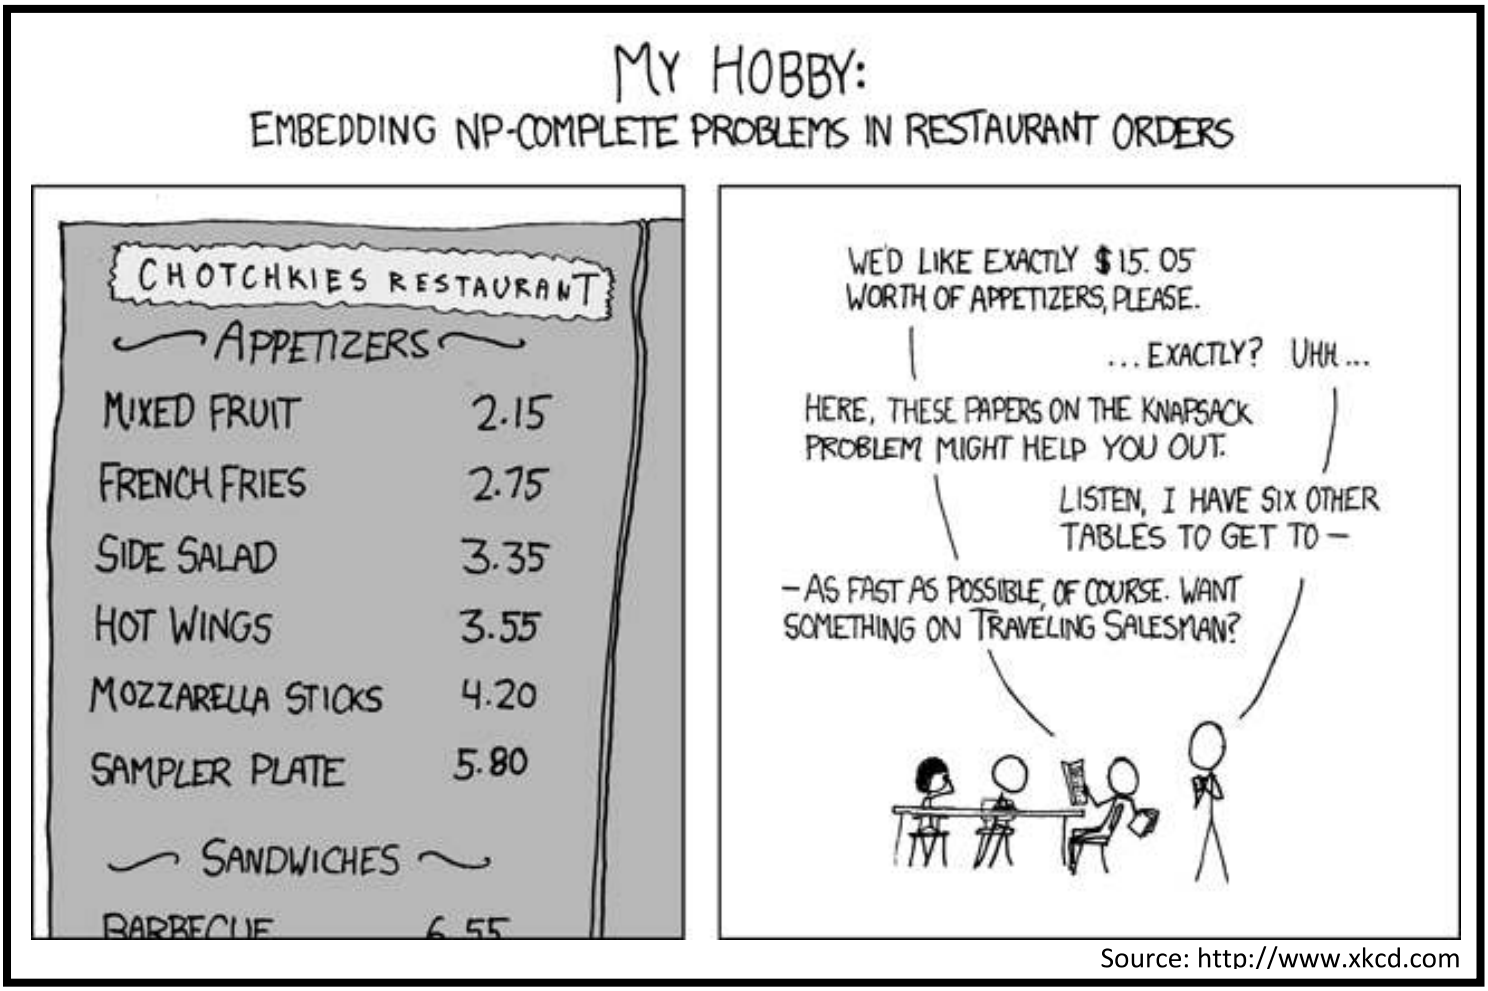

Source: http://www.xkcd.com

FROM THE PRESIDENT'S DESK

by Sarma Yadavalli (yadavalli@postino.up.ac.za)
ORSSA President

Sarma Yadavalli

Greetings Members. We are already into the first quarter of 2009! I hope this message finds you well and having made a great start to the year 2009. During the year 2008 we had a very successful IFORS 2008 conference, which was well attended and informative. ORSSA is proud of this achievement..

I welcome all new members of the executive, and the new executive will meet for the first time in March 2009. I am glad to announce that Dr Dave Evans has graciously accepted the appointment of Vice-President of our society. I am also pleased to see many young members coming forward to volunteer their services to the Society by joining the executive committee. This is a good sign.

However, with much regret, I note Prof Wim Gevers and Dr Johan Joubert's departure from the committee. Wim worked selflessly and tirelessly for the Society in the past, in various positions of the executive. On behalf of the society, I wish to thank Wim and Johan for their services and look forward to seeing them on the committee in the near future. Martie Harmse has served as Vice-President for the year 2008, and I sincerely thank her for her services. However, ORSSA welcomes her as a treasurer of the society and appreciates her acceptance to serve the society in this way.

Our Society (ORSSA) must be strong, energetic and healthy for us to continue with our hard work. It is also essential that we continue to grow and prosper. Also, we desire that Operations Research should play an effective role in building a better tomorrow.

I am aware that ORSSA runs chapters in certain cities and helps to organize regional seminars for Operations Researchers, where people get an opportunity to present their research findings. With the tremendous advances being made in Operations Research, I have no doubt that ORSSA will assume an even greater role in steering researchers in the country and so contribute to their development that can be at a par with the best in the world. It is indicative of our times that today, the science of Operations Research is becoming more and more inter-disciplinary and is collaborating with other branches of science. Therefore, at the top of the list of priorities one should aim to strengthen the ORSSA Chapters of the different regions.

NEW ORSSA EXECUTIVE COMMITTEE 2009

President:
Sarma Yadavalli

Vice-President:
Dave Evans

Secretary:
Isabella Nieuwoudt

Treasurer:
Marthi Harmse

ORiON editor:
Jan van Vuuren

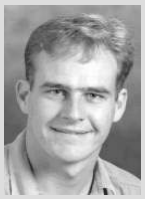

ORiON business manager:
Stephan Visagie

Newsletter editor:
Zane Simpson

Newsletter business manager:
Francois Bester

Webmaster:
Basie Kok

Additional members:
Ozias Ncube
Hennie Kruger
Maseka Lesaoana
Jacomine Grobler

Chapter Chairs:

Johannesburg:	Niel Manson
Pretoria:	Danie Payne
Vaal Traingle:	Andy Msiza
Western Cape:	Margarete Bester
KZN:	Carel Anthonissen

IFORS Representative:
Paul Fatti

Archivist:
Ilse du Plooy

MEMBER INTERVIEW: JACQUES DU TOIT

Contactable at: jacques@dip.sun.ac.za

Jacques du Toit

Jacques originally comes from a small town named Musina which is situated on the northern border of South Africa. Upon completing school he took three 'gap' years, spending most of his time in London and Tel Aviv. Upon returning to South Africa in 2001, the serious business of studying began. He began his university studies with the intention of doing computer science but found mathematics to be more enjoyable and so pursued a degree in Physical and Mathematical Analysis. He recently obtained a masters degree and is continuing with a PhD, also at the University of Stellenbosch.

How did you become involved/hear about OR, and what attracted you to it?

I first heard of OR during my undergraduate days when Prof Visagie and Prof de Kock visited one of our Applied Mathematics classes and told us of prospective OR courses. I was caught up in a degree already and was not keen to change so late in the game. It was only during my postgraduate studies that I came across it again and this time it was due to the fact that I had stumbled into an OR group doing research under Prof van Vuuren. The rest is history as, whoever they may be, say.

You are currently pursuing a PhD in OR. Is this correct, and what will your research be on?

I am conducting research within a persistent area surveillance system currently being investigated by the CSIR (Awarenet). This system will initially monitor illegal activities in the South African maritime exclusive economic zone (EEZ) in order to make information available to agencies, such as the National Parks Board, so that they may make informed decisions regarding activities in their area of responsibility.

I will be looking into the intention estimation component of the system using pattern recognition methodologies to make sense of and classify the observed behaviour of entities over time.

What led you to this decision instead of going directly into the workplace, and would you recommend this to other aspiring OR practitioners?

I was quite keen to get on with other things but found myself to have enjoyed the Masters so, that when the possibility of a PhD became a reality, I found it hard to ignore. I had toyed with the idea during my last year of Masters and had also considered going abroad for it. It seemed that the best decision was to continue immediately with the PhD (I fancied another gap year) as I am easily distracted and may not have managed to get back to studying. A number of PhDs at my local watering hole also stressed the importance of other factors, such as one's supervisor. I ultimately decided to continue my studies in Stellenbosch after careful deliberation.

I don't know if I would go so far as to recommend to others to study as long as I am, but I wish to learn more and studying further allows me to expand my horizons and investigate interesting problems.

Which specific area of expertise are you the most interested in?

That is a little difficult for me as my Masters was my first foray into OR work. I am mostly interested in fun mathematics. I am also a great fan of programming and so might try my hand at creating some tools myself in future. Presently, studying towards a PhD is exposing me to data mining and such notions as fuzzy set theoretic models, support vector machines and self organising maps. These ideas, and likely many more, will hold my interest over the next few years.

Have you been involved in any interesting OR projects or research, and what?

I was involved in basic research towards a threat evaluation and weapon assignment decision support system, currently being conducted at Stellenbosch University (this project featured in the September 2007 newsletter). I investigated a means to describe the smooth motion of a body featuring six degrees of freedom.

Interpolation and approximation methods were used to fit spline curves to a sequence of waypoints in three-dimensional space.

The work was wonderful in that it was very visual in nature and I am a big fan of pretty pictures.

What are your dreams and aspirations as an OR practitioner in South Africa?

Wow, I am a little unsure as to my dreams and aspirations as a human, nevermind as an OR practitioner --- I am more like a leaf in a river than a man in a boat. Upon completing my PhD, I expect I will

know a fair bit more about my capabilities as well as about OR. I will defer the question until then.

What would your commentary be on the role of ORSSA in promoting OR in general, and how do you think the society is doing in terms of achieving this goal locally and internationally?

Being the Operations Research Society of South Africa, I would expect that the society should play a central role in promoting OR. I believe this to be the case, although, I am unsure that I am qualified to answer this question as I am a 'back bencher'. I personally do my bit in social settings these days by telling people that I am studying towards an Operations Research degree; I have not yet met someone who knows what it is. Then again, people used to ask me during my masters what one might 'do with mathematics', as if it were an utterly useless pursuit.

Do you have a message for other young OR practitioners?

I have found the OR community to be most welcoming and I would encourage others to stick around and see for themselves. From the conferences I have attended and papers I have read, it is clear that OR straddles many different disciplines with something for everyone (everyone not being absolutely everyone in this case).

THIRD CALL FOR PAPERS

NATIONAL CONFERENCE OF THE OPERATIONS RESEARCH SOCIETY OF SOUTH AFRICA

University of Stellenbosch

20–23 September 2009

We are pleased to give third notice of the 38th Annual Conference of the Operations Research Society of South Africa (ORSSA) to be held at the Stellenbosch Institute of Advanced Study on the main campus of Stellenbosch University. During this conference we shall celebrate the 40th anniversary of ORSSA. Full details of the meeting, including a preliminary programme and the registration fee structure are available on the conference website:

www.orssaconf.co.za

Online registration opened on Monday March 2nd 2009.
Please make a note of the following important dates:

Friday 31 July 2009:	Deadline for Abstract Submission
Tuesday 30 June 2009:	Deadline for Early-bird Registration
Tuesday 15 September 2009:	Last day of Registration
Sunday 20 September 2009:	Welcome Reception
Monday 21 September 2009:	Start of Conference

Registration and abstract submission may be done online on the conference website.

HOW TO REDUCE CUSTOMER ATTRITION IN THE FINANCIAL SERVICES INDUSTRY?

The financial services industry is a maturing market that has seen a dramatic increase in competition over the last years. In times of strong market growth, a narrow focus on customer acquisition can be the key to ensuring business success in this highly competitive market.

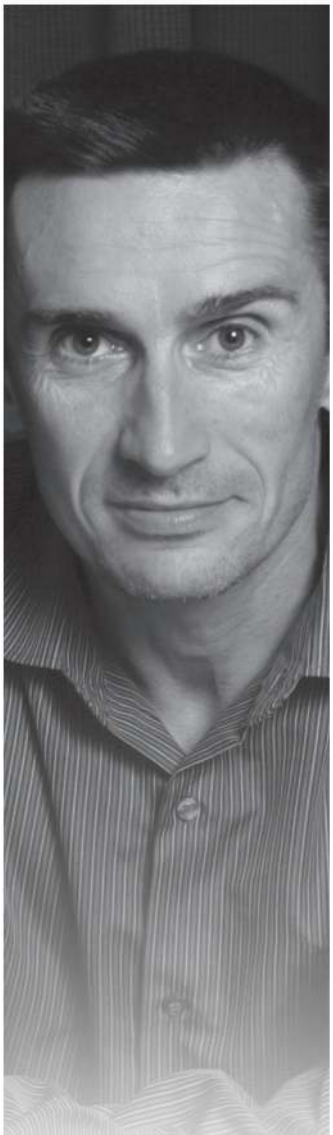

This is the view of Goran Dragosavac at SAS Institute SA, who explores how using analytical technologies to understand why your customers are leaving can stop potential attrition.

With the growing saturation of the market, financial service providers are beginning to re-evaluate ways to grow their business. Studies have shown that retaining existing customers has a bigger positive impact on the bottom line in the long-term than the unconsolidated acquisition of new customers.

If not through customer acquisition, how else can financial providers guarantee the growth of their business? In my opinion the fastest way to do this is to begin looking at how to start reducing the number of customers you are losing. The growth of your customer base is the difference between acquired and lost customers, and it is easy to see, that reducing the number of lost customers will have strong effects on the business. Financial organisations need to understand their customers on a more individual level, learn why customers have chosen to maintain a relationship with them, what is the present and future value and profitability of each customer, what are the distinct customer segments in respect to their purchasing behaviour and preferences – and of course, which of customers may seek to leave within a given period of time - and why?

Possibly even more importantly – they need to embrace new products and initiatives that will assist them to change the minds of these customers and in turn encourage them to stay. The answer is in adopting analytical technologies and solutions that would enable customer service organisations to create the knowledge they need to implement the right retention strategies to minimise defection of valuable customers.

Although, modern analytical technologies can help to reduce customer attrition, sometimes dramatically, it should be stressed that customer retention is best addressed if it is part of a bigger picture of customer lifecycle management. It is also only successful if an organisation can react quickly and positively, through the development of new products and services, campaigns, and opening the channels of communication with the client to ensure that the issues they want to deal with are dealt with timeously.

Successful retention starts with the first customer contact and continues throughout the entire lifetime of the relationship, and the most successful retention strategies are those that are in built on the pillars of other customer intelligence applications – such as customer's risk, value, as well as purchasing behaviour - and not in isolation.

To learn more about how to meet the requirements for real-time decision making, contact SAS on +27 11 713 3400 (Johannesburg and Pretoria) or +27 21 912 2420 (Cape Town) or visit our website, www.sas.com/sa.

sas | THE POWER TO KNOW.

Sugarcane Harvesting Schedules

by Jonas Stray (jonas.stray@hb.se)

PHD student in Operations Research at Stellenbosch University

Figure 1: *The author, Jonas Stray*

The project that this paper is about is funded by the University of Boras, Sweden, and is a multi-disciplinary collaboration project between myself, Professor Carel Bezuidenhout at UKZN, and Professor Jan van Vuuren at US. We have been warmly invited to share in the research of the South African Sugar Research Institute, the daily and seasonal operations at the Noodsberg and Sezela Mills in Kwa-Zulu Natal and accommodated very openly by four progressive and generous farmers in the Richmond area, south of Pietermaritzburg. The paper is an account of the modelling experience thus far, perhaps to be continued.

Figure 2: *Professor CN Bezuidenhout*

The project really began sometime during the spring of 2007 when a South African agricultural engineering professor met a colleague of mine in Cranfield, UK. The professor turned out to have been bitten by the operations research bug, and had come up with a scenario that he figured could be solved by means of some optimization tool. Now, I was in Sweden heading up an Industrial Engineering department that was (and is) in desperate need of lecturers with doctoral degrees, and I had wanted to pursue the degree for myself for quite some time. I invited the Professor, Carel Bezuidenhout of UKZN, to take the trip to Sweden and give a talk on the sugarcane planning problem that he had briefly described to my colleague. It turned out to be an interesting talk and an interesting problem in a very interesting country, so I packed my bags and headed for the Republic of South Africa to optimize sugarcane operations in some way or other. Upon landing in Cape Town, we took the obligatory tour around Cape Point, which was stunning, and visited some of the landmarks. Driving on the left (wrong) side of the road wasn't so hard for me since I'd been to the UK a few times, but my fiancé had a tough time. We headed to Stellenbosch to register and meet with Professor Jan van Vuuren who is my supervisor. I had been welcomed by Professor Bezuidenhout over the phone, but it was nice to meet someone face-to-face. Two weeks had passed, and neither Professor van Vuuren nor I knew enough about sugarcane to formulate any meaningful models, so off I went to Kwa-Zulu Natal and Pietermaritzburg to learn enough about sugarcane farming and milling to excuse the prolific variables and constraint production to come. I met with Professor Bezuidenhout who immediately hooked me up with numerous people from the industry; farmers, millers and researchers from every corner of the country were represented.

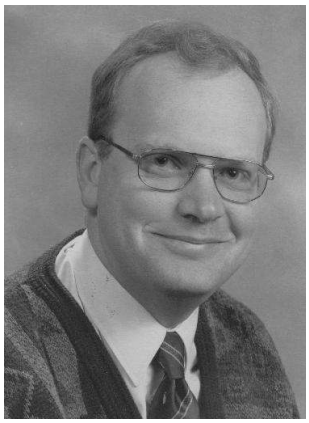

Figure 3: *Professor JH van Vuuren*

I didn't want my models to become great exercises in the art of constructing complicated constraints, but I also realized that I couldn't present a standard formulation of the transportation problem and use LINGO to solve it and present that methodology as my dissertation. I don't have the writing skills to make such a thing look unique enough to warrant the Bestowing of the PhD degree!! After reading some papers about optimization models in the sugar industry I was under the illusion that a linear program that determined the proportion of each field to harvest during every day of the season could be the way forward, so I spent weeks entertaining the idea, only to hit one merciless obstacle. It turned out that the decisions on the ground did not at all resemble my decision variable. A field is usually between 1 to 10 hectares in size and contains between 50 and 130 tons of cane, depending on the age of cane and the weather and many other factors. A typical South African sugarcane farmer always burns an entire field at a time, the main reason being that the fire is very hard to stop without some kind of firebreak. If the field is large, the farmer may then spend one or several days harvesting that field, involving very few management decisions, except things like providing direction for loading vehicles, cane cutters and trucks. Obviously, my fractional decision variable was useless,

since the decision was of an integer sort, and I now started to become very careful in choosing how to model the problem. I did what I used to tell my students back in Sweden, that is: carefully write down a problem statement, truly explaining what you perceive the problem to be, including every imaginable aspect. When I started to write down the problem, which ended up taking several days, I realized I had to go back to some of the people that I had met and I also had the distinct feeling of having skipped something or at least avoided something.

Besides making the mistake of starting out early on to formulate the problem as a linear program, I made the assumption that data was going to be the greatest obstacle. I knew that from other people's projects as well as from my own work and it is said that, in the end, often 90 percent of the work in a real world application of OR is data gathering and handling. This "fact" somehow made me imagine the model as a tool that is run once before the beginning of the season because then there would be enough time to update the data during the off-season. I thus made the gross error of letting the known difficulties of the modelling process dictate the design of the model itself. I wasn't comfortable with the model and I didn't know why. I simply had the sense that it wasn't going to be used by anyone involved in actual sugarcane operations. I suddenly realized that I had to be true to my own perception of reality. The problem wasn't only to maximize sucrose nor minimize operational costs nor just a combination of the two. It was *when* and *how* to do it!

After spending a day at the harvesting scene together with Shawn Kyle, a cane harvest operations manager in the Richmond region, I learned that a weekly harvest schedule that points out which fields to harvest that week, rarely lasts a week. Sometimes a cane fire inadvertently jumps to an adjacent field, which means that that field must be harvested while some field that was in the schedule has to stand back. Sometimes the wind blows in a direction that makes burning a particular field too risky, and that field's harvest is postponed. Some fields are accessible after rain, while some are not, and that might cause changes to the schedule. I realized that harvest schedules must be updated at least every week, not once a season.

I also realized that the problem is not just the one of maximizing sugar or minimizing cost, it is also one of staying within the permitted volumes per week and also to obey capacity constraints. The scheduling program must be able to assign fields to weeks and do so while maximizing the farmer's sugar output *compared* to the mill regions average[1]. It must minimize harvesting costs that depend on the risk of rain and the different effect that wet conditions have on the operations at different fields. It must estimate recoverable value (cane quality) and yield as far into the future as 9 months and must produce solutions that take tonnage restrictions into consideration.

The yield of a particular field depends on many factors, and which factors to include depends on what the model is required to do, but also what data are available. My scheduling tool should be able to predict yield as well as possible with as little variability as possible, but the data available to accomplish this is limited. I was more or less left (this might change, it has only been a year since the project started) with the choice of using historical data from individual farms or using weather-based growth models that assumed perfect farm management practice. I may resort to the yield matrices used by the fire insurance industry, but that would be very boring and would make the model quite lame (fire insurers assume that ratoon, variety and farm individuality does not matter). I am currently inclined to employ historical data from each farm and fit a regression model to be used as the prediction model within the scheduling tool. An example of the response surface generated from such a regression model can be seen in Figure 4.

Cane yield seems more difficult to predict at first, due to the larger variability, but the general trends are not completely obscured. In Figure 4, a response surface of yield together with the data points used to fit the underlying regression model shows the resulting imprecision.

The regression models themselves can be fitted quite easily with multiple linear regression analysis software, but the quality of the predictions for the coming season depends on how we adjust them to the current conditions. I am going to try to use a farmer's eye-estimates to select a model by minimizing the sum of squares of the residuals between his estimates and each of the previous year models that are available for his farm. In that way, I attempt to find a historical year that behaved like this year will behave, and thus end up with decent models of both RV and yield. The 95% prediction interval for the mean of the yield of field of a particular age, ratoon and time of year is usually ± 10 tons wide, while the prediction interval for an individual

[1] Farmers are today paid more for their cane if the cane quality is higher than the mill's weekly average, i.e. if they are better than the other farmers in the region. They are paid less than average if they have lower quality cane than the weekly average.

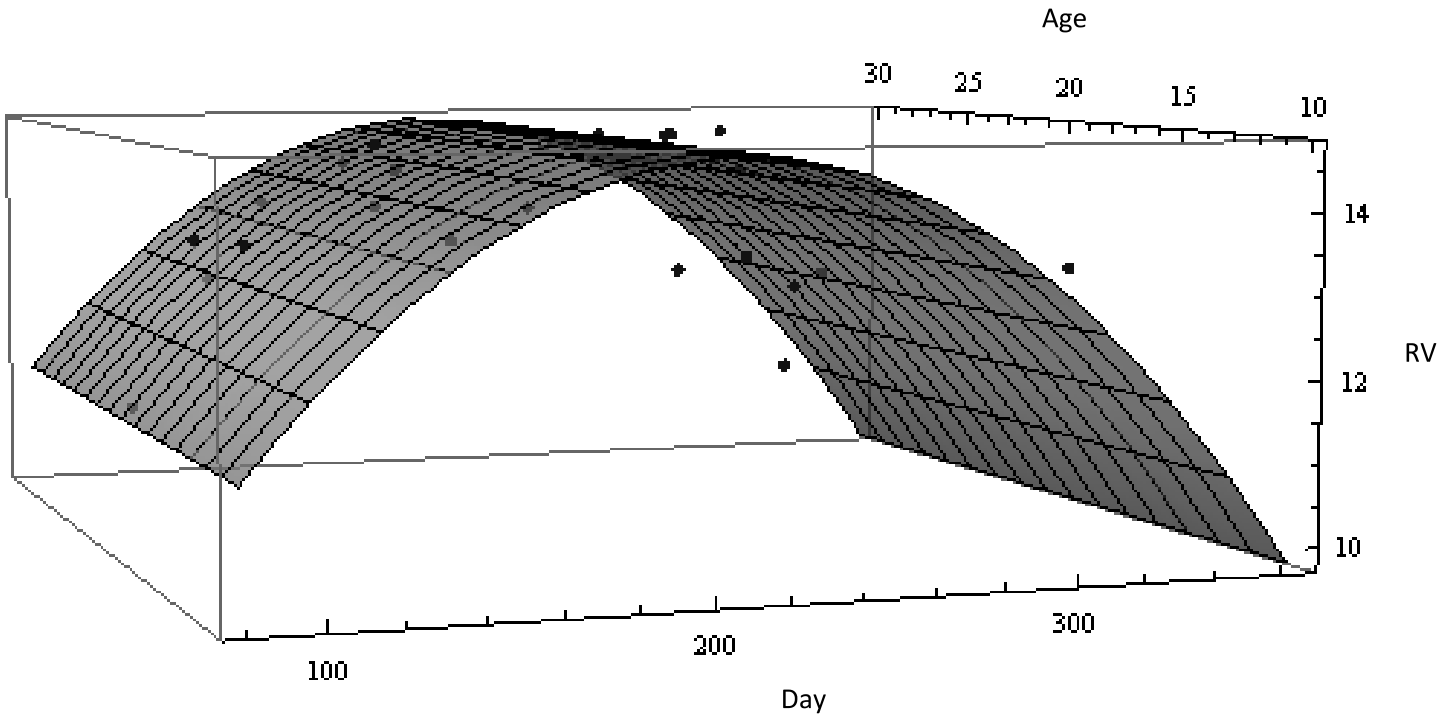

Figure 4: *RV (Recoverable value, or cane quality) is shown as a function of the day of the year counting from January 1st, and the age of the cane in the field in months. The points are the actual data points from a farm in Kwa-Zulu Natal during one season. The RV peaks in early July (around day 210) this year on this farm.*

field is greater than ± 30 tons (three times as bad!). Therefore, I am also using the farmer's estimates to correct the estimated value by field, not just by cane age, day of the year and ratoon, because the farmer actually sees and sometimes measures the state of each field.

In addition to the prediction models that were just discussed, I realized that if something out of the ordinary occurs, such as a frost or a pest infestation or a hail storm, one must account for the effects that the particular event might have on yield and cane quality. These phenomena have been researched in the literature, and subsequently modelled by, for each phenomenon, applying a function in addition to the regression models. The functions are straight lines that decrease the RV and yield respectively, by some constant value per day into the future. This solution, I hope, will represent the loss in value incurred by waiting to harvest fields that have been adversely affected by environmental anomalies.

In order to generate cash from the cane that stands on the field, the farmer must harvest it according to the rules and transport it to the mill. The costs incurred are sometimes expressed in a Rand per ton figure, but the farmers and I (I will tell you about the farmers soon) found it very interesting to see whether accounting for wet conditions would make a difference. The problem with wet conditions is that some fields may be harvested as usual, while some fields become harder to access or even impossible to access by machinery. There is a capacity drop across the entire midlands region during the wet season, and that is connected with the fact that nobody harvests during rain because the fields are very slippery and may also be damaged due to the top layer of the soil being scraped off by vehicles. A few hours after rain, many fields have dried enough to be entered by tractors with trailers, and some fields even by trucks. The fields that become harder to access when wet ought to be cheaper to harvest when dry, we thought. We therefore decided to differentiate between fields in terms of their accessibility during wet conditions and apply costs accordingly.

"We" are four farmers, their operations manager and I. The farmers are Clive Coulthard, Eric Lewis, Roger O'Neill and Malcolm Thompson, and the manager is Shawn Kyle. I was introduced to this genuine group (of friends to be) by the then Extension Officer to the Eston cane growers, Edgar Bruggeman. Professor Carel Bezuidenhout and I joined them on one of their meetings and, though a little sceptical at first, they agreed to help me develop and validate my scheduling tool. We decided to use the coming 2009 milling season as an experimental ground to develop and evaluate the ability of the scheduling tool to produce usable and meaningful schedules. I collected historical data from their farms, and built the models previously discussed. I designed a detailed return form to be filled in on a weekly basis as a reply to the schedule that the farmers

receive at the beginning of every week. I hope to catch every opinion and fact, and to be able to adjust the scheduling tool and its components accordingly, so as to make it truly applicable to almost every South African sugarcane farmer. I also designed a monthly evaluation form that asks a number of general questions to the farmers, such as "Does the schedule simplify your operations?" I hope to be able to say that "The scheduling tool was useful and, according to the growers, it saved them some money."

Figure 6 *From the left: Roger O'Neill, Shawn Kyle, Clive Coulthard, Malcolm Thompson and Eric Lewis.*

The engine in the scheduling tool is a mathematical programming model and a solver. The mathematical programming formulation is equivalent to the generalized assignment problem with the addition of a lower bound on the amount of resource that may be used. The solution procedure is a local search that starts with a random assignment and continues with one shift neighbourhood, one swap neighbourhood and one ejection chain type neighbourhood in succession. The shift neighbourhood is the set of solutions obtainable by taking the current solution and ejecting one field at a time from the week that the field is currently assigned to and inserting it into one week at a time, testing the objective for improvements at each insertion and if there is an improvement the field is inserted into the new week and the procedure starts again, setting the new solution to be the current solution. When the shift neighbourhood is exhausted, the algorithm takes the current solution to the swap neighbourhood. The swap neighbourhood is defined by the same procedural structure, but each test consists of two fields being swapped in terms of the weeks they are assigned to. The ejection chain procedure starts by ejecting one field and then finds a suitable field to insert into the ejected field's week whereupon the procedure creates a number of tentative solutions by inserting the first field into some of the weeks. The ejection chain then reverts back to the solution obtained by inserting the second field into the week of the first field, and then tries to find a suitable field to insert into the newly opened

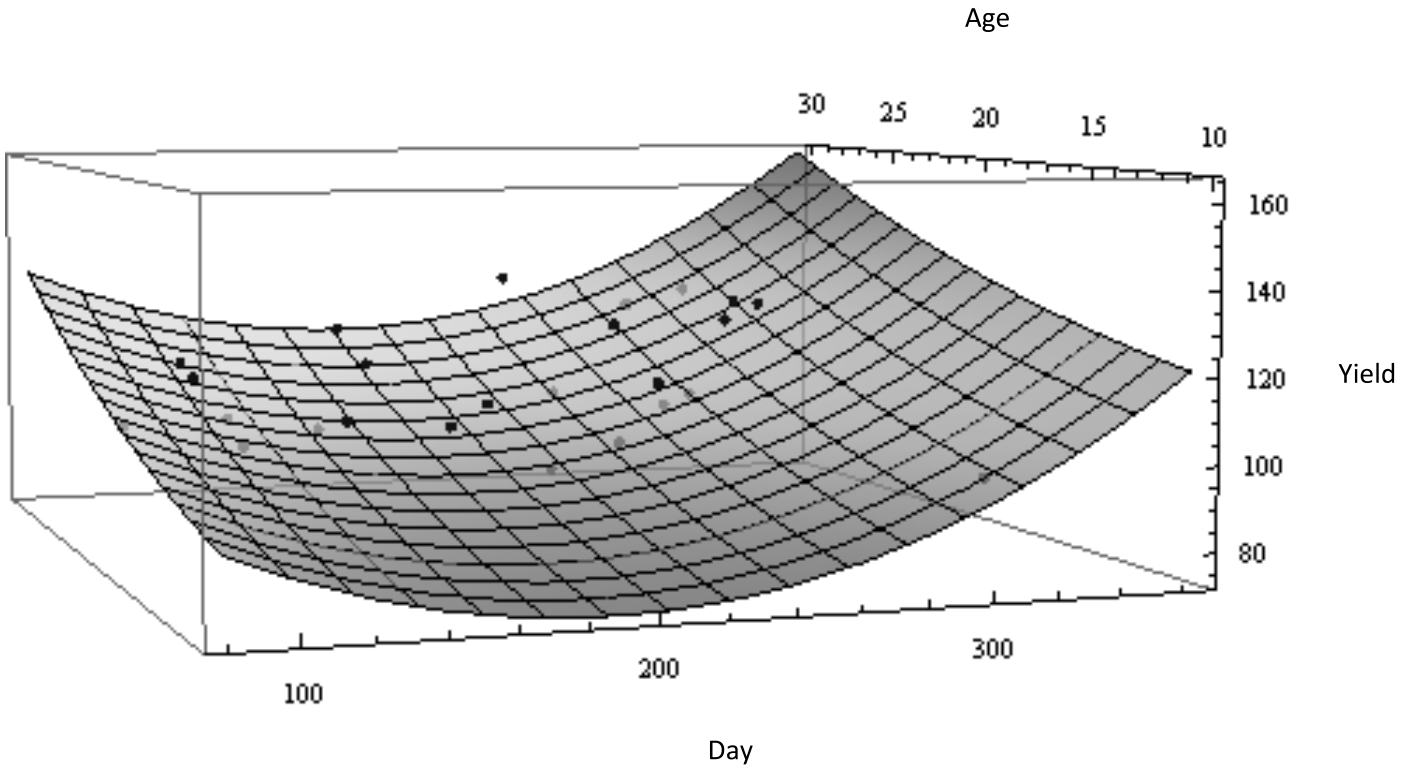

Figure 5: *This figure shows yield in tons per hectare as a function of the day of the year and the age of the cane. The model was generated using the same type of data as for the model in Figure 1. Notice that yields during the mid-season (June-August) for cane of a certain age tends to be lower. This is due to the slower growth generally observed during May to September.*

week (the second week). So the procedure continues until it has started by ejecting a field from every week.

If at any time during the procedure no suitable field is found to insert into the previously opened position, the procedure starts over by ejecting a field from the next week. Obviously, the best solution found so far is stored. When the ejection chain neighbourhood is exhausted, the algorithm takes the current solution to the shift neighbourhood. The search continues with penalty into the infeasible region, and solves my problems in less than 30 seconds, usually to within 2 percent from optimal on a fast PC, coded in Wolfram Mathematica 6.0. The stopping criteria are a time limit and an optimality criterion.

The scheduling tool has been built in Microsoft Excel with Visual Basic for Applications, but I have decided to rebuild it using Java so that the application can be run on any platform. I also think that programming it in Java will make it more user-friendly and robust, but for the 2009 validation, the Excel interface will suffice. The Excel interface will make sure that I treat the validation runs consistently and also provide information to the farmers in a readable, ergonomic fashion. An example of the output from the scheduling tool is shown in Figure 7.

Harvest Schedule

Farm ID	1			Solution alternative		3				
Date	09/02/23									
Present period					09/04/06	-	09/04/12			
Wk	15									
Field Name	**Exp. Tons**	**Age**	**Access**	**Event**	**Event Date**	**Notes**				
17	673	20	1	None	09/01/01					
32	137	22	1	None	09/01/01					
41	411	22	1	None	09/01/01					
42	298	19	1	None	09/01/01					
46	62	16	1	None	09/01/01					
Total	1581	Tons								

Following periods	Wk	Fields								
2	16	02a	04a	8						
3	25	16	23	31	34					
4	26	36	40							
5	35	01a	03c	04b	19					
6	36	01b	6	21	26					
7	45	12	13							

Figure 7: *This is a sample of the schedule that the manager will use (if he and the farmers approve of the schedule) as a guide to which fields to harvest each week. Each week I'll send them several of these schedules to choose from.*

DISCLAIMER

QUERIES AND CONTRIBUTIONS

Any queries and contributions to the newsletter are most welcome, especially article submissions. For any queries and contributions, please contact the newsletter editor:

Zane Simpson
Email: zane@sun.ac.za
Tel: 084 626 7588

Book Review

By Hans Ittmann (hittmann@csir.co.za)

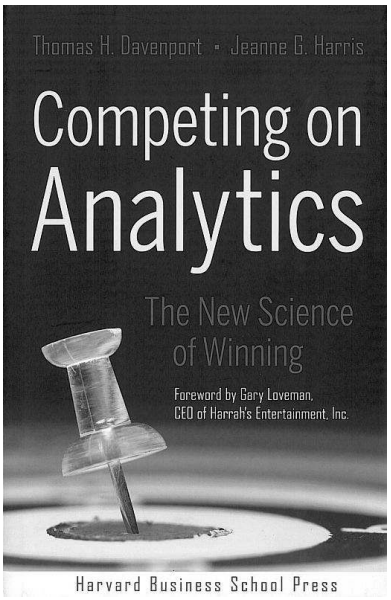

Competing on Analytics – The New Science of Winning by Thomas H. Davenport and Jeanne G. Harris, 2007, Harvard Business School Publishing Corporation, Massachusetts, USA, pp. 218. ISBN-10: 978-1-4221-0332-6. US $ 29.95.

Every few years there seems to be a new term added to the vocabulary of people working in the area of quantitative methods. *Analytics* is a term that appears to be used increasingly in the environment within which we as operations researchers operate. A number of papers in different journals contain the term; INFORMS, the OR society in the USA, started with a new online publication called *Analytics*; various articles in magazines refer to this new 'hot topic'. When I saw reference to *Competing on Analytics* I thought it would be a good idea to review the book to get a better idea of what analytics is all about.

Davenport and Harris define *analytics* as follows: "By *analytics* we mean the extensive use of data, statistical and quantitative analysis, explanatory and predictive models, and fact-based management to drive decisions and actions. The analytics may be input for human decisions or may drive fully automated decisions. Analytics are a subset of what has come to be called *business intelligence*: a set of technologies and processes that use data to understand and analyze business performance". Business intelligence according to them includes both data access and reporting, and analytics. Analytics, in turn, comprises analytical tools ranging from statistical analysis, forecasting/extrapolation, predictive modelling and optimisation. The authors argue quite strongly that organisations that use analytics extensively and systematically are able to out-think, outperform and out-execute their competition. Organisations compete in making the best decisions, which usually have systematically assembled data and analysis behind them.

The book has two distinct parts – the first focuses on the nature of analytical competition while the second emphasises the building of an analytic competency. The first part outlines how analytics is used in both internal and external processes. The second part provides a roadmap that describes how an organisation can become an analytical competitor and how to manage analytical people. This latter section also provides a quick overview of business intelligence architecture and offers some predictions of the future of analytics.

The common key attributes of typical analytical companies are, what the authors call, the four pillars of analytical competition, namely: a strategic, distinctive capability; an enterprise-level approach to and management of analytics; senior management commitment; and large-scale ambition, meaning aspiring to achieve. The five stages of analytical competition are laid out ranging from "analytical impaired" to "analytic competitor". The importance of experimentation is made clear as well as emphasising that companies need to be willing to run their business "by the numbers".

Various examples are presented of companies and government institutions of a wide variety that compete analytically. Many of these examples are typical successful operations research applications, the authors though seem to go out of their way to avoid using the term *operations research*. Those who compete successfully have analytical capabilities that are hard to duplicate, unique, adaptable to many situations, better than the competition and renewable. The final two chapters in the first part show how to compete in analytics with internal and external processes. The internal processes that are discussed include finances, manufacturing, R&D and human resource applications. The external processes focus on customer relationship management (CRM) and supply chain management. Acquiring and retaining customers are getting more difficult and expensive. Customers are harder to satisfy and more demanding. Marketing plays a critical role in retaining these clients. Some of the typical analytical applications in marketing are: CHAID, conjoint analysis, lifetime value analysis, market experiments, multiple regression analysis and price optimisation. On the other hand, the analytical applications in supply chains are: capacity planning, demand-supply matching, location analysis, modelling, routing and scheduling. The CRM

applications are very statistically oriented while the supply chain management applications are very OR related.

The second part of the book looks at developing an analytical capability. A road map for achieving this is presented through a five-stage process. The key elements of an analytical capability in terms of organisational, HR and technology capabilities are given. It is also necessary for an organisation to determine a strategic focus, for example, Wal-Mart in the USA focuses its analytical capability on supply chain management and marketing.

People are critical in the success of analytics because ultimately, it is people who make analytics work. Total buy-in and ownership are required throughout an organisation, starting at the CEO and senior executives. A few characteristics of such executives are:

- They should be passionate believers in analytical and fact-based decision-making,
- They should have some appreciation of analytical tools and methods,
- They should be willing to act on the results of analyses,
- They should be willing to manage a meritocracy.

The analytical professionals are equally important. UPS in the USA, for example, has a 60-person Operations Research Division made up largely of people with engineering and mathematical degrees, including several PhDs.

Finally, the technology and architecture of the business intelligence function are very important. Proper data management, tools for handling, processing and managing data and the appropriate analytical tools and applications are required. Much attention is devoted to data, i.e. are the data relevant, where does one source the data, how much data are required, what is the quality of the data and finally, how to manage the data. The analytical technologies range from spreadsheets, data mining tools, simulation tools, genetic algorithms, expert systems, etc. Clearly, a very comprehensive set of tools that reflects a combination of statistical and OR tools.

The final chapter of the book presents a great list of future trends and changes that can be expected. Some of these include:

- Pervasive business intelligence software,
- Increasing use of dedicated 'business intelligence appliances',
- More automated decisions,
- Consistent with the preceding trend,
- Greater use of alerts,
- More visual analytics,
- More prediction and less reporting,
- More mining of text.

In my opinion OR forms a critical component of analytics as described in this book. Operations research makes extensive use of data before any models or methods can be applied. There is the process of identifying what data are required, collecting the data, cleaning the data and making sure it is quality data. There is a specific skill attached to this. The second component of analytics is much more statistically oriented and maybe some OR people are not very well equipped to perform this part. Finally the quantitative analysis part fits ideally with what we as OR people do, namely, building mathematical models, etc. to add value to the data and ultimately to assist in making better decisions. Analytics is therefore a new topical term that in essence describes what statisticians and operations researchers do. It seems as if by using the term analytics, many people are able to convince CEOs, managers and decision makers of the importance of these analytical tools. We should also get on the bandwagon and use, even exploit, it to our own advantage in making OR more widely used and applied! In this way one can 'do good through doing good OR"! *Competing on Analytics* captures, through many examples, the real value that operations researchers and analytically minded people can contribute to the well being of organisations.