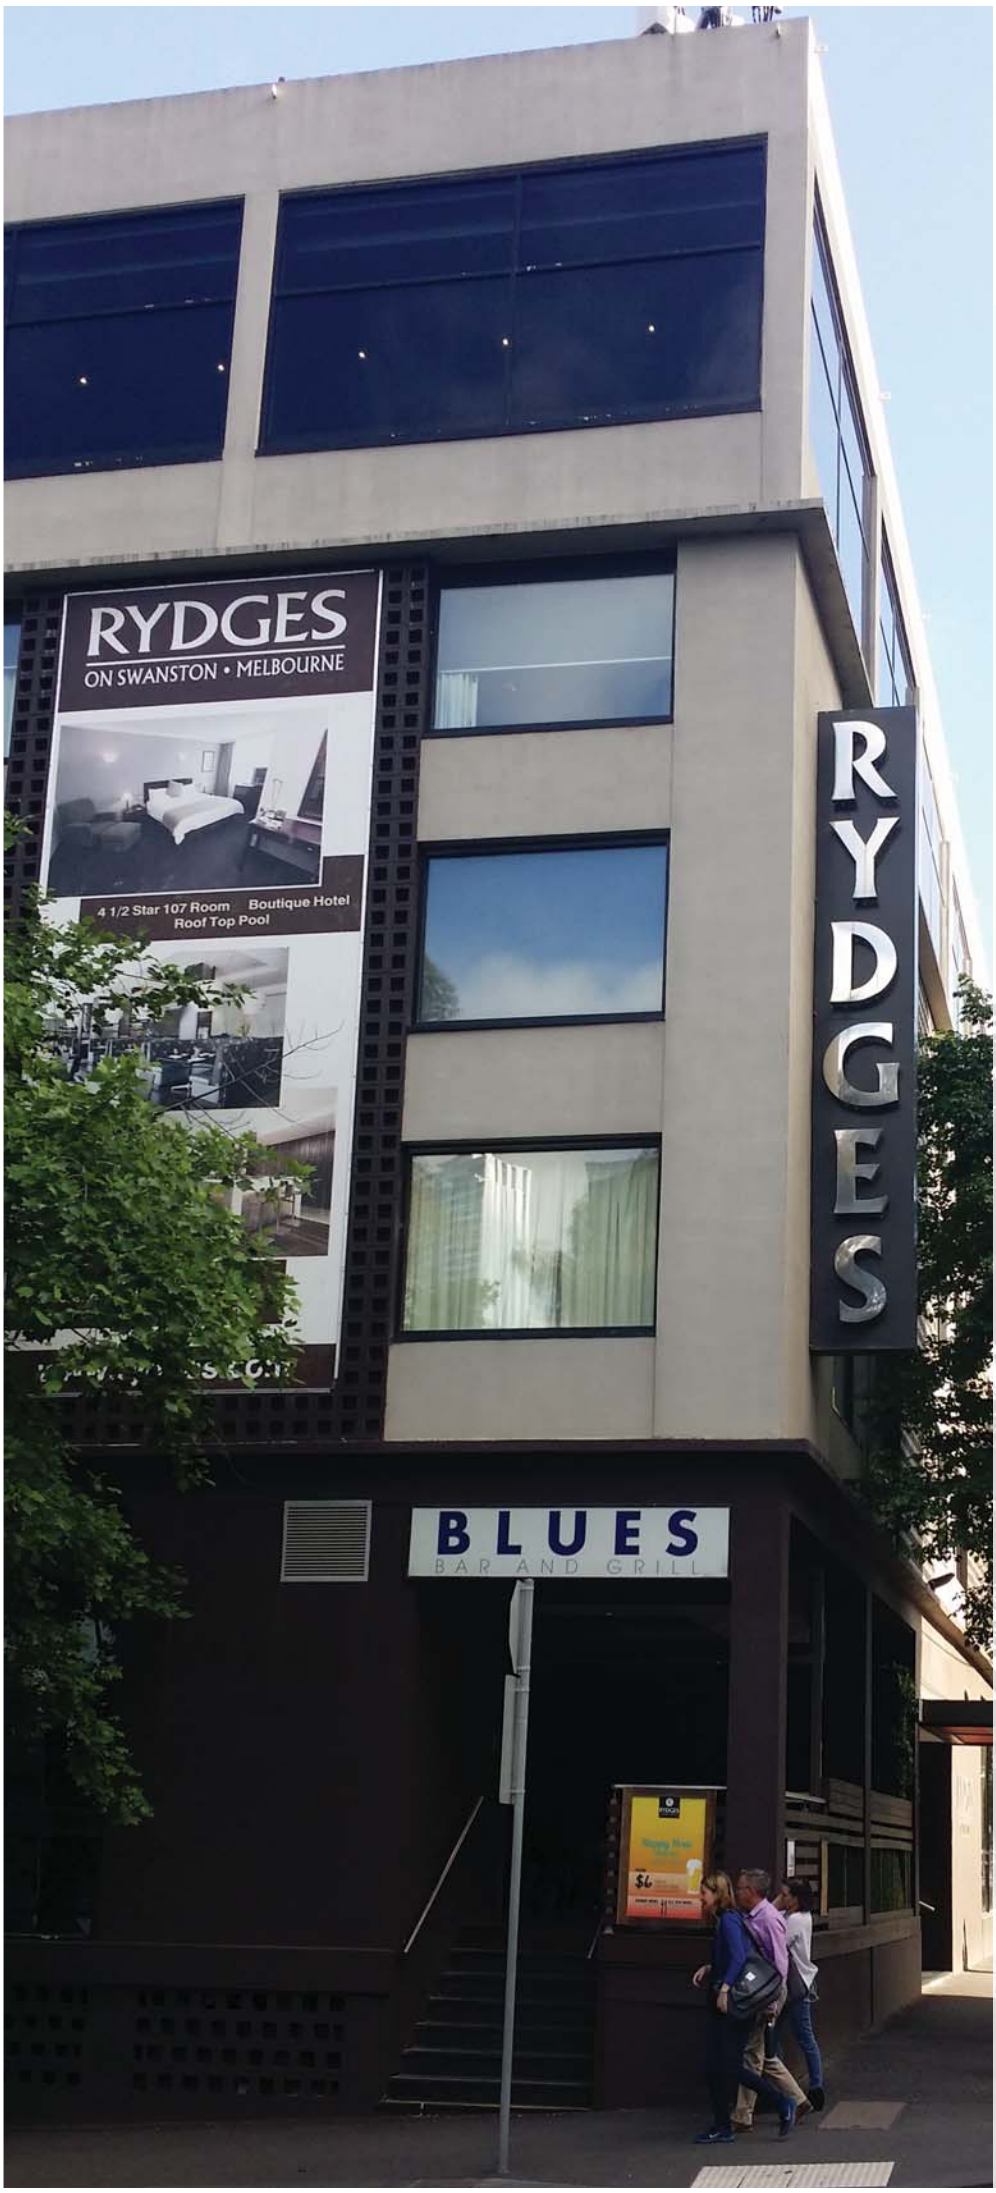

asor DORS

Book of Abstracts

ASOR-DORS Conference 2018

26^{th} National Conference for the Australian Society of Operations Research and Defence Operations Research Symposium

ASOR/DORS

2018

BOOK OF ABSTRACTS

December 4 – 6, 2018

Melbourne, Australia

Organized by

ASOR and DORS

Supported by

CSIRO Data61

Department Of Defence

ASOR-DORS 2018 Conference, Melbourne, Australia, 4-6 December 2018
Book of Abstracts
Prepared by Rodolfo García-Flores (CSIRO Data61)
30 November 2018

Organizing Committee

David Cox	DST
Simon Dunstall	CSIRO Data61/ASOR Melbourne
Rodolfo Garcia-Flores	CSIRO Data61/ASOR Melbourne
David Marlow	DST/ASOR Melbourne

Program Committee

Babak Abbasi	RMIT
Fred Bowden	DST
Thang Cao	DST
Clint Heinze	DST
Darryn Reid	DST

DORS Committee

Sanat Bista	Thang Cao
David Cox	David Marlow
Marc West	

Session Leaders

A1	Amanda Parker (Data61), Amanda Barnard (Data61/Swinburne), Ritaban Dutta (Data61), Babak Abbasi (RMIT)
A3	Carolyn Huston (Data61), Vincent Lemiale (Data61)
A4	Riccardo Mancini (U. Bologna), Rodolfo Garcia-Flores (Data61)
A5	Monica Chen (Data61)
A6	Andreas Ernst (Monash), Andreas Schutt (Data61), Reena Kapoor (Data61)
A7	Alan Dormer (Opturion), Philip Kilby (Data61)
A8	Vicky Mak (Deakin), Leonid Churilov (UniMelb / Florey Institute), Asef Nazari (Deakin)
A9	Saman Halgamuge (UniMelb), Peter Baumgartner (Data61), Melanie Ayre (Data61)

Published by: ASOR and DORS

ABSTRACTS

Keynote Seminar

Collaborative Models for Integrated Fire Control

Ted Goranson

Griffith University, Gold Coast, Australia; t.goranson@griffith.edu.au

Australia faces unprecedented defence challenges. Threats are likely to be more unpredictable and complex. The cost of systems is increasing rapidly. We have bought fifth generation systems, mostly from the US but the systems and doctrine need to be tailored for Australian use. We have a 'Plan Jericho' that at least for air operations concludes that we need systems of systems that are mission-centric rather than platform-centric. But neither we nor our allies have any of the fundamental technologies to enable this, even if the controlling organisations 'woke up'. The deficiency will likely have profound consequence.

We clearly need a paradigm shift, a radical advance in underlying science. We need systems that are intelligently adaptive, and we need them agile: the next war may be essentially over in 75 minutes. In response defence operations research has to expand from the comfortable role of advising force structure, to include force operation and restructuring. While AI and Autonomy are expected to save our bacon, there are scant signs that this time they will deliver.

This talk suggests that for several reasons Australia is the ideal place to develop the required paradigm shifting technology, and that the operations research community (vice the AI community) is more likely to succeed. The talk is organised as first describing an example scenario — integrated fire control — and a generic systems of systems solution.

Relatively exotic 'situation theory' can likely be integrated with existing OR/DOR techniques and military system legacies to support the required agile systems-of-systems, built on what we are buying. It may provide trustable, auditable 'narratives' because autonomy integrates humans. Some other benefits are outlined.

The case is made for a DST-led research program that partners with allied military research groups of course. But more significant are some commercial organisations that have far greater research budgets than allied militaries combined, and are years into technologies similar to those we need. These are indicated, together with promising examples of integration into state of the art machine learning and data mining systems. We simply need them.

Underscoring all of this is the need to reason about unknowns, a hard, hard problem distinct from the bookkeeping of uncertainty. A university research consortium is forming to conduct high risk, relatively low cost research, following promising early results on little known classified programs.

Keynote Seminar

Measuring truck service times at a container terminal

John Bartholdi

Stewart School of Industrial and Systems Engineering Georgia Institute of Technology, Atlanta, The United States of America; john.bartholdi@gatech.edu

How well do container terminals support truck operations? This is an important question for logistics hubs such as Panama, with five distinct container ports and a goal of being considered, functionally, a single port-of-call. Achieving this goal requires, among other things, making the movement of containers among ports fast and seamless. Truck operations are an important part of the land-side activities of any container port. A truck will typically bring a container to the terminal. After it enters the gate, the truck will visit various regions of the terminal ("yards") whose purpose and sequence depend on the type of container delivered. For example, an empty container will typically be dropped off in the appropriate yard depending on whether the container is dry or refrigerated. A full container for export may have to stop at a customs area for inspection and processing before being dropped off at the appropriate yard. After dropping a container, a truck will typically pick up a container to leave the terminal. The container may be empty or full, dry or refrigerated, and so retrieved from different yards of the terminal. Subsequently, the truck may have to take the container to customs for inspection and processing if it is for import; or it may take the container to another terminal for transshipment. Again, the sequence of services depends on the container type. Thus each truck will enter by a gate, visit some sequence of locations in the terminal, and finally leave by a gate. Container terminals typically judge the efficiency of truck operations by the gate-to-gate "turn time". But this is a coarse measurement, as it lumps all the different activities together, including time spent driving, queueing, or receiving service. It is used because it is easy to measure, because terminals typically require entry and exit from controlled gates and so can record arrivals and departures and later match them up. However, this fails to capture times in the various yards. It is technically possible to instrument cranes, etc. to measure service times, but such measurements would fail to capture the experience of the truck and container, which may have had to queue for service. It would also fail to capture times when the truck might have had to stop for traffic or congestion when traveling within the terminal. We show how to measure — unobtrusively — all these activities within the terminal. Moreover, our method requires no special hardware and imposes no reporting requirements on services within the terminal. And all this is done on data that is already collected. This analysis allows us to report on the truck turn-times; but, in addition, details such as total driving time within the terminal, total stop times within each service area, and total stop time within the terminal but outside service areas.

Furthermore, we can report times to drop off each type of container and times to pick up each type. We can also see how these times vary by time of day or day of the week, or in relation to the ship schedules. This level of detail can enable a port to focus efforts to improve processes where they will do the most good.

D1. Mathematical Modelling of Defence Systems and Operations

Ash modg thrakatuluk: One Model to bring them All: and in Warfighting bind them

Alex Kalloniatis[1], *Keeley Hoek*[2], *Mathew Zuparic*[3] *and Markus Brede*[4]

[1] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; alex.kalloniatis@dst.defence.gov.au
[2]Australian National University, Canberra, Australia; u5642917@anu.edu.au
[3]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; mathew.zuparic@dst.defence.gov.au
[4]University of Southampton, Interactions and Complexity Group, Agents, Southampton, United Kingdom; Markus.Brede@soton.ac.uk

If ever 'in-dated' other than in the catastrophic years of 1915-1917 of the European western front, the Lanchester equations of combat have long been regarded as out-dated representations of warfighting. Among other things they lack the dynamics of manoeuvre warfare, representations of Command and Control (C2), and the role of modern communications networks. In broader terms, apart from Force Application (or 'fires') the Lanchester model fails to bring together the further five doctrinal warfighting functions [1], Force Projection, Sustainment, Protection, C2 and Situational Understanding. For this reason the model remains poorly exploited for large scale Force Design optimisation while nevertheless remaining the subject of active research in military Operations Research even in this era of low-level warfare and counter-insurgency [2]. Any serious candidate for Modelling Complex Warfighting (as one of DST's Strategic Research Initiatives is named) should incorporate a major subset of these six functions. In this paper we bring together three of these functions with the Lanchester equations as the basis. The first step is to generalise the Lanchester system to a multi-networked — or multiplex — system of N Blue agents, M Red agents, which includes kinetic engagement between the adversaries, resource movement, or sharing, and resupply. Some initial linear forms of this belie simplicity once the requirement of the physicality of warfare is imposed, namely that entities once degraded below some threshold should not be able to recover. A version of the model whereby resources may be shared within one side, Blue or Red, via a form of diffusive coupling is also inadequate as this merely reinforces weak nodes, as initial optimisation studies of this model revealed. We formulate then a third version of this model where the pattern of resource sharing within one side takes into account relative strengths against specific adversaries with which entities are in kinetic engagement. The model exploits the formalism of multiplex networks. By using network rewiring techniques in stochastic optimisation, we numerically solve the model to identify the structure of optimised engagement and resource sharing networks for given organisation of the adversary.

We further propose that this model may be further embellished with networked C2 by adapting the famous Kuramoto model for phase synchrony. The new model, incorporating now three warfighting functions, represents Blue and Red forces distributed as resources at the nodes of separate networks. The Kuramoto dynamics governs how agents in the model may redistribute resources according to their degree of coordination, and battlefield damage inhibits the ability of agents to achieve such coordination by attenuating coupling across the resource network. We demonstrate this model with an illustrative scenario of two adversaries of different topology — one hierarchical, the other 'flat', with tactical agents engaged in combat and strategic agents facilitating coordination of combat resources. We show how the topology leads to diverse battlefield outcomes depending on coupling strengths, visualised as a fitness landscape. Furthermore, we show how the model may be further developed to incorporate Intelligence-Surveillance-Reconnaissance and logistics, and thus integrating five of the six warfighting functions.

Unsolved problems in land battle casualty estimation

Michael Lauren

NZDF, Defence Technology Agency, Auckland, New Zealand; m.lauren@dta.mil.nz

Estimation of likely casualty levels from battles has always been a difficult problem for defence analysts. Generally speaking, it is only possible to estimate an expected value, rather than forecast with any precision. While simulations using explicit representations of all combatants are often used to try to estimate battle outcomes, these can be unreliable, especially for cases where there may be many intangibles and other uncertain factors. Furthermore, there is a large overhead in trying to accurately represent the behaviour of every single combatant, particularly as the battle evolves. Owing to the difficulty of using such simulation methods, the approach of describing battles using a high degree of aggregation is often used instead, which saves the trouble of having to explicitly represent the actions of all combatants. Assuming one or other of these methods may provide some reasonable estimate of battle outcome and casualty level, an important problem in planning to deal with potential casualties is the rate at which they occur, and, as importantly, the variability in that rate. Work over the past two decades by our team of collaborators has demonstrated that this variability follows a very particular pattern. Specifically, that the temporal distribution of casualties tends to obey power laws. The hallmarks of this are high degrees of clustering in the time series, and a greater proportion of extreme events than occur in normal-type distributions. Examples of these characteristics have been found by ourselves and our co-workers for both conventional and irregular warfare. These observations support the hypothesis that warfare is an example of a so-called complex system. These systems generally display power-law statistics even though they include a wide range of phenomena, for which there is often little else in common — e.g. patterns of earthquakes, rain and turbulence data, and financial market data. In this paper we demonstrate how this sort of variability can be analysed, and discuss theories for understanding it. Furthermore, we present simulation results using simple agent-based models that demonstrate the difficulty of generating synthetic casualty time series that match the properties of recorded historical casualty time series. It is demonstrated that these model results deviate from those expected from the famous Lanchester differential equations, apparently due to the explicit representation of 2D geometry. This supports the idea that geometry plays a critical role in modelling combat outcomes, and is likely responsible for the clustering observed in the time series. However, while it appears that there may be a connection between the clustered nature of the casualty time series and the rate of attrition for each side, it remains unclear what this relationship is, and how to use it in a general combat metamodel that may be capable of completely replacing current aggregated modelling methods. Moreover, the inability to generate realistic synthetic time series points to there being important deficiencies in the way warfare is currently modelled, even for simulations using agent-based modelling techniques where the entities are much more capable of behaviour compatible with a complex system.

Complexity of decision making in warfighting with non-combatant actors: the Blue-Green-Red model

Mathew Zuparic[1], *Maia Angelova*[2], *Ye Zhu*[3] *and Alex Kalloniatis*[4]

[1]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; mathew.zuparic@dst.defence.gov.au
[2]Deakin University, School of Information Technology, Melbourne, Australia;
maia.a@deakin.edu.au
[3]Deakin University, School of Information Technology, Melbourne, Australia;
ye.zhu@deakin.edu.au
[4] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
alex.kalloniatis@dst.defence.gov.au

The networked Kuramoto model [1] has long been employed in the applied sciences as a means to understand complex mechanisms displayed in biological, social and physical systems. In DST, there has been a steadily growing body of work which applies the model as a means to understand how local properties affecting networked decision makers in Command and Control give emergent collective behaviours — both positive and negative. This application is achievable due to the underlying nature of the model's interactions; all things being equal, agents in an organisation governed by the same tempo who interact generally want to achieve a level of co-ordination in their decision making. Such a model can be used as the basis for representing adversarial interactions between two opposing networked forces [2], the 'Blue-vs-Red model', where each of Blue or Red seeks both internal decision synchronisation and for their respective tactical agents to stay ahead of the adversary's decision making cycle, a la John Boyd. This latter effect is encoded as parameter that defines the strategy of one side in relation to the other, how far ahead of the adversary decision making it aspires to. The conflicting goals embedded in the model results in cost/benefit trade-offs: for example, risky strategies, namely on side seeking to be too far ahead of an adversary's decision making, may lead to internal chaotic behaviour, undermining the coherence of internal decisions [3]. In this work, in seeking to extend the Blue-vs-Red model to complex battlespaces of modern warfighting, we include a third Green (or White) network of non-combatant decision makers. Here, members of this third group are not in an explicitly adversarial relationship with either Blue or Red. Nevertheless there is differentiation within Green, with its leaders seeking to synchronise with the leadership in Blue (for example in nation building, reconstruction or mentoring activities) while its tactical agents engage with Red (for example influenced by insurgent networks in the local population). Such differentiation creates further tensions and sources of instability that may be examined, and reflects some of the realities of experienced by the Australian Defence Force in its recent operations. Using computational exploration of the model we find that the addition of Green in this way leads to the appearance of more complex and nuanced model behaviours than that experienced with just two adversarial networks.

Some examples include intricate landscapes between steady-state and dynamic model regimes as we vary key model parameters. Additionally, through a remarkable second order effect, we find that Green may have its own emergent strategy, namely to find itself ahead of Blue and/or Red's decision making, without an explicit strategy parameter. Finally, through some well-placed approximations of the underlying network structures, and assumptions about each network's dynamics, we are able to significantly reduce the total dimensions of the defining equations. This dimensional reduction enables us to understand underlying model behaviours at a fundamental level, and offers a means to perform significant parameter sweeps and expose regions of interest using nothing more than desktop capabilities.

Game of Combat: Mathematical approaches to Modern Warfighting

Alex Kalloniatis[1], *Tansu Alpcan*[2], *Antonin Demazy*[3], *Tristan Perez*[4], *Molloy Timothy*[5] *and Hossein Seif*[6]

[1] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; alex.kalloniatis@dst.defence.gov.au
[2]University of Melbourne, Department of Electrical and Electronic Engineering, Melbourne, Australia; tansu.alpcan@unimelb.edu.au
[3]University of Melbourne, Department of Electrical and Electronic Engineering, Melbourne, Australia; ademazy@student.unimelb.edu.au
[4]Queensland University of Technology, Faculty of Science and Engineering, Brisbane, Australia; Tristan.Perez@qut.edu.au
[5]Queensland University of Technology, Faculty of Science and Engineering, Brisbane, Australia; t.molloy@qut.edu.au
[6];

Combat Models are dynamical representations of warfighting between at least two opposing sides. Game Theory studies the outcomes of decision making under conflict, where two or more decision-makers choose their actions in light of the choices of the other. In modern warfighting, these two concepts are inextricably linked. Surprisingly, insights about the integration of Game Theory with Combat Models are yet under-explored-though it has considerable pedigree in the use of the MiniMax Principle in Morse and Kimball's treatment of generalised Lanchester equations [1]. In this paper we present two initial approaches to generalising such treatments of warfighting. Both approaches incorporate in some form the role of networked decision-making, namely Command and Control (C2), as a means of enriching the traditional massed-force approach of the Lanchester equations [1]. Both approaches can therefore can be thought of as non-cooperative differential games, where the 'strategies' available to the two sides are encoded in dynamical systems represented by differential equations. Both approaches also use a representation of C2 as the dynamics of agents performing individual Orient-Observe-Decide-Act (OODA) loops and interacting within their respective networks with the aim of achieving synchronised decision-making. In the first model, two networked forces are engaged in a tension between seeking to synchronise their internal decision-making while simultaneously seeking to be in advance of the other in decision-making (in the sense of Boyd's "get ahead of the adversary OODA loop"). Combat, as in a representation of attrition or kinetic engagement, is abstracted away in this model. Consistent with non-cooperative Games, we find numerical evidence for a non-trivial Nash equilibrium after allowing the dynamics to reach equilibrium, whereby one side gains advantage over the other — even though both are playing their best strategies.

Indeed, we seek to derive analytical expressions for the dependence of the Nash equilibrium on aspects such as the network structure, individual decision speeds, and information exchange coupling strengths between agents. Surprisingly, the network structure plays a weak role as a consequence of the dependence on the fixed point structure of the dynamical model. In the second model, Lanchester combat is explicitly integrated with C2 such that well/poorly-synchronised decision making enhances/degrades the ability of one side to overpower the opposing side and also resupply itself. We again show that this model may be 'gamefied' and its Nash equilibria numerically determined, in this case for a time-dependent coupling strength as each side intensify or relax their C2 against the ongoing combat dynamics. Initial studies show some challenges for scalability of this approach to larger networks that may be overcome with more sophisticated optimisation techniques. We discuss the prospects for extending these approaches to more sophisticated mathematical models of warfighting. We hope that the initial work in this paper further motivates the development of novel tools to gain insight into the behaviours and C2 of adversarial networks in modern warfighting.

A1. Optimization and AI in discovery, control and design

Beyond cross-validation: Embracing a holistic approach to model validation

Thomas Keevers

Defenec Science and Technology Group, Sydney, Australia;
thomas.keevers@dst.defence.gov.au

Cross-validation is the de facto standard for assessing model quality in machine learning. One data set is used to train or calibrate a complex model, while a second independent data set is used to estimate its generalized performance. The elegance and simplicity of this approach is sometimes perceived to render traditional theoretical approaches unnecessary. This attitude is exemplified in Chris Andersen's provocative essay "The End of Theory: The Data Deluge Makes the Scientific Method Obsolete", 'Petabytes allow us to say: "Correlation is enough." We can stop looking for models. We can analyze the data without hypotheses about what it might show. We can throw the numbers into the biggest computing clusters the world has ever seen and let statistical algorithms find patterns where science cannot.' However, a number of recent studies call into question the sufficiency of cross-validation. Neural networks that achieve human-level accuracy in image recognition are vulnerable against adversarial examples in which images become misclassified after miniscule manipulation, despite strong performance in cross-validation. Google Flu Trends was able to accurately predict influenza outbreaks for several years before the model suddenly mispredicted outbreak timings and intensities. We argue that the data science community should embrace a holistic approach to model validation. This means using cross-validation in addition to assessing data quality and completeness, identifying subjective judgements in model assessment, and using an extended peer review process.

Organisational Constructions with Easy Synchronization

Richard Taylor[1], *Subhra Dey*[2], *Regina Burachik*[3], *Yalcin Kaya*[4] *and Kalloniatis Alex*[5]

[1] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; richard.taylor@dst.defence.gov.au
[2]University of South Australia, School of Information Technology and Mathematica Sciences, Adelaide, Australia; subhra.dey@unisa.edu.au
[3]University of South Australia, School of Information Technology and Mathematica Sciences, Adelaide, Australia; regina.burachik@unisa.edu.au
[4]University of South Australia, School of Information Technology and Mathematica Sciences, Adelaide, Australia; yalcin.kaya@unisa.edu.au
[5]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; alex.kalloniatis@dst.defence.gov.au

The Kuramoto model (Kuramoto, 1975) was originally motivated by the phenomenon of collective synchronization whereby a system of coupled oscillating nodes will sometimes lock on to a common frequency despite differences in the natural frequencies of the individual nodes. This model has since been applied to a wide variety of application fields including, biological, chemical, engineering, and social systems. In the Kuramoto model each node has an associated phase angle as well as its own natural frequency. The basic governing equation is the differential equation:

$$\frac{d\theta}{dt} = \omega_i + k\sum_{j=1}^{n} A_{ij} sin(\theta_j - \theta_i), \quad i = 1, ..., n \ . \tag{1}$$

where A is the adjacency matrix of the network and k is a coupling constant which determines the strength of the coupling. Define the critical coupling kc to be the least value of k for which a frequency fixed point occurs whereby each $\frac{d\theta}{dt} = 0$. In DST over many years we have developed the case of the applicability of this model to Command and Control (C2), where synchronising decision-makers are now the subject of the representation. An ideal network is then a C2 structure that enables collective decision-making at the lowest possible coupling while distributing network load equitably. This corresponds to efficient manageability when the network is considered as an organizational structure. The links in this network may correspond to a management structure where an individual has a link to their manager and the people they manage, or a peer relationship outside of the management chain. In this context synchronization corresponds to decision making across a distributed network of individuals based on iterated observation-processing-action loops (Kalloniatis2014).

In this paper we describe two approaches to derive an ideal network structure. The first begins with a hierarchy which may be seen as an efficient structure for coordinating many individuals with relatively few formal relationships while balancing the number of relationships required for each individual. However, as a model for decision making and coordination this structure requires a relatively large critical coupling for the number of edges. We shall show that by carefully adding just n/4 edges to the bottom layer of a binary tree the resulting network forms a family of well-studied networks with excellent connectivity properties called expander graphs which also have very low critical coupling. The second approach adopts optimisation where we develop objective functions that maximise sparseness of the network, namely most economical, while also building in maximisation of synchronisation. It turns out that the problem can be addressed using the Deflected Subgradient Method. We compare and contrast resulting networks from the two approaches, including structures with up to N=127, namely the size of a seven level complete binary tree. We discuss future applications in combat modelling where multiple warfighting functions, including C2, are brought together.

The Graphene Machine - A Virtual Platform for Computational Materials Design

Benyamin Motevalli[1], *Baichuan Sun*[2] *and Amanda Barnard*[3]

CSIRO Data61, Docklands, Australia
[1]b.motevalli@gmail.com
[2]baichuan.sun@data61.csiro.au
[3]amanda.barnard@data61.csiro.au

Industry 4.0 is the "intelligent revolution" of the manufacturing sector that uses smart systems, data and multivariate analytics to increase productivity. Part of this endeavour is to close the loop, drawing information from the factory floor to inform changes in the material processing parameters that determine the quality and quantity of products. Targeting specific markets requires a certain degree of control, as even low-tech applications that have considerable fault-tolerance to structural and property variability still require manufacturers to deliver a product with batch-to-batch consistency. Translating variations in a manufacturing process into modifications that can control product design requires a reference that understands the relationships between process, structure, properties and performance. The Graphene Machine is a digital science platform that fuses complex experimental and computational data sources and applies machine learning to create the first industry-specific materials informatics platform, aimed at supporting Australia's fledgling graphene-nanotechnology sector. Graphene is a two-dimensional sheet of carbon atoms discovered in 2004 by Nobel Prize winners Andre Geim and Konstantin Novoselov, which is the lightest, strongest, most electrically conductive substance on earth; exhibiting a range of unique and exceptional properties that will revolutionise everyday products. Since its discovery more than 12,000 graphene related patents have been filed, with research continuing to grow rapidly across the public and private sectors. Proposed graphene applications include real-time sensing of changes in stress, temperature, moisture and the concentration of gases (including toxins). However, after well over a decade of intense research and development graphene products are still almost non-existent, because the level of control exercised in laboratories cannot be delivered at scale. The structure and properties of graphene are heavily influenced by even small changes in temperature, pressure, synthesis, processing and storage conditions, and the quality of the raw material (often graphite from a range of different mines). The first stage of development of the Graphene Machine draws on over 50,000 virtual experiments, expanded using ab initio thermodynamics to over 500 mil, to ensure we have a comprehensive description of the material. Contemporary modelling and simulation algorithms, coupled with available supercomputing resources, are cost effective and can capture the diversity, complexity and scope of the corresponding conventional experiment. This database will be used to training machine learning models to correlate processing parameters and structural features with functional properties.

Future stages include fusing experimental data from conventional characterization methods, and the construction of a user-friendly interface for exploring hypothetical combinations of experimental settings and their impact on performance. Using this tool graphene production can be made robust against changes in raw materials, water quality (used during the reduction of graphene oxide), temperature or humidity on any given day, allowing companies to avoid batch-to-batch inconsistency and specialize in different qualities of product for applications where perfection is essential or where dispersivity can be tolerated. This capability would enable companies to refine their business model and become more strategic, and compete internationally.

AI in materials discovery and design

Tu Le

RMIT Univeristy, School of Engineering, Melbourne, Australia; tu.le@rmit.edu.au

Machine learning methods are efficient tools to develop novel materials for different applications. They are fast, resource efficient, versatile, and low cost. They produce reliable, repeatable results, uncover hidden insights through learning from trends in the data, and are essential synergistic elements of any experimental project, especially those that generate high volumes of multi-dimension data using high throughput experiments. This presentation will provide examples of how these efficient approaches, coupled with a strong understanding of molecular properties and interactions, can generate useful and robust models for diverse materials properties, from 'smart', self-assembling amphiphiles for targeted drug delivery, to advanced functional polymers. This approach also allows the determination of factors that significantly affect the materials properties and therefore allow the design of new materials or optimization of existing materials.

Artificial intelligence for directing experiments or simulations

Amanda Parker[1] *and Amanda Barnard*[2]

CSIRO Data61, Docklands, Australia
[1]amanda.parker@data61.phas.ubc.ca
[2]amanda.barnard@data61.csiro.au

Often in science there is a broad space of parameters which could be varied and the most difficult job for a scientist is to decide which experiments to run. Without intention, the scientific outcomes can be limited or biased when specialist knowledge is used to make these decisions. We are working to develop an implementation of artificial intelligence (AI) computational methods that directs which simulations (or experiments) are conducted in order to maximize the scientific outputs while minimizing their cost. This talk broaches some of the challenges to constructing such an AI implementation in a materials science context. However, this problem is one which crosses many domains. We take advantage of active learning wherein each new experiment helps refine the model and the model then directs which experiment is the next conducted. While simple in concept there many choices in an active learning implementation which will determine its efficacy. In this talk we discuss a map to appropriate choices for active learning, feature selection, and clustering methods. A key challenge is making use of available data effectively. It is expected that specialist domain knowledge is used to construct the features which uniquely describe each experiment. However, to select without bias which features are most appropriate for a predictive model is a cross domain data science challenge. Too many features can lead to overfitting, problems with covariance or the curse of dimensionality, while too few can fail to provide the desired predictive capability. Moreover, feature selection only becomes more challenging if the data set is expanding as more experiments are run. Clustering can help with the challenge of feature selection and also reveal hidden structure within multi-dimensional data. However, a cluster can be a difficult notion to define mathematically and to capture algorithmically. When the data has more than 3-dimensions we have no default way to verify results. Common clustering algorithms such as K-means can perform very well but define an ideal cluster in a specific way that can be limiting (ie. spherical, equally sized clusters). Often the means of validating the clustering result (eg. explained variance) can be convoluted with the clustering method (eg. kMeans) inhibiting comparisons across methods. We introduce iterative label spreading (ILS), a means to gain a priori insight into expected clustering outcomes, assess the results of any given clustering method and perform clustering.

A4. Food and beverage supply chains

A dynamic lot-sizing model under Vendor Managed Inventory (VMI)

Farshid Evazabadian[1], *Regina Berretta*[2] *and Mojtaba Heydar*[3]

[1]University of NEwcastle, Callaghan, Australia; farshid.evazabadian@uon.edu.au
[2]University of NEwcastle, Callaghan, Australia; Regina.Berretta@newcastle.edu.au
[3]University of Tasmania, Hobart, Australia; Mojtaba.heydar@newcastle.edu.au

This paper explores the impact of Vendor Managed Inventory (VMI) on a decentralised supply chain by proposing two Mixed Integer Linear Programming (MILP) models for a dynamic lot-sizing problem. The models describe the decision making for lot sizing before and after the implementation of VMI. The proposed models highlight that VMI takes advantage of centralised decision making, and can reduce the cost of the lot sizing by better synchronisation of the decisions. Numerical results, provided to compare the efficiency of VMI with the traditional decentralised lot sizing, indicate significant cost reduction under VMI. A set of experiments is also designed to determine the impact on retailers' cost (ordering and holding) as well as vendors cost (setup and holding) on the gap.

A Two-stage stochastic model for Selection of Processing Hubs to Avoid Broccoli Losses

Rodolfo Garcia-Flores[1], *Elaine Luc*[2], *Peerasak Sanguansri*[3] *and Pablo Juliano*[4]

[1]CSIRO Data61, Docklands, Australia; rodolfo.garcia-flores@csiro.au
[2]CSIRO Data61, Docklands, Australia; elaine.luc96@gmail.com
[3]CSIRO Agriculture and Food, Werribee, Australia; Peerasak.Sanguansri@csiro.au
[4]CSIRO Agriculture and Food, Werribee, Australia; Pablo.Juliano@csiro.au

It is estimated that, at present, around one third of all food produced is lost, either in production and distribution, or after retail. To further complicate matters, uncertainty and variability in the commercial and natural environments must also be taken into account when trying to reduce food losses. The objective of the present paper is to develop a decision support system to increase the efficiency of the Australian broccoli supply chain and reduce food losses considering uncertainty. To that end, we develop a two-stage stochastic mixed-integer linear programming model to assist Australian broccoli producers in taking the most cost-effective investment decisions and, at the same time, reduce the losses by producing novel, high value-added products from produce discarded on the field or during transportation. The stochastic model we propose selects, in the first stage, the optimal location of processing facilities to add value to the produce that would otherwise be considered as food loss, and suggests transportation operations as the recourse decisions. The model is solved using Lagrangian decomposition and the subgradient method. The data used to feed the model was collected on the field through a survey applied to broccoli growers nationwide. The model suggests near-optimal investment decisions that are far from the worst possible outcome, had the final market and environmental conditions turned out to be very adverse. Our results represent viable operations for the industry in the medium term.

D4 and D10. ML and Defence Decision Analysis

A Genetic Programming Framework for Novel Behaviour Discovery in Air Combat Scenarios

Martin Masek[1], *Chiou Peng*[2], *Luke Kelly*[3], *Lyndon Benke*[4] *and Papasimeon Michael*[5]

[1]Edith Cowan University, Perth, Australia; m.masek@ecu.edu.au
[2];
[3]Edith Cowan University, Perth, Australia; l.kelly@ecu.edu.au
[4]Defence Science Technology Group, Fishermans Bend, Australia;
lyndon.benke@dst.defence.gov.au
[5]Defence Science Technology Group, Fishermans Bend, Australia;
michael.papasimeon@dst.defence.gov.au

Behaviour trees offer a means to systematically decompose a behaviour into a set of steps within a tree structure. Genetic programming, which has at its core the evolution of tree-like structures, thus presents an ideal tool to identify novel behaviour patterns that emerge when the algorithm is guided by a set fitness function. In this paper we present our framework for novel behaviour discovery using evolved behaviour trees, with some examples from the beyond-visual range air combat domain where distinct strategies emerge in response to modelling the effects of electronic warfare.

Expanded basis sets for the manipulation of random forests

Thomas Keevers

Defenec Science and Technology Group, Sydney, Australia;
thomas.keevers@dst.defence.gov.au

While random forests is considered one of the best off-the-shelf data mining algorithms, it suffers from poor interpretability and an opaque decision structure. In this paper, we develop a method for generating an "expanded basis set" for a random forests model that captures every possible decision rule and vastly improves the transparency of random forests classifiers. The expanded basis set allows the structure of a random forests model to be algebraically manipulated and facilitates a number of operations, including inverse mapping from outputs to the domain of inputs, systematic identification of every decision boundary, and comparison of two random forests models. The expanded basis set also facilitates visualization of the global behaviour of a random forests classifier and a data set by combining parallel coordinates with a non-linear binning transformation. The global visualization allows classifier performance to be compared against domain expertise, and areas of underfitting and overfitting to be readily identified. Additionally, the expanded basis set underpins the generation of counterfactuals and anchors — combinations of variables that control the local outputs of a random forests model. The basis states can also be used to place bounds on the model stability in response to single or multi-feature perturbations. These stability bounds are especially useful when the model inputs may be uncertain or subject to variation over time.

Machine Learning for Military Wargaming

Glennn Moy[1], *Slava Shekh*[2] *and Shuang Yu*[3]

[1]DST Group, Brisbane, Australia; glennn.moy@dst.defence.gov.au
[2]DST Group, Edinburgh, Australia; Slava.Shekh@dst.defence.gov.au
[3]Monash University, Melbourne, Australia; shuang.yu@dst.defence.gov.au

Military wargames have been used for centuries as part of the military planning process. In the Australian military, the "Joint Military Appreciation Process" (JMAP) has wargaming as a central component of "Course of Action Analysis". Modern wargames have their origin in ancient strategy games, such as Chess. The recent success of machine learning techniques, applied to games like Chess, Go and Shogi, raises the question as to whether these techniques could have an equally successful role in analysing strategy for military wargames. Currently, military wargames are typically run with little or no automation. Analysis of wargaming results is commonly undertaken by human experts, based on their experience and (sometimes) limited support from simple analytic tools. In this work, we explore the use of machine learning to better analyse and develop strategy in military wargaming. The current state-of-the-art for achieving super-human performance on complex strategy problems are the "Alpha" family of techniques — AlphaGo [1], AlphaGo Zero [2] and AlphaZero [3]. These have demonstrated super-human performance on a range of games, including Chess, Go and Shogi, but are unproven for more realistic problems, such as military wargames. In particular, successful civilian games have properties that may not be present in realistic military scenarios. For instance, successful games are typically symmetric — players often have identical pieces and equal opportunity for success. By contrast, real wargames are often asymmetric — with each side having a different array of assets and capabilities. In addition, civilian games often have interesting strategies at all levels of competency — so that novice, intermediate and advanced players can compete and progressively advance their skills. Games that are impossible for novice players to grasp, or too simple for advanced players to enjoy are unlikely to become successful. This property may prove to be important in the success of techniques relying on self-play for incremental improvement. We discuss a number of other differences between realistic military wargames and civilian games that may be relevant in order to achieve super-human performance. These include increased game complexity, more complex board representations (for instance, if different pieces can occupy the same board cell) and the challenges of players having different goals. We present our implementation of an AlphaZero-like technique, along with heuristic and MCTS baselines, on an exemplar wargame played on a hex grid called "MicroWargame". MicroWargame is based on "Coral Sea 2042", a scenario developed in JOAD, DST Group and used as a teaching aid for wargaming. Finally, we present our findings on the use of these techniques in MicroWargame, with an emphasis on the differences between games like Go, Chess or Shogi and more realistic wargames.

We find that, in some cases, wargames may only have a small number of meaningful strategies available, making simpler techniques such as heuristics or hyper-heuristics more suitable and efficient. In other cases, we note that the inherent asymmetry in real wargames may make it difficult for a self-play technique to explore the solution space and present possible modifications to standard AlphaZero techniques to address these issues.

Analysis of Helicopter Risk Minimisation Against Ground Based Rifle Shots During Hovering

Kutluyil Dogancay[1], *Ziming Tu*[2], *Gokhan Ibal*[3], *Mina Shokr*[4] *and Stewart Chris*[5]

[1]University of South Australia, School of Information Technology and Mathematica Sciences, Adelaide, Australia; kutluyil.dogancay@unisa.edu.au
[2]Defence Science and Technology Group, Fishermans Bend, Australia;
Ziming.Tu@dst.defence.gov.au
[3]Defence Science and Technology Group, Fishermans Bend, Australia;
Gokhan.Ibal@dst.defence.gov.au
[4]Defence Science and Technology Group, Fishermans Bend, Australia;
mina.shokr@dst.defence.gov.au
[5]Defence Science and Technology Group, Fishermans Bend, Australia;
Christopher.Stewart@dst.defence.gov.au

Defence Science and Technology (DST) Group has been conducting experiments to model and simulate aiming errors for small-arms fire against a manoeuvring helicopter from spatially distributed positions. The difficultly with conducting a large number of experiments to collect consistent data for all firing positions has motivated the need for new methods for generalising shooter aiming errors obtained from a relatively small number of firing positions to arbitrary firing positions. Such a generalisation offers several benefits including risk assessment posed by a small-arms operator, and helicopter tactics analysis for risk minimisation. The particular scenario of hovering helicopter is of special interest as the aiming errors reveal a peculiar behaviour markedly different to general helicopter manoeuvres. This paper starts with a brief outline of the statistical approach to modelling a small-arms operator's firing performance utilising closest point of approach (CPA) for fired bullets. The CPA data collected in experiments are firstly pre-processed for data pruning. The pruned data is used to estimate first and second-order statistical parameters for the CPA in the 3D space for a limited number of shooter positions, under various helicopter manoeuvres and conditions such as day-time/night-time, with/without tracer. The CPA data are modelled as multivariate Gaussian with aiming errors characterised by 3x1 mean vector μ and 3x3 covariance matrix Σ. Once the Gaussian statistics μ and Σ are estimated for each shooter position using the experimental data, generalisation to any shooter position is accomplished by function generalisation using an artificial neural network (ANN). To ensure that function generalisation produces feasible statistics with symmetric and positive-semi-definite covariance matrices, the task of generalisation is formulated in terms of μ and Σ. The error ellipsoids are described by their semi-axis lengths and angular rotation of semi-axes (Euler angles) about the principal, x, y and z-axes. Using these parameters rather than covariance matrix entries ensures that the covariance matrix generated by the ANN agent will obey the properties of a covariance matrix.

Therefore, for function generalisation purposes, the aiming errors are described by μ, semi-axis lengths of error ellipsoid, and Euler angles for ellipsoid rotation. The ANN agent for the shooter is computed using a single generalised regression neural network (GRNN) with an input consisting of shooter position relative to the helicopter at $[0, 0, 0]$ heading North and helicopter speed, and a vector output comprising the nine statistical parameters that define the CPA error ellipsoid. A GRNN is comprised of a radial basis layer followed by a linear layer. An important consideration for helicopter vulnerability is the probability of hit, while the helicopter hovers, by a small-arms operator on the ground. Trained ANN-based shooter agents have been simulated to ascertain how the probability of hit varies for different geometries. The probability of hit information obtained from the ANN shooter agent can be exploited to construct a probability of hit map to determine optimal hovering manoeuvres for a helicopter.

Stochastic multi-criteria decision analysis of combat simulation data for selecting the best land combat vehicle option

Thang Cao[1] *and Dion Grieger*[2]

Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia
[1]thang.cao@dst.defence.gov.au
[2]dion.grieger@dst.defence.gov.au

Land Combat Vehicles (LCVs) are being acquired to improve the fighting capability of the Australian Defence Force (ADF). The mission effectiveness of a LCV is often modelled via combat simulation in which the multi-criteria metrics are measured from the simulation output. Consequently, it is desirable to develop a Multi-Criteria Decision Analysis (MCDA) methodology to support upcoming decisions for operational capability of future vehicle options. Uncertainties and imprecision are common both in criteria measurements in combat simulation and Decision Makers (DMs) preference; however, option ranking and selection procedure from simulation are normally limited to a single response metric or deterministic preference for the multiple metrics in the current literature. Therefore, the stochastic multi-criteria acceptability analysis (SMAA) model is explored in this paper for aiding this decision making problem. These uncertainties are treated better using a probability distribution function and Monte Carlo simulation in the model. Moreover, the idea of "feasible weight space (FWS)", which represents the union of all preference information from DMs, is applied in combination with other weighting techniques such as Analytical Hierarchy Process (AHP). The purpose of this paper is to describe the application of SMAA, FWS and AHP to the results generated in a close-loop combat simulation, such that the options with uncertain data can be analysed, ranked and the best option selected for a specific task or scenario. Up to our knowledge, this combined approach has been applied for the first time to deal with the defence decision analysis problems with uncertainty and interdependency.

A7. Transport and logistics

A Branch-and-Price Framework for the Maximum Covering and Patrol Routing Problem

Paul Chircop[1], *Timothy Surendonk*[2], *Menkes van*[3] *and Toby Walsh*[4]

[1]Defence Science and Technology Group, Sydney, Australia;
paul.chircop@dst.defence.gov.au
[2]Defence Science and Technology Group, Sydney, Australia;
timothy.surendonk@dst.defence.gov.au
[3];
[4]Technical University of Berlin, Berlin, Germany; toby.walsh@data61.csiro.au

The Maximum Covering and Patrol Routing Problem (MCPRP) is concerned with the allocation of police patrol cars to accident hotspots on a highway network. A hotspot is represented as a time window at a precise location on the network at which motor vehicle accidents have a high probability of occurring. The nature of these accidents may be due to speeding, driver fatigue, or blind-spots at intersections. The presence of police units at hotspots serves as an accident prevention strategy. In many practical applications, the number of available cars cannot cover all of the hotspots on the network. Hence, given a fleet of available units, an optimization problem can be designed which seeks to maximize the amount hotspot coverage. The cars must be routed in such a way as to avoid multiple contributions of the patrol effort to the same hotspot. Each police car is active over a predefined shift, beginning and ending the shift at a fleet station. In this paper, we introduce a method for constructing a time-space network of the MCPRP which is suitable for the application of a branch-and-price solution approach. We propose some large scale test problems and compare our approach to a state-of-the-art Minimum Cost Network Flow Problem (MCNFP) model. We show that our branch-and-price approach can outperform the MCNFP model on selected large scale networks for small to medium fleet sizes. We also identify problems which are too large for the MCNFP model to solve, but which can be easily handled by our approach.

Linear Complexity Algorithms for Visually Appealing Routesin the VRP

Philip Kilby[1] *and Dan C.*[2]

[1]Australian National University, Canberra, Australia; philip.kilby@csiro.au
[2];

In the Vehicle Routing Problem (VRP) we develop efficient routes to serve customers using a fleet of trucks. "Visually appealing" routes are often preferred as they allow more flexibility for reordering visits in reaction to unexpected events. Compactness of the routes is a key desirable feature, which can be achieved by using a penalty on the area enclosing the customers on each route. We present several efficient, alternative methods to directly minimising the convex hull. Our methods are based on the computation of features similar to the convex hull area, but have low computational complexity, O (n) instead of O(n log n), and are significantly easier to implement. We also show that these penalties achieve good guidance: insertion-based route construction methods can be guided towards more visually attractive route shapes. This is achieved at the cost of only minimal loss of optimality with respect to a primary objective function like the routes' total length. Embedded in an Adaptive Large Neighbourhood Search procedure, the penalties were able to achieve similar performance in terms of route compactness to a full convex hull calculation. We also achieve good separation between routes.

Prioritising autonomous supply - comparing selection by Marginal analysis and Neural nets

Gregory Sherman[1] *and Mitchell Mickan*[2]

[1]DST Group, Edinburgh, Australia; gregory.sherman@dst.defence.gov.au
[2]University of Adelaide, Adelaide, Australia; mitchell.mickan@gmail.com

When managing inventory or supply systems it is important to make good choices about which stock to prioritise over others. We can improve the overall availability of the supplied systems by making optimal choices on which inventory items should be allocated to meet demands. In this paper we will show how machine learning algorithms can be used to prioritise inventory. The developed algorithms were tested on a real data set and the improvement in inventory allocation measured. Machine learning is a powerful technique for transforming inputs to outputs in order to best achieve a set goal. It has many applications in areas where there is an abundance of data, and where the resulting decisions can be measured. As such inventory management is a suitable area of application, in particular the prioritisation of supply. Such an approach is even more relevant to those inventory models that represent autonomous processes. The models we are interested in are those relating to system availability and that use item backorder calculations. These models rely on a traditional prioritisation approach known as marginal analysis, otherwise known as a process of marginal allocation using a greedy algorithm. Because marginal analysis does not take into account performance over time, nor complex relationships in datasets, there may be potential for a machine learning algorithm to provide better results if it can learn to exploit both temporal and relationship data. The benefit of such an improvement is the value of availability generated and cost savings made in the supply network.

Sustainable Logistics Planner: Electric Vehicle Routing using Adaptive Large Neighbourhood Search

Paresh Chiramel[1], *Debdeep Banerjee*[2], *Isabella Henggie*[3], *Simranjit Kaur*[4] *and Johanna Eva*[5]

[1]Quintiq, Melbourne, Australia; Paresh.CHIRAMEL@3ds.com
[2]Quintiq, Melbourne, Australia; Debdeep.BANERJEE@3ds.com
[3]Quintiq, Melbourne, Australia; Isabella.HENGGIE@3ds.com
[4]Quintiq, Melbourne, Australia; Simranjit.KAUR@3ds.com
[5];

The vehicle routing problem (VRP) is a well known NP hard problem that appears in our daily life. Especially in this age of online shopping and delivery, it is one of the most relevant puzzles in the world of algorithms. In a typical case the objective of VRP is to deliver/pick up a set of customer orders (within given time windows) while minimizing cost of travel and lateness of service. Logistics companies are now looking for transportation solutions to cut emissions and contribute towards the Paris agreement (https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement). One solution, among others, is to use electric vehicles to replace conventional transportation options. Although electric vehicles are emission free, their range is limited compared to conventional vehicles. This range restriction adds another layer of complexity to the VRP problem, namely the recharging of the batteries. Recharging of batteries takes time and adds a significant amount of complexity to the VRP problem. In this presentation we describe Zero Emission Fleet planner, a commercial solution that solves the problem described above. This solution is based on Quintiq's Logistic Planner solution, which allows users to specify different types of electric vehicles with different on-board charger capacity, battery range as well as the existing charging network. It provides functionality to users to manually plan and verify routes. Zero Emission Fleet planner has an Adaptive Large Neighbourhood Search algorithm that automatically generates good solutions. The main decision points are which order (pick-up/delivery) would be serviced by which vehicle, when and where a vehicle will be recharged, and for how long. The large neighbourhood search starts with an initial solution. Neighbourhoods are created by selecting an anchor order, then un-planning the surrounding orders, replanning them back onto vehicles in proper sequence, and then looking for additional recharge actions needed for each vehicle. Once this process provides a better solution than the current solution, the new solution is accepted and updated. This iterative process terminates once it reaches the maximum number of iterations allowed or the maximum runtime. This algorithm can also be run in a continuous mode, planning new orders as they come in. Quintiq's VRP(TW) solver is a very efficient implementation of an adaptive large neighbourhood search algorithm, that holds several world records for standard benchmarks (https://www.quintiq.com/optimization/vrptw-world-records.html). The solution described in this presentation extends the existing solver with support for electric vehicles fleet. Currently we are working on creating a benchmark to track our progress.

Analytical BusPlus

Arthur Maheo[1] *and Micheal Forbes*[2]

[1]Australian National University, Canberra, Australia; arthur.maheo@anu.edu.au
[2]The University of Queensland, Brisbane, Australia; m.forbes@uq.edu.au

We present a framework for Benders Decomposition which incorporates a number of recent advances, and use it as a novel technique to solve a public network design problem arising in multi-modal public transport. Our main contribution is to derive Pareto-optimal Benders cuts using an analytical procedure inside a Branch-and-Cut scheme. Canberra is a planned city designed by American architect Walter Griffin in 1913. It features a large number of semi-autonomous towns separated by greenbelts. As a result, Canberra covers a wide geographic area, which makes public transportation particularly challenging. Bus routes are long, which results in low frequency, and low patronage, especially during off-peak periods. To address these limitations, the BusPlus project designed, optimised, and simulated a Hub-and-Shuttle Public Transit System (HSPTS). A Hub-and-Shuttle model consists of a combination of a few high-frequency buses and a large number of shuttles (or multi-hire taxis). Buses run on designated corridors between key hubs in the city, called "bus legs." Taxis supplement buses by bringing passengers from their origin to the closest hub and take them from their last bus stop to their destination. Designing such an HSPTS creates a series of interesting challenges, including:

1. How to connect a set of potential hubs using bus legs in order to minimise operational costs and maximise trip convenience?

2. How to satisfy a request using a journey that may comprise a shuttle ride, a number of bus legs, and a final shuttle ride?

This work focuses on the first problem and primarily on off-peak hours — evenings, week-ends — as they are the most challenging from a cost and service standpoint. We model this problem as a MIP and use a Benders decomposition approach to solve it. The Benders decomposition operates on MIPs; it proceeds by a sequence of projection, relaxation, and outer approximation steps. The HSPTS can be seen as a two-level decision problem: first, determine on which legs buses will run; second, route the demand on the new network. This leads to a natural decomposition: the bus legs selection will be the master problem, and the detailed routing will be the sub-problem. A class of Benders cuts, called Pareto-optimal cuts, are known to out-perform standard cuts, especially in network problems. However, to derive such cuts we need to solve two linear problems: the regular Benders sub-problem and a modified Pareto sub-problem. We demonstrate how to derive Pareto-optimal cuts without using a solver. First, we solve the sub-problem using a dedicated algorithm, in our case a shortest path algorithm. Then, we demonstrate how to derive dual costs from this solution.

Finally, we prove that the cut generated from these costs are Pareto- optimal. We provide results of three different setups: modeling and solving the HSPTS as a single MIP; solving the HSPTS using a tailored Benders decomposition; and, solving the HSPTS using our analytical Benders framework.

D5. Force Design Analysis

Evolutionary Algorithms for Force Structure Options

Connor Hicks[1], *Elizabeth Kohn*[2] *and Thitima Pitinanondha*[3]

[1]Department of Defence, Canberra, Australia; connor.hicks@defence.gov.au
[2] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; Elizabeth.Kohn@dst.defence.gov.au
[3] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; thitima.pitinanondha@dst.defence.gov.au

A modern Defence Force consists of a diverse range of capabilities to support missions at the tactical, operational and strategic levels. Designing a balanced and affordable force structure to meet Government strategic objectives and assure national security has always been a challenge. Force design is a centralised and enduring process within the Australian Defence Organisation that seeks to translate Government strategic objectives into a coherent force structure within specified time and budget envelopes. This process increasingly relies on analytical approaches and tools such as wargaming, simulation and optimisation techniques. This paper investigates evolutionary algorithms (EAs) as a potential tool for generating and evaluating force structure options. EAs can evaluate an extremely large solution space of force mixes at a much faster rate than human cognition to determine a balanced and affordable force structure option according to an objective function. This paper also discusses the implementation of a software framework, dubbed "FORCESIGHT", which can be customised by developers to model any scenario where the use of EAs is appropriate. Based on the outcomes of a trial of FORCESIGHT, it is clear that the EAs approach could provide a result of respectable quality. It is demonstrated that EAs can lead to large increases in efficient evaluation of potential improvements to the Force-in-Being and Future Force.

Defining Project Target Benefits to Enable Portfolio Design in Defence

Ivan Garanovich[1], *Terence Weir*[2] *and Ofer Zwikael*[3]

[1]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; ivan.garanovich@dst.defence.gov.au
[2]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; terence.weir@dst.defence.gov.au
[3]Australian National University, College of Business and Economics, Canberra, Australia; ofer.zwikael@anu.edu.au

The Department of Defence requires analytical methods and tools that can support force design, and can provide the required evidence base for major Defence capability investment decisions. While operations research can provide a range of sophisticated techniques to enable analysis and visualisation of the investment portfolio, most analytical methods rely on the availability of metrics for Defence capability to enable mathematical modelling and optimisation. It is observed that the force design process translates Government and Defence policy guidance into a vision of the future force structure that can be articulated to government. The output of the force design process is expressed in the Integrated Investment Program (IIP) which represents Defence's acquisition and capital investment plan in major equipment, information and communications technology, facilities and workforce that forms the basis for the future force structure. It is suggested that Defence can borrow the latest best practices from the project management discipline to define measures and metrics for Defence capability investment projects to enable rigorous portfolio design. This paper outlines key principles for the definition of target investment benefits in Defence context which would link Defence capability investment to the pre-defined strategic goals. Key attributes of target benefits are defined for Defence investment projects which are aligned with the current Defence Capability Assessment Program and Gaps and Opportunities process. By better formulating target benefits for capability investments in the IIP, Defence will be able to enhance the quality of force design decisions and improve resource and budget allocation. Benefit and value measures for Defence investment projects will enable analysis for capability maximisation, portfolio cost minimisation and trade-off analysis between cost and capability.

Towards Whole-of-Force Data Farming Analysis for Force Design

Peter Hoek[1], *Andrew Au*[2] *and Edward Lo*[3]

[1]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; peter.hoek@dst.defence.gov.au
[2] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; andrew.au@dst.defence.gov.au
[3] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; edward.lo@dst.defence.gov.au

Increasing complexity of modern military operations necessitates the development of techniques to supplement traditional force design approaches that rely solely on intuition. Joint Future OpeRating Concept Explorer (JFOrCE) is an agent-based simulation capability being developed in NetLogo to support options analysis for the Australian Defence Force (ADF) force structure review. This paper describes the design concept using a top-down approach of realising future operating concepts in simulation to support force design analysis. The implementation of JFOrCE and its application to a data farming approach is presented. The proposed approach shows potential for JFOrCE to become a valuable decision-support tool, identifying not only the most likely outcome, but the possible range of outputs including the location of outliers that may have significant implications.

A pragmatic approach to modelling and managing a defence force system

Donald Lowe[1] *and Mark Unewisse*[2]

[1]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Canberra, Australia; donald.lowe@dst.defence.gov.au
[2]Department of Defence, Defence Science and Technology Group, Joint and Operations Analysis Division, Edinburgh, Australia; mark.unewisse@dst.defence.gov.au

One of the key challenges for Australia is the design, delivery and management of its defence force which is both sizeable and complex. Despite a relatively small population estimated as the 53 rd largest in the world, Australia has a Defence budget that is about the 13 th largest in absolute terms. It is required to be constantly available for operations and consists of many systems and systems-of-systems. Moreover, in order to meet evolving and future operational needs, the force is continually changing via the management of multiple capability development projects that deliver new capabilities over time. This is a complex network of linked decisions performed by many different decision-makers at different levels of authority in an asynchronous manner through time. Ultimately, this activity must maintain alignment with, and aim to fulfil, the Australian government's strategic guidance. A systems thinking method for modelling the complex capability dependencies has been devised by DST researchers called the SCMILE Services-Based Approach (where SCMILE is an acronym derived from Sensing, Command and Control, Physical Mobility, Information Mobility, LogisticsandEnablers- Engagement) [3, 4]. Its simplicity and universality make it suitable for the high-level capture and tracking of the system design of a military force in its current state, as it evolves and in potential future states. This approach is already starting to be exploited within the Australian Department of Defence via the current development of an enterprise software tool, called the Force Integration Management System (FIMS). A different framework called the Strategic Lines of Capability (SLC) was introduced in 2016 which was designed to represent the linkages between strategic guidance, the contributions of the different capabilities to undertake operations to achieve strategic tasking, through to the projects that deliver these capabilities through time [6]. In this way, the framework is comprised of a series of layers that that provides a systematic and gradated progression from strategic effects to capability investments. In this paper, the SLC framework is further elaborated, listing the basic steps of its construction, showing its relationship to the Defence Integrated Investment Program (IIP) and its associated Portfolio-Program-Project (P3) Management philosophy and explaining how it conforms with systems design principles already enunciated in such domains as systems engineering, software design and Capability-Based Planning. This systems thinking framework is then combined with the aforementioned SCMILE approach to produce a larger, coherent, consistent and more complete methodology that not only transparently traces from strategic tasks down to the multiple capabilities that contribute to them, but also provides a systematic method of modelling the systems-of-systems that comprise these capabilities.

It is shown how together, they have the potential to provide a pragmatic and systematic method of coherently designing and managing the Australian Defence Force in a continual fashion.

On the Set-Union Budget Scenario Problem with Alternative Initiatives

Richard Taylor[1] *and Jerzy Jagiello*[2]

Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia
[1]richard.taylor@dst.defence.gov.au
[2]Jerzy.Jagiello@dst.defence.gov.au

At MODSIM 2015 and MODSIM 2017, (Taylor 2015) and (Jagiello and Taylor 2017) formalized the Budget Scenario Problem presented by (Order, 2007, 2009). In the Budget Scenario Problem a list of initiatives is provided, each with an anticipated cost. Each initiative is scored against a number of scenarios with a value indicating how useful the initiative is against that scenario. For a collection of initiatives the total score is calculated as a sum of best initiative scores within the collection for each scenario. This paper extends this idea to take into account the sum of the best l initiative scores for some fixed l. This is designed to give a more robust indication of the value of a collection of initiatives. Since this extension is analogous to the Set-Union Knapsack Problem (Goldsmith, 1994) we call this formulation the Set-Union Budget Scenario Problem with Alternative Initiatives. In mathematical terms the Set-Union Budget Scenario Problem with Alternative Initiatives can be expressed as:

$$\max \sum_{|\psi|=l,s=1}^{r} p_s \sum_{j\in\psi} (\delta_{js}(\alpha_j + \sum_{k=1}^{n} \beta_{jk} x_{jk}) x_j) \tag{2}$$

Subject to

$$\sum_{j=1} c_j x_j \leq B \ , \tag{3}$$

where: r is the number of scenarios, s is the index of scenario, $x_j \in 0, 1$ for initiative j, c_j the cost of initiative j, p_s the probability of scenarios, ψ is the subset of l initiatives, j is the index of initiative, l is a fixed number of initiatives, δ_{js} is the value of score for initiative j and scenario s, α_j is the measure of independence for initiative j, β_{jk} is the value of interdependency between initiative j and k, B is the cost bound. In simple terms this problem can be thought of as determining the best collection of tools (toolbox) that can be bought for a given cost, where the best tools in the toolbox (collection) are used for a particular job (scenario), while the total value reflects the expected ability of the toolbox (initiative collection) to address any single job (scenario). The above formulation was tested on the data provided by (Order, 2009) and the following results, showing score vs cost bound for two extreme dependency cases Alpha = 1 (initiatives are independent) and Alpha = 0 (initiatives are fully dependent) were obtained (see Figures 1 and 2). With the growing number of alternative initiatives differences in scores between independent and dependent initiatives also grow. Note that scores are higher for the independent case since any set of initiatives I within the cost bound contribute the full values of δ_{js} to the expression (1).

In the dependent case however the value δ_{js} is spread over all the links β_{jk} that attach to j and may include some where one end of the link are not within I, and so do not contribute to (1). The novelty of this paper is the extension of the budget scenario problem given by (Jagiello and Taylor, 2017) to take account of multiple alternative initiatives in the value calculation. If we again use the analogy of the toolbox for a collection of initiatives, this approach may be thought of as a more robust measure of total value in that some of the best tools (initiatives) in the toolbox may not be available at a particular time and the other tools in the toolbox could be used. Further research will be required to improve our optimizer to solve or approximate solutions to larger problems.

A4. Food and beverage supply chains

Integrating shelf life constraints in capacitated lot sizing and scheduling for perishable products

Shuo Chen[1], *Regina Berretta*[2], *Alexandre Mendes*[3] *and Alistair Clark*[4]

[1]The University of Newcastle, School of Electrical Engineering and Computing, Newcastle, Australia; sylvia.chen@uon.edu.au
[2]The University of Newcastle, School of Electrical Engineering and Computing, Newcastle, Australia; Regina.Berretta@newcastle.edu.au
[3]The University of Newcastle, School of Electrical Engineering and Computing, Newcastle, Australia; alexandre.mendes@newcastle.edu.au
[4]University of the West of England, Bristol, United Kingdom; Alistair.Clark@uwe.ac.uk

In this study, we consider a multi-item capacitated lot sizing and scheduling problem for perishable food products, which have a fixed shelf life period due to depreciation issues, such as physical deterioration or perceived value loss. We incorporate shelf life constraints within a classical lot sizing and scheduling problem that considers lot sizing and partial sequencing of production on a single machine over a finite planning horizon associated with demand in each period. The model includes setup times and setup costs. Moreover, it considers that disposal occurs when products reach their shelf life in inventory. We present two variants of this lot sizing and scheduling problem integrating shelf life constraints and also with and without disposal. We test the performance of the proposed formulations on a set of instances from the literature.

An application of vessel route planning models to the import/export of seasonal products

Riccardo Accorsi[1], *Emilio Ferrari*[2], *Riccardo Manzini*[3] *and Alessandro Tufano*[4]

Department of Industrial Engineering - University of Bologna, Bologna, Italy
[1]riccardo.accorsi@unibo.it
[2]emilio.ferrari@unibo.it
[3]riccardo.manzini@unibo.it
[4]alessandro.tufano2@unibo.it

The distribution of food products represents a great trade opportunity for maritime carriers and shipping companies, especially within the Mediterra-nean basin which concentrates many important food processing and consuming countries. Fresh products, as fruit and vegetables, are characterized by seasonal and climate-driven volumes, and logistics networks and distribution (i.e. ship-ping) operations should designed and planned in agreement with such trends. In this working paper an application of the Vessel Routing Problem with Selective Pickups and Deliveries (VRPSPD) to the maritime import/export of food sea-sonal product is illustrated. The VRPSPD belongs to well-known class of vehi-cle routing problems intended to plan the routes of the maritime distribution of commodities between sources and destinations. A time-dependent formulation of the VRPSPD is applied in this paper to maximize the profit of a maritime carrier involved in the import/export of fruits among Mediterranean ports. A simple numerical example is used to validate the model and to identify oppor-tunities for future problem investigations in the seaborne trade of seasonal and perishable products.

A hierarchical facility location problem in dairy industry under demand uncertainty

Rasul Esmaeilbeigi[1], *Richard Middleton*[2], *Rodolfo Garcia-Flores*[3] *and Mojtaba Heydar*[4]

[1]The University of Newcastle, School of Mathematical and Physical Sciences, Newcastle, Australia; rasul.esmaeilbeigi@uon.edu.au
[2]The University of Newcastle, School of Electrical Engineering and Computing, Newcastle, Australia; richard.middleton@newcastle.edu.au
[3]CSIRO Data-61, Optimization Applications and Platforms, Docklands, Australia; rodolfo.garcia-flores@data61.csiro.au
[4]University of Tasmania, School of Natural Sciences, Discipline of Mathematics, Hobart, Australia; mojtaba.heydar@utas.edu.au

We consider a special case of the Hierarchical Facility Location Problem (HFLP) in which demand is a random variable. This study is motivated by a recovery network design problem in dairy industry. Consider a set of cheese makers that produce raw whey. Whey is a by-product of cheese production that should be processed due to economical and/or competitive advantages. We are interested in collecting raw whey produced by each cheese maker (demand) in order to convert it into a commercial product called 40% demineralized whey powder (40DWP). Designing a whey recovery network is an important Reverse Logistics (RL) problem. Whey RL network design is first addressed by Garciía-Flores et al. (2015) in which the authors study an actual cluster of cheese makers in Minas Gerais, Brazil. The authors formulate the problem as a HFLP with two levels of facilities: collection centers and plants. Given a set of whey producing cheese makers, the problem is to select some of them to open these facilities and to allocate other cheese makers to the open facilities such that the total transportation and fixed costs are minimized. Esmaeilbeigi et al. (2017) study a variant of the problem in which demands are random variables and propose the two-stage stochastic programming approach to address the problem. They present the so-called extensive formulation of the problem that can be solved by a standard linear solver.

A sustainable multi-product inventory control model for fresh fruits considering deterioration

parichehr paam[1], *Regina Berretta*[2] *and Rodolfo Garcia-Flores*[3]

[1]The University of Newcastle, Newcastle, Australia; parichehr.paam@uon.edu.au
[2]The University of Newcastle, Newcastle, Australia;
Regina.Berretta@newcastle.edu.au
[3]CSIRO Data61, Docklands, Australia; rodolfo.garcia-flores@data61.csiro.au

The increase in energy price, on one hand, and environmental concerns, on the other hand, have urged stakeholders of fruit supply chains to rethink the operation of their warehouses with a sustainability-based view. One of the most significant innovations to store a wide range of fruits and vegetables is the regular atmosphere (RA) and controlled atmosphere (CA) warehouses. In CA mode, because of more precise adjustment of atmosphere elements, food deterioration is slowed down. In addition, warehouses are closed during this mode and there is no fruit flow. So, the demand can only be satisfied from a warehouse operating under RA mode. Warehouses can switch their mode between RA and CA. Obviously, a warehouse is in RA mode at the beginning, since it has fruit input. In this study, we propose a multi-product Mixed Integer Linear Programming (MILP) model to optimize the inventory stage of fruit supply chains using CA and RA warehouses. We optimize not only the inventory flow of each product, but also the number of active warehouses, operational length and mode of warehouses (RA/CA), while satisfying the demand. In other words, the model decides on the optimal inventory flow, the optimal time of when to turn warehouses on/off and when to switch their mode from RA to CA or vice versa. Besides, the model considers fruit deterioration over storage with a fixed deterioration rate per period, which differs for RA and CA warehouses, and incurs a penalty cost for deteriorated fruits to the objective function. The objective function of the model is to minimize total inventory costs consisting of handling cost, energy cost, setup cost when a warehouse starts to operate, and penalty cost of deteriorated fruits. To test the model, we have collected a data set describing actual operations at a large apple company in Australia. In our model, the minimization of total costs, which comprise energy costs as well as the penalty cost of fruit deterioration, depends on the way the warehouses are set up by means of switches that determine if a warehouse operates in CA or RA mode. The decision maker has control over the modes of operation of the warehouses, which are represented by decision variables in our model. In addition, several numerical examples are provided to test the applicability of the model using Gurobi optimization solver 7.0. In order to minimize the computational time, we determine and fix the used warehouses for each fruit product, as the warehouses are all the same at the beginning. The traditional inventory control models do not consider energy expenditures and environmental impacts in the objective function, and do not make decisions on the setup and operational mode of warehouses. Therefore, the proposed multi-product MILP model can be used as a novel tool by managers and stakeholders of all fresh food supply chain, using RA and CA warehouses, to move toward the sustainable development.

D5 and D6. Force Design Analysis and Modelling Unknowns

Preparing the Australian Army for a Fourth Industrial Revolution

Scott Holmes

UNSW, Canberra, Australia; scotth183@gmail.com

The Fourth Industrial Revolution represents a period of transformational change. Army's organisational design, combat forces and approach to future warfare preparation are dominated by the technologies and structures created by the last military transformation to 'modern warfare'. Three-dimensional modern warfare developed during the First World War by combining technologies from the First and Second Industrial ages. Subsequent improvements have evolved modern warfare through the inclusion of digital and electronic technologies of the Third Industrial Era. Optimising a military for dated norms becomes increasingly risky as a transition to the Fourth Industrial Era transforms society and the basis for military success. My presentation will contend that Army needs to focus on organisational design in addressing the magnitude of change accompanying the realisation of a Fourth Industrial Revolution.

Autonomous Analyst for Force Design

Damion Glassborow[1], *Sanat Bista*[2], *Daniel Bongiorno*[3], *Darryn Reid*[4] *and Wark Steven*[5]

[1]Defence Science and Technology Group, Perth, Australia;
damion.glassborow@dst.defence.gov.au
[2]Defence Science and Technology group, Canberra, Australia;
sanat.bista@dst.defence.gov.au
[3]Defence Science and Technology group, Canberra, Australia;
daniel.bongiorno@dst.defence.gov.au
[4]Defence Science and Technology Group, Adelaide, Australia;
darryn.reid@dst.defence.gov.au
[5]Defence Science and Technology Group, Adelaide, Australia;
steven.wark@dst.defence.gov.au

Force Design is a cyclic process to translate Government's strategic policy and guide the development of an affordable and capable future force for Australia [1]. It specifically involves the following activities:

1. Baselining the current force, strategy and investment plan, and assessing gaps and opportunities;
2. Collating and understanding the gaps and opportunities;
3. Prioritising the gaps and opportunities, and developing options;
4. Recommending responses to the gaps and opportunities; and
5. Synthesising the responses and developing future force options for consideration by Government.

The Defence Science and Technology Group (DSTG) is interested in researching and developing an intelligent decision support system capability to support Force Design in handling less-structured application areas [2]. This work is inspired by a concept known as the Autonomous Analyst. Autonomous Analyst is not a particular technology, but a problem choice to drive technology and analysis capability development, comprising of two interlocking components: firstly, and most obviously, we wish to employ emerging autonomous machine technologies to analytical advantage; secondly, because this is neither as straightforward nor as comprehensively beneficial as often assumed, Operations Analysis has to drive artificial intelligence and machine learning forward in the first place in order to obtain sufficiently capable algorithms.

Moreover, the latter component is really part of a broader requirement for Operations Analysis to enable the development of autonomous systems that are properly robust within known limits under normal conditions of typical environmental and operational uncertainty. A use case of an autonomous system is in knowledge synthesis using the large quantities of information within Defence, Industry and open literature. Defence literature could include a wide variety ranging from publications and reports to data from discussion exercises, analytical wargames and simulations. This would require:

- A Defence domain of knowledge corpus containing relevant data that are discoverable and reusable,
- Analysis and visualisation of the corpus data and uncertainty visualisation [3] using emerging technologies,
- Trained machine and deep learning models that will output answers with precision, confidence and justification

Research and Development Questions:

- How could data from a variety of sources be collected and collated? Data may come in the form of quantitative data from Defence assets or qualitative in the form of subjective matter expert (SME) text data, workshop data, open and closed source literature. Data may also be inherently coded such as coming from a wargame.
- How might we deal with 'continuous' time based data such as audio and video recordings? How would we fuse this data with types mentioned previously?
- How would we analyse the context in which data is collected? Is there a way to establish Defence context taxonomy in a semi/fully autonomous manner?
- Defence personnel and experts have background knowledge and context behind their language and terminology. How might we disambiguate different Defence domain language and terminology? This is a broadly-applicable research question about understanding domain specific language through learning context.
- Similarly if we are to ask a question of our autonomous analyst, how might we disambiguate the question to understand the intention? And like-wise return an answer in a similar domain specific language.
- How might we establish themes and topics amongst this varied and ambiguous data?
- How could we determine a quality or trustworthiness of the knowledge behind the qualitative data? Is trust model based analysis [4] a possibility?
- How might we shape the evolution of the autonomous analyst as a trusted system?

- How might we determine the quality of our own analysis? How do we validate data for a hypothetical event or situation?
- How might we use data visualisation to convey a message from an autonomous system?
- How could we establish trust in the autonomous system?
- When automation is introduced to handle big data challenges, the human analyst needs to understand the relationships and context within the data in order to understand their significance. What visualisation and data storytelling approaches are suitable for an autonomous analyst to explain its reasoning and results to the human analyst?

Seeding an Analysis Revolution

Darryn Reid[1] *and Tim McKay*[2]

Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia
[1]darryn.reid@dst.defence.gov.au
[2]tim.mckay@dst.defence.gov.au

Like its military theory host before it, Operations Analysis coalesced around the much older notion of rational knowledge as convergence towards verified truth, with decision-making consequently taken as making the right decision. This viewpoint manifests through ergodic bases for quantitative analysis, whereby the future is asserted to be symmetrical with the past. Yet nobody really believes the future and past to be symmetrical. The assumption is almost always tacit and motivated by the strong invariance of statistical properties of the system that follow: thus, for instance, we have problem tidy formulations in terms of cumulative expected utility maximisation. Unfortunately, the inability to deal adequately with problem settings involving deep uncertainty also follows. The constricting effects of this conception subsequently spawned a qualitative branch, built from imported social science and psychology to address what in military theory had already been split off as the "art" of war: the divide between 'hard' qualitative methods and the 'soft' qualitative approaches remains a defining feature of the field to this day. As exemplified by Force Design and Employment, decision-making problems featuring fundamental uncertainty, complexity and terminal failure potential press heavily on the limits of contemporary military and defence Operations Analysis. Fundamental uncertainty is the manifestation of underlying paradoxes that arise systems that allow self-reference, and these paradoxes generate limits to knowing within the problem context; ergodic assumptions exclude such phenomena to yield the intuitively attractive but overly strong time invariance whereby sampling from the inaccessible future is equivalent to sampling from the past. It is true that we cannot address arbitrary uncertainty, and that a bounding assumption of some form is necessary to progress; however, the defence of ergodicity as necessary though empirically false dissolves with the realisation that beyond this degenerate heart there are any number of weaker conditions from which to choose, each with a corresponding invariant, that admit different degrees of deep uncertainty. In choosing an invariant, we face a fundamental trade-off: the weaker the invariant, the higher the piecewise efficiency for obtaining decision-making reliability within broader bounds, so the goal in dealing with deep uncertainty in analysis settings amounts to identifying the strongest invariance that is nonetheless empirically reasonable in the applied problem. Defence Science and Technology (DST) has thus initiated its Modelling Complex Warfighting (MCW) Strategic Research Investment (SRI) with the goal of seeding a wider revolution in Operations Analysis around deep uncertainty, complexity and failure exposure. The ambition of its goal necessitates a deeply collaborative structure; self-referentially, it is an exercise in dealing with the kind of uncertainty that it investigates. Thus this programme has been constructed on the basis of an economic analysis informing both its structure and, in light of the inimical effects incumbent in collaboration itself, its collaboration approach.

The revolution will be an inflationary expansion of problem conceptions, rather than incremental improvements in solution methods to problems we already know about, and this expansion will yield the unification of the quantitative and qualitative branches. Behind the scenes, what will have shifted is the base conception of knowledge and decision-making: knowledge is now about reliability within known limits rather than convergence to verified truths, and decision-making is now about making decisions that are robust with respect to precisely characterised types of deep uncertainty rather than making the right decision.

Operational Research and Science Diplomacy

Clinton Heinze[1], *Ben Francis*[2], *Torgny Josefsson*[3], *Todd Mansell*[4], *McCann Linda*[5], *Jamie Watson*[6] *and Yi Yue*[7]

[1]DSTG, Melbourne, Australia; clinton.heinze@dst.defence.gov.au
[2]DSTG, Melbourne, Australia; ben.francis@dst.defence.gov.au
[3]DSTG, Melbourne, Australia; torgny.josefsson@dst.defence.gov.au
[4]DSTG, Canberra, Australia; todd.mansell@dst.defence.gov.au
[5]DSTG, Canberra, Australia; linda.mccann@dst.defence.gov.au
[6]DSTG, Canberra, Australia; jamie.watson@dst.defence.gov.au
[7]DSTG, London, United Kingdom; yi.yue@dst.defence.gov.au

The Defence Science and Technology Group (DST) maintain a significant science diplomacy presence in North America, Europe, and Asia aligned to Defence and broader Australian Government strategy. Recent trends have seen proportionately more operational research (OR) professionals filling these science diplomacy roles. A brief exploration of the emerging field of science diplomacy in an Australian context precedes a preliminary investigation of the overlap between the skills of OR generally and military OR specifically and those required by science diplomats. This view is supported by a series of personal accounts by the authors of their experiences and achievements in science diplomacy. "At first glance, scientists and diplomats are not obvious bedfellows. While science is a quest for truth, Sir Henry Wotton, the 17th century English diplomat, famously pegged an ambassador as 'an honest man sent to lie abroad for the good of his country.' " — Ian Chubb, (2012) The relatively recent emergence of science diplomacy as a professional endeavour is informed by a consideration of the (at least) three points at which science and diplomacy interact:

1. Science for diplomacy: scientific cooperation to further national political interests and improve international relations;
2. Diplomacy for science: diplomacy for facilitating international scientific endeavours;
3. Science in diplomacy: the use of science and technology in the conduct of diplomatic relations.

Diplomats tend to be concerned with the first of these and scientists the second, but there are an increasing number of professional science diplomats (or diplomat scientists) who blur the lines between these areas. OR is, at its heart, an analytical pursuit and the overlap with diplomacy is made clear by Blackwill in his 'ideal qualities of a successful diplomat' where he lists "Demonstrate an analytical temperament" as a critical capability a diplomat should possess. OR is also inherently multi-disciplinary and in many cases it is the OR professional who provides the glue that unifies the multidisciplinary team and this too is a highly sought after capability in science diplomacy.

OR, more than the pure sciences, prefers solutions that trade accuracy for timeliness or politics or other pragmatics of the domain. This approach, strongly found in OR, is of particular relevance to diplomacy. War and diplomacy are closely linked; Clausewitz's famous aphorism captures this interconnection and was entertainingly, or perhaps ominously, restated by Zhou Enlai as "All diplomacy is a continuation of war by other means". It should be no surprise then that those with a background in military theory have some of the attributes that are particularly suited to diplomacy and so it might be concluded that military OR in particular provides a great training ground for those seeking a career in science diplomacy. With relevant and recent roles in science diplomacy combined with careers in defence operational research the authors provide an argument for the particular advantages that OR offers to science diplomacy and provide accounts of the role of science diplomat that advocate for its importance, argue that it is of significant national and international importance and is personally and professionally rewarding.

A8. Healthcare and service industries

The operating room scheduling problem based on patient priority

Omolbanin Mashkani[1], *F.J. Hwang*[2] *and Amir Salehipour*[3]

University of Technology Sydney, School of Mathematical and Physical Sciences, Sydney, Australia
[1]omolbanin.mashkani@student.uts.edu.au
[2]feng-jang.hwang@uts.edu.au
[3]amir.salehipour@uts.edu.au

An effcient operating theatre schedule contributes significantly to enhancing efficiency of hospital operation management and plays a critical financial role for most hospital setting. In this paper, an operating room scheduling problem based on patient priority is investigated at tactical and operational levels subject to specific strategic decisions. At tactical level the main goal is to generate a cyclic time table, known as the master surgical schedule (MSS) and can be repeated over the planning horizon of several months to years. Furthermore, operational level concerns about allocating patients to operating rooms and determining the day of surgeries, called surgical case assignment problem (SCAP). To handle the problems at both decision levels simultaneously as MSS-SCAP problem, a bi-level integer linear programming (ILP) model and a heuristic approach are proposed. The objective function is to maximize the total priority scores of the patients assigned to the surgical scheduling blocks over a given planning horizon. An adaptive ILP model also is proposed to solve the SCAP taking into consideration the dynamics of waiting list. The computational experiments are conducted using a set of random data to evaluate the performance of the proposed bi-level ILP model and heuristic, in terms of solution quality and computation time. Our numerical results indicate that the proposed ILP is capable of yielded optimal solutions for the small-scale instances and near-optimal solutions for medium-size instances within 3,600 seconds. The developed heuristic algorithm can generate quality solutions within 2 seconds for large-scale instances.

Post-disaster volatility of blood donations in an unsteady blood supply chain

Babak Abbasi[1], *Zahra Hosseinifard*[2], *Masih Fadaki*[3] *and Nigel Clay*[4]

[1]RMIT University, School of Business IT and Logistics, Melbourne, Australia;
babak.abbasi@rmit.edu.au
[2]Deakin University, Faculty of Business and Law, Melbourne, Australia;
zahra.hosseinifard@deakin.edu.au
[3]RMIT University, School of Business IT and Logistics, Melbourne, Australia;
Masih.fadaki@rmit.edu.au
[4]RMIT University, school of Science, Melbourne, Australia; nclay67@gmail.com

The blood supply chain triggers an urgent call for additional blood when faced with shortages in response to surged demand as a consequence of catastrophic events. Past studies show that the response to a call for blood after a disaster is substantive. Yet the consequential impact on the supply chain is not well understood. This is due to the perishability of blood and the fact that donors are not eligible to give blood for a certain period after a donation has been made. In this study the donation process is modelled with a discrete Markov chain and the impact of a call for blood resulting from a disaster is studied. This paper highlights new insights that aid planners to mitigate the negative impacts of a substantial response to a call for blood.

D10. Defence Decision Analysis

A Systems Approach to Analysing Organisational-level Adaptability: Review of the Australian Army Lessons Network as a Case Study

Matthew Richmond[1], *Amina Omarova*[2] *and Vernon Ireland*[3]

[1]Defence Science and Technology Group, Edinburgh, Australia;
matthewkevinrichmond@outlook.com
[2]University of Adelaide, Adelaide, Australia; amina@aurecongroup.com
[3]University of Adelaide, Adelaide, Australia; vernon.ireland@adelaide.edu.au

This paper describes a methodology for reviewing organisational-level adaptability from a systems perspective. Taking an action learning approach, we reviewed the Army Lessons Network (ALN) as a case study in order to: (i) develop practical options for improvement of the ALN, (ii) reflect on the review methodology and identify options for improving its effectiveness in subsequent reviews, and (iii) demonstrate the utility of adaptive review and lay the foundations for its further application.

The DoDAF view of the SCMILE services framework

Nicholas Tay[1] *and Donald Lowe*[2]

Dept of Defence, ACT, Australia
[1]nicholas.tay@dst.defence.gov.au
[2]donald.lowe@dst.defence.gov.au

The SCMILE Services Framework (SSF) is a service-based approach being used to model and articulate ADF capability dependencies and their integration and interoperability (I2) implications. First proposed in the last decade (Lowe et al, 2006), the SSF is being employed by the Joint Integration Concepts and Assurance (JICA) Directorate for I2 assurance and validation. There are three key strengths of the SSF:

1. SIMPLE - The SSF is based on six simple functions which describe how a system or entity can completely exchange services with another system.

2. SCALABLE - The services can be applied to systems, sub-systems or system-of-systems. Scalability permits flexibility in choosing the appropriate level of a system to consider dependencies.

3. SYSTEMATIC - The codification of agreements require that the systems become standardised. This encourages managers to consider systems consistently using the same terms and characteristics.

Another major advantage of the SSF is that it is similar to several frameworks currently employed within the ADF, including the Joint War-fighting functions, represented in the following table. Previously, Lui et al. (2017) explored several examples where the SSF could be applied within the United States' Department of Defense Architecture Framework (DoDAF). This presentation will refine the options based on JICA's employment of SCMILE for I2 dependencies. At first glance, architects may presume that SCMILE is similar to need-lines or information exchanges in an operational view. In actual fact, the information mobility function (i.e. communications and information networks) and the widespread use of systems indicate that SCMILE is best placed as a DoDAF resource exchange element. Furthermore, the SCMILE matrices are improved extensions of the SV-3 (Systems-Systems Matrix), demonstrated by the following figure.

The articulation of the SSF as a DoDAF resource exchange element within DoDAF provides several advantages.

- Firstly, it can provide an improved understanding of the SSF for architects and coalition partners who are familiar with DoDAF.
- Secondly, it allows the SSF to access a range of systems modelling tools that have been developed around DoDAF and other compatible architecture such as UPDM (Unified Profile of DoDAF and MODAF1) and the NAF (NATO Architecture Framework).
- Lastly, it provides a connection between the dependency data collected from the SSF with existing data (such as information exchange requirements) that have been collected within DoDAF.

Confidence in Estimation

Chelsea Kramer[1] *and Niem Tri*[2]

DSTG, Fishermans Bend, Australia
[1]chelsea.kramer@icloud.com
[2]niem.tri@defence.gov.au

Subject Matter Experts (SMEs) are used throughout the Australian Defence Force (ADF) to examine needs for future capabilities and to develop future warfighting concepts. It is imperative that SMEs give objective and reliable estimates which can give proper insight into the consequences of choices and actions. However, most estimates are not precise and rational but influenced by internal and cognitive biases. The ability to understand these influences can lead to more objective and reliable estimates.

Keynote Seminar

Metaheuristics for Optimization Problems with Routing and Location Characteristics

Paolo Toth

Department of Computer Science and Engineering - University of Bologna, Bologna, Italy;
paolo.toth@unibo.it

The family of optimization problems including both "routing" and "location" characteristics will be consider. After a short description of the main features of the "routing problems" and of the "location problems", the Capacitated Location-Routing Problem (CLRP) and the Generalized Traveling Salesman Problem (GTSP will be addressed in more detail. In the Capacitated Location-Routing Problem (CLRP), we are given a set of available depots (each depot is located at a node of a network, and has an associate cost, a capacity and a set of identical vehicles stationed at the depot) and a set of customers (each customer is located at a node of the network and has a positive demand). All the demands of the customers must be satisfied through a set of routes (each route starts from a depot, visits a subset of customers whose global demand must not exceed the capacity of the associated vehicle, and returns to the starting depot). In addition, each customer must be visited by exactly one route, and the global demand of the customers visited by the vehicles associated with a depot must not exceed the capacity of the depot. The aim of the CLRP is to determine the subset of depots to be opened, the customers to be assigned to each open depot, and the routes to be performed to satisfy the demands of the customers. The objective is to minimize the sum of the global cost of the open depots, of the global cost of the used vehicles, and of the global traveling cost associated with the performed routes. The CLRP is NP-hard, since it generalizes two well-known NP-hard problems: the Capacitated Facility Location Problem (CFLP) and the Multi Depot Vehicle Routing Problem (MDVRP). The Generalized Traveling Salesman Problem (GTSP) is a generalization of the well-known Traveling Salesman Problem (TSP) in which the set of nodes is partitioned into clusters. In the GTSP the aim is to find a minimum-cost simple cycle visiting exactly one node for each cluster. The GTSP is NP-hard, since it generalizes the TSP. The most effective metaheuristics proposed for the solution of the CLRP and of the GTSP will be described, and experimentally compared on the benchmark instances from the literature, by taking into account both the quality of the solutions found and the CPU times required to obtain the solutions.

D9. Simulation of combat operations

Exploring the Impact of Alternative Random Number Generators within a Stochastic Combat Simulation

Kevin Clark[1], *Darren Lohmeyer*[2] *and Dion Grieger*[3]

Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia
[1]kevin.clark@dst.defence.gov.au
[2]Darren.Lohmeyer@dst.defence.gov.au
[3]dion.grieger@dst.defence.gov.au

Stochastic simulations rely on random number generators (RNG) to represent real world variation and uncertainty. Within the context of operations analysis there is a desire to select RNGs that provide statistically consistent, robust and reliable results. Our investigation examined the impact of both RNG and initial seed selection on results within the Combined Arms Analysis Tool for the 21st Century (COMBATXXI) closed loop combat simulation model. COMBATXXI uses multiple random number streams for each simulation replication, with each stream being called up to 700,000 times in the scenarios examined in this study. Each stream is used for a different function within the model, such as observations, direct fire engagements and communication. Further, these are not necessarily drawn from with the same frequency as any other stream. Two different RNG algorithms implemented in COMBATXXI were tested: Mersenne Twister and a Linear Congruential Generator. The Linear Congruential Generator has a substantially shorter period than the Mersenne Twister, reducing how many replications can be run before random number streams from one replication overlap with streams from another. COMBATXXI provides sets of seeds for both random number generators however these seeds were not guaranteed to generate streams of random numbers that did not overlap. Consequently one new set of block separated seeds was produced for each RNG, enforcing non-overlap between streams and between replications. The RNGs were initially tested outside of COMBATXXI by generating streams of random numbers using both the original and revised seeds. These streams were tested for distribution and correlation to ensure that the statistical properties were comparable. In each case, the number streams provided similar results. To test the impact of RNG and seed choice within COMBATXXI, an analysis of three alternative military vehicles was conducted. 200 replications of the simulation were run for each of the three alternatives and with each of the RNG and initial seed combinations. The results were compared across six primary metrics to test for differences between the three alternatives as well as differences between each of the RNG-initial seed combinations. There were no statistical differences between the four RNG-initial seed combinations for any of the three alternatives. However, when the three alternatives were compared against each other, the number of replications required to identify statistically significant differences between the alternatives was different, dependant on the initial RNG and seed selection.

In five of the six metrics RNG-seed combinations converged towards the same statistical finding as the number of replications increased, albeit at different rates. For one of the metrics one Mersenne Twister seed set rejected the null hypothesis while the others did not. These results suggest both RNG choice and initial seed selection may need to be considered in the design of combat simulation studies. These decisions may provide greater efficiency of simulation computing resources via fewer replications without sacrificing the statistical confidence of study findings.

Aircraft in Warfare

Clinton Heinze

Defence Science and Technology Group, Melbourne, Australia;
clinton.heinze@dst.defence.gov.au

There are arguments about when operational research (OR) began, but with a career at the conjunction of air combat and OR it will not surprise you to learn that I hold fast to the view that it originated in 1916 with the publication of Aircraft in Warfare by Frederick Lanchester. In September of this year, after searching for decades, I finally acquired a first edition copy. Surviving examples are not particularly difficult to find, but they are often expensive or in relatively poor condition and, as with many OR professionals, the tendency to optimise engenders fussiness. It is now 102 years since Lanchester published this remarkable book and yet it remains relatively unknown; a miscarriage of academic justice that I will attempt to partly redress. After briefly summarising Lanchester's life and work the focus will turn to Aircraft in Warfare, reviewing each of the chapters in turn. The book is best known for the introduction of Lanchester's Power Laws but is replete with gems of wisdom and insights and occasionally, it must be said, folly. Even the preface by Major General Sir David Henderson is timeless and includes the self-effacing comment that "I am strongly entrenched behind a barricade of military prejudice, with some dim recollections of early scientific training..." but is then followed with a critique of Lanchester's consideration of the numbers of aircraft that might be acquired by an army where he writes "The only safe line on which to proceed is to consider, first, what are the services which the aircraft are required to perform? Second, how much of our available resources are we justified in devoting to these services?" This is the very essence of force design as it now exists. Lanchester's innovations included the glide bomb, powered air-to-surface weapons, air-to-air missiles, joint personnel recovery, armour plate on aircraft, a principled approach to calculating aircraft stability following stores-release, the first example of joint doctrine for the employment of air power, the advantages of a single-type fleet, the problems of mis-identification of civil air traffic and dozens of other deeply interesting insights. There are, it must be admitted, occasional mis-steps. For example, "we may dismiss from our minds any general usage of the air as a commercial highway; the traffic in merchandise which will be airborne will never become a great percentage of the world's total." Or the complete discounting of air to air missiles "Again, we are confronted with the fact that any such weapon would be of little service apart from attack on an airship, and so may be looked on as useless lumber." Concluding Chapter 18 is an account of the strategic employment of air power; Lanchester states, "presuming the continued supremacy of our air fleets, no resistance or defence by the enemy of a permanent character can be maintained". The idea that air supremacy is a necessary precondition to military victory was dismissed as nonsense at the time but would eventually become an accepted maxim of military thought. And thus it remains.

A5. Decision making under uncertainty

Dynamic Relocation of Aerial Firefighting Resources to Reduce Expected Wildfire Damage

Nicholas Davey[1], *Simon Dunstall*[2] *and Saman Halgamuge*[3]

[1]The University of Melbourne, Melbourne, ; ndavey@student.unimelb.edu.au
[2]CSIRO Data61, Docklands, Australia; simon.dunstall@data61.csiro.au
[3]The University of Melbourne, Melbourne, Australia; saman@unimelb.edu.au

Aerial firefighting resources are an integral part of modern wildfire suppression strategies. In many locations around the world where wildfires pose a serious threat, firefighting authorities have access to fleets of different aircraft. These can be used to provide support to land-based resources during extended attack of existing fires or to quickly suppress recent spark events during the initial attack phase. As the amount of time that a fire has been burning is a predictor of the amount of damage it causes, fast aerial response times are critical. Therefore, there is significant value in dynamically repositioning aircraft to airbases and fires over the course of a fire day or fire season. In this paper, we devise one such approach based on Model-Predictive Control to make relocation decisions at various times over a single day. These relocation decisions are based on solving an underlying Mixed Integer Linear Program (MILP) so as to minimise expected damage over a lookahead horizon. The inputs to this program are updated at each of these decision times based on prevailing stochastic weather conditions, the current state of fires in the region, and the current assignment of aircraft to bases and fires. The expected fire damage profiles used in this model are based on empirical data that is pre-computed for the region of interest. We apply our model to a scenario in Central Chile and show that with careful parameter selections it is possible to make improved relocation decisions to reduce the expected fire damage in a region using this approach.

A6. Methodologies, platforms and novel applications

Single-machine coupled task scheduling with time-dependent processing times

Mostafa Khatami[1], *Amir Salehipour*[2] *and F.J. Hwang*[3]

University of Technology Sydney, School of Mathematical and Physical Sciences, Sydney, Australia
[1]mostafa.khatami@student.uts.edu.au
[2]amir.salehipour@uts.edu.au
[3]feng-jang.hwang@uts.edu.au

In this paper we study the single machine coupled task scheduling problem with time-dependent processing times. The generic coupled task scheduling problem deals with the two-task jobs where the second task of a job must be processed after the completion of the first task and with an exact delay between the tasks. We study a variant, in which all jobs have identical processing times for the first tasks, as well as identical values for delays, and the second tasks of jobs have time-dependent processing time. The processing times of the second tasks follow the simple linear processing times. We investigate an optimal case of the problem, and propose a heuristic algorithm for the general case. Over a set of randomly generated instances we show promising result for the heuristic algorithm, particularly, for large instances.

Keynote Seminar

Integer Linear Programming in the Age of Cloud Computing

Robert Dell

Naval Posgraduate School, Monterey, The United States of America; Dell@nps.edu

Cloud computing allows us to temporarily leverage massive computing power. Harnessing this power creates opportunities to quickly solve many (e.g., hundreds of thousands) of integer linear programming (ILP) instances that each provide an optimal plan to a possible future. At this scale, we find the inclusion of features such as: solution persistence (ensuring the model is not ignorant of its own prior advice); elastic constraints (allowing constraint violations when this is advantageous); and secondary objectives to be key to successful gleaning true insights. This talk presents ILP models selected from diverse military applications with an emphasis on how these features enabled impactful analysis and how the age of cloud computing provides new challenges and opportunities.
Collaborators on this work include many officer-students and colleagues at the Naval Postgraduate School. References are supplied for following up on the details.

A5. Decision making under uncertainty

Real Option Evaluation of a Tailings Dam Failure

Nicolas Langrené[1], *Margaret Armstrong*[2], *Wen Chen*[3], *Renato Petter*[4] *and Petter Carlos*[5]

[1]CSIRO Data61, Docklands, Australia; nicolas.langrene@data61.csiro.au
[2]Fundacao Getulio Vargas, Escola de Matemática Aplicada, Rio de Janeiro, Brazil; margaret.armstrong@fgv.br
[3]CSIRO Data61, Docklands, Australia; wen.chen@data61.csiro.au
[4]Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil; rapetter@gmail.com
[5]Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil; cpetter@ufrgs.br

Traditional mining projects store the byproducts of mining operations behind a tailings dam. When the wall of a tailings dam is breached, it sends a flood of often toxic material into the surrounding countryside, causing serious environmental damage and often loss of life. For example 19 people were killed in the recent Samarco disaster in Brazil and the Rio Doce River was polluted for more than 660 kilometers down to its mouth and out into the Atlantic Ocean (Fernandes et al, 2016). Although these dam failures are often considered extremely rare events, they are far more common than is realized. In a study sponsored by the European Union, Rico, Benito, Salgueiro et al (2008) identified 147 cases worldwide up to 2008; more recently 36 such cases were identified in the 18 year period from 2000 to mid-2017, an average of two per year (Wise-Uranium, accessed 2017). In the literature, most studies on tailings dam failures focus on the technical causes of the dam failure, or on the loss of life and the environmental damage. From the environmental and safety point of view, investing in the safety of tailings dams or in alternative mine designs devoid of such dams should be the preferred option. In this work, we show that this preference still holds from a pure financial point of view, once the actual financial impact of such disasters is taken into account, paving the way for a stricter enforcement of safety regulations. This work focuses on the evaluation of the financial impact of such disasters within a real option framework from the point of view of mining companies. We establish a closed-form formula for the expected value of a conventional mining project subject to the risk of tailings dam failure. The stochastic components considered are the metal price, the occurrence of a tailings dam failure, as well as the penalty cost and recovery period after such a failure. These components are calibrated to the available statistical data on tailings dam failures. Beyond this assessment of the financial impact of tailings dam failures on the value of mining projects, we analyse two dynamic options available to mining companies: firstly, the option to perform preventive maintenance and temporary repairs to the dams to postpone the risk of failure; and secondly, retrofitting the mine with an alternative design devoid of tailings ponds, such as dry processing, to remove the risk of failure altogether.

We obtain a semi-analytic value for these two real options by a simple dynamic programming numerical scheme combined with Monte Carlo simulations of the dynamic risk factors. We analyse the attractiveness of these two options from the point of view of mining companies, and discuss the factors that could lead to a growing recourse to dry processing in the future.

Better management of production incidents in mining using multistage stochastic programming

Margaret Armstrong[1], *Lorenzo Reus*[2] *and Bernardo Pagnoncelli*[3]

[1]CNRS UMR 9217, i3, CERNA Centre for industrial economy, PSL Research University, MINES Paristech, Paris, France; margaret.armstrong171@gmail.com
[2]University Adolfo Ibanez, School of Business, Santiago, Chile; lorenzo.reus@uai.cl
[3]University Adolfo Ibanez, School of Business, Santiago, Chile;
bernardo.pagnoncelli@uai.cl

Among the many sources of uncertainty in mining are production incidents: these can be strikes, environmental issues, accidents, or any kind of event that disrupts production. In this work, we present a strategical mine planning model that takes account of these types of incidents. When confronted by production difficulties, mines which have contracts to supply customers have a range of flexibility options including buying on the spot market, or taking material from a stockpile if they have one. Earlier work on the same subject (Armstrong et al, 2011 and 2012; Armstrong and Galli, 2013) was limited in that the optimization could only be carried out for a few time periods (up to 5 years). By using decomposition schemes we are now able to efficiently solve large-scale versions of the model, with a horizon of up to 30 years. We consider decision trees with 30^{12} scenarios and perform a sensitivity analysis on some parameters of the model, such as minimum and maximum buying limits, in order to understand their effect on the optimal policy. The results provide a "roadmap" for mine management as to the optimal decisions, taking future possibilities into account. We present extensive numerical results using the new sddd.jl library, using Julia language, and discuss policy implications of our findings.

Compete or collaborate? Pricing strategies under a multinomial logit model with non-linear network effects.

Felipe Maldonado[1], *Gerardo Berbeglia*[2] *and Pascal Van*[3]

[1]CSIRO-Data61, Canberra, Australia; felipe.maldonado@anu.edu.au
[2]The University of Melbourne, Melbourne, Australia; g.berbeglia@mbs.edu
[3]Stewart School of Industrial and Systems Engineering Georgia Institute of Technology, Atlanta, ; pascal.vanhentenryck@isye.gatech.edu

Consumer choice models are a key component in Operations Research, where the demands for certain type of products are assumed to follow a particular form, and sellers adjust their strategies accordingly to maximise their revenues. Among these model the multinomial logit (MNL) model is one of the most widely used in the literature due to its analytical tractability. In this research we consider a MNL model where we also incorporate network effects over the consumer's decisions, meaning that a certain product is more attractive (increases its utility) when more people purchase it. We aim to study seller's pricing strategies based on this model, where the purchasing decisions are affected by past consumption. In general terms we assume that the willingness to purchase is influenced by the known quality of the products, their prices, and network effects as a function of consumption history. The consumers can purchase the product $i \in 1, \ldots, n$ that maximises their expected utility or they can choose to leave the market without making any purchase, what is called the no purchase option. Unlike most of the previous research in this area, we consider that the no purchase option also presents network effects, representing the ineffectiveness of the market (e.g., if products are too expensive, even when they have good quality, they may not attract enough people), capturing in that way, the effect of consumers buying similar products somewhere else.

The main contributions of this research can be summarised in the following way:

- Non-linear network effects in a consumer choice model: We propose a variation of the Multinomial Logit Model for consumer choice where we incorporate non-linear network effects, representing in this way, market interactions where consumers only see a score function of the past consumption. Since the probability of choosing the available products (or the no purchase option) dynamically changes over time due to the network effects, we apply stochastic approximations techniques to prove that such probability converges almost surely to an asymptotic stationary distribution, that represents the market share of each product in the long run.
- Monopolistic and competitive pricing are analysed: For a market with n sellers, we model their expected revenues based on the asymptotic market share distribution and the displayed prices. First, we study the case where the sellers act collaboratively, adopting a monopolistic pricing strategy to maximise the overall expected revenue, we show that the market share of the no purchase option is decreasing in terms of the network parameter $r : 0 < r < 1$, and that the overall expected revenue is increasing in that parameter as long as r is large enough. We then study the case where the sellers compete, inducing a price competition game that has a unique pure Nash Equilibrium (we also provide an algorithm to compute it). We finally compare experimentally and analytically both cases, incorporating into the analysis, the consumers' perspective.

Considering Uncertainty in the Optimization of Hybrid Ground Source Heat Pump Systems

Hansani Weeratunge[1], *Simon Dunstall*[2], *Julian de Hoog*[3], *Guillermo Narsilio*[4] *and Halgamuge Saman*[5]

[1]The University of Melbourne, Melbourne, ; hweeratunge@student.unimelb.edu.au
[2]CSIRO Data61, Docklands, Australia; simon.dunstall@data61.csiro.au
[3]IBM Research Australia, Melbourne, ; juliandehoog@au1.ibm.com
[4]University of Melbourne, Parkville, Australia; narsilio@unimelb.edu.au
[5]University of Melbourne, Parkville, Australia; saman@unimelb.edu.au

Heating, Ventilation, and Air Conditioning (HVAC) systems are receiving an increased interest for efficiency improvements since they consume approximately 40Ground source heat pump systems exploit the large thermal mass of the ground and utilize the ground as a heat source in winter and heat sink in summer to provide heating and cooling. However, the continuous long-term operation may affect the ground temperature if the amount of heat extracted and rejected is not balanced annually. As a result, the system performance may decline and ultimately result in failure of the system. A possible solution to this problem may be to integrate a supplementary heat source or sink to prevent this annual energy imbalance. For a heating dominant system, solar thermal energy can be integrated as an auxiliary heat source to maintain the energy balance. This solar assisted ground source heat pump systems (SAGSHP) can achieve higher system efficiency than conventional systems. Energy storage plays an important role in synchronizing the intermittent solar generation and the heating demand. In addition, thermal capacity of the building can also be employed to store thermal energy and shift the consumption of the heating system. These methods can be used to achieve peak shaving and take advantage of off-peak electricity prices. In general, HVAC systems are designed to accommodate peak demand, which typically occurs only on a few days per year. Due to the high initial cost of GSHP systems, this might lead to an uneconomic, oversized heating system. Therefore, it may be more economical to design a hybrid heating system combining a conventional heating method to provide the peak heating requirement. This way, the investment cost can be minimized by reducing the capacity needed by each energy source. An intelligent control strategy plays a critical role in meeting the demand efficiently. In addition, a significant amount of energy and cost can be saved by integrating the occupancy and weather predictions into the controller. Furthermore, incorporating the time of use electricity price can minimize the cost by shifting the electricity consumption into periods with low electricity price. However, for a system with a significant energy generation and consumption, an efficient control system must consider the uncertainties involved in the occupant activities and environmental conditions to successfully respond to uncertain conditions. Real options valuation is a technique that can be used to valuate operating decisions under uncertain conditions. This approach considers multiple decision pathways as a consequence of uncertainty involved and the flexibility of decisions.

This study aims to develop an optimized control model of a hybrid ground source heat pump system using least squares Monte Carlo simulation (Longstaff and Schwartz 2001), considering the uncertainties involved in the environmental conditions and occupant behaviours.

Exact Analysis of Periodic Review Inventory Control Systems with Lost Sales and Zero Lead Time

Nazanin Esmaili[1], *Bryan Norman*[2] *and Jayant Rajgopal*[3]

[1]University of Technology Sydney, Faculty of Engineering and IT, Sydney, Australia; nazanin.esmaili@uts.edu.au
[2]Lubbock, Texas Tech University, Texas Tech University, Manufacturing and Systems Engineering, Department of Industrial, Texas, The United States of America; bryan.norman@ttu.edu
[3]Pittsburgh, University of Pittsburgh, Department of Industrial Engineering, PA, The United States of America; rajgopal@pitt.edu

We study periodic review inventory control policies with lost sales, stochastic demand, zero lead-time and a target service level to be satisfied. The objective is to minimize the total expected inventory management effort, subject to a meeting a specified minimum alpha-service level. In order to obtain structural results for this inventory control system, it is modeled as a discrete time Markov chain. We present a novel approach to derive exact closed-form solutions for the limiting distribution of the on-hand inventory level at the end of a review period, given values of s and S. We then establish a relationship between the limiting distributions for adjacent values of the reorder point (s), and use this to develop an efficient recursive algorithm for determining the limiting distribution for every feasible reorder point. There is no need to solve a system of balance equations, and as input, the algorithm only needs a simple vector of probabilities, which is a one-time calculation for the given value of S. The value of the algorithm lies in the fact it can be applied to optimize any policy that can be stated as a function of these limiting probabilities, and because of the ease with which we can derive these probabilities, it allows us to address a system with many different items that are controlled independently. We illustrate this using an application in healthcare inventory management, and show that optimal policies for a particular objective can be readily derived for several thousand items across a number of locations, and these yield sizeable savings over the policies currently in use.

An AHP Trade-off Study

Richard Taylor

Defence Science and Technology group, Canberra, Australia;
richard.taylor@dst.defence.gov.au

The Analytical Hierarchy Process (AHP) is a multi-criteria decision making (MCDM) method using pairwise comparisons in order to rank a number of options against a collection of criteria (Saaty 1980). In order to deal with the data demands of many criteria and options a number of researchers have investigated the use of incomplete matrices whereby only a subset of all possible pairwise comparisons are required. In using this approach a sacrifice is made in the form of increasing the sensitivity of the rankings produced to individual comparisons. Given that these comparisons are subjective it is desirable to limit or reduce this sensitivity level. We present an analysis of optimal graph designs corresponding to the comparison matrices for the purpose of AHP analysis that find the best balance between the number of edges in the graph, corresponding to the number of input comparisons required by subject matter experts, and the sensitivity of the output rankings produced to changes in the comparisons. To do this we generate regular graphs of different degrees in which the ratio of the largest and the second smallest eigenvalues are as small as possible. This produces graphs which are highly connected and in which the edge expansion number is high for the number of edges. Assume that the logarithm of the pairwise comparisons are random variables with standard deviation. By using the Geometric Mean, also called the Logarithmic Least Squares Method, we are able to analytically derive expressions for the standard deviation of the option rankings. We then compute the maximum standard deviation across the option rankings as a measure of the sensitivity of the graph design. By appropriately calibrating the number of edges and the sensitivity measure we find the trade-off point between regularity degree and sensitivity. Below we show this trade-off for the regular graphs with 50 nodes. The Best Worse Method (BWM) was recently proposed as a way of minimising the pairwise comparisons by first identifying the most and least important options, and then using only those pairwise comparisons which involve the most and least important options (Rezaei 2015). The author suggests that comparisons of this kind are easier to make than comparisons between more equally ranked options. The BWM leads to the construction of graph designs with average degree $4-8/n$. This graph is shown on the figure below, illustrating its average degree and its sensitivity. We note that it appears to have good sensitivity properties comparable to the best regular graphs of degree 4 that we have been able to find.

Competing for cognitive resources: Measuring workload in a time pressured dual-taskenvironment

Hector Palada[1], *Andrew Neal*[2], *David Strayer*[3], *Tim Ballard*[4]*and Heathcote Andrew*[5]

[1]Univeristy of Queensland, Brisbane, Australia; hector.palada@uqconnect.edu.au
[2]The University of Queensland, Brisbane, Australia; a.neal@psy.uq.edu.au
[3]The University of Utah, Salt Lake City, The United States of America; david.strayer@utah.edu
[4]Univeristy of Queensland, Brisbane, Australia; t.ballard@uq.edu.au
[5]The University of Tasmania, Hobart, Australia; andrew.heathcote@utas.edu.au

Measuring workload is critical for systems design to ensure that individuals are not overloaded and can make decisions and perform tasks effectively and safely. Understanding why and how a particular measurement method is sensitive to workload is important for mitigating overload. The Detection Response Task (DRT) is an international standard (ISO, 2016) for objective real-time assessment of workload that has minimal effects on primary task performance, making it an attractive option for workload measurement in many settings. The DRT is suitable for use in active tasks as driving, requiring operators to make a simple button-press in response to a stimulus (e.g., an LED on a head-mounted wand or a vibrotactile device taped to the clavicle) that occurs randomly every 3-5 seconds. An increase in DRT response times or failures to respond as primary task load increases is thought to occur due to competing resources being reallocated to the primary task. However, alternative processes could account for these effects, including changes in response caution, response bias and non-decision processes. We examine how people respond to changes in task demands in a dual-task environment with the aim of identifying what psychological processes the DRT is measuring. The primary task required complex decision making under time pressure, and we modeled it, and the DRT responses, using the linear ballistic accumulator (Brown and Heathcote, 2008) and a single-bound diffusion (Wald) model (Heathcote, 2004). As time pressure increased, the rate of information processing increased on the primary task while response caution decreased. In contrast, the rate of information processing in the DRT declined with greater time pressure. These results are consistent with the hypothesis that the DRT's sensitivity to workload is due to the reallocation of cognitive resources as demands increase on the primary task.

Assortment Optimisation Under non-conventional customer choice models

Alvaro Flores[1], *Gerardo Berbeglia*[2] *and Pascal Van Hentenryck*[3]

[1]Australian National University, Canberra, Australia; alvaro.flores@anu.edu.au
[2]The University of Melbourne, Melbourne, Australia; g.berbeglia@mbs.edu
[3]Stewart School of Industrial and Systems Engineering Georgia Institute of Technology, Atlanta, The United States of America; pascal.vanhentenryck@isye.gatech.edu

Revenue management is an extensive area of research nowadays, with applications across every industry. In the centre of this area lays the assortment problem, which amounts to find a subset of products to offer in order to maximize revenue, provided that customers follow a certain model of choice. Most studied models satisfy the following property: whenever the offered set is enlarged, then the probability of selecting a specific product decreases. This property is called regularity in the literature. However, customer behaviour often shows violations of this condition (such as the compromise effect, where adding extreme options sometimes leads to a positive effect for compromise products, whose probabilities of being selected increase in relative terms compared to other products). We study two models of customer choice where regularity violations can be accommodated, and show that the assortment optimisation problem can still be solved in polynomial time. First we analyse the Sequential Multinomial Logit (SML). Under the SML model, products are partitioned into two levels, to capture differences in attractiveness, brand awareness and, or visibility of the products in the market. When a consumer is presented with an assortment of products, she first considers products on the first level and, if none of them is purchased, products in the second level are considered. This model is a special case of the Perception-Adjusted Luce Model (PALM) recently proposed by Echenique, Saito, and Tserenjigmid (2018). It can explain many behavioural phenomena such as the attraction, compromise, similarity effects and choice overload which cannot be explained by the Multinomial Logit (MNL) model or any discrete choice model based on random utility. We show that the seminal concept of revenue-ordered assortment sets, which contain an optimal assortment under the MNL model, can be generalized to the SML model. More precisely, we show that all optimal assortments under the SML are revenue-ordered by level, a natural generalization of revenue-ordered assortments that contains, at most, a quadratic number of assortments. As a corollary, assortment optimization under the SML is polynomial-time solvable On the other hand, the Two-Stage Luce model (2SLM), is a discrete choice model introduced by Echenique and Saito (2018) that generalizes the standard multinomial logit model (MNL). The 2SLM does not satisfy the Independence of Irrelevant Alternatives (IIA) property nor regularity, and to model customer behaviour, each product has an intrinsic utility and uses a dominance relation between products. Given a proposed assortment S, consumers first discard all dominated products in S before using an MNL model on the remaining products. As a result, the model can capture behaviour that cannot be replicated by any discrete choice model based on random utilities.

We show that the assortment problem under the 2SLM is polynomially-solvable. Moreover, it proves that the capacitated assortment optimization problem is NP-hard and presents polynomial-time algorithms for the cases where (1) the dominance relation is attractiveness correlated and (2) its transitive reduction is a forest. The proofs exploit a strong connection between assortments under the 2SLM and independent sets in comparability graphs.

A6. Methodologies, platforms and novel applications

The Relax-and-Solve framework for solving difficult scheduling problems

Amir Salehipour

University of Technology Sydney, Sydney, Australia; amir.salehipour@uts.edu.au

This study develops a novel framework for solving difficult scheduling problems. The framework, which we call the Relax-and-Solve, performs a set of "relax" and "solve" operations, aiming to destruct a solution and re-construct an improved solution. We demonstrate performance of the proposed framework by solving two challenging scheduling problems, namely the aircraft landing problem and the just-in-time job shop scheduling problem. By comparing our Relax-and-Solve method with the state-of-the-art methods we demonstrate both efficiency and the quality of solutions delivered by the Relax-and-Solve method, including obtaining new best solutions. The quality of solutions and the short computation time of the Relax-and-Solve method characterizes it for real-time and online applications. This is very important because due to the typical short time window available for producing real-time and online schedules, effective and efficient algorithms are therefore paramount in order to deliver high quality solutions quickly.

Analysing Fantasy Sport Competitions with Mixed Integer Programming

Steven Edwards

Monash University, Melbourne, Australia; steven.edwards@monash.edu

A fantasy sport competition is an online competition in which participants act as the coach and selector of their own fantasy team of real players. These competitions are remarkably popular with currently over 5.5 million teams in the major competition of the English Premier League. A fantasy team scores points based on the statistical performances of the team's players in their real-world sporting matches. The objective for each coach is to finish the season with the highest total number of points. During the season, coaches must manage a budget as well as trade players in and out of the team subject to a number of constraints. Due to their well-defined nature, as well as the simple objective function, these competitions lend themselves very naturally to analysis by Mixed Integer Programming (MIP). In this paper we consider three different problems for the 2018 season of the AFL SuperCoach competition, modelling and solving each with MIP. The aim of each problem is to highlight the gap between what was achieved by real players and what was theoretically possible. The first problem is to determine all the decisions that a coach should have made to obtain the highest score possible. The second problem is to determine the lowest starting budget from which it would have been possible to win the competition. The third problem is to determine whether it would have been possible for a team that was set up at the start of the competition and completely forgotten about to win the competition.

Tight MIP formulations for bounded length cyclic sequences

Thomas Kalinowski[1], *Tomas Liden*[2] *and Hamish Waterer*[3]

[1]University of New England, Armidale, Australia; tkalinow@une.edu.au
[2]Linkoping University, Norrkoping, Sweden; tomas.liden@liu.se
[3]University of Newcastle, Newcastle, Australia; hamish.waterer@newcastle.edu.au

We study cyclic binary strings with bounds on the lengths of the intervals of consecutive ones and zeros. This is motivated by scheduling problems where such binary strings can be used to represent the state (on/off) of a machine. In this context the bounds correspond to minimum and maximum lengths of on- or off-intervals, and cyclic strings can be used to model periodic schedules. Extending results for non-cyclic strings is not straight forward. We present a non-trivial tight compact extended network flow formulation, as well as valid inequalities in the space of the state and start-up variables some of which are shown to be facet-defining. Applying a result from disjunctive programming, we also convert the extended network flow formulation into an extended formulation over the space of the state and start-up variables.

D2. Data Analytics for Defence Analysis

Towards the Identification and Visualisation of Causal Events to Support the Analysis of Closed-loop Combat Simulations

Dion Grieger[1], *Martin Wong*[2], *Antonio Giardina*[3], *Marco Tamassia*[4], *Torres Luis*[5], *Rajesh Vasa*[6] *and Kon Mouzakis*[7]

[1]Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia;
dion.grieger@dst.defence.gov.au
[2]Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia;
Martin.Wong@dst.defence.gov.au
[3]Deakin University, Applied Artificial Intelligence Institute (A2I2), Geelong, Australia;
antonio.giardina@deakin.edu.au
[4]Deakin University, Applied Artificial Intelligence Institute (A2I2), Geelong, Australia;
marco.tamassia@deakin.edu.au
[5]Deakin University, Applied Artificial Intelligence Institute (A2I2), Geelong, Australia;
luis.torresandrade@deakin.edu.au
[6]Deakin University, Applied Artificial Intelligence Institute (A2I2), Geelong, Australia;
rajesh.vasa@deakin.edu.au
[7]Deakin University, Applied Artificial Intelligence Institute (A2I2), Geelong, Australia;
kon.mouzakis@deakin.edu.au

Analysis of output data from closed-loop combat simulations can provide insights into the relationships between model inputs and outputs. However, analyses that only consider end-of-run output data may not be sufficient to explain why those relationships exist. In the case of stochastic models, where multiple replications of the same scenario are conducted, the presence of outliers and multi-modal results also needs to be accounted for.

Decision Support Tool for Whole-of-Force Assessment

Li Jiang[1] *and Thitima Pitinanondha*[2]

DST, JOAD, CANBERRA, Australia
[1]li.jiang@dst.defence.gov.au
[2]thitima.pitinanondha@dst.defence.gov.au

Capturing data from force assessment activities has long been considered a challenging task due to the volume and variety of the data required for evidence-based analysis and decision-making processes. The pressure of designing the data collection methodology, developing the data collection tools, presenting the collected data in real-time and delivering an analysis result in short-time is a big challenge for the analysts involved. To deal with those challenges effectively, a number of rigorous data collection methods have been studied to explore the best available technologies and tools to provide fast and reliable data collection, processing and visualisation capability. This paper reports experiences and findings in designing a decision support tool for the 2018 Baseline Assessment activity as part of the Defence Capability Assessment Program (DCAP). The rationale for and experience gained during the tool design to enable effective and robust data collection and in-depth analysis of the collected data subsequently are also discussed. The paper recommends: (1) leveraging the strengths of various tools and the latest technologies to support whole-of-force assessment; and (2) completion of a robust data design process before designing and developing a database to collect and process data. In support of these recommendations, the paper proposes a blueprint for a future database that can offer enduring and broader data analytics for force design and analysis, using an advanced database system and Artificial Intelligence (AI) technologies.

Towards An Effective Force Visualisation

Li Jiang[1] *and Sanat Bista*[2]

DST, JOAD, CANBERRA, Australia
[1]li.jiang@dst.defence.gov.au
[2]sanat.bista@dst.defence.gov.au

Force Design is a complex and critical process in Defence. Making good decisions that can address capability issues in the future within budgetary constraints is considered a wicked problem in the wider Defence community. It requires not only sound methodologies but also necessitates integrating considerable information and data across the entire Defence portfolio. Prolific, interdependent and temporal-sensitive data are used in the Defence decision-making process. This calls for dedicated research on new approaches with sufficient scientific rigour that can address the increasing challenges within the force design process. A review of the literature and the status quo of current force design tools and practices has shown the existing visualisation approaches (1) provide only fractional information for particular activities in the force design process; (2) lack sound data design solutions for supporting force design; and (3) provide at best limited back-end database management support. This situation often leads to fragmented data and information being captured within the analysis and visualisation tools, thus offering limited and short-term support for only a few activities in the force design decision making process. Moreover, lack of a sound database design makes it challenging to store data in a format that is fit for purpose and readily accessible to conduct analysis. Based on research on the intricate data sets required by the force design process, this paper proposes a conceptual framework for visualisation that can provide (1) the critical information to support decision-making with increased scientific rigor; and (2) offer advanced data analytics capabilities with effective force visualisation options.

Keynote Seminar

Understanding information and trust: From individuals to populations

Amy Perfors

Computational Cognitive Science Lab, University of Melbourne, Melbourne, Australia;
amy.perfors@unimelb.edu.au

In our increasingly connected and uncertain world, modern warfare requires an understanding of how information flows across social networks and how it can be used (or misused) to shape the behaviour and beliefs of people in those populations. Trust and influence are essential elements for this understanding: whether (and how) they are created, maintained, and destroyed has a large impact on which information is believed and shared. As a computational cognitive scientist, my work addresses these questions using two main tools: controlled experiments with individuals and small groups, combined with mathematical and computational models that predict and analyse how these individual behaviours might scale up to larger populations. In this talk I will describe how this approach is shedding new light on critically important defence issues, especially misinformation and polarisation.

D8. Information warfare analysis

Understanding the Emergent Cognitive Battlespace

Don DeBats[1], *Maryanne Kelton*[2], *Emily Bienvenue*[3], *Zac Rogers*[4] *and Troath Sian*[5]

[1]Adelaide, Flinders University, South Australia, Australia; don.debats@flinders.edu.au
[2]Adelaide, Flinders University, South Australia, Australia;
maryanne.kelton@flinders.edu.au
[3]Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia;
Emily.Bienvenue@dst.defence.gov.au
[4]Adelaide, Flinders University, South Australia, Australia; Zac.Rogers@flinders.edu.au
[5]Adelaide, Flinders University, South Australia, Australia; Sian.Troath@flinders.edu.au

Total warfare in the 1940s mobilized whole societies but total warfare in our times will be waged within societies. Digital technologies flowing through the informational domains, including social media platforms and blockchain networks, have created new channels of influence. The transmission of disinformation through these increasingly powerful and inclusive domains creates uncertainty that erodes societal trust and threatens democratic control. Technological change grows exponentially more powerful but in a non-linear pattern that is characteristically uncertain and cannot be foreseen or predicted. The resulting uncertainty in what can be understood as the cognitive battlespace exposes societies and institutions as never before to both external and internal manipulation. The Future Operating Environment: 2035 recognizes that the greatest disruptions to the strategic landscape for both civilian and military functioning over the next two decades are likely to occur in cyber domains. Advantage in the physical battlespace in which digital systems contend, whether in the immediate or other regions, cannot be assured outside the strategic war of information. This is a war that cannot be won and must be fought. The authors of this paper contend that a fundamental shift is taking place in the nature of the Defence operating environment from the physical battlespace toward the cognitive. In arguing this core point, the authors first present a social science analysis of the emergent non-linear patterns of the cognitive battlespace at the strategic level, the characteristics of which include: the unrestricted and unbound competition below the threshold of the traditional use of force; the continuous presence of targeted trust attacks to delegitimise liberal democratic institutions and democratic civil society and disrupt the liberal international order and; the increasing use of peer networks as primary mediums of disinformation transmission and dislocation of the state in this process. Here we observe that power is now diffused amongst a diversity of actors within network structures from which the state, including both government and military, is displaced from its traditional and more comprehensive hierarchical position. Power no longer flows neatly through vertical institutions but flows increasingly in horizontal networks in which the role of the state is diminishing.

Digital technologies, from the development of the internet to social media and DLTs, are central in any analysis of this issue. Despite common perceptions of disintermediation, organisations sit atop the peer-to-peer networks, controlling large consolidated data storage systems. In this landscape, the western liberal state's ability to exercise power essential for its strategic position is lessened. In the second section of the paper, the initial steps toward the development of a framework of analysis for cognitive warfare are identified, mapped, and visualized. The initial categories of analysis and the integral constituent components discussed in this research include the strategic intent of the adversarial actors and the nature of the disruptive narrative.

Modelling Sharing of Information in the Australian Political Twittersphere

Robert Ackland

Australian National University, Canberra, Australia; robert.ackland@anu.edu.au

There is currently a lot of interest in the spread of information on social media. One strand of research is focused on how 'echo chambers' (the tendency of people to connect with like-minded others) and 'filter bubbles' (algorithms selecting news content based on previous behaviour) can result in people being exposed to information that reinforces previously held beliefs. A second area of research focuses on the role and impact of 'socialbots' — social media accounts programmed to appear to be human-operated — with particular emphasis on measuring impact as contribution to information diffusion. A final related body of work looks at the association between the veracity of information and its diffusion, attempting to establish the patterns of diffusion of true and false news stories. Research into the spread of information on social media involves significant methodological challenges relating to, for example, how to construct social networks over which the speed and extent of spread of information can be measured, and how to develop scalable methods to label network nodes (for example as conservative/liberal) so that political and social dimensions of the spread of information spread on social media can be studied. This presentation summarises ongoing research into the spread of news stories on Twitter. The stories tweeted by three Australian news brands on three randomly sampled days were collected. We then collected all the retweets of these stories over the following seven days, and for those Twitter users who were retweeting politics-related stories (we refer to these as "political retweeters"), we further collected who they followed on Twitter at that time, and we also collected all other retweets they made during this time period. Our dataset therefore contains detailed dynamic data on the production and sharing of news stories on Twitter. The first type of analysis involves mapping the trajectories of news stories as they are shared by political retweeters. Multiple correspondence analysis is used to estimate the political affiliation (right/left) of the political retweeters using information on what federal politicians they followed. We are then able to answer questions such as: how does the type of news story affect the trajectories of its diffusion in terms of political diversity of people sharing the news, breadth (calculated using follower network shortest path information between retweeters) and speed of diffusion? The second type of analysis involves the use economic decision theory to model political retweeting as consumption of information. Tests of revealed preference allow us to identify sets of retweeters who are deemed to share common preferences for consuming information, and this leads to the construction of related index number measures of attention.

Operational design within complex and ambiguous environments

Guy Duczynski

Edith Cowan University, Joondalup WA, Australia; g.duczynski@ecu.edu.au

The complexity and ambiguity associated with military operations has increased over recent years and continues to put pressure on the Joint Military Appreciation Process (JMAP) and other planning methods devised to guide actions in less chaotic environments. Operational design is receiving interest as a companion method to JMAP and invites a more detailed investigation of the area of interest before courses of action are developed and analysed. This paper sets out a software-supported operational design planning method with clear steps and outcomes culminating in an agreed end-state and lines of operation. The method encourages planners and others to comprehend more thoroughly the interactions within the area of interest and devise purposeful actions that progressively reset the conditions to achieve the end-state.

Biased Information Transmission and Proliferation

nicolas fay[1], *Amy Perfors*[2] *and Bradley Walker*[3]

[1]University of Western Australia, Perth, Australia; nicolas.fay@uwa.edu.au
[2]The University of Melbourne, Melbourne, Australia; amy.perfors@unimelb.edu.au
[3]University of Western Australia, Perth, Australia; bradley.walker@uwa.edu.au

Knowledge is often acquired through information that is passed from person to person, as opposed to via first-hand experience. This type of information acquisition is known as social learning (Bandura, 1977). However, not all information is created equally — human cognitive biases lead to a preference for particular types of information. Evolutionary theory suggests that, because of its greater survival value, negative information (e.g., location of a predator) will be preferred over positive information (e.g., location of a food source) (Nettle, 2018). A human preference for negative information was demonstrated in a recent large-scale experiment run by colleagues and I (Bebbington, MacLeod, Ellison, and Fay, 2017). The present paper reports a similarly large-scale experiment ($N = 425$) that examines the mechanisms driving the preferential transmission and survival of negative information. Participants read a text that contained positive and negative information, plus ambiguous information that could be resolved either positively or negatively. This information was then communicated to a second person, and this process was repeated across four persons to simulate multiple information transmission episodes. Sociality was manipulated to determine the extent to which a preference to transmit negative information arises due to an individual-level cognitive bias, or due to the interplay between a cognitive bias and the social-interactional processes inherent to typical human-to-human information exchange. Consistent with Bebbington et al. (2017), our findings supported a negativity bias: (unambiguous) negative information survived multiple transmission episodes better than (unambiguous) positive information, and, when ambiguous information was resolved, negative resolutions were more common than positive resolutions. In addition, the social context modulated information transmission: information fidelity decreased as sociality increased. Furthermore, highly interactive social exchange attenuated the negativity bias relative to the less social conditions, although a preference for negative information remained. In conclusion, negative information is more likely to be retained and transmitted than positive information. We speculate that this negativity bias is adaptive, because it prioritizes costlier hazard avoidance over reward. Hence, negative information, on the whole, is more valuable. Our findings represent a crucial first step in understanding why particular types of information proliferate (e.g., on social media). In addition, the experimental paradigm developed can be used to test the most effective way to gain social influence, and how best to guard against information warfare.

A7. Transport and logistics

Capacity Alignment Planning for a Coal Chain: A Case Study

Saman Eskandarzadeh[1], *Thomas Kalinowski*[2] *and Hamish Waterer*[3]

[1]The University of Newcastle, Callaghan, Australia;
saman.eskandarzadeh@newcastle.edu.au
[2]University of New England, Armidale, Australia; tkalinow@une.edu.au
[3]The University of Newcastle, Newcastle, Australia; hamish.waterer@newcastle.edu.au

We study a capacity alignment planning problem for a coal chain. Given a set of train operators, a set of train paths, and a terminal comprising of a dump station and a set of routes from the dump station to the stockyard, we seek a feasi- ble assignment of train operators to train paths, to time slots at the dump station, and to routes. The assignment must maximize the number of system paths in the resulting schedule and the schedule should perform well with respect to various performance criteria. We model the problem as a mixed-integer conic programme (MICP) with multiple objectives which we solve using a hierarchical optimization procedure. In each stage of this procedure we solve a single objective MICP. De- pending upon whether we evaluate the associated performance criteria under a 2- or 1-norm we reformulate the MICP as either a mixed-integer second-order cone programme or as a mixed-integer linear programme respectively, and can stream- line the hierarchical optimization procedure by exploiting properties of the model or observed behaviour on practical instances. We compare the performance of the procedure under the different norms on a real instance of the problem and find that the quality of the solutions found by the faster 1-norm procedure compare well to the solution found under the 2-norm.

On the importance of Stock Reallocation in Service Parts Resupply Network: An Approximate Policy

Aghil Rezaei[1], *Sobhan Asian*[2] *and Ramtin Rabiee*[3]

[1];
[2]La Trobe University, Melbourne, Australia; s.asian@latrobe.edu.au
[3]Nanyang Technological University, Singapore, Singapore; rrabiee@ntu.edu.sg

Over recent years, increased competition in the airline industry and higher cost of labor and fuel expenses have put extra pressure on this sector. In such a capital-intensive industry, it is critical to employ flexible inventory control policies that can efficiently manage the on-hand inventory within the network, so the system can operate with minimum delay and operation costs. Stock reallocation policy can improve the performance of repairable service parts inventory systems by enabling the movement of service parts among several holding locations or bases. Comparing the costs of stock reallocation with that of backorder, this policy can potentially lead to significant cost reduction specially for industries with expensive service parts and crucial uptime. This paper studies a repairable service parts inventory system of an MRO with a central repair facility and several bases, where stock reallocation policy is used to transship a service part from a base with sufficient stock to another base with an urgent need for inventory. We use Markov decision process to model the stock reallocation problem which consists of determining the optimal time of a lateral transshipment and identifying the optimal origin and destination bases. The optimal stock reallocation problem was formulated as a dynamic program. Since using the stock allocation policy results in a state-space collapse, the multi-dimensional queue length process can be approximated through a single queue system. Hence, an approximation is proposed in which, the relative value function of a K base model is approximated by the sum of the relative value functions of the aggregate queue and individual single queues. In other words, the K base inventory system is treated as K + 1 independent queues, with a queue representing the aggregate behavior of the inventory system and K queues representing the behavior of individual bases. We conduct intensive numerical experiments to evaluate the accuracy of our proposed solution method. It is observed that including stock reallocation capability in the inventory system results in a 4.871% cost reduction. Additionally, the approximate stock reallocation policy performed very close to the optimal policy with maximum and average optimality gaps, of 0.682% and 0.090%, respectively. In addition to the performance evaluation of the proposed stock reallocation policy, a sensitivity analysis is performed to better characterize the effects of varying the key parameters. Following observations were made in the sensitivity analysis: 1) As the demand ratio increases, the optimality gap of the approximate stock reallocation policy increases, and the cost reduction amount decreases. 2) as server utilization increases, the cost reduction amount decreases. 3) As the total initial stock increases, the number of parts available for inventory pooling between bases increases, and hence, increasing the base-stock level magnifies cost reduction.

4) As the stock reallocation cost and rate increases, cost reduction due to using the reallocation policy decreases and increases, respectively. 5) Increasing the number of bases in the resupply network leads to increased cost reduction, which results from using the stock reallocation policy.

The fixed route vehicle-refueling problem with variable fuel consumption rate

Juan Calle[1], *Vicky Mak-Hau*[2] *and Roberto Perez-Franco*[3]

[1]Deakin University, Centre for Supply Chain and Logistics, Melbourne, Australia;
jcallesalazar@deakin.edu.au
[2]Deakin University, School of Info Technology, Melbourne, Australia;
vicky.mak@deakin.edu.au
[3]Deakin University, Centre for Supply Chain and Logistics, Melbourne, Australia;
roberto.perezfranco@deakin.edu.au

Operations research techniques have been widely used to solve planning problems in the transportation industry. The range of the problems varies from the operational, tactical and strategical level. This works tackle an operational planning problem in the trucking industry such as is the problem of deciding where and how much to refuel a truck along a route (Gorman, Clarke et al. 2014). We divide the literature that has tackled the problem of optimizing refueling process in the trucking industry into two flows. One flow of work has studied the problem of integrally deciding the route and the stops to refuel. This problem appears in the Less than Truck Load context. The basic idea behind is that heavy items that will be delivered should be dropped at the beginning of the route and that lighter items should be delivered at the end of the route. This problem is known as the traveling salesman with refueling (TSPR) (Suzuki 2012). The other flow concerns to the Truck Load industry, in which the route is being already defined from a specific origin to a destination and the amount of load carried by truck is constant. The key behind this is to avoid truck stops that are far from the route and choose stops that have good prices. This problem is called the fixed route vehicle-refueling problem (FRVRP) (Suzuki and Lan 2018). Our work concerns this second flow of work. However, the contribution we make could still be useful for scholars working in the first flow, as one approach to solve the TSPTW is to make the routing decision first and then the refueling decisions. Although the first software to solve the FRVRP in the industry appeared in the 90's, the first reports of scholars who have worked in this problem date just from 2007 (Suzuki 2008). Since then, there have been almost a dozen articles dealing with this topic. The most common approach to solve this problem has been to model it as a mixed integer linear program (MILP), in which the fuel consumption rate is constant along the route and the distance for going to the service station is the same as the distance for coming back to the main route. Our purposed model enhances the standard formulation in two aspects: one is that we add the concept of segments to the standard model in order to represent the different fuel consumption rates. The second is that we add the concept of branching nodes for representing the asymmetry of cost to and from a refueling stop, which means that the distance and fuel consumption for going to one stop may be different from the distance and fuel consumption to return to the main route. We test the model in instances of different sizes to determine the computational time to solve it.

Cyclic Freight Train Planning; Integrated Problems

Ata Elmi[1], *Andreas Ernst*[2] *and Mohan Krishnamoorthy*[3]

[1]Monash University, School of Mathematical Sciences, Melbourne, Australia;
Atabak.Elmi@monash.edu
[2]Monash University, School of Mathematical Sciences, Melbourne, Australia;
andreas.ernst@monash.edu
[3]The University of Queensland, School of IT and Electrical Engineering, Brisbane, Australia;
m.krishnamoorthy@uq.edu.au

The railway is known to be the best mode of land transport in terms of energy consumption and land use per ton - km transported; and also in terms of economic efficiency for freight transportation. Thus, making an optimal transportation plan becomes an important issue for railway networks. This research studies the freight train planning problems for Australian transportation companies. The commonly adopted planning process by railway freight transportation companies in Australia are as follows:

- Timetabling: The amount of freight for each service demand is known, and thus the timetabling determines the number of required train itineraries. The general aim of the timetabling problem is to construct a train schedule that matches the required number of itineraries while satisfying the service demands delivering time-windows.
- Rolling stock assignment: There are multiple types of wagons and locomotives utilised in railway freight transportation. The rolling stock assignment aims to determine the wagons and locomotives required to cover the transportation of service demands for each itinerary due to the physical restrictions and capacities of railway tracks.
- Rolling stock balancing: The main aim is guaranteeing the rolling stock rotation and their availabilities for the timetabled itineraries. Empty wagons and inactive locomotives have to be moved to ensure rolling stock availability because of a limited number of resources while considering terminal capacities.

The cyclic freight train planning problem considered in this research is modelled based on the network flow structure. The proposed mathematical model contains a multi-layer network structure. The first set of layers are regarding freight services, and each freight service has to flow through the network from the departure node to the arrival node. The other sets of layers are ensuring the flow of rolling stocks through the network to satisfy the requirements of train itineraries. Thereupon, the second set of layers are regarding the container wagons. There are different types of wagons regarding weight, length, loading capacity and double stacking capabilities.

The second set of layers guarantee the flow for each type of wagons separately. The third set of layers are for the locomotives, and as there are different types with differ- ent hauling capacities on different tracks, therefore each layer within the third set en- sure the flow of each locomotive type in the network due to the hauling requirements for the itineraries. The freight services are forced to flow due to the delivery demand within the provided time-window for each service. Besides, the rolling stocks flow are in direct relation with the freight services flows in the network. The considered problem in this research is purposed to integrate the timetabling, assignment and balancing concepts of tactical and operational level problems in train planning. The proposed mathematical model is applied to real-life data set in the Australian con- text. The proposed approach efficiently decreased the number of train itineraries. Moreover, the number of active rolling stock resources decreased due to efficient assignments and balancing.

D10. Defence Decision Analysis

The Wheels vs Tracks Problem for Armoured Fighting Vehicles in the Australian Context

Nikoleta Tomecko[1] *and Kasia Krysiak*[2]

[1]Defence Science and Technology Group, Edinburgh, Australia;
nikoleta.tomecko@dst.defence.gov.au
[2]Defence Science and Technology Group, Fishermans Bend, Australia;
Kasia.Krysiak@dst.defence.gov.au

In armoured fighting vehicle design the Iron Triangle concept describes the design tensions that exist between the three primary characteristics of these vehicles: Mobility, Survivability and Lethality. Traditionally, wheels and tracks represent two different trade-off instances between different aspects of these three factors and are suited to different operational conditions. To provide some clarity to the wheels vs tracks argument for the ADF, a wheels vs tracks study was undertaken in the Australian context. This study collated results of previous studies and performed a meta-analysis, synthesising the results to produce an understanding of the impacts of wheels and tracks on operational outcomes, analysing the current evidence of the strengths and weaknesses of wheels and tracks, and interprets these in different contexts characterised by environmental and operational variables. The results of the meta-analysis show that overall a tracked vehicle will offer a greater operational capability advantage more often. In particular:

- Out of the 72 different contexts defined, 62 show an operational advantage for tracked vehicles. Only 9 contexts had an overall utility skewed towards a wheeled vehicle, and in one context wheeled and tracked vehicles were judged as equal.
- The analysis identified 14 contexts with an intensity rating of extreme, and in all of those contexts, tracked vehicles had an operational advantage over wheeled vehicles.
- In 15 of the 20 contexts judged to be most likely, tracked vehicles had an operational advantage over wheeled vehicles, while the remaining five showed an operational advantage for wheeled vehicles

A Tool for Management and Measurement of VfM Indicators

Minh Nguyen[1], *Anh Pham*[2], *Richard Taylor*[3], *Yue-Jin Wang*[4] *and Jagiello Jerzy*[5]

[1]Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia;
minh-tuan.nguyen@dst.defence.gov.au
[2]Australian Department of Defence, ACT, Australia;
[3] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
richard.taylor@dst.defence.gov.au
[4] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
Yue-Jin.Wang@dst.defence.gov.au
[5] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
Jerzy.Jagiello@dst.defence.gov.au

One of the key deliverables for the newly created Contestability Division of the Australian Department of Defence (DoD) is to take a leading role in helping to identify systemic, recurring issues for resolution, and in developing standards, tools and models to support project contestability and to ensure Defence makes the best decisions. The Australian National Audit Office (ANAO) has issued a number of reports recommending DoD assess the Value for Money (VFM) for projects (see e.g. (ANAO 2009, 2016)). Accordingly, Contestability Division has proposed a method to analyse VFM (Pham Waddell 2017, 2018) and collaborated with the Joint Force Analysis (JFA) team, Defence Science and Technology, to develop a tool for managing and measuring the VFM indicators. The notion of VFM is associated with the concepts of whole life service value provided by project capabilities. Common criteria to assess the VFM are Economy, Efficiency, Effectiveness and Equity. Good VFM is the optimal use of resources to achieve the intended outcomes by spending less, spending well, spending wisely and spending fairly (National Audit Office 2010, Department for International Development 2011). Our developed VFM tool has evolved from a simple scoring scheme to Multi-Criteria Decision Making (MCDM) covering all aspects of the four criteria above. The VFM structure has three hierarchical levels arranged in major criteria and sub-criteria. The evaluation results (detailed and summarised) are presented in both numerical tabular and graphical form where calculation and assumption details are hyper-linked for transparency and back tracking. By applying MCDM, we evaluate the influence of multiple criteria by calculating a utility value where each criterion takes a certain weight. Both the scoring/rating method and pairwise comparison method (see e.g. (Nguyen 2003) and references therein) are used in this work and implemented in our decision support tool. The pairwise comparison method is preferred for use because it allows measurement of evaluation consistency and consensus in group decision making. Two pairwise comparison methods included in the VFM tool are the Analytic Hierarchy Process (AHP) (Saaty 1980) and Best-Worst Method (BWM) (Rezaei 2015).

Multi-level sensitivity analysis can be done directly from the tool to see any impact of weighting value changes. In summary, our developed tool, organises and presents the VFM indicators as MCDM problems. It can conveniently examine various proposed project options and be compared and visualised in both tabular and graphical form. Contestability Division has used the tool in their pilot studies. Results from the tool compared with the current assessment methods indicate that the MCDM evaluation approach and the implemented VFM tool are very promising and useful. The proposed model provides assessment structures which assist in making explicit Decision Maker's preferences as well as providing a tool for automating and documenting some of the currently time-consuming evaluation processes.

Strategic Risk Management in Practice

Filinkov No[1]*, Terence Weir*[2]*, Steven Lord*[3]*and Hossein Seif*[4]

[1]DST, JOAD, Canberra, ; Alexei.Filinkov@dst.defence.gov.au
[2]DST Group, Canberra, Australia; terence.weir@dst.defence.gov.au
[3]DST Group, Canberra, Australia; Steven.Lor@dst.defence.gov.au
[4]DST, JOAD, Canberra, ;

Contemporary risk management methodologies are typically used for identification and prioritisation of strategic risks. The International Risk Management standard, ISO 31000:2009, is the world-wide basis for best practice in strategic level risk processes. However, due to the qualitative and subjective nature of strategic risk, its analysis requires a more nuanced approach than that used in more tactical or operational settings and this paper discusses the need to understand the range and nature of strategic threats, and how to represent risk assessments. As such, a particular focus of this work is on how to incorporate best practices in strategic risk analysis, and operations research into the design and application of strategic risk management in the Defence context. A number of steps are recommended incorporating international risk management best practices within the context and uncertainties unique to strategic risk management for Defence (as opposed to tactical or engineering risk management).

A6. Methodologies, platforms and novel applications

DEVELOPMENT OF A FUNCTIONAL-STRUCTURAL GROWTH MODEL FOR Thalassia hemprichii using GAMA

Therese Rollan[1], *Edgardo Macatulad*[2], *Leorey Marquez*[3], *Rene Rollon*[4] *and Shaira Mia*[5]

[1]Diliman, University of the Philippines, Quezon City, The Philippines;
[2]Diliman, University of the Philippines, Quezon City, The Philippines;
egmacatulad@up.edu.ph
[3]CSIRO Data61, Clayton, The Philippines; leorey.marquez@csiro.au
[4]Diliman, University of the Philippines, Quezon City, The Philippines;
rrollon@iesm.upd.edu.ph
[5]Diliman, University of the Philippines, Australia;

Seagrass are flowering plants submerged in shallow marine waters, usually located on semi-enclosed lagoons and along coastlines, and co-existing with intertidal mangroves and corals. Seagrass play an important role in providing food and shelter for various marine species, stabilizing the sea bottom, maintaining water quality and supporting the livelihood of local economies. These also help terrestrial life through it's massive harness of carbon which is more than thirty times the capacity of terrestrial greenery. Thalassia hemprichii is one of the dominant species of seagrass in the Philippines which grows primarily via asexual reproduction (clonal) and exhibits horizontal expansion typically in a span of 11 days. The Philippines has over 36,000 km of coastline with large portions surrounded by seagrass beds which are located mostly in the northwestern, western and southern parts of the country. However, this vast coverage has been decreasing in the past several years due mostly to anthropogenic factors. The major long-term threats to seagrass habitat include coastal eutrophication as a result of overfeeding in fish cages, effluent from industrial, commercial and domestic facilities, inadequate septic systems, human wastes, and storm drain runoff carrying organic waste and fertilizers. In addition, extreme events such as typhoons damage, and in some cases completely wipe out, wide areas of seagrass beds. To address this problem, a study is being conducted that aims to identify and understand the various factors and processes involved in the growth and/or decline of seagrass in order to analyze and understand its natural behavior and estimate its reaction to various external factors. This study will assist local communities, government authorities, environmentalists, ecologists and other researchers to formulate effective strategies for seagrass protection, rehabilitation and conservation. This paper describes the development of a planning support and scenario evaluation tool that performs geosimulations of seagrass (Thalassia hemprichii) growth at specified sites in Bolinao, Pangasinan, Philippines.

The tool is implemented using the GAMA (GIS Agent-based Modeling Architecture) Framework which, in the final complete version, will provide a user interface that will feature user-controlled action panels, multi-layer displays, statistical reporting and agent inspection. The geosimulation uses agents to represent the main components of seagrass growth, namely the apical meristem or apex, the internode and the shoot. The interactions between agents and their environment are governed by complex equations, constraints, and thresholds using species-specific parameters obtained from intensive field studies. These parameters include apex density, plastochrone interval (number of days within which internode is produced), horizontal elongation rate, branching rate, horizontal internodes between shoots, shoot spacing along rhizome and median maximum age of shoot and internode. The geosimulation conducts growth using a time step of 25 days which represents the duration of internode growth, and a threshold of 58 apices denoting the maximum number in a 1 square meter plot. A simulation run starts with randomly distributed pairs of apex and internodes in a number of specified plots and terminates with meadows of seagrass covering wide areas. The paper presents initial results from the study and discusses future work to improve the model.

Exposing performance differences between branching strategies using evolutionary generation

Simon Bowly

The University of Melbourne, Melbourne, Australia; sbowly@student.unimelb.edu.au

Recent work in the optimisation literature has focused on both online and offline learning methods to improve Mixed Integer Programming (MIP) solver performance. Particular emphasis has been placed on learning efficient branching rules and node selection strategies [1, 2]. The proliferation of these and similar machine learning-based methods underscores the importance of curating data sets which effectively challenge a range of strategies. While it is preferable to benchmark new algorithm developments using well-studied instances, the focus of reported results is often on comparisons with the current state of the art, as opposed to identifying the underlying reasons for improved performance [3]. Benchmark data sets commonly used for algorithm evaluation may not contain the requisite diversity to effectively demonstrate the strengths and weaknesses of a new solution strategy. As such, we propose using methods for generating new targeted instances to gain confidence in an algorithm's performance on future problems. In this work a multiobjective evolutionary algorithm is used to generate a spectrum of MIP instances which expose performance differences between different branching strategies. Similar single-objective approaches have been applied in the past to generate test cases which challenge heuristics [4] or have specific properties [5]. The method maintains diversity across a large population of instances, producing a benchmark set suitable for exploratory analysis. The results are focused on producing instances which vary the tree size required to prove optimality for multiple branching strategies implemented in the SCIP solver [6]. Using this method to search the problem space allows small, fast-to-solve instances to be generated which illustrate performance variation between branching rules. In comparison with existing benchmark sets, commonly composed of very large problems which require significant time to solve and are hard to analyse, this data may prove useful for gaining insights and further tuning of MIP algorithms.

Solver Independent Rotating Workforce Scheduling

Nysret Musliu[1], *Andreas Schutt*[2] *and Peter Stuckey*[3]

[1]Vienna University of Technology, Vienna, Austria; musliu@dbai.tuwien.ac.at
[2]The University of Melbourne, Melbourne, Australia; andreas.schutt@data61.csiro.au
[3]CSIRO, Data61, Melbourne, Australia; peter.stuckey@monash.edu

Rotating workforce scheduling is a specific personnel scheduling problem arising in many spheres of life such as, e.g., industrial plants, hospitals, public institutions, and airline companies. The aim is to schedule workers, so that workforce requirements for shifts and days, minimal and maximal length of shifts, and shift transition constraints are met. The schedule horizon is split over a number of weeks equal to the number of workers. Each day of the week has a number of shifts requiring a pre-defined number of workers on duty. For each worker, one has to decide whether they are on duty on one shift or have an off duty day. The schedule is rotating, which means that the first, second, ..., second last, last worker cover the first, second, ..., second last, last week of the schedule in the first week and then the second, third, ..., last, first week in the second week and so on. Finding a feasible solution for the rotating workforce scheduling problem is NP-complete. The state-of-the-art method is a heuristic approach based on min-conflicts heuristic and tabu search (MC-T), which is also deployed in a commercial software FCS. MC-T even outperforms the best complete solution approach, which uses the Satisfiability Modulo Theory (SMT) solver MathSAT. We investigate different models for the rotating workforce scheduling problem using the high-level modelling language MiniZinc. For instance, in one of our models we directly express the constraints in linear constraints as one would do for Mixed-Integer Programming (MIP). In another model, we use the global constraints "gcc" and "regular" for modelling the problem in a Constraint Programming (CP) way. In addition, we look at the impact on redundant and symmetry breaking constraints. Since MiniZinc is solver-independent as AMPL, we solved the same models with the MIP solver Gurobi and the CP solver Chuffed in order to find the best model for each solver and to compare them. To our surprise, Gurobi performed at best on the CP-like model mainly due to a transformation of the "regular" constraint to network flow, whereas Chuffed performed at best on the MIP-like model due to exposure of important literals to learn on. We empirically compared Gurobi and Chuffed using their best model to MC-T and MathSAT approach on the standard benchmark set and newly generated benchmark set. The results are consistent on both sets. Chuffed and Gurobi outperformed the MathSAT and MC-T in terms of average runtime and the number of solved problems, i.e., finding a feasible solution or proving the infeasibility of the problem. Thus, they improve the state-of-the-art for solving rotating workforce scheduling. Chuffed was the best for feasible problems, whereas Gurobi the best for infeasible ones. Overall, Gurobi was the most robust approach. Note that this line of work was published at the 15th International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR 2018) (see https://link.springer.com/chapter/10.1007/978-3-319-93031-2_31).

Integer programming approaches to course schedule assembly and student allocation problem in helicopter aircrew training

Vicky Mak-Hau[1], *Brendan Hill*[2], *David Kirszenblat*[3], *Bill Moran*[4], *Nguyen Vivian*[5] *and Ana Novak*[6]

[1]Deakin University, School of Information Technology, Melbourne, Australia;
vicky.mak@deakin.edu.au
[2]Deakin University, School of Information Technology, Melbourne, Australia;
brendanh@deakin.edu.au
[3]Parkville, University of Melbourne, Department of Mathematics and Statistics, Victoria, Australia;
d.Kirszenblat@unimelb.edu.au
[4]University of Melbourne, School of Engineering, Parkville, Australia;
wmoran@unimelb.edu.au
[5] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
Vivian.Nguyen@dst.defence.org.au
[6] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
ana.novak@dst.defence.gov.au

In this paper, we study a timetabling problem that arose from helicopter aircrew training for the Royal Australian Navy. A student (trainee pilot) has to complete a syllabus, i.e., a sequence of courses. A course is equivalent to a subject or a unit in tertiary education. The courses have a pre-requisite structure. Each course is associated with a pass rate, obtained from historical data, and a number of repeated sessions spanning the same duration, but occupying different (and possibly overlapping) time slots. A feasible schedule is a sequence of course sessions such that each course in the syllabus is covered by exactly one session, and that all pre-requisites requirements are observed. The optimisation problem is to assemble course sessions into feasible schedules and allocate students to these schedules such that the class size restrictions for each course session is not exceeded. The objective is to minimize the sum of time-span in completing the syllabus over all students. We have developed and numerically tested a number of integer linear programming (ILP) formulations for the combinatorial optimisation problem described above. The first two formulations are based on the model where a directed graph is used with vertices representing courses and arcs representing connections of the courses. An optimal solution will construct Hamiltonian paths on the digraph, each representing an ordered set of the courses to be taken hence (in other words, a schedule). The first ILP model uses binary variables to create schedules. There will be multiple copies of the same schedule with different indices, but the number of schedules is precisely the number of students, hence no upper bounding on the number of schedules is required. The second ILP model uses general variables to represent the number of students assigned to each schedule. However, upper bounding on the number of schedules as well as phantom schedules and course sessions are required.

The objective function is quadratic, and therefore a two-stage method is proposed as a heuristic. We then present a third integer programming model that uses a binary variable for each possible schedule, although these schedules must be generated in advance. Thus, we propose two methodologies for generating these feasible schedules: 1) a Revised Knuth's Dancing Link Algorithm, and 2) column generation. Our column generation approach is heuristic: we generate all columns for the root node linear programming relaxation, and use these columns to solve the integer program to obtain a feasible solution. In most cases we have tested, however, the root node linear programming relaxation produced naturally integral solutions, and is therefore optimal to the integer program. We also propose a branch-and-price approach for exact solutions.

It depends - how to kill a project with data

Melanie Ayre[1], *Simon Dunstall*[2], *Reena Kapoor*[3] *and Andreas Schutt*[4]

CSIRO Data61, Docklands, Australia
[1]melanie.ayre@data61.csiro.au
[2]simon.dunstall@data61.csiro.au
[3]reena.kapoor@data61.csiro.au
[4]andreas.schutt@data61.csiro.au

Manufacturing is once more a growth sector in OECD economies. 20th century trends in off-shoring to lower-labour-cost countries have eased, with re-shoring increasingly popular. While the main drivers for this are in quality standards and bespoke products, the adoption of digital technologies that reduce the total cost of production while enabling customisation will be critical to its continuation. While there are a number of industries that have fully-automated production, there are many that either need a gradual transition through collaborative automation ('co-bots') or where 'hand-made' is a critical value proposition. In these, digitalisation often means decision-support, replacing tacit-knowledge-derived decision-making with a process that uses forecast evaluation and situational awareness data to provide greater control and transparency. Drawing on our experience with a number of manufacturing businesses who have sought to explore development and deployment of such decision support tools in planning and scheduling in manual and mixed build processes, we will describe the fundamental issue that affects progress in such projects — which is codifying the tacit knowledge. Actual operational-level data is often spread across multiple access-controlled systems, and being able to extract it, correlate it, and know what to trust and act on in real time is a skill in itself. We will present a case of a dynamic re-scheduling system, designed to provide decision support in disaster recovery scenarios for customised production. At first glance, this seems to be a relatively straightforward combinatorial optimisation problem, albeit multi-objective and with numerous side-constraints, hence challenging to solve. However, the real problem arises when trying to validate solutions with respect to real production operations. Creating a tractable version of the problem requires defining constraints, whereas design to maintain throughput requires keeping options open. The answer to every data-led question becomes 'It depends.'

D10. Defence Decision Analysis

Weighting Methods for Multi-Criteria Decision Analysis with Application to a Defence Prioritising Problem: A Comparative Study

Yue-Jin Wang

Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia;
Yue-Jin.Wang@dst.defence.gov.au

Multi-Criteria Decision Analysis (MCDA) deals with prioritising, ranking or selecting alternatives based on human judgement from a finite set of alternatives in terms of multiple, usually conflicting criteria. In Defence applications, an alternative can be an issue, scenario, option, strategy that is under consideration. In the majority of MCDA models, deciding on the criteria and assigning their weights play a significant role. Once the weight of each criterion is decided, the MCDA aggregated scores are calculated and the rankings of alternatives are determined. Various weighting methods have been proposed in the literature and applied for solving different MCDA problems. Among them, direct ranking, ranking order derived weighting, and pairwise comparison methods are some popular weighting methods. Each weighting method, however, differs in terms of accuracy, complexity and theoretical foundations, and produces different sets of weights. This paper presents a comparative study that aims to prioritise issues in a real example of a Defence decision analysis problem. The MCDA formulation consists of four criteria and fourteen issues. Numerical evaluations of the issues against each criterion were identified by twenty stakeholders using a "point-allocation" tool individually. Direct ranking and the Analytical Hierarchy Processes (AHP) method were applied to eliciting the weights of criterion importance. Based on the data collected, several criteria weight determination methods are tested, including the direct ranking (DR), the reciprocal rating (RR), the rank-order centroid (ROC), AHP and Best-Worst Modelling (BWM). The influences of the selection of weighting methods on the final prioritisation results are analysed and the advantages and limitations of those popular weighting methods are further discussed. For the selected MCDA example, it shows that RR weighting, ROC weighting, and AHP or BWM weighting schemes produce identical top five ranking orders for the fourteen issues, while the DR weighting gives quite different ranking results. The results also indicate that the normalised weights of the most important criteria from the DR weighting tend to be lower than the weights from AHP or BWM. An overlap between definitions of criteria might be a possible explanation.

Autonomous Command and Control of a Multi-UAV Asset Search

Slava Shekh[1] *and Glennn Moy*[2]

[1]DST Group, Edinburgh, Australia; Slava.Shekh@dst.defence.gov.au
[2]DST Group, Brisbane, Australia; glennn.moy@dst.defence.gov.au

Recent advances in platform autonomy are increasingly having an impact on military capability. One consequence of the increased availability and reduced cost of tactical autonomous systems is the increasing number of platforms that need to be managed in the modern battlespace. Decision support systems can assist in reducing the cognitive load on decision makers through high-level recommendations combined with autonomous decision making. We consider a specific problem in high-level Command and Control — using multiple UAVs to search for a threat. In general, the most effective UAV search strategies are a complex function of environmental factors, friendly capabilities, threat properties and cost. We explore a simplified version of this problem to understand the performance of a range of techniques in this problem space. In our simplified problem, a series of static blue agents must be positioned in four grid locations in order to intercept red agents that are attempting to reach a goal location. Each blue algorithm aims to maximise the chance of detecting red agents, given specified red start and goal probability distributions. We explore three broad classes of algorithm — simulation algorithms, path analysis and topography analysis. Simulation techniques are based on running multiple simulations of the underlying scenario, parameterised by the problem state information. In subsequent runs, each simulation approach employs a different method for choosing (based on earlier results) where to place blue assets. Path analysis algorithms, by contrast, use a probability distribution representing likely red paths to directly estimate the frequency of particular cells occurring in the search space. This is a more efficient technique for obtaining optimal blue behaviour in situations where the red paths can be effectively estimated, and where there is little dependence between red and blue behaviours. A third class of approaches identifies blue behaviour based on the topography of the space, and focuses on the identification of choke points to recommend blue agent search locations. Topography analysis can be a very efficient approach, but achieves this efficiency at the expense of ignoring all aspects of the red behaviour other than expected start and goal locations. We compare the performance of these three classes of algorithm across different scenarios. We find that topography analysis outperforms all others for problem domains where it is applicable. The path analysis approach provides the best alternative where topography analysis cannot be used, as long as there is sufficient computational time budget. If there is limited computation time, bandit-based and evolutionary approaches provide good (but often non-optimal) solutions. In future, we aim to consider more complex problem domains in which red agents cannot be defeated by simply placing blue agents in fixed locations. In domains where red agents continuously adapt their behaviour, there is a need for more advanced blue detection approaches, which also adapt and learn through repeated interaction with red agents.

In this regard, machine learning techniques and self-play reinforcement learning, where blue agents can find weaknesses and iteratively improve their own strategy, are techniques worth considering in the domain of high-level Command and Control.

How to setup successful military-university collaborations: Initial research insights from working with the Australian Defence Forces

Matthias Guertler[1] *and Giti Datt*[2]

University of Technology Sydney, Sydney, Australia
[1]matthias.guertler@uts.edu.au
[2]giti.datt@uts.edu.au

Defence forces around the globe face continuously increasing challenges and uncertainties resulting from fast technological advancements, demographic changes, global politics and asymmetric threats (Tomes, 2004; Australian Government, 2016). Military-university collaboration (MUC) can help defence to successfully tackle those. MUC allows access to latest research insights and expertise, including novel technologies, processes and tools to address various socio-technical issues (Hernandez et al., 2017). Thus, there has been an on-going call for increased MUCs (e.g. Australian Government, 2016). However, initiating and managing MUCs are not trivial. Often tacit differences in culture, values and interests can cause conflicting needs and requirement, such as protecting sensitive knowledge versus scientifically publishing project outcomes. If insufficiently considered, they can result in limited outcomes or even failures of MUCs. Therefore, MUCs need to be carefully and systematically planned (Hernandez et al., 2017). Despite this relevance, research on MUC is limited (Hernandez et al., 2017) and focuses mainly on military education, mental health and medical collaborations. Only a few articles address the planning, initiation and management of MUCs, such as MUCs to fight environmental disasters, building a knowledge management system for MUC (Burita et al., 2015), or MUCs to improve mental health services (Hernandez et al., 2017). To achieve the overarching goal of systematic planning and managing of MUCs, the resulting research question is: What are success factors of initiating and managing university-military collaborations? To answer this question, we retrospectively analysed research projects between the Royal Australian Navy and the University of Technology Sydney. Based on the interviews with our Navy contacts, we developed an initial MUC framework. It was evaluated through feedback from our Navy contacts and contacts of other services to avoid a Navy bias. Our MUC framework comprises concrete success factors and recommendations for planning, initiating and managing MUCs. These 12 factors are structured in three clusters: context (e.g. overarching objectives, organisational constraints), culture (e.g. organisational culture, traditions), and project management (e.g. team continuity, communication). This includes knowledge protection tools like NDAs or review opportunities before submitting academic publications, definition of communication frequency or on-boarding of new team members. The framework addresses a critical research gap and sets the basis for future research that explores and develops practical guidelines, processes and tools to efficiently plan, initiate and manage MUCs. Along with better project outcomes, it can strengthen national eco-systems combining different sectors, such as academia, defence, industry and public sector. In this respect, this research also contributes to a better understanding of cross-sector collaboration in general.

A security-based petroleum trading model

Greg Calbert

Defence Science and Technology, Joint and Operations Analysis Division, Edinburgh, Australia;
calbertgregy@gmail.com

Without liquid petroleum, either jet fuel, diesel, or other products such as lubricants and fuel oil, the Defence force ceases to function. Furthermore, the national support base, the bed-rock of the Defence force stops as well. Petroleum products will continue to form an energy source of choice for Defence because of superior energy density for decades to come. While Australia is seen as a regional power, its energy resilience is in a state of change and is generally seen as declining. Currently the Government sees market forces as providing petroleum supply security, given the nation stores 50-55 days of stocks which is below the mandatory 90 days required by the International Energy Agency. With this context, Defence Science and Technology (DST) Group has developed a security-based global petroleum simulation, called SPECULA, in order to model the effects of regional and global changes in oil production, refining, shipping or distribution of petroleum products during conflict or significant environmental events. This paper briefly describes the security context of the petroleum supply chain in the Asian region. Then previous economic models of Australia's petroleum supply security are reviewed and critiqued in terms of their ability to model conflict scenarios. The SPECULA model is then described along with model parameters and outputs. Finally, the future challenges of this model are addressed in the discussion/conclusion.

Putting the science behind our advice

Niem Tri

DSTG, Fishermans Bend, Australia; niem.tri@defence.gov.au

Many publicly funded research programs in the areas of bias, knowledge elicitation, prioritisation and scenario analysis are not interested in investigating these themes in the military context. For example, knowledge elicitation works best when using diverse subject matter experts, but this condition is often compromised when access to experts is limited with many experts sharing similar educational and work experiences. This paper proposes using the DSTG's summer vacation project (SVP) program as a vehicle to develop our research capability by exploring these topics within the military context. This paper reviews the student projects to date, outlining their research question, the approach as well as reporting on preliminary results. We also discuss the implications of using this model and end with some recommendations for improving outcomes for DSTG while supporting the development of the STEM pipeline. In conclusion, student projects are an effective mechanism to extend public research and make the findings applicable in the Defence setting. Both DSTG and the students benefit from this arrangement and I suggest that more researchers put effort into supporting the STEM pipeline.

A9. Analytics in Industry4.0, IoT and Mechatronics

Development of Extraterrestrial Food Production Systems: An Application of Mechatronics and the Internet of Things

Kurt Benke[1], *Garry O'Leary*[2], *Glenn Hale*[3] *and Richard Watson*[4]

[1]Centre for AgriBiosciences, AgriBio, Bundoora, Australia; kbenke@unimelb.edu.au
[2]Agriculture Victoria Research (AVR), Horsham, Australia;
Garry.Oleary@ecodev.vic.gov.au
[3]Centre for AgriBiosciences, AgriBio, Bundoora, Australia;
glenn.hale@ecode.vic.gov.au
[4]Research Scientist and IT Consultant, Vermont, Australia; richardwatson@tpg.com.au

Significant resources are being directed towards the development of bio-regenerative food systems by NASA and others to support long duration space flights and to decrease the mass and volume of pre-packaged supplies (Perchonok et al, 2012). The space mission is supported by research in controlled-environment agriculture (CEA) based on indoor sealed environments that make extensive use of light emitting diodes (LEDs) for illumination, renewable energy systems, recycled water, together with automatic control of temperature, humidity and CO2 levels with natural and artificial growth media and synthetic biology (Benke and Tomkins, 2017; O'Leary and Benke, 2018). A research agenda is described for CEA for use on Earth and extra-terrestrial food production systems planned for future space missions to the planet Mars. Potential applications of CEA include food production in remote and inhospitable environments such as deserts, spacecraft and submarines and on other planets where endurance is constrained by human factors rather than the range provided by propulsion systems. The task is difficult and not to be underestimated, but it is not impossible. Bioengineering and OR can provide a baseline and methodology to support research originating from terrestrial CEA systems, such as so-called Vertical Farms, which are now appearing in smart cities and highly connected urban clusters around the world (Benke and Tomkins, 2017). Design and development of CEA systems requires software and hardware support from many disciplines, including mechanical engineering, electronics, biosystems engineering and information and communications technology (ICT). The combination of these disciplines has been termed mechatronics, which reflects its origins in hybrid systems based on integration of electronics with mechanics. In the specialised field of extra-terrestrial CEA, to assure fail-safe operation, additional disciplines contribute to the mix, including space medicine and aerospace engineering. Associated with the field of mechatronics is the emerging field of Internet of Things (IoT), which involves network connectivity of sensors, software and actuators that support data processing and the operation of robots and adaptive systems. The IoT has many applications, including smart homes, smart cities, manufacturing industry and food production systems (Tzounis et al, 2017).

There are many problems and research challenges presented by mechatronic and IoT systems, some of which can be addressed by OR using computer simulation with support from artificial intelligence (AI) at the systems level (O'Leary and Benke 2018). The challenges include latency issues for communication with Earth together with the need for fail-safe reliability engineering. A recent study reviewed progress in applying OR methods to IoT systems and noted that OR methods are now being applied to some of the major IoT research challenges, particularly Big Data, in the context of data management and predictive analytics (Ryan and Watson 2017). In agriculture, the application of Big Data offers support to farming operations, facilitation of real-time operational decisions and the redesign of legacy business processes (Wolfert et al, 2017).

A 5-parameter complexity classification of the two-stage flow shop scheduling problem with job dependent storage requirements

Yakov Zinder

Sydney , University of Technology, Sydney, Australia; yakov.zinder@uts.edu.au

The talk presents NP-hardness results and polynomial-time algorithms for the two-stage flow shop scheduling problem with a buffer. In the majority of publications, the buffer only restricts the number of jobs that have completed the first operation and are waiting for the start of the next one. Much less is known about the more general model where different jobs have different buffer requirements. In particular, this model with the additional assumption that each job needs the buffer for the entire duration of its processing is important for multimedia systems, where data files share the same memory, [5] and for supply chains, where the change of the mode of transportation involves unloading and loading [1]. To the best of the authors' knowledge, only publications [2], [3] and [4] studied the computational complexity of such scheduling problems. The talk addresses this gap in the literature by presenting a classification that establishes the borderline between NP-hard and polynomially solvable cases of the makespan problem. According to this classification, all instances form families, each associated with a vector $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5)$. For $i \in 1, 2, \alpha_i = 1$, if all processing times on machine i are the same. Otherwise, $\alpha_i = 0$. If $\alpha_3 = a$, the buffer requirement of each job is proportional to its processing time on the first machine. Otherwise, $\alpha_3 = 0$. The parameter $\alpha_4 \in =, \leq, \geq, \emptyset$ and specifies the relation between the processing times of each job on the first and second machine. Thus, $\leq$ indicates that, for each job, its processing on the first machine does not take more time than on the second. The parameter α_5 indicates whether or not any two jobs can occupy the buffer simultaneously.

Cyclic job shop scheduling problem with sequence-dependent setups; mixed integer-programming

Ata Elmi[1], *Mohan Krishnamoorthy*[2] *and Andreas Ernst*[3]

[1]Monash University, Department of Mechanical and Aerospace Engineering, Melbourne, Australia;
Atabak.Elmi@monash.edu
[2]The University of Queensland, School of IT and Electrical Engineering, Brisbane, Australia;
m.krishnamoorthy@uq.edu.au
[3]Monash University, School of Mathematical Sciences, Melbourne, Australia;
andreas.ernst@monash.edu

Cyclic scheduling has been an effective scheduling method for the repetitive discrete manufacturing environment. The technologic manufacturing cells like robotic cells are tried to eliminate or decrease the number of buffers. Thereupon, manufacturing systems without buffers are encountered in real life. A production line, structured as a job shop, requires that: (i) the operations of each job of the Multiple Part Set are assigned to the machines in advance and (ii) the routing (order of processing) of each job passing through the machines, not necessarily the same for each job, is known and fixed. The cyclic job shop scheduling problem is to find the processing order in which the operations are repetitively processed on each machine (infinitely many times). Such an order, together with the starting times of all the operations, is called a cyclic or periodical schedule. A regular time interval in which all the operations are repeated is called a period or a cycle time. This study investigates the problem where the machines have no buffers that cause the blocking conditions, which rapidly decrease the number of feasible solutions and, therefore, makes it a lot harder to find feasible solutions. We also considered the sequence-dependent setups, which frequently appear in different manufacturing environments. The sequence-dependent setups are also encountered in two types, anticipatory setups and non-anticipatory setups. This research provides a novel mixed integer programming model for the cyclic job shop scheduling where the operations are scheduled within a cycle. In addition, due to the blocking condition considered in this research, the anticipatory and non-anticipatory sequence dependent setups have different effects on the complexity and character of the problem. The characteristics of the blocking cyclic job shop scheduling problem and/with anticipatory and non-anticipatory sequence dependent setup times are discussed, and an example is used to demonstrate the effectiveness of the proposed method.

The advantages of this study and achievements can be outlined as follows:

- The previous studies on modelling the cyclic job shop scheduling problem are based on performing the operations in iterative cycles but the proposed model in this research schedules all the operations within a single cycle.
- Due to the scheduling within a single cycle, the proposed model could be extended simply to different resource constrained versions.
- Two kinds of sequence-dependent setups are considered based on anticipatory and non-anticipatory concepts.
- The non-anticipatory setups do not affect the blocking condition as much as anticipatory setups as the related part does not have to wait on the previous machine because of related setup.
- The anticipatory setups directly affect the blocking conditions as any part may need to wait on the previous machine until the related setup is completed.

As future research, an efficient solution method will be proposed for the considered problem. Moreover, this research could be extended to the manufacturing systems with transporting robots or the servers that are performing the setups.

Situational Awareness for Industrial Operations

Peter Baumgartner[1] *and Patrik Haslum*[2]

[1]Data61 and Australian National University, Canberra, Australia;
peter.baumgartner@data61.csiro.au
[2]Australian National University and Data61, Canberra, Australia;
Patrik.Haslum@anu.edu.au

The smooth operation of industrial or business enterprises rests on constantly monitoring, evaluating and projecting their current state into the near future. Such situational awareness problems are not well supported by today's software solutions, which often lack higher-level analytic capabilities. To address these issues we propose a modular and re-usable system architecture for monitoring systems in terms of their state evolution. States are explicitly represented and externally analyzable. In addition, different state trajectories can be explored simultaneously, e.g., for assessing their plausibilities and what-if reasoning. The need for that arises when dealing with noisy or unreliable input sources. In the paper we describe the system architecture with a focus on a core component, the state inference engine, and its implementation as a shallow embedding in Scala.

D7 and D3. Combat enablers and Advances in Wargaming

Female Engagement Teams - Will They Work For Us?

Erica Dill-Russell

NZDF, Defence Technology Agency, Auckland, New Zealand;
erica.dill-russell@nzdf.mil.nz

The New Zealand Defence Force (NZDF) uniformed population is 84% male and 17% female. In contrast, the environments that the NZDF deploys into are gender equal. The NZDF has recently established a Female Engagement Team (FET) whose primary role is — and will be — to engage with local female and adolescent population, where such engagement would be culturally unacceptable for male soldiers. There is relatively little published information on FET capabilities and what little there is, is retrospective analysis of FET experiences. With the increasing prevalence of population centric conflict and extant issues with low levels of female participation in the military, gendered capabilities will become more important. As such, researchers from the NZDF's Defence Technology Agency (DTA) are conducting a 3 year longitudinal cohort study of the NZDF FET. This research aims to provide empirical data on the development and operational employment of FET capabilities, building on the mainly retrospective research available on the development of FET by providing analysis of the capability as it is developed. This paper will describe Phase 1 of the study, which evaluated the development and selection procedure for the FET capability within 1 NZSAS Regiment. Results are presented along with potential implications for the FET capability going forward. The study design used mixed methods (interviews, surveys and informal observation) to analyse research participants' perceptions and experiences. Research participants included FET candidates along with unit members where the FET is being developed. Participants' interview responses were coded thematically for prevalence using NViVO analysis software. The responses were then disaggregated to compare FET aspirants and Regiment member responses. These techniques yielded rich data indicating the different views held by the different groups. Results highlighted FET candidates' reasons for applying to join the FET and indicated that the development of the FET capability filled a gap in providing a perceived elite role for women in the NZDF to apply for, much like the NZSAS for men. Interviewees were almost uniformly positive about the development and necessity for a FET type capability. Perceptions of the benefits of the capability focussed on the ability to search, gather intelligence, and engage with local women, where it was culturally inappropriate for male soldiers to do so. Whilst most participants (both Regiment members and FET candidates) thought there was no additional risk in deploying a FET capability, there remained concerns associated with the deployment of women in general such as physical capability, vulnerability and male behaviour in combat.

Potential barriers to the success of the FET capability were identified and solutions proposed; namely a) ensuring the role of the capability is messaged correctly, as ambiguity could affect the perception of the capability, and b) ensuring the right training is provided to FET members, so they can deploy effectively in their articulated role. This research will continue with its second cohort beginning in November 2018. Further research analysis will focus on the training and integration of the FET and its eventual operational deployment.

A Fast Heuristic for Generating Medium-Term Military Aircraft Fleet Schedules

Nicholas Joseph[1], *David Marlow*[2] *and Robert Dell*[3]

[1];
[2]DST Group, Fishermans Bend, Australia; david.marlow@dst.defence.gov.au
[3]The Naval Postgraduate School, Monterey, The United States of America; dell@nps.edu

We develop an algorithm to generate monthly schedules for a fleet of military aircraft over a fixed horizon (e.g. three years). The fleet may consist of a number of squadrons with some aircraft deployed to operational areas. All aircraft must enter depot maintenance at regular intervals and need to achieve a certain number of flying hours prior to their induction. These schedules must satisfy many complex requirements to ensure that the right type of aircraft are available when and where they are needed over the planning horizon. For this reason, schedules that merely maximise flying hours may not be useful in practice. Unfortunately, incorporating other requirements such as aircraft availability among different squadrons, balancing aircraft usage, and entering planned maintenance at set times greatly increases the complexity of the problem. When optimised with modern mixed-integer programming (MIP) solvers, this added complexity results in long computation times that may not be practical for end users. We therefore develop an alternative heuristic-based solution that generates feasible schedules of reasonable quality in a fraction of the time required by full MIP-based optimisation approaches. Our method generates aircraft schedules using the current squadron assignments, the number of months the aircraft have been in their respective squadrons, and the accumulated hours each aircraft has flown in each squadron to date. Our algorithm progresses from month 0, only allowing aircraft to swap out of a squadron once a minimum number of months have been achieved. Additionally, maximum durations within squadrons are also enforced. For aircraft that are available to swap into different squadrons, their next squadron is chosen based on 1. whether the aircraft is due for scheduled maintenance, 2. the target number of aircraft required for each squadron, and 3. the balance of hours the aircraft have flown in each squadron. At each month, the aircraft hours are updated using the average flying rate of their assigned squadrons. Once all aircraft assignments are complete, the optimal flying hours and positions are computed by solving a reduced MIP with far fewer binary variables than the full problem. We run our approach for six scenarios with differing fleet sizes and squadron numbers and compare them with the optimal solutions provided from CPLEX. For all cases, the heuristic is able to generate feasible squadron assignments in a matter of seconds. Furthermore, as the subsequent MIP contains far fewer binary variables, the aircraft hours and positions are solved within a few minutes. In contrast, solving the full MIP to optimality takes up to 2 days for the largest problem sizes. More importantly, CPLEX takes a very long time to find incumbent solutions with objective values at least as good as those generated by the heuristic (at least two hours for medium-to-large instances).

Our heuristic is a viable approach for rapidly finding feasible fleet schedules of reasonable quality and therefore has promise for military fleet scheduling in practice. Finally, it can be used to provide good quality initial solutions for warm-starting other solution approaches when optimality concerns become more important.

Strategic Wargaming for Whole-of-Force Design

Thitima Pitinanondha[1] *and Nitin Thakur*[2]

[1] Defence Science and Technology, Joint and Operations Analysis Division, Canberra, Australia; thitima.pitinanondha@dst.defence.gov.au
[2]Providence Consulting Group, Canberra, Australia; nitin.thakur1@defence.gov.au

Force design is a critical strategic decision-making process in the Australian Defence Organisation that aims to produce a plan for future capabilities that would operate effectively to achieve desired strategic objectives under likely future threats and challenges. As the strategic environment changes rapidly, and uncertainties and unforeseen events appear at all time, these require the force design process to be flexible and robust. Wargaming has been a method and tool used in support of Defence planning at operational and tactical levels for many years. This paper discusses a wargaming approach applied in support of strategic decision-making at strategic level to assess the capability of the force so as to plan for future capabilities. It also lays a foundation for a strategic wargaming approach that could be applied for future research and associated activities to support strategic decision-making for whole-of-force design.

A3. Natural hazards and emergency management

Stochastic multi-resources scheduling and sequencing decisions for emergency operations

Shahrooz Shahparvari[1] *and Behrooz Bodaghi*[2]

RMIT University, School of Business IT and Logistics, Melbourne, Australia
[1]shahrooz.shahparvari@rmit.edu.au
[2]behrooz.bodaghi@gmail.com

The optimization of scheduling and sequencing aspects relating to multi-resources is significant during disaster management due to the uncertainty about the magnitude of the disaster, a lack of reliable information, the availability of resources, and the needs of the victims. Particularly during the Black Saturday Bushfire, Victoria, Australia, the investigations of Victorian Bushfires Royal Commission has revealed that resources such as medical teams and medical supplies were poorly coordinated during the initial disaster emergency response.

Synthetic populations in bushfire scenarios

Joel Robertson

RMIT University School of Mathematics/Computer Science, Melbourne, Australia;
joel.robertson@student.rmit.edu.au

Bushfires pose a significant threat to Australia's regional areas. To minimise risk and increase resilience, communities need robust evacuation strategies that account for people's likely behaviour both before and during a bushfire. Agent-based modelling (ABM) offers a practical way to simulate a range of bushfire evacuation scenarios in which the individual interactions and responses of agents allow collective behaviours to emerge. A person's response when they become aware of a bushfire threat will largely depend on where and why they are in the region, so the creation of a synthetic population of agents that accurately describes the makeup of the community throughout the day is a crucial requirement for an ABM bushfire simulation. Although there are established techniques for synthesising a population, these usually rely on existing sample data sets to incrementally generate a full population. In regions where there is a large transient demographic that might vary depending on the season, the time of week or the weather conditions on the day, this data is not likely to exist for each given bushfire scenario. The population can instead be represented by a specific set of inputs that correspond to the conditions on the day. These inputs will capture the proportion of different groups in the population, the types of activities people are likely to engage in, how these activities are distributed throughout the day, and the potential locations of activities in the region. This set of inputs is designed in conjunction with emergency services personnel and allows the representation of a scenario population using simple and attainable data. Applying the evacuation scenario with this synthetic population in an ABM bushfire simulation will demonstrate the possible impact of population structure on outcomes. It is hoped that this will give an improved understanding of the risks associated with evacuation, and contribute to the development of tailored evacuation plans for each community to help them prepare for and respond to bushfire.

The story of CSIRO Spark, from a science idea to a technology suite blending fire science, computation and OR

Simon Dunstall

CSIRO Data61, Melbourne, Australia; Simon.Dunstall@data61.csiro.au

In 2011, a research team specializing in bushfire science, and a research team specializing in numerical simulation methods, decided to work together to construct a GPU-based approach for bushfire spread simulation. The goal was to create a bushfire modelling platform that was modular and openly-developed, so that it could be rapidly customized, be modified to incorporate the latest findings from fire science, be applicable to a range of real-world use cases,and bring many researchers' model development and validation efforts together in a common implementation with proven and expanding functionality. Seven years later, this platform known as Spark has been a success relative to each of the goals, and has been embraced by a number of research institutions and agencies. Of particular interest here is the origins and progress of Spark between 2011 and 2018 as a dataset-generating tool for risk analytics. The initial 2011 commitment to build Spark partly stemmed from experience gained in powerline bushfire risk analytics work in Western Australia, where a power utility was seeking to quantify its maximum foreseeable loss, and a suitable modern fire simulation tool was not available to the research team. At the present time, Spark has now been used multiple times in Australia for risk analytics work associated with optimising asset refurbishment and replacement strategies to maximally reduce bushfire likelihood within budget constraints. The architecture of Spark enables very fast and highly parallel ensemble simulation runs on Australia's major GPU-based computing cluster. This makes it practical to execute risk analytics studies that utilise up to two million bushfire simulations in order to build datasets of fire consequence under various fire weather conditions. This consequence data is combined with likelihood data to estimate risk at the level of individual power poles or spans, or as is occurring in work currently underway Chile, at a resolution of quite small land parcels (one hectare) over very large spatial scales (hundreds of square kilometres). Using the risk estimates, either mathematically-based optimisation approaches, or human-centric optimisation processes, are applied to identify interventions and projects that can maximise fire risk reduction.

INDEX OF AUTHORS

Index of Authors

SUPPORTED BY