Number 105 April 2016

ERCIM NEWS

www.ercim.eu

Special theme:

Planning and Logistics

Also in this issue:

Keynote:
Trends and Challenges in Logistics and Supply Chain Management
by Henk Zijm, Professor of Production and Supply Chain Management, University of Twente

Research and Innovation:
Business Process Execution Analysis through Coverage-based Monitoring
by Antonello Calabrò, Francesca Lonetti, Eda Marchetti, ISTI-CNR

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 6,000 printed copies and is also available online.

ERCIM News is published by ERCIM EEIG
BP 93, F-06902 Sophia Antipolis Cedex, France
Tel: +33 4 9238 5010, E-mail: contact@ercim.eu
Director: Jérôme Chailloux
ISSN 0926-4981

Contributions
Contributions should be submitted to the local editor of your country

Copyright notice
All authors, as identified in each article, retain copyright of their work
ERCIM News is licensed under a Creative Commons Attribution 4.0 International License (CC-BY).

Advertising
For current advertising rates and conditions, see http://ercim-news.ercim.eu/ or contact peter.kunz@ercim.eu

ERCIM News online edition
The online edition is published at http://ercim-news.ercim.eu/

Next issue
July 2016, Special theme: Cybersecurity

Subscription
Subscribe to ERCIM News by sending an email to en-subscriptions@ercim.eu or by filling out the form at the ERCIM News website: http://ercim-news.ercim.eu/

Editorial Board:
Central editor:
Peter Kunz, ERCIM office (peter.kunz@ercim.eu)

Local Editors:
Austria: Erwin Schoitsch (erwin.schoitsch@ait.ac.at)
Belgium:Benoît Michel (benoit.michel@uclouvain.be)
Cyprus: Ioannis Krikidis (krikidis.ioannis@ucy.ac.cy)
Czech Republic:Michal Haindl (haindl@utia.cas.cz)
France: Steve Kremer (steve.kremer@inria.fr)
Germany: Michael Krapp (michael.krapp@scai.fraunhofer.de)
Greece: Eleni Orphanoudakis (eleni@ics.forth.gr),
Artemios Voyiatzis (bogart@isi.gr)
Hungary: Erzsébet Csuhaj-Varjú (csuhaj@inf.elte.hu)
Italy: Carol Peters (carol.peters@isti.cnr.it)
Luxembourg: Thomas Tamisier (thomas.tamisier@list.lu)
Norway: Poul Heegaard (poul.heegaard@item.ntnu.no)
Poland: Hung Son Nguyen (son@mimuw.edu.pl)
Portugal: Joaquim Jorge (jorgej@tecnico.ulisboa.pt)
Spain: Silvia Abrahão (sabrahao@dsic.upv.es)
Sweden: Kersti Hedman (kersti@sics.se)
Switzerland: Harry Rudin (hrudin@smile.ch)
The Netherlands: Annette Kik (Annette.Kik@cwi.nl)
W3C: Marie-Claire Forgue (mcf@w3.org)

Cover photo by courtesy of Electrocomponents plc.

12th EUROPEAN COMPUTER SCIENCE SUMMIT

www.informatics-europe.org/ecss/ecss-2016.html

ECSS 2016
BUDAPEST - HUNGARY
24-26 OCTOBER

© FreeImages.com/Rayko Svensson

INFORMATICS EUROPE

EÖTVÖS LORÁND UNIVERSITY
FACULTY OF INFORMATICS

INFORMATICS DRIVING THE DIGITAL WORLD

ECSS is the forum to debate the trends and issues that impact the future of Information and Computer Sciences in Europe. Join leading decision makers in research and education and hear from renowned speakers in the field how Informatics is driving the digital world.

Conference Chairs:
Lynda Hardman, President, Informatics Europe
Zoltán Horváth, Dean of Faculty of Informatics, Eötvös Loránd University, Hungary

Program Chairs:
Gregor Engels, Department of Computer Science, University of Paderborn, Germany
László Monostori, MTA SZTAKI, Hungary Academy of Sciences

Pre-Summit Workshop For Deans, Department Chairs and Research Directors

Trends and Challenges in Logistics and Supply Chain Management

by Henk Zijm

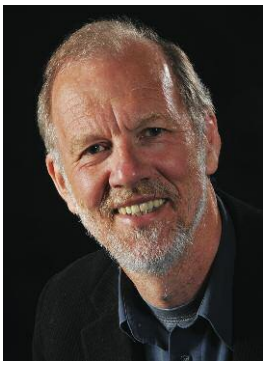

Henk Zijm, Professor of Production and Supply Chain Management, University of Twente

In today's global economies, logistics is a key facilitator of trade, and hence an important factor in rising prosperity and welfare. Natural resources are scarce and not evenly distributed in terms of type and geographical location in the world. Logistic chains enable the distribution of materials, food and products from the locations where they are extracted, harvested or produced to people's homes and nearby stores. At the same time, current logistics systems are fundamentally unsustainable, due to the emission of hazardous materials (CO_2, NO_x, particulate matter), congestion, stench, noise and the high price that has to be paid in terms of infrastructural load. Things are even getting worse: while the European Commission has set (not achieved) targets to reduce Greenhouse Gas Emissions (GGE) in 2015 to 60 % as compared to 1990, the percentage of transport related GGE increased from 25 % in 1990 to 36 % today.

The still growing world population stresses the need to further increase productivity while at the same time diminishing the ecological and societal footprint. This requires a quality upgrade of the human resource pool by better education and training, including lifelong learning programs. Productivity can also be improved by better support tools, easier access to relevant information, and further automation of both technical processes (i.e. robotics) and decision making (artificial intelligence). The same tools might also help to reduce border-crossing logistics systems' vulnerability to crime and illicit acts, such as theft, organized immigration crime (human trafficking) and customs law violations.

The continuing urbanization poses a further challenge. The development of wealth in Asia and Latin America has resulted in a huge shift from agricultural and nomadic forms of living to urban life. More and more cities with over ten million inhabitants have emerged, requiring different modes of transport and logistics systems than available today. There is an increasing interdependency between supply chain design or management and urban planning or land-use management. In addition, due to both political conflicts and natural disasters, the importance of humanitarian logistics can hardly be overestimated.

But also consumer behaviour is changing rapidly, demonstrated for instance by the rapid advance of e-commerce, with a profound impact on both forward and reverse logistics and supply chains. Clearly, meeting the continuous pressure on fast delivery is only possible by an excellently functioning logistics network.

Fortunately, technological innovations are expected to at least partially address some of these challenges. The design of new and lightweight (bio-)materials, miniaturization and de-materialization of products helps to diminish both their costs and ecological footprint. Technologies like 3D-printing and micro-machining are also a step forward towards mass-customization but in addition have a profound logistic impact, for instance in stimulating "local for local" production, thereby also reducing so-called anticipation (safety) stocks, because they allow production at the place and time needed.

But also the impact of robotics will change the logistics landscape considerably, as it did already in automotive assembly lines and automatic storage and retrieval systems, assisted by digital dynamic identification systems such as RFID, and all controlled by innovative warehouse management systems. Similar developments are found at container terminal sites in both seaports and inland harbors. Without exception, such systems rely heavily on smart sensor and actuator systems, evolving towards the so-called Internet of Things (IoT). The same IoT is currently innovating both passenger and freight transport rapidly; vehicle transportation in 2050 is foreseen to be largely unmanned transportation.

But technological innovation is only a part of the story; at least equally important is the development of smart business models based on joint responsibilities and fair allocation of revenues instead of on individual profit maximization. Complex modern supply chains are first and foremost characterized by the fact that many stakeholders are involved in shaping their ultimate manifestation, not only shippers and logistic service providers but also the financial sector and governmental agencies, and ultimately the customer. Such systems require adequate planning and control mechanisms, including distributed architectures, cloud computing solutions, cognitive computing and agent-based decision support systems. The recent attention for data driven models (big data analytics) marks an important further step towards full-blown automated decision architectures.

Multi-stakeholder systems aiming at cooperation between essentially autonomous companies require tools that basically draw on game-theoretical concepts. But the key idea - established in the Nash equilibrium theory - that players may have to give up their individual optimal solution in order to achieve an overall stable equilibrium solution is still hard to accept, in particular for private companies that were used to concentrate on their individual profits. This is perhaps the biggest hurdle to be overcome to arrive at sustainable logistics; it involves not only smart business solutions but more importantly a change of mind. and indeed trust in the value of collaboration.

A similar change of mind is requested to implement ideas of re-use of products or components, both via (electronic) second markets or directly from dismantling disposed products in closed loop supply chains, as an example of the circular economy. Also, the rising attention for sharing economy concepts (as e.g. in car sharing, cloud computing, music streaming) may have important consequences for supply chain design, planning and control in that the focus will at least partially switch from delivering products to delivering services.

KEYNOTE

JOINT ERCIM ACTIONS

SPECIAL THEME

The special theme section "Planning and Logistics" has been coordinated by Rob van der Mei, CWI and Ariona Shashaj, SICS.

Introduction to the Special Theme

Simulation Models and Test Beds

Case Studies

ManyVal 2015 - Workshop on Many-Valued Logics

by Carles Noguera

ManyVal is a series of international workshops on the logical and algebraic aspects of many-valued reasoning. The 2015 workshop served as the official meeting of the ERCIM Working Group of the same name.

The aim of the workshop series is to bring together both established and young researchers sharing an interest for a specific topic. Accordingly, each edition has a sharp focus. The attendance is limited in order to facilitate close and informal interaction. There are no parallel sessions and contributed talks are allocated ample time.

The chosen topic this year was modal and first-order many-valued logics. The workshop took place at Hotel Les Sources in Les Diablerets, Switzerland, from 11-13 December 2015. The program committee was chaired by George Metcalfe (University of Bern) and the organizing committee was chaired by Denisa Diaconescu (University of Bern), with assistance from other members of the Logic group at the Mathematical Institute.

There were 28 participants at the meeting from 10 countries. The program included invited talks by Itaï Ben Yaacov (Université Claude Bernard, Lyon 1, France), Marta Bílková (Charles University, Prague, Czech Republic), Xavier Caicedo (Universidad de los Andes, Colombia), Petr Cintula (Academy of Sciences, Prague, Czech Republic), Rafael Peñaloza (Free University of Bozen-Bolzano, Italy). It also included eleven selected contributed talks of half an hour.

Sponsors of ManyVal 2015 included the Swiss National Science Foundation (SNF) and the European Research Consortium for Informatics and Mathematics (ERCIM).

Past editions of ManyVal took place in Gargnano (2006), Milan (2008), Varese (2010), Salerno (2012), and Prague (2013).

Link: http://mathsites.unibe.ch/manyval2015/

Please contact:
Carles Noguera, Institute of Information Theory and Automation (UTIA), Czech Academy of Sciences/CRCIM
Chair of the ERCIM ManyVal Working Group
Tel: +420 26605 2563
noguera@utia.cas.cz

Call for Nominations

Cor Baayen Award 2016

The Cor Baayen Award is given each year to a promising young researcher in computer science and applied mathematics. The award was created in 1995 to honour the first ERCIM President.

The award consists of a cheque for 5000 Euro together with an award certificate. The winner of the 2016 Award will be invited to present her or his work at the European Computer Science Summit 2016 in Budapest, co-located with the annual ERCIM and Informatics Europe fall meetings.

Eligibility

Nominees must have carried out their work in one of the "ERCIM countries": Austria, Belgium, Cyprus, Czech Republic, Finland, France, Germany, Greece, Hungary, Italy, Luxembourg, Norway, Poland, Portugal, Spain, Sweden, The Netherlands and the United Kingdom. Nominees must have been awarded their PhD (or equivalent) after 30 April 2013. A person can only be nominated once for the Cor Baayen Award.

Submitting a nomination

Nominations should be made by a staff member of an ERCIM member institute. Self nominations are not accepted. Nominees must have performed their research at any research organisation from the country of the nominating institution. Nominations must be submitted online. Deadline is 15 May 2016.

Cor Baayen Award Selection Committee

The selection of the Cor Baayen award winner is the responsibility of ERCIM's Human Capital Task Group, who may consult expert opinion in reaching its decision.

The Cor Baayen Award is named after the first president of ERCIM and the ERCIM 'president d'honneur'. Cor Baayen played an important role in its foundation. Cor Baayen was scientific director of the Centrum Wiskunde & Informatica (CWI) in the Netherlands from 1980 to 1994.

More information and submission form:
http://www.ercim.eu/activity/cor-baayen-award

Please contact:
The ERCIM representative of your country
(see http://www.ercim.eu/about/member-representation)
or
Claude Kirchner, Inria
ERCIM Cor Baayen Award coordinator
Claude.Kirchner@inria.fr

How Science can Improve Marine Ecosystem Predictions: the BlueBRIDGE Case

by Sara Garavelli

BlueBRIDGE - Building Research environments fostering innovation, decision making, governance and education - is a European project aimed at providing innovative data services to scientists, researchers and data managers to address key challenges related to sustainable growth in the marine and maritime sectors.

Imagine yourself wanting to build a wonderful new three storey house in the countryside. You go to the local authorities and they say that unfortunately you can only have two storeys, due to birds nesting nearby and having a higher building will affect their habitat. Or imagine that a local council wants to move a flood defence system seaward but it's not possible because the beach and environment in the surrounding areas will be damaged. Human decisions are highly influenced by ecosystems.

A central problem in ecology is determining the processes that shape the complex networks known as food webs formed by species and their feeding relationships. The topology of these networks is a major determinant of ecosystems' dynamics and is ultimately responsible for their responses to human impacts.

A real example is documented by the Vancouver harbour case. The Robert Bank Terminal 2 (RBT2), Vancouver, BC, project proposed to increase the size of the port, adapting to the increase in demand for container shipping traffic which is expected to triple by 2030. To understand the impact of this change on the ecosystem, the project leaders enlisted help from food web models, in particular of the Ecopath with Ecosim (EwE) model.

Jeroen Steenbeek, Head Technical Committee at Ecopath Research and Development Consortium, explains that "Making predictions of the future state of the marine ecosystem of the Vancouver port was an exciting experience because we were able to see how ecosystem models could impact real world decisions. Single-species fisheries models, which are usually adopted for management evaluation, must be complemented as they are unable to capture interactions between species and information on their spatial distribution. The ecosystem modelling approach of EwE offers a means to incorporate interactions and spatial constraints into a useful tool for ecosystem-based fisheries management".

"Through the EwE models available via the software provided by the Ecopath Consortium the project leaders of the Vancouver port were able to understand in advance the impact of the extension of the port on the marine resource dynamics", continues Steenbeek, "EwE infact has shown that the building of the extension to the port may not be that environmentally impactful and that not all of the counter measures foreseen by the initial project may be needed".

EwE has proved how useful it can be displaying its potential in the research toward the impacts of climate change and human activities on marine ecosystems. However, today there are still a set of technical limitations that prevent a wider adoption of the EwE approach. The key for good predictions are, in addition to thorough understanding of a marine ecosystem, good data. Different data in different formats, which correspond to a majority of the situations in which the users of EwE work, requires a huge effort and resources just too initially set up the simulation models. In addition, running ecosystem simulation models requires large computational resources to which nor everyone has access.

That's why the Ecopath Consortium has joined the BlueBRIDGE project. BlueBRIDGE - Building Research environments fostering innovation, decision making, governance and education - is the newly funded H2020 project aimed at providing innovative data services to scientists, researchers and data managers to address key challenges related to sustainable growth in the marine and maritime sectors referred to as the "Blue Growth long term strategy".

BlueBRIDGE can support the process of Ecopath model creation by making available for Ecopath a seamless access to different species, fisheries and environmental data sources.

In the doing this, users of Ecopath do not need to worry about data harmonization or conversion and can focus on the ecosystem modelling, which inevitably speeds up their research. In addition, the BlueBRIDGE underlying infrastructure D4Science, can equip Ecopath with cloud capabilities on which they can run their model (at the moment EwE is only available via a desktop version). BlueBRIDGE can enhance the EwE capabilities and expand its usage to a wider audience, contributing to the creation of more powerful and accurate instruments for the prediction of human activities on seas and ecosystems. The results of the BlueBRIDGE and EwE collaboration will be available for all the researchers worldwide. Keep your eye on BlueBRIDGE!

The project is coordinated by ISTI-CNR, the ERCIM Office contributes with expertise for administrative and financial management.

Links:
http://www.bluebridge-vres.eu
@BlueBridgeVREs
http://www.ecopath.org
http://www.d4science.org

Please contact:
Donatella Castelli, ISTI-CNR, Italy
donatella.castelli@isti.cnr.it

Introduction to the Special Theme

Visions of the Future: Towards the Next Generation of Logistics and Planning Systems

by Rob van der Mei, CWI, and Ariona Shashaj, SICS

In December 2013, one of the biggest online retailers announced its intention to use drones to deliver products to consumers. Perhaps the next step towards future logistics will be the use of teleportation beams! Although we are under no immediate threat of being pelted by delivery drone rain while dodging teleporting beams of products, logistics systems have undoubtedly become more complex in order to face new challenges imposed by the growth of the global market, mass urbanization, and the move towards sustainability. This Special Theme of the ERCIM News is dedicated to recent advancements in logistics, planning, scheduling and supply chain optimization.

Logistics is the science that orchestrates the flow of resources through supply chains in terms of transportation mode, warehousing and third-party organization. Modern logistics and service systems need to provide efficiency, correctness and robustness in the process of planning and controlling the raw materials, products and people flows and the related information. Some current challenges within logistics include: the globally dispersed nature of companies; the increase in diversity of storage and transportation modes; the mass urbanization that has occurred over the last decade; and the requirements for flexibility, transparency and sustainability.

These challenges notwithstanding, the recent advances in ICT technologies make it possible to have the right resources in the right place at the right time. There is a growing interest and increased research efforts in data science and big data, which are set to become powerful tools for future logistics. The development of machine-learning and data-mining algorithms and their application to complex and voluminous data in order to extract knowledge will improve logistics operations and supply chain management systems by achieving transparency and control over entire systems, improving predictive analysis and risk assessment as well as real-time adjustment and responses to environmental conditions.

By enriching the physical world with contextual information, Augmented Reality (AR) is an emergent technology which will play a fundamental role in the future of industrial processes. The benefits of AR applications in the

fields of logistics and planning operations range from sensory integrative models to intelligent transportation and execution of maintenance/warehouse operations.

A selection of articles in this special theme discusses the potential use of these cutting-edge technologies to logistics operations, such as data-drive models derived through big data techniques and the potential use of machine-learning approaches for intelligent transportation systems, as well as remote maintenance systems through AR applications.

Although the improvement of emerging technologies will lead us to the future of logistics, what stands behind today's efficient logistics operations and planning is the application of mathematical tools. This special theme includes a selection of publications that provide an overview of these tools and their application. In particular, an open source optimization framework is discussed and used within specific case studies in logistics, such as risk assessment and vehicle route optimization. Stochastic optimization is discussed in two different contexts: firstly, in solving placement and packing problems in scenarios involving complex industrial objects, and secondly to efficiently plan complex supply network and logistics maintenance operations. Mixed integer programming techniques are used to optimize train timetables and reduce losses on railway infrastructure capacities while satisfying time constraints and quality of service.

A selection of four articles describe the simulation models and test beds in the field of logistics and planning operations: application of simulation-based optimization approach in order to improve the distribution network of grocery retail stores, description of a 3D modelling software for production factory planning, development of a test bed infrastructure for logistics and transportation, and the study of new models towards global collaborative and robust production networks.

Finally, some real world case studies are included in this special theme: modelling and validation of a system for shipping lane management; study of the impact of patient logistics management on breast cancer treatment; improving the planning and distribution of fire-fighter stations and vehicles; study of the challenges related to the cash supply chain network; and optimizing the planning of railway shunting yards.

Please ocntact:
Rob van der Mei, CWI
R.D.van.der.Mei@cwi.nl

Ariona Shashaj, SICS
ariona@sics.se

Defining the Best Distribution Network for Grocery Retail Stores

by Pedro Amorim, Sara Martins, Eduardo Curcio and Bernardo Almada-Lobo, INESC TEC

Large food retailers have to deal with a complex distribution network with multiple distribution centres, different temperature requirements, and a vast range of store formats. This project used an optimization-simulation approach to help food retailer Sonae MC make the best decisions regarding product-warehouse-outlet assignment, product delivery modes planning and fleet sizing.

In the food retail sector, maintaining food quality across the supply chain is of vital importance. Product quality is dependent on storage and transportation conditions. Compared with other types of retailers the supply chain is very complex, and is best considered as an amalgamation of three types of intertwined food supply chains: frozen, chilled and ambient [1].

Store formats vary among food retailers, in terms of size, product ranges and sales volumes. To be closer to and more convenient for their customers, retailers are opening new stores every year in new locations. The allocation of new stores to the distribution centres is frequently defined ad hoc. These allocations affect the distribution of the different product categories (temperatures) and the operational capacity of the warehouses [2]. Furthermore, poor distribution planning might result in low vehicle utilization and consequently in an overestimated fleet size and/or a fleet mix that is unnecessarily costly. Since product-warehouse-outlet assignment (assignment to distribution centres), product delivery mode planning (direct shipping and/or hub-and-spoke) and fleet sizing (types and dimensions) are interrelated, solving them separately results in oversimplification. However, tackling these decisions simultaneously requires extremely complex mathematical programming which is unfeasible for the real world. Furthermore, existing literature in this field fails to include critical ingredients, such as the consistency of the operations throughout the year, rendering the proposed approaches not applicable.

A real-world case

Sonae MC is the leading grocery retailer in Portugal. Its store outlet is divided in three segments: hypermarkets, supermarkets and convenience stores. To supply its 210 stores (franchising stores are not considered here) with 60,000 products, it relies on two hubs and two specialized warehouses (for fish and meat). Stores are supplied every day with a dedicated, heterogeneous fleet and transportation related annual costs amount to around 40 million euros.

Recently, thanks to an optimization-simulation driven approach developed at INESC TEC, Sonae MC has improved its distribution network and it was able to cut transportation related costs by more than 750 thousand euros, while maintaining the same service level to stores. After several iterations, Sonae MC selected a solution that simplifies and segregates the distribution of the different product categories. This allowed for a reduction of the fleet size and a change in the fleet mix. Additionally, it has performed some adjustments in the warehouse-outlet assignment that has resulted in reduced warehouse operational costs.

Hybridize to simplify

A hybrid approach was used to simplify the problem. This was necessary in order to cope with the complexity and interdependencies of the decisions presented by this real-world scenario. First, several what-if scenarios were defined, based on real-world practice, to identify a distribution framework that achieves transportation savings. At this stage the different store allocation and product delivery modes are analysed. Secondly, an optimization algorithm proposed by Amorim et al. (2012) for routing problems consid-

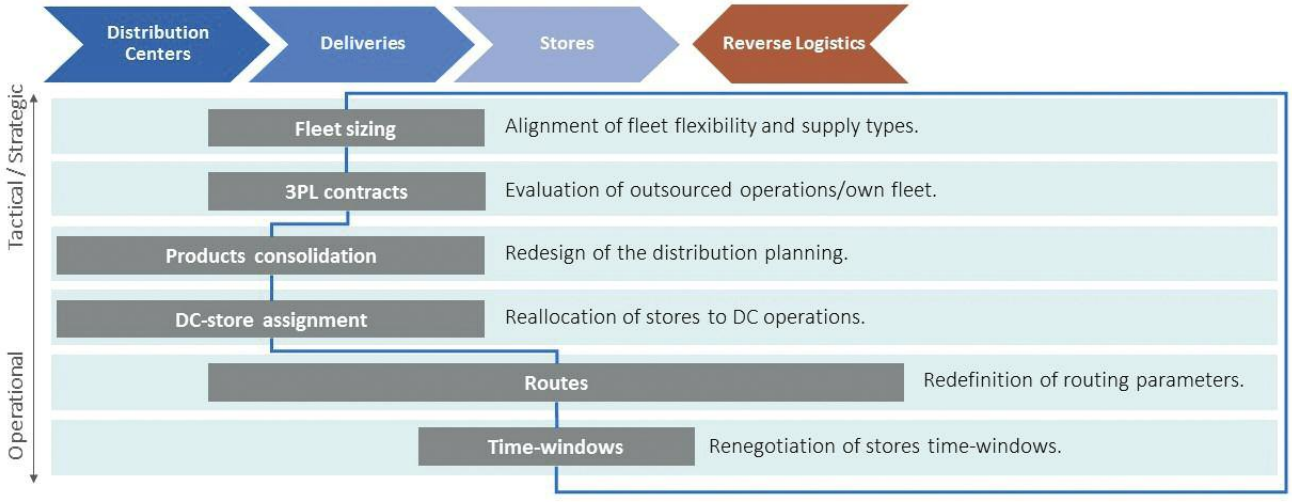

Figure 1: Overview of the opportunities tackled in the project.

ering heterogeneous fleet and time-windows has been adapted to define the routes [3]. The different processes and decisions give rise to considerable operational complexity, so the evaluation of the solution is best done in a simulation environment. This allows operational and financial Key Performance Indicators (KPIs) to be obtained and any readjustments to be made before making any real implementation. The simulation model, at the final step, provides valuable insights on the implementation of strategic and tactical decisions.

Managers at Sonae MC, impressed by the quality of the solutions provided, the user friendly outputs and the automatically generated inputs, decided to put these new solutions into practice. To ensure a smooth transition, adoption was progressive: the new distribution planning process was implemented gradually, and with the new operational performance the number of vehicles was reduced. The implications of the project go beyond better decisions: the project highlights the advantages of looking at operations from a different perspective - with a focus on cost optimization - that may yield solutions that would not be easily devised empirically.

This project was developed by the Centre for Industrial Engineering and Management of INESC TEC, in Portugal. It started in November 2014 and concluded in April 2015. The estimated savings of the project have already been achieved by the company, ensuring a pay-back within a few months. This research is being extended towards a more quantitative scenario generation step.

References:
[1] Akkerman, Renzo et al.: "Quality, safety and sustainability in food distribution: a review of quantitative operations management approaches and challenges", Or Spectrum 32.4 (2010): 863-904.
[2] A. H. Hübner et al.: "Demand and supply chain planning in grocery retail: an operations planning framework", International Journal of Retail & Distribution Management 41.7 (2013): 512-530.
[3] Amorim, Pedro, et al: "A rich vehicle routing problem dealing with perishable food: a case study", Top 22.2 (2014): 489-508

Please contact:
Sara Sofia B. Martins
INESC TEC, Porto, Portugal
Tel: +351 22 209 4370
sara.martins@fe.up.pt

Planning Production Efficiently

by Andreas Halm, Fraunhofer Austria

Many parameters and variants have to be taken into account when planning new production buildings. How do you find the optimal positioning for the production equipment and the most efficient transport paths? Is there a way to find bottlenecks before they hinder your production? The new GrAPPA software makes planning the production a lot easier.

Planning a new production facility or reorganizing an existing production is a difficult task with many factors that must be considered: the positioning of the different production stages, for example, or the size of temporary storages. To run the production at optimal capacity, the workload needs to be equally distributed across all transport paths. While the actual planning is usually done using a two-dimensional floor plan, a three dimensional model offers considerable advantages in presentations, when making estimates involving height or vertical distance, or in granting an intuitive overview of the factory building and the planning scenario. A three dimensional model usually has to be done in a separate step in a different software.

Half the work with a new approach

The software GrAPPA [L1] has been developed by the computer scientists of the Fraunhofer Austria visual computing group together with engineers from the logistics and production planning group. This package makes the planning process easier and more efficient. While drawing, placing and moving workstations on the floor plan, the engineer is also creating a simple 3D model of the factory. This can later be refined by adding detailed and polished models, but even in the early stages it be used to gain quite intuitive insights into the layout of the factory.

The embedded material graph editor allows the user to define relationships and dependencies between the materials being processed, combined or created in the factory. It defines which materials come into the factory and which go out. Numerous properties need to be taken into account. Not only does the size and weight of the materials matter, but it may also be useful to know, for example, how many items of a given material can fit into one transport container to be transported within the factory.

Despite the complexity of such a graph, the embedded editor is quite easy and intuitive to work with. In many cases, this data can also be imported from pre-existing SAP databases, if an existing production is to be analysed. Once the transport paths have been placed in the floor plan, the possible efficiency of the scenario can be analysed. Since it may be more costly to transport large or heavy items, transport paths from one workstation to another should be as short as possible for these materials. Additionally, some paths may be used more frequently than others. If two heavily used paths cross, it might slow production and hinder efficiency. All the information gathered from analysis is displayed clearly in different diagrams, also serving numbers and tables for thorough reliable evaluation. The amount of data displayed to the user is reduced by compressing the data into its most important subsets (as also shown in [1, 2]). Additionally, coloured highlights in the three-dimensional view point out where the planning engineers should focus.

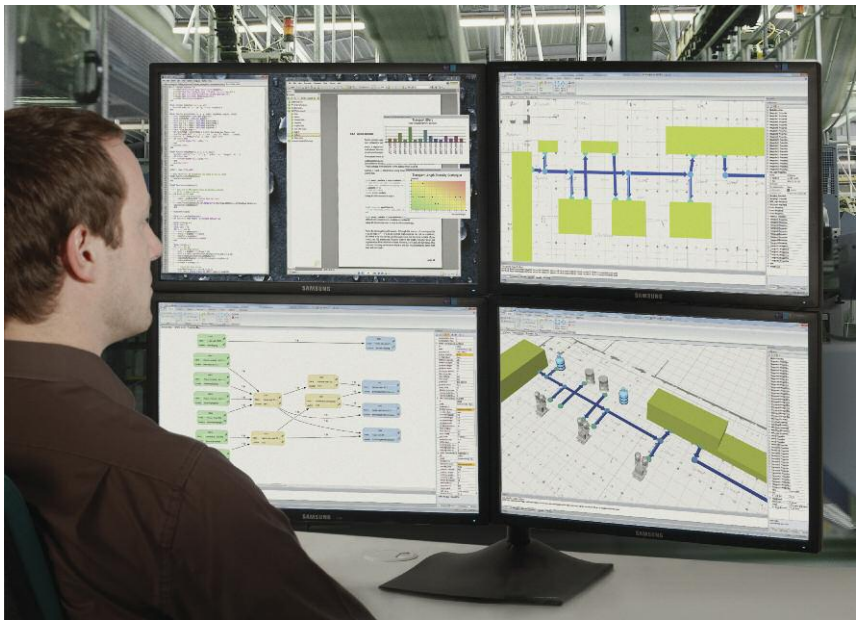

Figure 1: production planning with the GrAPPA software.

Easily adaptable and extendable

Most of the algorithms, diagrams and analysis tools in GrAPPA are programmed in a specially tailored script language. Whenever part of the factory layout or the data changes, the scripts get notified of the change and may recalculate results, modify information or rebuild diagrams. This has the advantage that all the diagrams, computations and the data layout can be easily adapted to fit the needs of a given factory. In the simplest case, this may be just a colouring issue, to get the diagrams to fit into the corporate identity of the business. Modifying the algorithms is more complex, for example adding a new transportation system would change the outcome of the whole analysis. One key advantage is that it also means that the data layout, by which data is imported into the program, is not fixed but can be adapted to the layout that is already in place in the business's databases. Also implementing new algorithms or integrating different data combinations is an easy task, thereby making this a perfect test platform for new developments in the area.

Link:

[L1] http://www.fraunhofer.at/grappa

References:

[1] A. Halm et al.: "Energy Balance: A web-based Visualization of Energy for Automotive Engineering using X3DOM", in "Proceeding of the International Conference CONTENT", 2013, pp. 1 - 10

[2] A. Halm et al.: "Time-based Visualization of Large Data-Sets -- An Example in the Context of Automotive Engineering", in "International journal on advances in software", 2014, pp. 139 - 149

Please contact:

Andreas Halm
Fraunhofer Austria Visual Computing
Tel +43 316 873 5410
andreas.halm@fraunhofer.at

ICT for a Logistics Demonstration Centre

by Miguel A. Barcelona, Aragón Institute of Technology

Information and Communication Technologies (ICT) have become a key element in improving collaboration and decision-making in the field of transport and logistics. The Spanish ICT4Logistics Demonstration Centre is a test-bed infrastructure where providers or potential customers may run pilots and demonstrations.

Today's business environment is characterised by complex and dynamic supply chains (SC) [1]: an organisation's decisions are increasingly influenced by the decisions of others who are linked in a value chain [2]. Consequently, Information and Communication Technologies (ICT) have become a key element in improving collaboration and decision-making in the field of transport and logistics.

The Spanish ICT4Logistics Demonstration Centre [L1] is a public initiative promoted by Red.es [2], the public corporate entity attached to the Ministry of Industry, Energy and Tourism which is responsible for promoting the development of the Information Society in Spain, in collaboration with the Aragon regional Government and the Aragon Institute of Technology (ITAINNOVA) [L3].

A public demonstration centre is a physical space in which ICT companies offer practical demonstrations of products and services designed to improve productivity and competitiveness, to other companies, which may be potential users of technology. Each demonstration centre also acts as a meeting point for regional innovation and the needs of small and medium enterprises (SMEs) in their respective demarcation, and fosters intersectoral collaboration and business development.

The main objectives of the ICT4Logistics Demonstration Centre are:

- To promote the creation of meeting spaces between suppliers in the ICT sector and the potential demand from logistics companies.
- To disseminate the benefits of incorporating ICT into production processes among companies in the logistics sector.
- To facilitate the transfer of technology, specialised services and knowl-

edge between the ICT sector and logistics companies.
- To advise, train, develop and transfer technology into companies, in particular SMEs.
- To provide logistics companies with the necessary infrastructure for testing the innovation technologies.

To this end the centre is composed of four elements:
- Smart warehouse: includes technologies applied to reception, picking, shipping, storage and automatic characterisation for freight, as well as to create a test-bed infrastructure for an intelligent warehouse scenario.
- Smart store: composed of technologies for the identification and collection of products, for real-time information about the existing stock of each product on the shelf and methods to efficiently replace products.
- Smart transport: contains equipment designed for the simulation of vehicle traffic allowing static and dynamic mapping.
- Smart supply chain: used to evaluate the potential application of identification, monitoring and control systems into the management of flow products in order to simulate and coordinate a better decision-making process.

Since the official opening in April 2013, more than 1000 people and 350 enterprises have attended more than 100 events, workshops, training sessions, demonstrations and sectorial events hosted in this centre. In particular, it hosted the 6th European Conference on ICT for Transport Logistics in 2013 and promotes the annual Spanish ICT4Log event.

The centre is being used to promote the transfer of knowledge into SMEs in practice. As a success story, in the Model and Inference Driven - Automated testing of services architectures (MIDAS) EU Project [L4], the ICT4Logistics Centre has been used to develop a test-bed infrastructure to guarantee the successful integration of service-based supply chain management systems according to GS1 Logistics Interoperability Model standard [3].

In the future we will work to extend this test-bed infrastructure by adding electronics, sensors and robotics as an Internet of Things (IoT) laboratory where providers or purchasers may run pilots and practical demonstrations.

Figure 1: Point of sale infrastructure.

Figure 2: Warehouse infrastructure.

Links:
[L1] http://www.cdlogistica.es/en
[L2] http://www.red.es
[L3] http://www.ict4log.eu
[L4] http://www.midas-project.eu

References:
[1] Chopra, S. & Meindl, P. Supply chain management: Strategy, planning & operation. Springer, 2007.
[2] Jardim-Goncalves, R., Agostinho, C., Sarraipa, J., Grilo, A., & Mendonça, J. P. Reference framework for enhanced interoperable collaborative networks in industrial organisations. International Journal of Computer Integrated Manufacturing,2013.
[3] GS1, GS1 Logistics Interoperability Model v1.0, http://www.gs1.org/lim

Please contact:
Miguel Ángel Barcelona
ITAINNOVA (Aragón Institute of Technology)
C/ María de Luna 7, 50018, Zaragoza, Spain
mabarcelona@itainnova.es

Production Planning on Supply Network and Plant Levels: The RobustPlaNet Approach

by Péter Egri, Dávid Gyulai, Botond Kádár and László Monostori, SZTAKI

In today's competitive manufacturing industry, efficient production planning is of crucial importance to match the supply with the customer order stream. The RobustPlaNet approach aims to support the decision makers at different hierarchical levels of production enterprises to react or even to proactively plan and manage their activities in a more robust way than it used to be in the traditional practice.

European manufacturing industry would benefit by transforming the current hierarchical supply chains into networks, based on partnerships with a new understanding of business approaches and services provided to the partners (B2B) and to the final customers (B2C). To this end, we investigated innovative decision-making methods that integrate production planning at single plants into the management of the whole global network. The research is funded by the European Union within the RobustPlaNet 7th Framework project, involving four academic and seven industrial partners from Hungary, Germany, Italy and The Netherlands.

Robustness aspects of production planning

In general, production planning methods rely on deterministic input data, hence fail to cope with the dynamic effects of the execution environment and the considerable uncertainty of the underlying planning information. Robust techniques that can provide feasible production plans are required to address these challenges. Robustness in production planning involves refined approaches that can handle predictable or unpredictable changes and disturbances, respond to the occurrence of uncertain events (reactive approaches) or protect the performance of the plan by anticipating uncertain events (proactive approaches) [1]. In our methodology, robust planning is the logical layer of the robust production. A production plan is considered robust if it achieves an acceptable level of the selected performance indicators even when unpredictable disruptions occur during the execution of the plan. The robustness of the system often works against other efficiency criteria, hence a trade-off is required if the objective is to increase system's robustness. In the RobustPlaNet project, robust planning methods are defined both at the network and plant levels of the production hierarchy (Figure 1).

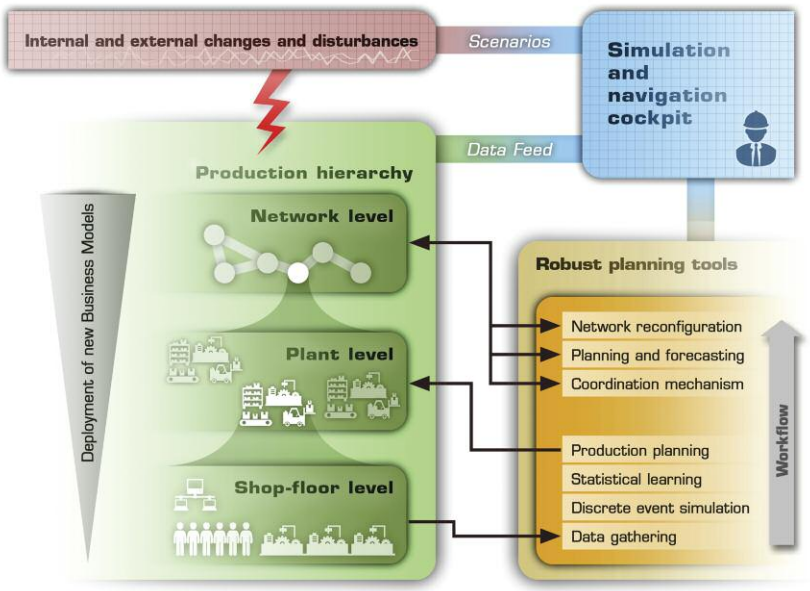

Figure 1: Robust planning methods of the RobustPlaNet project, defined at network and plant levels.

Planning in supply networks

The main goal on the network level is to enable robust and system-wide efficient production in a dynamic, unpredictable and global environment. Performance can be improved by implementing three complementary tasks. Firstly, the network design deals with configuration and re-configuration in order to provide a structure that is robust against diverse, largely unpredictable supply chain risks. Secondly, the supply chain planning matches supply and demand by considering statistically predictable fluctuations and disturbances. Finally, the coordination mechanism harmonizes the decisions made in a distributed manner, and supplements the risk management toolkit with risk sharing schemes.

The proposed methodology for supply network planning starts with identifying and analyzing the main sources of disturbances. Possible ways of managing these problems are then investigated, including estimated cost analysis. Owing to the complexity of a networked production system, the resultant alternative action plans are then evaluated in a multi-method simulation environment, enabling the combination of multi-agent, discrete event, and system dynamic modelling aspects in a general framework. This planning process will enable the network to respond to various circumstances, supporting efficient operations on the lower levels of the production hierarchy.

Robust production at the plant level

At the plant level in the production hierarchy, planning methods are aimed at calculating robust production plans that respect the possible uncertainties of a real production environment, maintaining the target level of key performance indicators, although disturbances occur during the execution of the plan. The robustness of the plans is provided

by proactive, simulation-optimization approaches that mostly rely on near real time algorithms processing actual historical production data that is used in discrete-event simulation models to analyze possible future production scenarios. By applying statistical learning methods on the dataset provided by the simulation experiments, the outcome of future scenarios can be predicted analytically, even considering the stochastic parameters and random events that were represented in the simulation model. The prediction models are integrated into the mathematical optimization models that are responsible for calculating the robust production plans [2].

IT platform to support the planning activities

These novel planning tools and models are all integrated in a common software platform called Simulation and Navigation Cockpit. Each of the functional modules can be triggered independently directly from this platform which employs the modules as black-boxes in the background and offers an intuitive web-based GUI on role-basis. Additionally, the integrated platform allows the user to define custom, problem-specific workflow mechanisms where the modules are chained sequentially, each operating on the same database. Depending on the organizational requirements of a production company, different roles can be granted with different data access levels within the integrated platform, ensuring that users can only access data that they are permitted to.

This research is supported the by the European Union 7th Framework Programme Project No: NMP2013-609087, Shock-robust Design of Plants and their Supply Chain Networks (RobustPlaNet).

Links:
www.robustplanet.eu
https://www.linkedin.com/groups/8137114

References:
[1] T. Tolio, M. Urgo, J. Váncza: "Robust production control against propagation of disruptions" CIRP Annals-Manufacturing Technology, vol. 60, no. 1, 2011.
[2] D. Gyulai, B. Kádár and L. Monostori: "Robust production planning and capacity control for flexible assembly lines" IFAC-PapersOnLine, vol. 48, no. 3, 2015.

Please contact:
Botond Kádár
Fraunhofer PMI, SZTAKI, Hungarian Academy of Sciences, Hungary
Tel: + 36 1 279 6176
botond.kadar@sztaki.mta.hu

Modelling and Validating an Import/Export Shipping Process

by Giorgio O. Spagnolo, Eda Marchetti, Alessandro Coco and Stefania Gnesi, ISTI-CNR

In recent years, business process management has become increasingly popular in many industrial contexts and application domains. This is mainly because it facilitates the modelling of process specifications and the development of an executable framework, while providing concise definitions and taxonomies. The data acquired during the business process execution phase can be used for quality analysis and to demonstrate compliance to specifications. We describe an experience in the real world context of the Livorno Port Authority.

Business process management usually relies on a Business Process Model (BPM), specified using one of the available Business Process Modelling Notations (BPMN). Working with the Livorno Port Authority, we successfully adopted BPM data analysis to find commonalities and discrepancies between the mandatory guidelines for controlling the arrival and departure of goods and the current implementation of the shipping systems known as the Tuscan Port Community System (TPCS) [3].

TPCS is a web-services based platform with multilevel access control and data recovery facilities, which aims to support and strengthen the management of shipping lanes according to Italian regulations. TPCS processes a huge amount of information, enabling a reduction in costs and the streamlining of administrative procedures. It can be used for the complete management of all connected applications handling cargo movements and import/export operations, and involving all the actors interested in the information flow, such as shipping agencies, custom forwarders, freight forwarders, terminals, hauliers, and also the Control Authorities.

The users can generate and manage Cargo Manifests during export, and Unloading Lists during import operations. In both operations, users can interact with the platform in order to know the status of goods (e.g. loaded or unloaded, cleared for customs etc.). They can also request and receive goods certificates and authorizations (e.g. phytosanitary authorization) for dangerous goods via the SUD (Sportello Unico Doganale - Customs Single Window) interface.

A storytelling approach has been used to create the business process model for the TPCS. The method is summarised in Figure 2, where three main stakeholders carry out the tasks proposed by the storytelling methodology: the Tellers(Port Community Experts), the Facilitator and the Modellers(BPM Experts). The Tellers are those individuals who participate in the process and therefore have domain knowledge. They are asked to describe their activities explicitly through a story. The Facilitator is an experienced professional in the application domain who helps the story tellers to produce coherent stories and the first abstraction of the models. The

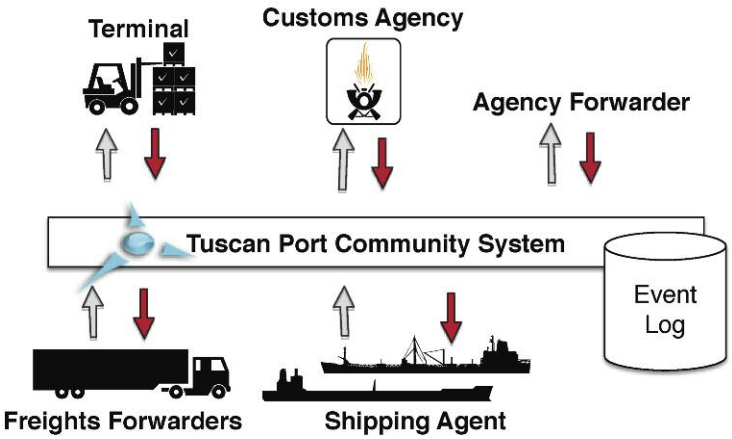

Figure 1: The actors involved in the TPCS and their interaction.

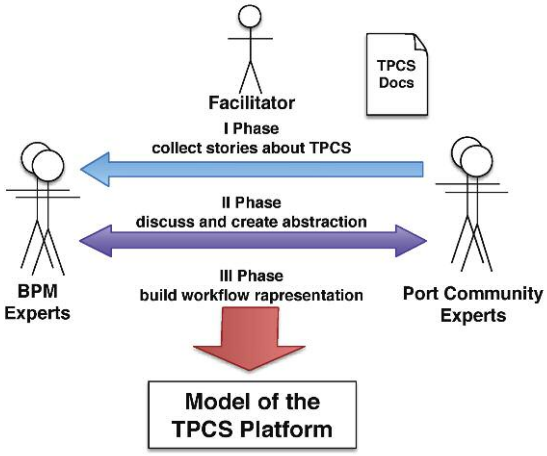

Figure 2: Methodology applied to create the model.

Modellers are process analysts who refine the graphical model developed based on abstractions extracted from the stories.

An event logger has been integrated into the TPCS code system in order to monitor the information exchange within the platform and between the different web-services, and to collect data to assess conformance to the BPM specification.

Conformance checking techniques have been used to relate events in the event log to activities in the process model and compare them [1-2]. The experience highlights important challenges in the application of process mining techniques and enables inconsistencies detected during the process execution to be promptly corrected. The mining activity has been confirmed to be a useful means for quality assurance and the checking of software in operation. However, this experiment has also shown that the identification of the BPM elements can be a key factor for the final results. The identification of the rules, the policies, the roles, and the responsibilities, as well as the interactions between users and the platform represent the main criticalities when deriving an accurate business processes model and when correctly managing it.

Once the business process of the TPCS platform was derived using the story telling techniques, the platform was then equipped with the activity process log. Next, using process mining techniques on the log of the platform with the process model derived, we performed conformance analysis. The results obtained were used to investigate the behaviour of the platform and validate the process model.

This work is a contribution to the industrial application of formal techniques for the monitoring of business processes.

References:

[1] W. van der Aalst, et al.: "Replaying history on process models for conformance checking and performance analysis", Wiley Int. Rev. Data Min. and Knowl. Disc., March 2012.
[2] R. M. Dijkman, et al.: "Semantics and analysis of business process models in BPMN", Inf. Softw. Technol., November 2008.
[3] G. O. Spagnolo, et al.: "An experience on applying process mining techniques to the Tuscan Port Community System", LNBI Proc., 238, 2016.

Please contact:

Giorgio O. Spagnolo, ISTI-CNR, Italy
Spagnolo@isti.cnr.it

Where is the Money? - Optimizing Cash Supply Chain Networks

by Leendert Kok and Joaquim Gromicho, ORTEC

"German CIT truck robbed with Bazooka" 12-12-2015, "Attack on CIT in Amsterdam" 4-06-2015, "Failed robbery on the Italian highways" 15-05-2015. These are three examples of the motivation of one part of a joint research project of ORTEC, Geldservice Nederland (GSN), and the Vrije Universiteit (VU) Amsterdam called 'Optimizing cash supply chain networks'.

One of the goals in cash supply chains is to manage the inventory of ATMs. GSN, a cooperation of banks, manages the inventory of the majority of ATMs in the Netherlands. To this end, they generate money pickup and delivery orders which they outsource to Cash in Transit (CIT) companies to deliver the money to and from the ATMs. Since trucks carrying money are attractive to robbers, an important goal in planning the delivery routes (besides being efficient to lower transport costs), is to make the routes unpredictable.

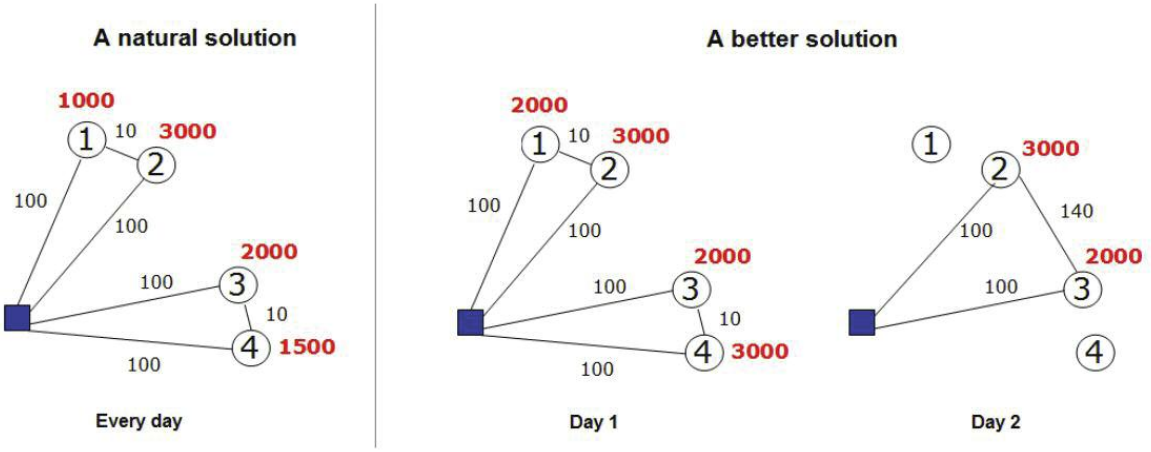

Figure 1: Example of improving the long term transport costs per delivered volume ([3]).

In 2015, three PhD students at the VU started working on this problem. The problem can be seen as a special case of the inventory routing problem, see [1], where the goal is to manage the inventory of fuel for a set of customers, for example, by making sure that they do not run out of stock. The long term goal is to minimize transport costs. The problem is a generalization of the well-known vehicle routing problem (VRP), by adding two extra decisions to the problem: when to deliver and how much.

Within the cash replenishment business, an extra requirement appears for these delivery routes: they should be unpredictable. In a sense, the problem is the opposite of the consistent VRP [2], where the goal is to deliver to the same customer at the same time of the day and the same day of the week. There are different measures for the unpredictability of vehicle routes: the time of replenishment, the sequence of replenishments, and the routes driven between two replenishments (K-shortest paths).

The PhD students developed a model for varying the time of delivery by generating multiple non-overlapping time windows with the last delivery for each specific ATM. The VRP with multiple time windows has received little attention in the literature. A special modelling characteristic in cash supply is that waiting time at customers (ATMs in this case) is not allowed. Especially when time windows are tight, this extra constraint has a major impact on solution methods for this vehicle routing problem.

As mentioned above, a key decision within inventory routing is the moment of delivery and the amount to deliver. To minimize the long run transport costs, an objective was introduced by [3] which minimizes transport costs by delivered volume. Figure 1 illustrates how a larger transport cost still leads to a better plan with respect to the long term transport costs.

ORTEC, one of the world's largest suppliers of advanced planning software, has extensive experience in solving Inventory Routing Problems and, with the involvement of many master's students, is continuously updating its innovative solutions. One of the current master's students is looking at ways to improve the inventory routing solutions by making a clever preselection in which customers are selected for delivery. The student developed a mechanism that decides which so-called may-go orders (orders that may be postponed until the next day) should be included in the batch for the planning day. The results are very promising and indicate that with minimal computational effort, the long term transportation costs can be reduced by more than 5%.

The final part of the research project focuses on contract design. Currently, GSN decides on the orders, and third parties are responsible for delivering the money. This leads to suboptimal solutions (the preferred amount to deliver may depend on the driven vehicle routes). Moreover, service level agreements are done between GSN and the banks, but the logistics companies are crucial for fulfilling them. This research aims to develop incentives to better deal with all these challenges.

References:

[1] L. Bertazzi, S. Bertazzi, M. Grazia: "Inventory routing problems: an introduction", EURO Journal on Transportation and Logistics, Vol. 1, No. 4, pp. 307-326; 2012

[2] C. Groër, B. Golden, E. Wasil: "The Consistent Vehicle Routing Problem", Manufacturing & Service Operations Management, Vol. 11, No. 4, pp. 630–643, 2009.

[3] W. J. Bell, et al.: "Improving the distribution of industrial gases with an online computerized routing and scheduling optimizer", Interfaces, 13, pp. 4–23, 1983.

Please contact:

Leendert Kok
ORTEC, The Netherlands
leendert.kok@ortec.com

An Industrial Take on Breast Cancer Treatment

by Sara Gestrelius and Martin Aronsson, SICS Swedish ICT

When a woman is diagnosed with breast cancer she is drawn into a network of examinations, tests and treatments with interdependencies in several steps. Because the prognosis for recovery is tied to how quickly she gets help, it is important to reduce the time from the first doctor's appointment to the start of treatment.

In early 2013 the health authorities of Region Östergötland in Sweden were investigating why the time between the first doctor's appointment and the start of treatment was longer than the goal of 28 days for breast cancer patients at the hospital in Linköping. To help answer their question, researchers from SICS Swedish ICT were tasked with proposing and testing operational research methods that could elucidate how the different process steps contributed to the overall lead time. The methods proposed should also allow for different process changes to be investigated.

The pilot study began with an investigation step where the researchers learned about all parts of the care process. This included interviews with doctors and administrative staff, and also analysis of data that the department had gathered over the years. The initial analysis resulted in some important insights about the process, e.g. that the patient inflow varied depending on the time of the year, and that the process had a number of discrete steps that negatively affect the lead time. In fact, based on the initial analysis it could be proven that the goal lead time of 28 days or less was impossible to live up to for many patients, given the current process and methods. Patients arriving on certain week-days were out of sync with the discrete process-steps, which resulted in long lead-times.

The discrete steps are often a consequence of doctors having certain days allocated to certain tasks, or operating rooms only being available on certain days. This is because operating rooms are shared among departments, and set-up times prevent doctors from performing all tasks every day. For example, the doctors do not have time to hold multidisciplinary conferences every day. The focus on high resource efficiency negatively affected the patient flow.

Based on the understanding acquired in the investigation step we constructed a discrete event simulation. The arrival process was modelled as an inhomogeneous Poisson process based on the historic data, and the basic rules of the simulation model were reviewed by the interviewed staff. The purpose of the simulation was to investigate whether there was potential for shorter lead times, and if so, where to invest in resources. First of all the lead times from the simulated standard process and the historic data were compared, and the discrepancies were discussed with the hospital staff. Then different proposed process changes were tested in the simulation to provide insights that were otherwise difficult to gain. For instance, the simulation could detect when increasing resources in a particular process step did not result in time gains because the patient still had to wait for other parts of the process. The simulation model was also used to estimate how congested the different process steps were. Being able to explore the process, and predict and quantify in advance how changes affect the lead time and the load curve, stimulated great interest among the hospital staff.

In addition, the hospital identified a few departments with resources that were likely to benefit from optimizing planning support systems. Operating rooms, as well as other shared resources such as pathology analysis competence, are interesting subjects for optimization and planning as they deal with demands from many patient streams and categories.

Shorter lead times are closely linked to the question of good planning. Of course you always have to be particularly considerate when working with people. Perhaps the patient needs a few extra days, for instance, to take in the news that she has cancer before she is ready to go into surgery. Nevertheless, SICS believes that a more industrial approach to the planning process would benefit everybody involved. Apart from using simulation there are other operational research techniques that are apt for healthcare systems, such as optimization of schedules [1].

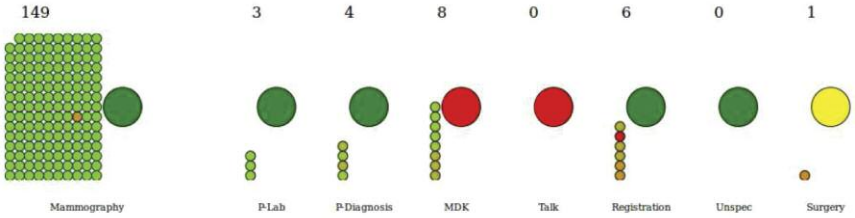

Figure 1: Example of a simulation of a care process. The small circles are patients, and the colour represents time since their first doctor's appointment. The large circles are process-steps that are open (green), open but working at full capacity (yellow) or closed (red).

The SICS researchers involved with the project consider healthcare a perfect candidate for their planning methods. They think better decision support tools would enable a care that is more adapted and coordinated, both for the patient and for the health professionals.

Link:
http://www.sics.se/projects/flecs-flow-efficiency-in-cancer-treatment

Reference:
[1] M. Persson: "On the Improvement of Healthcare Management Using Simulation and Optimisation" Blekinge Institute of Technology, Karlskrona, doctoral thesis, 2010.

Please contact:
Sara Gestrelius, Martin Aronsson
SICS Swedish ICT
sarag@sics.se, martin@sics.se

Evaluating Operational Feasibility before Investing: Shunting Yards in Sweden

by Sara Gestrelius, SICS Swedish ICT

The cost of building a new shunting yard is in the order of €100 million and the expected lifespan is over 50 years. Discussions about new railway infrastructure investments are currently ongoing in Sweden, and SICS Swedish ICT are using mathematical optimization to solve NP-hard shunting yard problems to help the Swedish Infrastructure Manager choose between different yard layouts.

Shunting yards are the hubs in the network that distributes freight cars all over Sweden. Inbound trains arrive at the yard where their cars are sorted into new outbound trains. Shunting yards often consist of different sub-yards: an arrival yard where the trains arrive and are decoupled, a classification bowl where the cars are sorted into new outbound trains, and a departure yard where finished outbound trains can wait for their departure time. The system can be likened to a factory with incoming stock (the arrival yard), a production facility (the classification yard), and outgoing stock (the departure yard). The arrival yard is connected to the classification bowl via a hump, and the cars roll from the arrival yard to the classification bowl by means of gravity. However, not all shunting yards have a departure yard, and the number of tracks in the various sub-yards vary. Figure 1 shows a graphical representation of a typical shunting yard with all three sub-yards.

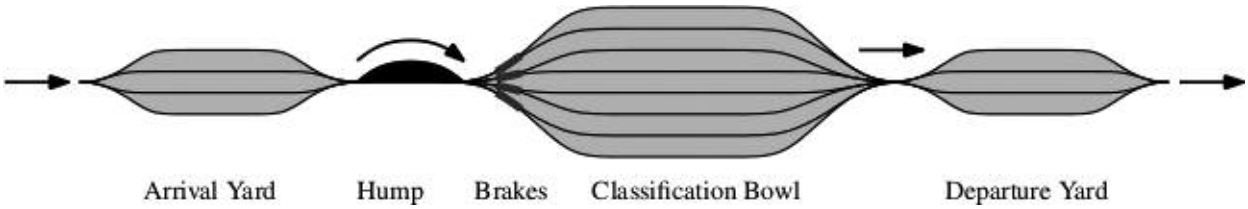

Figure 1: Graphical representation of the typical lay-out of a shunting yard.

The tracks required in the various sub-yards for successfully building all outbound trains on time are highly dependent on the trains' arrival and departure times, the car bookings and the operation of the shunting yard. In Sweden, shunting yards are operated in a particular way and the classification bowl tracks are divided into train formation tracks and mixing tracks. Train formation tracks are used for compounding outbound trains while mixing tracks are used as a storage place for early cars whose outbound trains are not yet being built. The mixing tracks are emptied by pulling the cars on the mixing tracks back to the arrival yard and then rolling them over the hump again. It is necessary to do detailed operational planning to know the track requirements, but unfortunately scheduling a shunting yard using Swedish operational practices is an NP-complete problem [1].

SICS Swedish ICT has previously worked with RWTH Aachen University and ETH Zürich to develop mathematical optimization models and methods for planning shunting yards [1]. A heuristic is used for scheduling the arrival yard, the departure yard, and the times when the mixing tracks are emptied. The output from the heuristic defines an optimization problem for the car sorting in the classification bowl, where the objective is to minimize the number of extra car shunt moves brought about by mixing. In this new project the models are used to generate shunting schedules for different yard layouts. The schedules can be analysed to estimate how much (or little) work that is required to process a certain set of trains given a specific shunting yard layout. Fewer track resources will lead to more shunting work and vice versa. By varying the number of tracks in each sub-yard the trade-off between the number of tracks and the shunting work can be analysed, and the goal is to provide the Infrastructure Manager with data that helps them find a good balance between the cost of building tracks and the cost of shunting. In particular, it is important for the Swedish Infrastructure Manager to identify when the track resources are too scarce to operate the intended traffic in a robust and proper manner.

The first project aiming at investigating shunting yard layouts using optimization was started during spring 2015, and its successor is due to finish at the end of 2016. Within the second project two shunting yards are currently being investigated, Hallsberg Rangerbangård and Sävenäs Rangerbangård. In the Hallsberg case study the departure yard usage is in focus, and in particular questions regarding the trade-off between arrival yard usage and the departure yard usage are being evaluated. Methods for investigating flat shunting are also being developed as a complement to the existing models. It is possible that some freight trains only require car swaps rather than the full resorting offered by shunting yards.

In Sävenäs there are plans for building a completely new yard, and therefore a few different, sometimes completely new and inventive, yard set-ups are assessed. The traffic in Sävenäs also requires some further sorting with respect to the car-order in the outbound trains as there are groups of cars that

should be dropped off along the train's journey. Methods for this extra sorting are being developed.

Another pertinent question regarding freight trains is the routing of cars and models for optimizing the routing are currently being developed. Experimenting with different routings can be useful for improving the overall efficiency of the shunting system, and also for investigating the effect of closing one shunting yard partially or completely. The latter is relevant when discussing re-investment in shunting yards as it may require existing shunting yards to be closed down, which makes the economic cost of the disturbed traffic an important factor when deciding between different investment alternatives.

As stated above, heuristics are part of the current schedule generation. As part of future work the project will therefore also look into the possibility of constructing an optimization model that can handle all three sub-yards.

Link:
https://www.sics.se/projects/pragge

Reference:
[1] M. Bohlin, S. Gestrelius, F. Dahms, M. Mihalák and H. Flier, "Optimization Methods for Multistage Freight Train Formation", Transportation Science, Articles in Advance, 2015.

Please contact:
Martin Aronsson, Zohreh Ranjbar, Martin Joborn, Sara Gestrelius, Markus Bohlin,
SICS Swedish ICT
martin@sics.se, Zohreh.Ranjbar@sics.se, mjoborn@sics.se, sarag@sics.se, bola@sics.se

Boosting the Responsiveness of Firefighter Services with Mathematical Programming

by Pieter van den Berg, TU Delft, Guido Legemaate, Amsterdam Fire Department, and Rob van der Mei, CWI

In life-threatening situations where every second counts, the timely arrival of firefighter services can make the difference between survival and death. Motivated by this, the Stochastics Department at CWI in collaboration with the TU Delft and the Fire Department Amsterdam-Amstelland in the Netherlands have developed a mathematical programming model for determining the optimal locations of the vehicle base stations, and for optimally distributing the different firefighter vehicle types over the base stations. Extensive analysis of a large data set for the Amsterdam area demonstrates that, and how, response time can be improved by relocating only few base locations and redistributing different vehicle types over the base locations.

The Dutch capital Amsterdam was the first city in the Netherlands with a professional fire service, which was established in 1874. With 144 personnel and nine fire stations covering 30 square kilometres, it ensured fire protection for approximately 285,000 inhabitants. Today, it is the regionally organized fire department Amsterdam-Amstelland that, with 1,150 personnel and 19 fire stations covering 354 square kilometres, is responsible for over 1,000,000 inhabitants. Obviously, over time the questions of how many fire stations were needed for a given coverage and where to locate them had to be answered numerous times as new needs and means for fire protection emerged.

In this study, we introduce a model to (1) determine optimal locations for the fire stations, and (2) to find an optimal distribution of the vehicles over the selected locations.

We divide the region into a set of demand points from which calls can arise, and define a subset of locations that can be used as a potential base location. We consider different types of vehicles that are used for a different type of call. Each type can have different response time targets. We limit the number of vehicles that we may use of each type. For each demand point and for each vehicle type, we have a weight that indicates the importance of covering a certain demand point by a vehicle of a specific type. The expected number of calls is commonly used for the weights, but other risk measures could be included as well. To determine the coverage, we introduce a subset of the base locations that contains all locations that are sufficiently close to cover our demand point by the specific vehicle. Whether a base location is close enough depends on three variables: the travel time from the base to a location, the pre-trip delay, and the response time target. We assume that the travel times are the same for each vehicle type and that the pre-trip delay, which is the time elapsed before the vehicle starts driving to the scene, is fixed. This pre-trip delay also consists of two parts: (1) triage and dispatch, and (2) chute time. The triage and dispatch time is the time spent in the call centre to assess the importance of the call and assign a vehicle. The chute time is the amount of time that elapses between the assignment of a call and the crew's departure from the base. The response time target can differ for each vehicle type and for each demand point. For our model, the goal is to maximize the number of calls that are covered by the appropriate vehicle type, while minimizing the number of base locations used [3].

To apply the model to the region of Amsterdam, we defined a set of 2,643 demand points and 2,223 potential base locations. Travel times between potential base locations and demand points are provided by the fire department and are based on estimated travel times on the road network between each location. In our analysis, we include the four most common types of vehicles used at Dutch fire departments: fire apparatus (FA), aerial apparatus (AA), rescue

Figure 1: Overview of different types of firefighter vehicles: fire apparatus (top left), aerial apparatus (top right), rescue apparatus (bottom left) and marine rescue units (bottom right). Photos: Jeffrey Koper.

apparatus (RA) and marine rescue units (MR), see Figure 1 for an illustration of the different vehicle types.

The number of available vehicles of each type is respectively 22, nine, three and two. In the current configuration, we have 19 bases. We take the absolute number of calls per vehicle that occurred in 2011 to model the recurring risk. In this period, the number of calls per vehicle type is 29,016, 9,182, 615 and 703, respectively. As an illustration, Figure 2 shows a fictitious relocation of four fire stations.

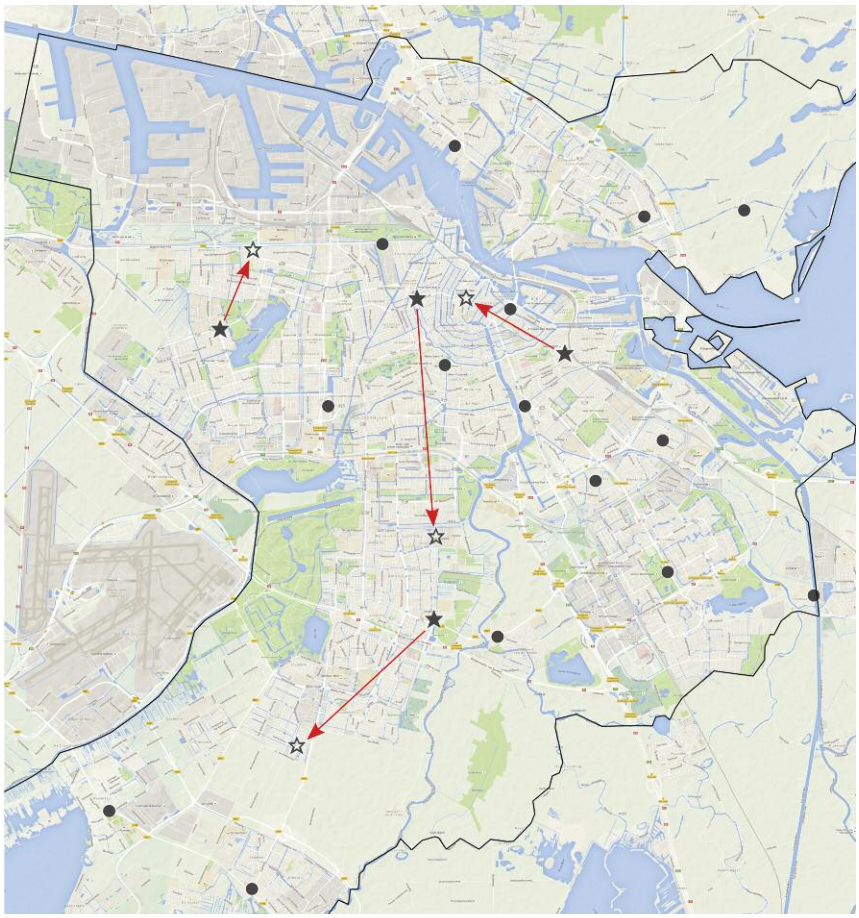

Figure 2: Possible distribution of fire departments where four movements are allowed. Source: G.A.G. Legemaate.

Extensive analysis of a large dataset of historical incidents demonstrates:

- that, and how, a response time can be improved by simply relocating only three out of 19 base locations and redistribution of the different vehicle types over the base locations;
- that there is no need to add new base locations to improve performance: optimization of the locations of the current base stations is just as effective.

The results show an enormous potential for improving the arrival time for firefighter services worldwide. Furthermore, because relocating existing base locations might be considered costly, with only minor adaptations the model can also be used to relocate fire trucks and personnel to temporary locations during large scale incidents to improve coverage.

Link: http://repro.project.cwi.nl

References:

[1] P. Kolesar, W.E. Walker: "An algorithm for the dynamic relocation of fire companies", Operations Research 22 (2), 249-27, 1974

[2] A.J. Swersey: "A Markovian decision model for deciding how many fire companies to dispatch" Management Science 28, 352-365, 1982

[3] P.L. van den Berg, G.A.G. Legemaate, R.D. van der Mei: "Boosting the responsiveness of firefighter services by relocating few base stations in Amsterdam", 2015. Under review (a preprint available from the authors upon request)

Please contact:

Rob van der Mei
CWI, The Netherlands
mei@cwi.nl

Pieter van den Berg
TU Delft, The Netherlands
P.L.vandenBerg@tudelft.nl

Guido Legemaate
Fire Department Amsterdam-Amstelland, CWI, The Netherlands
g.legemaate@brandweeraa.nl

Predicting the Demand for Charging Stations for Electric Vehicles

by Merel Steenbrink, Elenna Dugundji and Rob van der Mei, CWI

At the Stochastics research group at CWI we use socio-economic features of neighbourhoods to predict the demand for charging stations for electric vehicles. Based on a large set of behavioural data, a discrete choice model is estimated, and the utility of charging stations in every neighbourhood is calculated. In this way the demand can be predicted and the municipality of Amsterdam can proceed proactively.

The municipality of Amsterdam would like to stimulate the use of electric vehicles in the city. To this end, a substantial number of charging stations [L1] have been constructed in the city. Over 1,000 stations are already in operation, but the city wants to further increase this number to 4,000 by 2018 [L2, L3, L4]. The policy is quite encouraging: if you buy an electric vehicle, and there's no charging station within a 200 metre radius of your residence, you can apply for one in front of your door. However, the process of investigating suitable places is a time-consuming task. If it were possible to predict where the demand for charging stations will be, the scaling-up process could take place more effectively.

In this research a Discrete Choice Model was estimated to predict the demand for charging stations. Every charging activity can be seen as a choice for a certain station over other stations. In this way the utility of every station can be determined as a function of characteristics of the stations. When this function is estimated, the utility of future stations can be determined, based on their characteristics. In this way the municipality can focus on potential neighbourhoods, where the utility is high, and already investigate the appropriateness of these locations even before any requests have been made.

To estimate the model, behavioural data from 2012, 2013 and 2014 were used. The data describe the charging activity at every station, which reveals how often, how much and how long a poll has been used. To be able to use different characteristics, and to decrease the size of the choice set, the choices have been aggregated to the area of neighbourhoods. Every choice for a charging station is a choice for a neighbourhood. Socio-economic data from the Central Bureau for Statistics (CBS) have been incorporated into the model. The model takes into account: average number of cars per household, percentage "Western" inhabitants (persons with nationality from Europe, North America, Indonesia or Japan), the percentage of apartments. Furthermore it considers the mean percentage of time per day that the stations in each neighbourhood were used. Because the data were aggregated there was a compensating term with the number of stations per neighbourhood.

We estimated a Discrete Choice Model which calculates the probability of an electric car driver choosing a charging station in a certain neighbourhood [1],[2],[3]. We calculated this probability for every neighbourhood and in this way obtained a ranked list of neighbourhoods with high and low probabilities. We thereby are able to advise the municipality in which neighbourhoods to focus their expansion activities.

Until now we have incorporated into our model socio-economic characteristics for which data were readily available from the CBS. Further work coupling with other sources of data is likely to identify other relevant variables - for instance, proximity to malls or large office buildings may turn out to be important. These factors will be added to the model as we develop it. The model also does not currently differentiate between day and night time charging sessions. This consideration could reveal additional variation in the usage patterns of the charging stations in each neighbourhood.

Some neighbourhoods with low population density have no aggregate socio-economic statistics reported by CBS to protect the privacy of these residents. These neighborhoods have been currently excluded from this investigation due to lack of descriptive data. However, many of these neighbourhoods – largely industrial areas or sporting facilities – may require charging stations. Additional data on land use and the built environment may allow us to include these neighbourhoods.

Finally, when we have ascertained which neighbourhoods need charging stations, the next step is to determine the exact locations of these charging stations in each neighbourhood.

Links:

[L1] https://chargemap.com/city/amsterdam
[L2] http://kwz.me/U3
[L3] http://kwz.me/U5
[L3] http://kwz.me/U6

Refernces:

[1] F. Goetzke, R. Gerike, A. Páez, E.R. Dugundji: "Social Interactions in Transportation: Analyzing Groups and Spatial Networks", Transportation, 42 (5), 2015, 723-731. http://dx.doi.org/10.1007/s11116-015-9643-9

[2] M. Maness, C. Cirillo and E.R. Dugundji: "Generalized Behavioral Framework for Choice Models of Social Influence: Behavioral and Data Concerns in Travel Behavior", Journal of Transport Geography, 46, 2015, 137–150. http://dx.doi.org/10.1016/j.jtrangeo.2015.06.005

[3] E.R. Dugundji, J.L. Walker: "Discrete Choice with Social and Spatial Network Interdependencies: An Empirical example using Mixed Generalized Extreme Value Models with Field and Panel Effects", Transportation Research Record, No. 1921, 2005, 70-78. http://dx.doi.org/10.3141/1921-09

Please contact:

Elenna Dugundji
CWI, The Netherlands
e.r.dugundji@cwi.nl

Data-driven Optimization for Intelligent and Efficient Transport

by Björn Bjurling and Fehmi Ben Abdesslem, SICS Swedish ICT

Data-driven models derived with big data techniques can be used to improve and automate strategic and tactical decision making in heavy-duty road transportation.

Tactical decision making at carriers (e.g. assignment planning and vehicle allocation), relies largely on the skills and experience of individual dispatchers and decision makers. Decisions are typically supported by both commercial route planning software and physical models and statistics [1]. For larger carriers, substantial cost savings in terms of fuel consumption and driver scheduling may be achieved by improving decision support models and by considering larger data sets [2]. Recent developments in intelligent data analysis and computing technologies have made it possible to extract valuable information from big data, which in many domains has revolutionalized business models and increased the competitiveness of early adopters. Recently, there has also been intense activity in recording and collecting data in the automotive and transportation domains. The Data-driven Optimization for Intelligent and Efficient Transport (DOIT) project takes advantage of the advances in emerging big data technologies to analyse this data with the aim of improving decision making in the transportation industry.

One of DOIT's use-cases is concerned with improving assignment planning: to find the most cost-effective combination of vehicles, drivers, and routes for a given set of transport assignments (with specification of pick-up/delivery location and time, as well as information about the cargo). This can be modelled as a variant of the Vehicle Routing Problem (VRP). The main issues to address in this use case are: developing data-driven cost models and devising ways of using these new models together with existing commercial planning tools.

The DOIT develops data-driven models for fuel consumption and travel time estimations with the aim of improving decision making in the transportation industry.

DOIT develops data-driven models for fuel consumption and travel time estimations as input to the VRP. These models exploit big datasets generated from vehicles, carriers, and manufacturers. Thousands of trucks generate data about their fuel consumption and travelling time, and these data are correlated with external attributes such as: time of day, route taken, specifications of the truck, freight weight, and even the weather and road conditions on that day.

With data-driven models, the solution to the VRP can thus take into account a larger set of attributes in estimating the corresponding costs, and thus gain in accuracy timeliness. There is, however, no obvious way to extract sufficiently certain and exhaustive assignment data as input to the VRP. Although basic assignments can be detected from collected data with some certainty, it is not possible to satisfactorily infer critical and realistic assignment information (e.g. earliest pick-up time or routing preferences) without making use of exogenous statistical models. DOIT will instead make use of data-driven cost models based on collected trip data to infer estimated costs for given assignment, where a trip is defined as the time between a vehicle starting and its first stop.

Link:
https://www.sics.se/projects/doit

References:
[1] D. Macias: "Estimation of Fuel Consumption for Real Time Implementation", Master Thesis, KTH, Automatic Control, 2012.
[2] Winckler: "BMW Eco Navigation" BMW Group Technology Office, Presentation Notes, 2011.

Please contact:
Björn Bjurling,
SICS Swedish ICT
bgb@sics.se

Adopting a Machine Learning Approach in the Design of Smart Transportation Systems

by Davide Bacciu, Antonio Carta, Stefania Gnesi and Laura Semini, ISTI-CNR

We have applied a machine learning approach to both implement and assess new services for the users of a bike-sharing system. The aim is to predict the destination station of a bike in use, given information on its pick up details.

Bike-sharing systems (BSSs) are a means of smart transportation in urban environments with the benefit of a positive impact on urban mobility: the availability of a public bike, permitting point to point displacement can reduce the use of private cars in town.

The design of a BSS is multi-faceted and complex: a BSS has many components – including the human users as well as the bikes and stations. The users are an intrinsic part of any BSS, and their individual patterns of behaviour and specific preferences and needs have a decisive impact on the collective usability and performance of the system. Other issues that must be taken into consideration include the costs of installing, maintaining and running the BSS, and the particular characteristics of the city. Municipalities obviously wish to offer a successful and convenient service, so that people continue using it, reducing traffic congestion.

One of the main problems of a BSS is balancing the bicycles among the different stations. Some stations tend to fill up while others are empty, with these situations varying dynamically during the day. Studies have shown that user satisfaction can be improved by providing information on the status of the stations at run time. Most current systems supply this information in terms of number of bicycles parked in each docking station, through services available via the Internet. When the arrival station is full or the departure station empty, the user needs to know whether this situation is likely to change soon. A service is needed that can predict whether, at a given moment, there is high probability that someone will take a bike, making a place available in a full station, or that someone will return a bike to an empty station.

Machine Learning (ML) methodologies could be used to learn computational models of prediction features from BSS usage data [1]. Machine learning can be applied to realize data-driven adaptive approaches to predictive modelling of input-output data relationships. Through these methodologies, it is possible to learn the (unknown) relationship between a feature and its inputs by exploiting historical data providing examples of such an input-output map. A trained model can then be used to provide predictions on future values of the feature in response to new input information, i.e. providing an implementation of the feature component. This approach has two advantages. On the one hand, it provides a powerful and general means to realize a wide choice of predictive features for which there exist sufficient (and significant) historical data. On the other hand, trained ML models are provided with a measure of predictive performance that can be used as a metric to assess the cost-performance trade-off of the feature. This provides a principled way to assess the runtime behaviour of different components before putting them into operation.

A CicloPi Bike Station in Pisa.

We are using this machine learning approach to both implement and assess new predictive services for the users of a BSS. One example is LocationPreview, a service that predicts the destination station of a bike in circulation given information on its pick up details. This feature is implemented using one of the sub-features: StationProfile and UserProfile. These subfeatures are both based on training machine learning models, and are paradigmatic of a class of learning tasks dealing with static vectorial data. They differ in the data they use: the first is built on a feature that maintains the history (more than 280,000 entries) of all movements station-to-station with date and time, but user unaware (to train the models), the latter exploits user profiles by training personalized predictors.

Part of this research has been done in the context of the European FP7 project QUANTICOL [L1], on the quantitative analysis of smart transportation systems. In addition, we have begun collaboration with PisaMo S.p.A., an in-house public mobility company of the Municipality of Pisa, which introduced the BSS CicloPi in Pisa a few years ago.

Link:
[L1] http://www.quanticol.eu

Reference:
[1] D. Bacciu, S. Gnesi, L. Semini: "Using a Machine Learning Approach to Implement and Evaluate Product Line Features", WWV 2015, El. Proc. in Theoretical Computer Science, vol. 188. EPTCS, 2015, p. 75-83.

Please contact:
Stefania Gnesi, ISTI-CNR, Italy
stefania.gnesi@isti.cnr.it

Remote Service Using Augmented Reality

by Björn Löfvendahl, SICS Swedish ICT

For a global industrial company like ABB, sending maintenance specialists around the world costs a lot of money and is neither time efficient nor environmentally friendly. Augmented reality allows the specialist to stay in the office and guide a local employee remotely.

ABB FACTS runs large power installations in many parts of the world, often located far from cities and airports. Usually the sites have their own service personnel, but sometimes when unexpected problems occur it is necessary to send in maintenance specialists, incurring significant expenses for both ABB and its customers. Furthermore, flying in specialists is neither time efficient nor environmentally friendly. This raises the question as to whether the problems could be solved remotely without having to fly in external personnel.

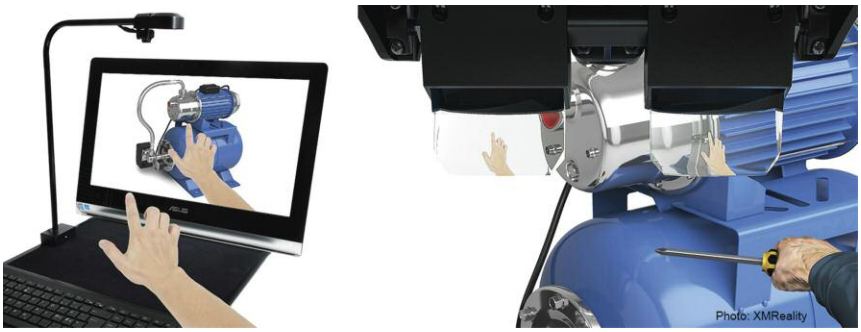

Figure 1: The AR equipment. A screen together with a camera pointing downwards for the guide (left) and a pair of HUD goggles with a front facing camera for the follower (right).

The project INCODE [L1] (Information and Competence on Demand) investigates how modern technology such as augmented reality (AR) and virtual reality (VR) can be used to facilitate the transfer of knowledge within a company. INCODE is run by SICS Swedish ICT Västerås together with Interactive Institute Swedish ICT Piteå and several companies, as part of the Process Industrial IT and Automation programme (PiiA), a Swedish strategic innovation programme. The project will run from summer 2015 to spring 2017.

AR is about superimposing digital objects like 3D models or text onto the real world around you, as opposed to VR where you are totally immersed in a virtual world. By using AR equipment developed by XMReality, one of the project partners, the project aims at solving ABB's problem with dispersed service. With this equipment it is possible to guide another person, a "follower", remotely using both voice and hands (see Figure 1). The guide sees what the follower sees and can use her hands to show how certain task should be performed. The follower wears AR goggles and can both hear the guide's voice and see her hands superimposed in front of him. Using AR for remote guidance has been done before [1], but what is interesting about this equipment is the potential to use gestures in a natural way.

In a pilot study prior to INCODE the AR equipment was evaluated. Tests were performed in an ABB laboratory where two different ABB technicians were given a maintenance task to complete. One user performed the task guided with the AR equipment and another user was only guided using a mobile phone. Although this was only a small test there were some interesting observations. The technician with the AR equipment perceived a much higher sense of safety compared to the other, as the guide could see what the technician using AR was doing and correct him if he did something wrong. A third test with a user, equipped with AR gear but with no connection to or knowledge of ABB's equipment, was also performed. This user had no problems performing the same task as the more experienced technicians. In fact, some parts of the test went even faster compared to the other since after a while this user started acting like a robot, doing what the guide told her to do without questioning.

INCODE continues the work started in the pilot and plans to do tests in real industrial environments. There are several conditions that have to be studied further, including:

- What happens if the connection between the guide and the follower is broken?
- Is it possible to integrate XMReality's and a customer's systems and what advantages would this have?
- What routines must be implemented to make sure that the follower will not expose himself to danger just because the guide tells him to perform a certain task?

Being able to quickly perform remote service has many advantages over having to travel a long distance and perform the service in place. For ABB's customer it would mean that the installation would be up and running much faster, which would save them a lot of money. ABB would be able to coordinate maintenance tasks much more easily and the specialists would not have to waste time on travelling. Reduced travelling would also mean reduced travel costs and a more environmentally friendly way of working.

There are still several questions and issues to solve, but we believe that AR has the potential to revolutionize how industrial maintenance will be performed in the future.

Link:
[L1] https://www.sics.se/projects/incode

Reference:
[1] J. Wang et al.: "An Augmented Reality Based System for Remote Collaborative Maintenance Instruction of Complex Products", Proc. IEEE CASE 2014, Taipei, Taiwan, 2014.

Please contact:
Björn Löfvendahl, SICS Swedish ICT
bjorn.lofvendahl@sics.se

Value Stream Mapping with VASCO - From Reducing Lead Time to Sustainable Production Management

by René Berndt and Alexander Sunk, Fraunhofer Austria

Value stream mapping (VSM) is a lean management methodology for analyzing and optimizing a series of events for production or services.

Even now, the first step in value stream analysis - the acquisition of the current state - is still created using pen and paper by physically visiting the production site (see Figure 1). VASCO is a tool that contributes to all parts of value stream analysis - from data acquisition on the shop floor, detailed analysis, over planning, through to simulation of possible future state maps (FSM) - always taking economic and ecological factors into consideration (see Figure 2).

When a new product is manufactured - starting with the raw materials at the incoming goods department right through to the end product in the hands of the customer - multiple activities are required, for instance, assembly, transport and temporary storage. The aggregation of these various steps is called "value stream". More efficient planning and implementation of the value stream means more profit for the business. Maximizing the efficiency of a value stream is becoming increasingly difficult owing to the flexible nature of modern production. Enterprises need to adopt their business to the requirements of the market - e.g. how to cope with demand fluctuations or to fulfil individual customer requests. This is one of the major challenges for value stream planning.

Based on the experience of many projects in the field of value stream optimization – and in cooperation with partners from industry - Fraunhofer Austria has been developing the software tool VASCO. This solution is tailored for flexible and sustainable value stream planning. Even the visualization and analysis of complex value streams can be performed efficiently.

VASCO is implemented as a Microsoft Visio Add-In and supports the user in creating intelligent value stream maps in a fast and user-friendly manner. Its main advantages are the automatic linking of processes and configurable logistic tasks, as well as the (optional) display of five "economic" data-lines (production time, transport distances, space area usage, energy consumption, production costs) directly underneath the value stream. Furthermore, "ecological" data-lines are implemented for assessing disposal and carbon footprint per part produced.

One of the significant features of VASCO is that all VSM related symbols are fully customizable by a configuration file. This configuration file defines which properties are added to the symbol. These properties can be classified into two major categories: manual input values or calculated values. The manual values are entered by the user, whereas the calculated values depend on a formula consisting of manually entered or other calculated values. The formula definition is also part of the configuration file and can be modified even at run-time. As a result, each company can customize VASCO to their individual needs.

Another key feature of VASCO is extensibility. While VASCO is a Microsoft Visio Add-In, it can be customized by plugins itself. The basic version of VASCO is shipped with three plugins, extending the basic functionality of the tool:

- KPI-Plugin The KPI-Plugin adds additional visual features (see Figure 4) to the Visio page. This shape clearly displays the key performance indicators of the factory. Once a VASCO graph is complete and VASCO itself is in calculation mode, the values are calculated and auto-

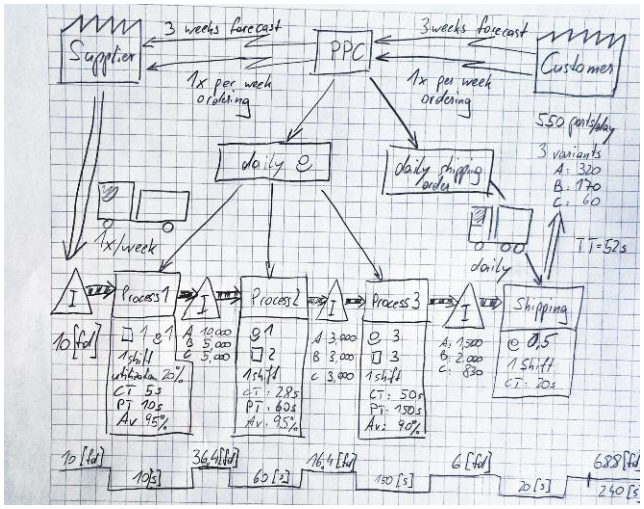

Figure 1: Typical hand-drawn VSM created at the production place.

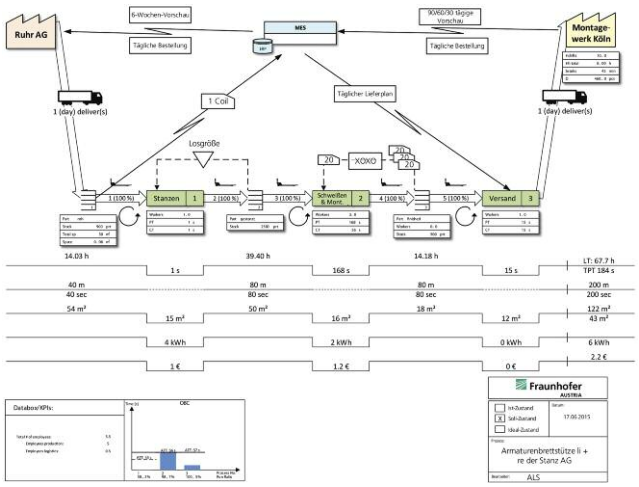

Figure 2: VSM diagram created with VASCO.

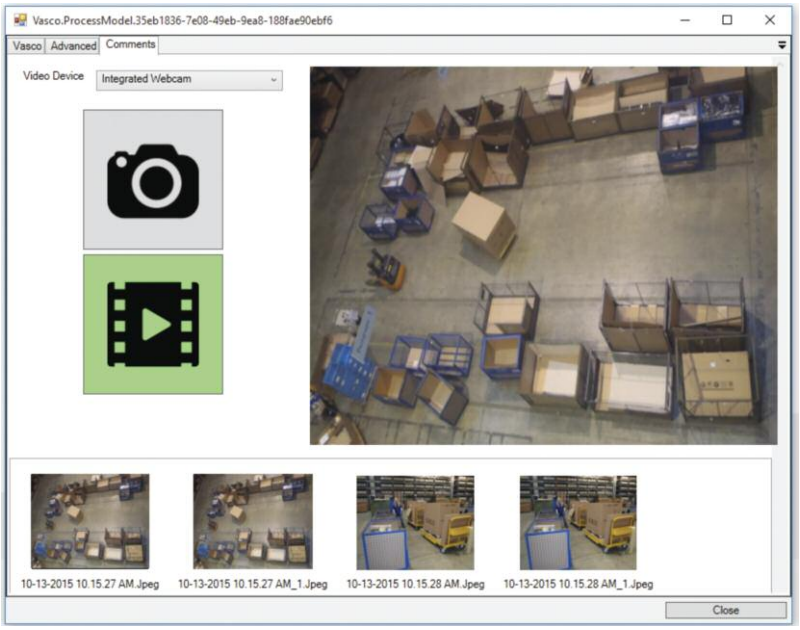

Figure 3: Annotating the VSM with images/videos/comments taken during the acquisition at the production place.

matically updated when a value in the graph changes.

- OBC-Plugin The operator balance chart (OBC) visualizes the total amount of work of each process compared to the takt time. The takt time can be considered as an average external rate at which customer requires goods produced. Therefore, the OBC visualizes how the cycle times of processes in the considered value stream fulfill this need. An OBC also supports optimal workload balancing between processes by making the amount of work for each operator very nearly equal to, but slightly less than, takt time. Figure 5 shows the OBC chart of the given example.

- Comment-Plugin VASCO was designed to make the acquisition and calculation of a new value stream easier and to replace the pen and paper acquisition. With the pen and paper method it is always possible to add different comments to the different symbols. In order to give the VASCO user a similar feature during the acquisition a comment plugin was created. This comment plugin enhances every symbol on a VASCO page with a comment tab (see Figure 3). When we observed that users sometimes only copied key figures from a machine into this comment tab during the data acquisition process, we further enhanced the comment-plugin with a snapshot ability. This feature enables the user to take a snapshot with the tablet instead of copying the values. It is also possible to record a video with the comment plugin. This can be used to record different views of the machine or to record the voice of the person who does the acquisition so that there is not even the

VASCO KPIs		
Flow rate / Flussrate	1.39	[%]
Lead time / Durchlaufzeit	36.51	[fd]
Processtime / Prozesszeit	3.38	[min]
Total number of employees / Anzahl Mitarbeiter ges.	6	[#/day]
Production / Produktion	5	[#/day]
Logistics / Logistik	1	[#/day]
Employee productivity / Mitarbeiterproduktivität	75	[parts/pers.]

Figure 4: The Key Performance Indicator giving an overview of the operating numbers in the factory.

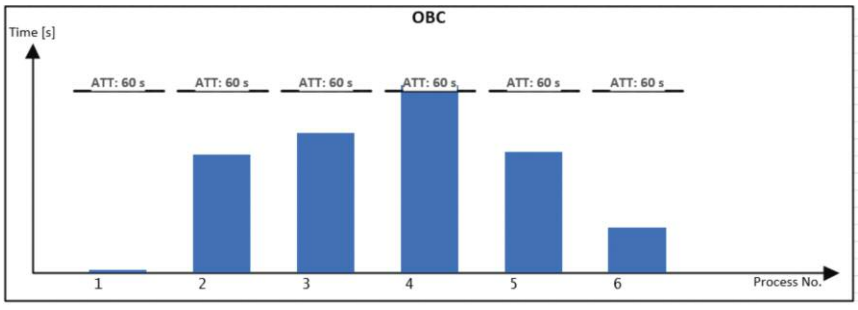

Figure 5: OBC showing the process time in relation to the takt time.

Current work-in-progress is the integration of sustainability criteria within a VSM. Ecological aspects are an increasingly important factor in addition to traditional economic considerations. "Sustainability" is the keyword that is causing enterprises to a change the way they operate. Responsible entrepreneurs and managers in Austria and throughout Europe also have responsible customers. Products with a very good environmental balance provide a competitive advantage by fulfilling the customer's needs and the demand of efficient resource use.

Link: http://www.fraunhofer.at/vasco

Reference:

[1] R. Berndt, et al.: "Vasco - Mastering the Shoals of Value Stream Mapping", in Proc. of the Sixth International Conference on Creative Content Technologies, 2016.

Please contact:

Thomas Edtmayr
Fraunhofer Austria
thomas.edtmayr@fraunhofer.at

Risk Analysis for a Synchro-modal Supply Chain Combined with Smart Steaming Concepts

by Denise Holfeld and Axel Simroth, Fraunhofer IVI

SYNCHRO-NET is an EU Horizon 2020 funded research project that is developing a powerful and innovative synchro-modal supply chain platform. Based on three demonstrators SYNCHRO-NET will show that a cloud-based solution can support the deployment and application of the synchro-modality concept, guaranteeing cost-effective robust solutions that de-stress the supply chain to reduce emissions and costs for logistics operations while simultaneously increasing reliability and service levels for logistics users.

Synchro-modality is a novel paradigm in logistics management which focusses on highly flexible transport operations: In a synchro-modal scenario the shipper agrees with the logistics service provider (LSP) on the delivery of products at specified costs, quality and sustainability, but leaves the decision on how to deliver according to these specifications to the LSP. This freedom gives the LSP the possibility to deploy different modes of transportation flexibly. The decision to switch to different modes of transportation may depend on actual circumstances such as traffic information or availability of resources. The core of the SYNCHRO-NET solution will be an integrated optimisation and simulation platform, incorporating: real-time synchro-modal logistics optimisation; smart steaming ship simulation and control systems; synchro-modal risk/benefit analysis modelling; dynamic stakeholder impact assessment solution; and a synchro-operability communications and governance architecture, in order to enable such an approach. Each module in itself represents a massive step forward compared to the current state of the art. The key innovation, however, will be the integration of these modules into a collaborative platform that ensures that all aspects of the supply chain are included.

Furthermore SYNCHRO-NET will introduce synchro-collaborative business models that allow (often sceptical) logistics operators to share economic and efficiency benefits, releasing commercially sensitive data in a controlled way without eroding their competitive advantage. For example, with a single press of a button, the user will be able to see how the consequences of a particular smart steaming strategy affects: ship bunker fuel usage AND multimodal hinterland logistics AND service level risk AND stakeholder perception AND end-to-end costs AND total emissions AND congestion in key nodes/corridors, etc. i.e. the decisions will be taken based on the global picture. SYNCHRO-NET will also provide the operational capability needed to manage the complexities of a synchro-modal supply chain. Thus, when the preferred strategy has been identified, SYNCHRO-NET will allow it to be implemented by providing powerful real-time multimodal logistics scheduling, supply chain optimization and ship scheduling/routing systems, governed by enterprise-collaborative business models.

Perhaps the most important output of SYNCHRO-NET will be the demonstration that steaming, coupled with synchro-modal logistics optimisation delivers amazing benefits to all stakeholders in the supply chain: massive reduction in emissions for shipping and land-based transport due to modal shift to greener modes and the optimized planning processes leading to a reduction in empty kilometres for trucks and fewer wasted repositioning movements.

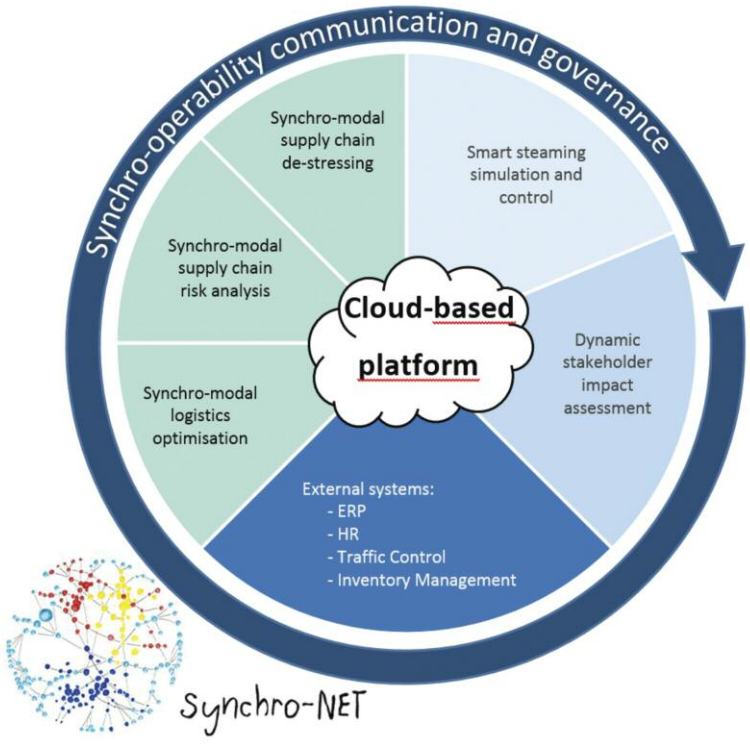

Figure 1: the Synchro-Net concept.

This will lead to lower costs for ALL stakeholders – shipping companies and logistics operators will benefit from massive reduction in fuel usage, faster turnaround times in ports and terminals and increased resource utilization/efficiency. Customers and end users will have greater control of their supply chain, leading to more reliable replenishment activity and therefore reduced safety stocks and expensive warehousing. Authorities and governmental organizations will benefit from a

smoother, more controlled flow of goods through busy terminals, and reduction of congestion on major roads, thus maximizing the use of current infrastructure and making the resourcing of vital activities such as import/export control, policing and border security less costly.

With a specific attitude to risk, three modules work in an integrated manner to determine a concrete door-to-door route for a user. The 'Synchro-modal Real-time Optimization' module uses "classical" objective values like costs, punctuality, resource utilization etc. to determine a solution. This is offset by a solution generated by the 'Supply Chain De-stress' module, which creates reliability and introduces flexibility into solutions, e.g. by flattening peaks in resource usage, including buffers in scheduling etc. For this purpose, extended objectives are necessary. In addition to monetary objectives, effectiveness-driven goals should be included, such as customer satisfaction and supplier reliability. The 'Synchro-modal Supply Chain Risk Analysis' module is intended to balance both types of objective functions depending on the risk attitude of the user. This module is to be developed as meta-process [1] with novel risk measures for complex supply chains [2].

SYNCHRO-NET had its kick-off in June 2015 and will last for three and a half years. It is co-funded under the Horizon 2020 Framework Programme of the European Union and brings together an industry-driven group of companies and research organizations, involving top stakeholders and developers from 10 countries and 20 organizations. The consortium currently carries out a requirements analysis from the perspective of all stakeholders in the smart steaming synchro-modal supply chain.

Links:
http://www.synchronet.eu
http://cordis.europa.eu/project/rcn/193399_en.html

References:
[1] D. Holfeld, A. Simroth: "Monte Carlo Rollout Method for Optimization under Uncertainty", 15th International Conference on Civil, Structural and Environmental Engineering Computing CC 2015, Prague, Czech Republic, September 1-4, 2015
[2] I. Heckmann, T. Comes, S. Nickel: "A critical review on supply chain risk – definition, measure and modeling", OMEGA Volume 52, p. 119-132, 2015

Please contact:
Denise Holfeld, Axel Simroth
Fraunhofer Institute for Transportation and Infrastructure Systems IVI
Tel: +49 351 4640 651
denise.holfeld@ivi.fraunhofer.de

Designing Sustainable Last-Mile Delivery Services in Online Retailing

by Niels Agatz, Leo Kroon, Remy Spliet and Albert Wagelmans, Erasmus University Rotterdam

The continuous growth of online sales together with the current inefficiency of delivery services puts a lot of pressure on urban areas in terms of congestion, emissions and pollution. Researchers at Erasmus University Rotterdam are developing decision support tools to facilitate efficient delivery services for products purchased online.

Online retail sales are continuing to grow at a fast pace, even in times of global economic downturn. Online sales represent the only growth sector in the declining retail market. However, despite the potential of the internet as a sales channel, online retailers face many logistical challenges in fulfilling online demand. Internet order fulfillment, also called e-fulfillment, is generally considered the most challenging and critical part of the operations of companies selling physical goods online. Handling small individual orders and shipping them to customers' homes in a timely and cost-efficient manner has proven difficult. This is particularly the case for the "last-mile" delivery, i.e., the last leg of the e-fulfillment supply chain in which the delivery is made to the customer. Owing to the large number of small individual orders, the last-mile often covers many stops in urban areas. Consequently, the

last-mile leg of the delivery is disproportionately expensive.

The efficiency of the last-mile not only impacts the profitability of online retailing but also affects environmental and social performance criteria such as emissions and traffic congestion. For instance, vehicle-miles are directly related to emissions and the required number of vehicles for delivery has an impact on congestion.
The last-mile logistics project started in 2015. The primary goal of the project is to develop decision support tools to facilitate innovative operating strategies and to analyze the design of retail networks including an online channel that provides the most benefits in terms of various sustainability criteria. The research is carried out at the Erasmus University Rotterdam and is funded by the Dutch research funding body TKI-Logistiek, NWO, the largest Dutch online supermarket AH.nl and the logistics consultancy company ORTEC. The ambition is to have prototype decision support tools available for the industry by 2019. The current research focusses on two themes: optimizing delivery operations and the design of a multi-channel retail network.

The logistical challenges in delivery operations are especially apparent when the customer has to be home to receive the goods [1]. To minimize the risk of not finding a customer at home for

delivery, and thus creating unnecessary vehicle movements, the retailer and the customer may agree on a delivery time window [2]. The design of such a time window menu involves decisions on the length and the number of the offered time windows and on the associated delivery fees. Moreover, the time window selected by each customer directly affects the delivery routes, and hence the transportation costs, emissions and additional strain on congestion. Our project develops different mathematical model formulations and subsequently focusses on the design of algorithms to optimize the time windows on offer. These algorithms can be used as decision support tools for online retailers but also as evaluation tools to assess the impact of different operating strategies on various sustainability criteria.

The second theme focusses on the support of strategic design decisions associated with multi-channel retailing. One example of such a decision is where to place pick-up point locations that allow customers to pick-up goods ordered online at a convenient time and place. This decision needs to take into account both customer preferences and operational costs. Another example relates to the consolidation of different transportation flows in a retail network that operates both online and offline channels, see also [3]. To support these strategic design decisions, we are also formulating mathematical models and developing optimization approaches. Note that decisions on the network design have an impact on the delivery operations and vice versa. This connects the two themes in the project.

Overall, our project aims to improve the efficiency and sustainability of the last-mile of online retailing. We believe this will become increasingly relevant given the growth of online sales and the increasing logistics challenges in urban areas.

Link:
http://www.erim.eur.nl/Last-Mile

References:
[1] A.M. Campbell, M.W.P. Savelsbergh: "Decision support for consumer direct grocery initiatives", Transportation Science 39(3): 313-327, 2005.
[2] N.A.H Agatz, et al.: "Time Slot Management in Attended Home Delivery", Transportation Science 45(3): 435-449, 2011.
[3] N.A.H Agatz, et al.: "E-fulfillment and multi-channel distribution - A review", European Journal of Operational Research 187(2): 339-356, 2008.

Please contact:
Niels Agatz , Leo Kroon, Albert Wagelmans, Remy Spliet
Erasmus University Rotterdam, The Netherlands
nagatz@rsm.nl, lkroon@rsm.nl , spliet@ese.eur.nl, wagelmans@ese.eur.nl

Spare Parts Stocking and Expediting in a Fluctuating Demand Environment

by Joachim Arts, TU Eindhoven

Capital goods, such as trains and railway infrastructure that facilitate our public transport, are an important part of our daily lives. Maintenance operations are necessary to ensure safety and prevent disruptive failures. To make these operations run smoothly, it is crucial to have the right amount of spare parts available. Eindhoven University of Technology and NedTrain collaborate to optimize the spare parts supply chain using stochastic modelling and optimization.

Interchangeable parts have revolutionized modern manufacturing, but the idea was originally born as a maintenance logistics innovation by a French general by the name of Jean Baptiste de Gribeauval. When artillery or other equipment broke down during battle, performing the repair would usually take longer than the battle itself, rendering equipment useless for the remainder of a battle. Gribeauval made a plan to construct equipment from interchangeable parts such that repair could be performed quickly by replacing the broken part. The broken part is then repaired offline to be used in a future exchange of parts. This method of repair-by-replacement is now common in the maintenance logistics operations of many asset owners in the rail, aerospace, medical devices, defence, and other industries.

To successfully run a repair-by-replacement system, logistics managers and engineers need to make several decisions, the most important of which are:
- How many spare parts of each type do you need?
- How do we decide which broken parts need to be repaired with priority?

Smart methods to make these decisions have large financial and societal consequences. Expenditure on spare parts in the USA in 2003 and 2006 constituted 8% of GDP (see e.g. [1]). Temporary unavailability of a spare part can lead to delayed or cancelled trains and flights, standstill of large manufacturing plants, and rescheduling of medical procedures that need specialized equipment, among many other adverse societal consequences. Therefore Eindhoven University of Technology and NedTrain set up a collaboration to answer these questions within a larger research program of NedTrain in rolling stock lifecycle logistics.

Fluctuating demand and expediting repair

NedTrain, a subsidiary of the Dutch railways, is responsible for the maintenance of most of the rolling stock in the Netherlands. Their core business is to run a repair-by-replacement system so that rolling stock is always available to carry passengers according to the time table. The spare part supply chain of NedTrain and many other companies has two important characteristics:
- Demand for spare parts is a non-stationary stochastic process that results from degradation processes and

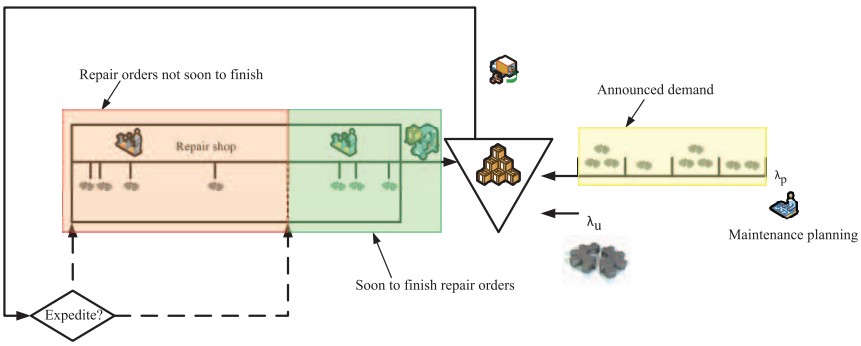

Figure 1: Schematic representation of the spare parts supply chain.

maintenance planning. This means that demand intensities fluctuate over time.

- The repair lead time of a broken part can be shortened but only for a limited number of all the repairs handled by a repair shop. Typically this is done when the stock of ready-for-use parts is critically low.

Natural questions that arise in this setting are:

1) When should the repair of a spare part be expedited?
2) How do optimal expediting policies and stocking decisions depend on the way demand fluctuates over time?
3) Can we reduce the spare parts inventory investment by using the answers to the previous questions in a smart way?

Structural results for optimal expediting policy

We constructed the model shown in Figure 1 to answer questions 1) and 2) above. It works as follows: When a part fails, the broken part is immediately sent to a repair shop. At this point in time, the supply chain manager can choose to request an expedited repair lead time for this part. The repair lead time will be shorter when an expedited repair lead time is requested, but this can only be done for a limited number of repairs per time unit. Using a Markov decision process model, we can show that the optimal expediting rule exhibits the following structure: When the number of parts whose remaining repair lead time is longer than the expedited repair lead time exceeds a certain threshold, the repair of the next broken part will be expedited. The size of this threshold depends on the total number of spare parts in the system, and the forecast of demand in the near future.

Multiple items, fleets, and repair resources

Decisions on how many parts to stock and when to expedite the repair of a part need to be taken for many spare parts that belong to different (train) fleets, but that may use one of several shared repair resources. We devised an algorithm based on decomposition through column generation. To answer question (3) above, we conducted a numerical study that shows that spare part investments can be reduced by 25% on average compared to the state of the art.

Practical decision support tool

Users that do actual spare parts planning and control can interact with our models through an application we developed. This application is called TRAINcontrol and a screenshot is shown in Figure 2. This app helps users to run sensitivity analyses and answer "what if?"-type questions.

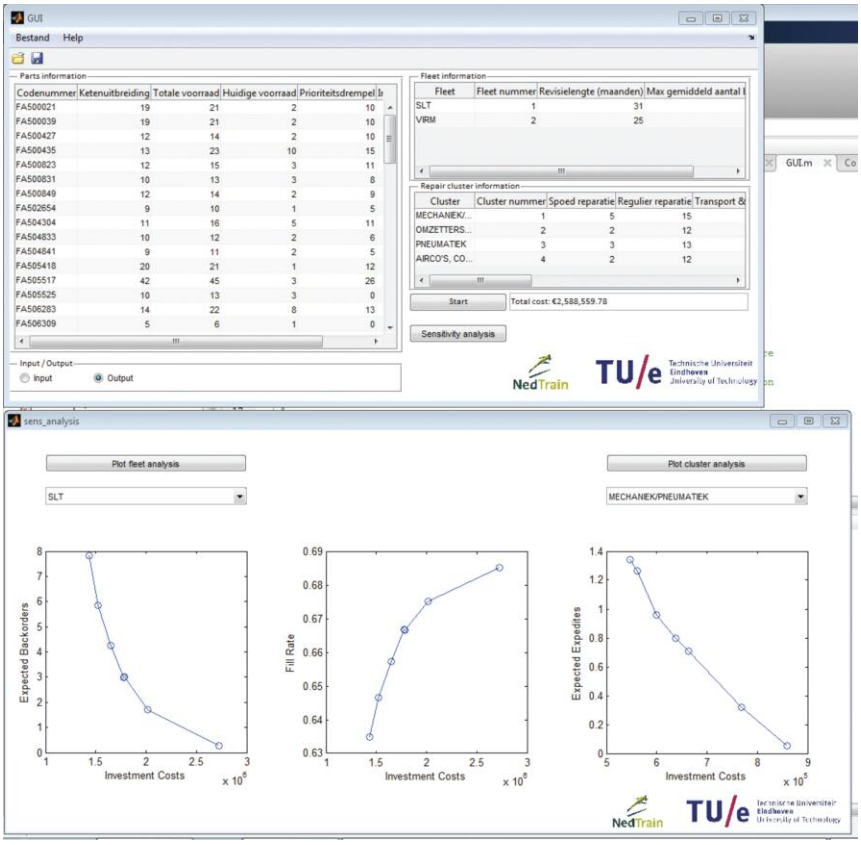

Figure 2: Screenshot of the TRAINcontrol application.

Link:

TRAINcontrol MATLAB app:
http://www.mathworks.com/matlabcentral/fileexchange/54498-traincontrol

References:

[1] J.B. Jasper: "Quick Response Solutions: Fedex Critical Inventory Logistics Revitalized", Fedex White Paper, 2006, http://www.fedex.com/us/supplychain/casestudies/quick_response.pdf.

[2] J. J. Arts: "Spare Parts Planning and Control for Maintenance Operations", PhD Thesis, Eindhoven University of Technology, http://alexandria.tue.nl/extra2/760116.pdf.

Please contact:

Joachim Arts
TU Eindhoven, The Netherlands
j.j.arts@tue.nl

Packing with Complex Shapes

by Abderrahmane Aggoun, KLS OPTIM, Nicolas Beldiceanu, Gilles Chabert, École des Mines de Nantes and François Fages, Inria

Constraint programming methods can be used to solve hard geometric placement and packing problems such as those encountered in the operational logistics and in the conception of packing plans.

Modern warehouse management systems (WMS) provide advanced management features and monitoring of movements of goods within a warehouse. They do not, however, satisfy the growing demand for digital processing - especially the processing required to achieve an agile supply chain model capable of handling custom commands whilst minimising cost and environmental impacts.

This subject was investigated within the French ANR funded Net-WMS-2 project [L1] (2011-2015) - an industrial project involving Inria Saclay, Ecole des Mines de Nantes and SME KLS OPTIM. This project was a follow-up of the EU FP6 STREP project Net-WMS which showed how constraint programming methods can substantially improve industrial box placement and packing problems, and identified problems including: the difficulties of placement with rotations, and the industrial need to handle complex shapes and curved objects.

The daily use of the application Optim Pallet (3D), which was developed from the earlier project (Net-WMS), resulted in an up to 15% reduction in transportation costs for one company. As a consequence of these encouraging results, our industrial partners were keen to engage in research on new innovative optimisation applications to assist with handling complex shapes.

Our previous studies demonstrated the feasibility and the realisation of prototypes for handling complex shapes: mixtures of boxes, polyhedrons, cylinders, and curved objects defined by Bézier curves, with the development of new methods of modelling and solving discrete-continuous hybrid constraints. Constraint programming methods can address NP-hard optimisation problems and provide solutions that are often effective and easy to maintain. They are based firstly on a model of the problem as a set of constraints to satisfy, and secondly on a resolution that uses predefined constraint solvers and a generally heuristic enumerative search procedure.

As part of this project, the constraint solver Choco over discrete domains has been interfaced with the continuous constraint solver Ibex in order to deal with complex shapes in placement problems. We have shown that it is possible to define the search strategies themselves as search constraints [2], and have extended the MiniZinc modelling language, which is now a standard in the field, with Horn clauses for this purpose, and have shown that the best placement strategies for squares could be modelled declaratively in ClpZinc without loss of efficiency. However, these methods have proved inadequate for curved shapes, and the use of stochastic optimisation methods was necessary, in this case using the CMAES solver. We have therefore shown that the constraints of placement and non-overlap of complex objects can be modelled by error functions (e.g. penetration depths [3]) to minimise, and developed an original and general interface MiniZinc-CMAES for solving such problems over the reals [1].

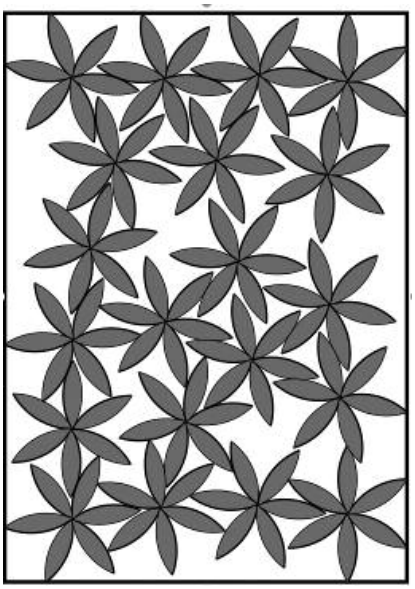

Figure 1: ClpZinc model and approximate resolution using our CMAES backend for the placement of 20 rosettes defined by union and intersections of circles.

Using these declarative modelling techniques, combined in an original way to stochastic optimisation methods, we were able to solve in an approximate way both mixed curved-square shapes placement problems, such as boxes and cylinders for truck loading problems, and complex shapes defined by Bézier curves for packing problems. We have also shown, on the problem of Korf for placing squares in a rectangle, that the exact resolution square shapes placement problems was done without significant loss of performance compared to the state of the art, but with a fully declarative modelling/programming in ClpZinc of the search strategy [2] which is very new. The software developed by academic partners is distributed with sources as free software. The project results have been developed as demonstration prototypes by KLS OPTIM company.

Link:

[L1] http://lifeware.inria.fr/wiki/Net-WMS-2

References:

[1] T. Martinez, F. Fages, A. Aggoun: "A Stochastic Continuous Optimization Backend for MiniZinc with Applications to Geometrical Placement Problems", in Proc. of CPAIOR'16, LNCS, 2016.

[2] T. Martinez, F. Fages, S. Soliman: "Search by Constraint Propagation", in Proc. of PPDP'15, ACM, 2015.

[3] I. Salas and G. Chabert: "Packing Curved Objects", in Proc. of IJCAI, 2015.

Please contact:

François Fages, Inria, France
Tel: +33 (0)682822937
Francois.Fages@inria.fr

Abder Aggoun, KLS OPTIM, France
Tél : + 33(0)608864555
Abder.Aggoun@klsoptim.com

Allocating Railway Tracks Using Market Mechanisms and Optimization

by Victoria Svedberg, SICS Swedish ICT

New rules call for new methods. The deregularisation of the railway market in Sweden has not only opened the market for new railway operators, but also posed challenges to the old methods. SICS Swedish ICT is developing a process to make timetabling transparent, easy and fair, while still satisfying societal needs.

How can track time be distributed between competing companies and needs in a deregulated railway market? The basic rule is to squeeze in as many trains as possible while at the same time maintaining the quality of the traffic for society as a whole. But what to do when the railway network is already heavily congested and many trains need the same track time? The European Union advocates a fair market on the tracks by the EU-railway directive (2001/14/EC). For a well-functioning competition on the tracks, the decision about which competing railway operator or maintenance entrepreneur will obtain a specific track time has to be fully transparent and non-discriminatory. There is currently no method for this, but SICS, together with the Swedish Transport Administration, is developing one.

A project at SICS investigates how track time can be distributed between competing companies and needs in a deregulated railway market (Photo: Skånetrafiken).

The project "Cost-benefit efficient track allocation" investigates a new timetabling approach in which track time is sold to an operator at a market based price, which mirrors the highest price the operator is willing to pay for it. Giving the track time to the operator who values it the most provides transparent competition. The operator wants to maximize the amount of passengers and/or transported goods, and this ensures that societal service is maximized. The only trick is how to get the operators to reveal how much they are willing to pay for track time.

The solution is a timetabling process consisting of three parts: framework agreements, auctions, and a spot market. The auctions take place one year in advance and only very sought-after track timeslots are included. The resulting prices of the auction become the foundation of the pricing mechanism in the spot market. The prices in the spot market are set by a kind of revenue management and the operators can apply for their desired track timeslots and pay the price. The auctions and the spot market will provide an opportunity for the operators themselves to settle disputes over track time. The goal is to develop a fully functional, transparent method to rule out competitors on the same market segments and to add flexibility in an otherwise long process.

Non-commercial traffic, which is subsidised by the state and thus financially stronger will not take part in the auction and the slot market. Instead the non-commercial traffic enters framework agreements which last for a longer time period, typically five years. These framework agreements should contract capacity utilization rather than a fixed timetable. As an example, a regional operator can sign a framework agreement where the infrastructure planner (the Swedish Transport Administration) agrees to set up the timetable so that trains of the regional operator can depart from a station every 15 minutes daily.

How do we then make sure that neither too much nor too little capacity is booked in the framework agreements during the negotiation?

To compare the possible other trains and the non-commercial traffic we intend to use a cost-benefit analysis. The basis is the generalized cost and operating cost of the non-commercial traffic. By linearizing the expression for the costs, modelling it as an MIP problem and solving it, the minimum cost is found. Since the timetable is not known yet, the minimum value is an estimation of the best possible timetable. The resulting timetable will work as an input to calculate a consumer and producer surplus when conducting a cost-benefit analysis [1].

This work is still at an early stage. In the near future we will further investigate the slot market and its combined operational and economic issues.

Link:
http://www.sics.se/projects/sameff-trackallocation

Reference:
[1] V. Svedberg, M. Aronsson, M. Joborn: "Timetabling based on generalised cost", SICS Technical report T2015:05.

Please contact:
Victoria Svedberg, Martin Joborn, SICS Swedish ICT
vicsve@sics.se, mjoborn@sics.se, martin@sics.se

OscaR, an Open Source Toolbox for Optimising Logistics and Supply Chain Systems

by Renaud De Landtsheer, Christophe Ponsard and Yoann Guyot, CETIC

The efficient operation of logistics and supply chain systems requires businesses to solve large and complex optimisation problems such as production planning, fleet management and order picking. The OscaR Open Source framework supports the optimisation of such problems through a variety of powerful solvers offering various expressivity, optimality, and scalability trade-offs.

Optimisation problems often arise in logistics and supply chain systems. Although a number of "off-the-shelf" commercial solutions exist, companies may have problems using them for various reasons: their problem might be too specific to be managed within a closed solution, or its cost might be too high for smaller structures. The multiplicity of problems might also require the use of different techniques or even a hybrid approach based on a combination of methods, which is also difficult to achieve in a closed framework.

OscaR is an Open Source software framework [L1] developed jointly by the CETIC research centre, the University of Louvain and the N-SIDE company. CETIC is contributing through the SimQRi (Simulative Quantification of procurement induced Risk consequences and treatment impact in complex process chains) [L2] and TANGO (Transparent heterogeneous hardware Architecture deployment for eNergy Gain in Operation) projects [L3] . The framework notably supports constraint programming, constraint-based local search, mixed integer programming and discrete event simulation, and is especially tailored for combinatorial optimisation. Here we illustrate the use of two relevant OscaR modules on specific case studies in logistics.

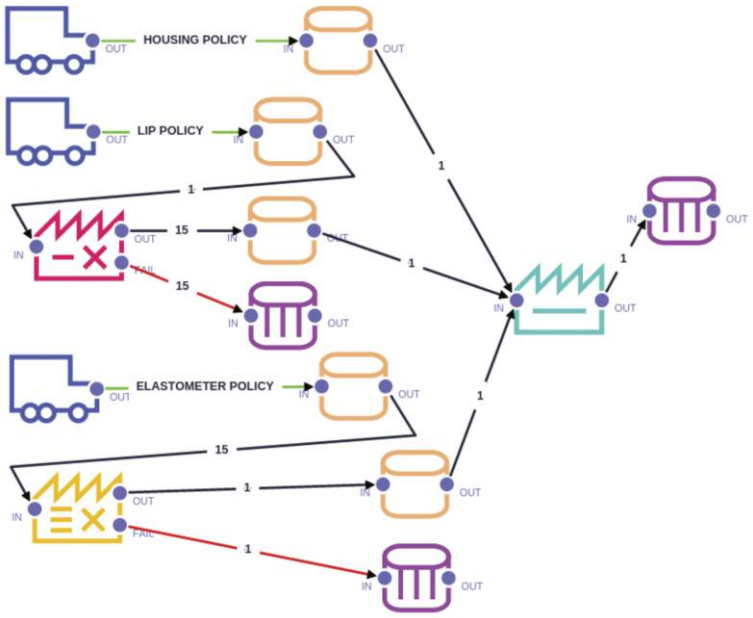

Figure 1: Supply chain model for a complex assembly.

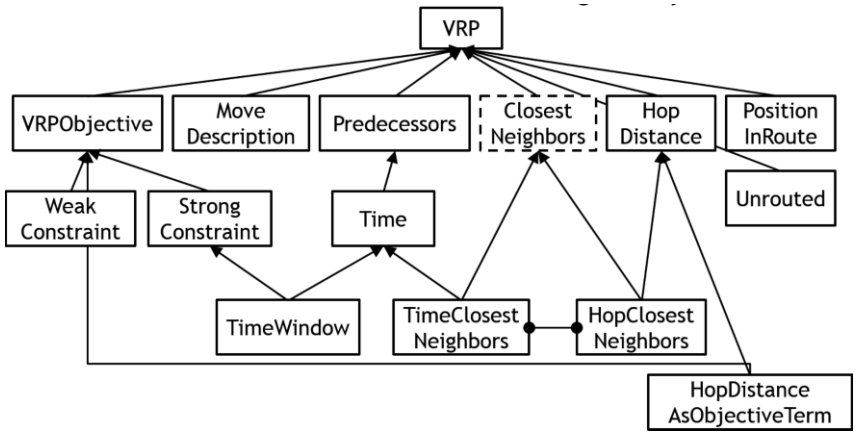

Figure 2: Library of reusable VRP problem elements.

Managing procurement risks through Discrete Event Simulation (DES)

Decision makers in the field of procurement and production logistics need to design cost optimal processes taking into consideration the firm's exposure to supply risks. Assessing risks in logistics processes requires quantification of the impacts at the delivery side (e.g. order delay, quality, quantity) of several factors, preferably in a statistical way. This requires the efficient simulation and computation of key indicators on the global logistics system. To address these issues, OscaR.DES is able to capture the problem directly at the logistics domain using both textual and graphical primitives (see Figure 1) such as suppliers, storage, production processes, order policies, etc. Based on this, the DES engine can very efficiently compute the system evolution by evaluating changes only when they occur and updating complex risk indicators combining basic properties (process delays, storages volume…) using a rich set of operations (logical, arithmetic, temporal,…). Large sets of histories can then be aggregated through Monte-Carlo techniques, e.g. to explore the procurement policies that are able to best address the identified risks. Figure 1 illustrates a typical risk-oriented model for a complex assembly process.

Optimising vehicle fleet routing

Local search techniques are known to be efficient at solving such problems - especially large problems. The CBLS engine of OscaR supports the modular specification of routing problems by assembling problem elements (such as objective functions, strong and weak constraints, time windows, traffic jams, SLA as shown in Figure 2). On the other hand, it provides a rich domain specific language for working out powerful search procedures combining a rich set of neighbourhoods (e.g. insertPoint, onePointMove, threeOpt) in a declarative style using movement, acceptor and meta-heuristic operators (e.g. tabu search, simulated annealing). This not only results in very efficient search procedures expressed at the logistics domain level but also reduces development time and facilitates maintenance when a problem evolves.

Links:

[L1] http://oscarlib.bitbucket.org
[L2] http://www.simqri.com
[L3] http://tango-project.eu

References:

[1] L. M. P. Van Hentenryck: "Constraint-based Local Search", MIT Press, 2009.
[2] R. De Landtsheer et al.: "A Discrete Event Simulation Approach for Quantifying Risks in Manufacturing Processes", Int. Conf. on Operations Research and Enterprise Systems, February 2016.
[3] R. De Landtsheer et al.: "Combining Neighborhoods into Local Search Strategies", 11th MetaHeuristics International Conference, June 2015.

Please contact:

Renaud De Landtsheer
CETIC, Belgium
Tel: +32 472 56 90 99
renaud.delandtsheer@cetic.be

Integrated Resource Planning in Maintenance Logistics

by Ahmad Al Hanbali, Sajjad Rahimi-Ghahroodi, and Henk Zijm

The MLOG (Optimal Exploitation of Resources in Maintenance Logistics) project, executed jointly by the University of Twente and the University of Qatar, focuses on an integrated planning of resources needed for asset maintenance.

In many Western countries, maintenance and overhaul of capital assets constitute some 15 % of their GDP. Smart sensor and data gathering techniques, as well as advanced decision support systems, are exploited to design integrated logistics support (ILS) systems. In the MLOG project we focus on an integrated planning of resources needed for asset maintenance rather than the piecemeal approach usually considered in the literature and in practice. The goal is to minimize overall resource investments subject to agreed service level constraints. Exact methods are only computationally feasible for very small problems, hence we have developed fast but highly accurate approximation methods. In an actual case study, our results show that integrated planning can achieve an overall cost reduction of up to 27 % without sacrificing the offered service quality, when compared with the common practice solution used in companies. Combining our approach with smart sensor and condition monitoring data (Internet of Things) may further enhance asset availability and hence industrial production both in Qatar and in the Netherlands.

Maintenance logistics has received considerable attention in recent decades owing firstly to the significant investments associated with capital-intensive assets, which in turn require a high operational availability, and secondly to the need to prevent environmental damage – for example, the BP accident in 2010 in the Gulf of Mexico - or safety incidents – for example in medical equipment and aircraft. Unplanned downtime of advanced capital equipment can be extremely expensive; for an average aircraft it is estimated at some S10,000 per hour and for a high-tech lithography system in the semiconductor industry it may amount up to $100,000 per hour. Consequently, unplanned downtime should be prevented as much as possible, for example, by exploiting advanced condition monitoring techniques and preventive maintenance policies, and if they occur, they should be kept as short as possible by using optimal corrective maintenance policies [1]. The latter implies that malfunctioning parts or components causing the system breakdown are immediately replaced by ready-for-use ones, since repair of the complete system on site induces unacceptable long downtimes. Typically, a system failure induces a set of actions as depicted in Figure 1.

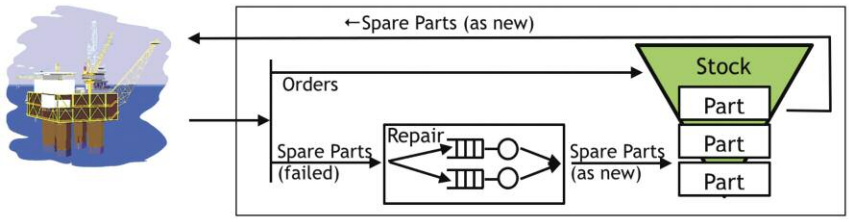

Figure 1: In the MLOG project, we are developing fast approximation algorithms to optimize service supply chains, while reducing the overall costs with 27 % in an actual case study.

The availability of resources needed for repair (spare parts, service engineers, and tools) in fact defines the operational availability of the complete asset. However, both spare parts and skilled service engineers often require high investments (parts worth $50,000 are no exception for capital-intensive assets), hence there is a trade-off between resource investments and service provided. Much research has been carried out on spare parts management see [2]. So far the planning of resources has largely been performed independently (per resource) while their simultaneous availability is essential to minimize downtimes.

Contribution

As part of this project, we have been studying a service region in which a local spare parts inventory supplies different types of parts. Failures occur randomly and are often due to one malfunctioning part that needs replacement. In addition, a skilled service engineer is needed to complete the repair and replace the defective part. Malfunctioning parts are repaired off-line and subsequently stocked for future use. The focus here is on the integrated, multi-resource approach, for which different service policies are available: (i) complete backlog in case of unavailability of either parts or engineers, (ii) an emergency service fulfillment from an external source in case of unavailability of either parts or engineers, (iii) and the heterogeneous policy with backlogging for one resource and emergency for the other.

The most realistic heterogeneous policy would involve having parts supplied by an emergency service and engineers backlogging. This is due to the long spare parts replenishment time that is often encountered compared with the engineers' service time. To evaluate such a heterogeneous policy, we have developed accurate methods using Mean Value Analysis and Laplace Transform techniques that can handle practical problems with a high number of different types of parts [3]. Our methods can tackle these problems quickly as opposed to the standard Matrix-Geometric approach, which is only computationally feasible for very small problems. This is due to the detailed description of the system state needed in the standard approach. In addition, we also consider the overall logistic support system optimization, i.e., we determine optimal stock levels and service engineer crew size, such that the average total costs (including spare parts holding costs, hiring cost of service engineers, and emergency costs) are minimized subject to pre-specified service level constraints. For real problems with high number of part types, our methods yield solutions very fast while the total cost error is negligible when compared with an exact method (only verifiable for small problems). The optimization algorithm demonstrated an overall cost reduction of 27% in an actual case study, when compared with the separated optimization, which is common practice in companies. Future research may include advance information based on condition monitoring data (smart sensors, IoT) to move to preventive maintenance, thereby further increasing asset productivity.

MLOG was awarded by the NPRP program, 7th cycle in 2014, and funded by the Qatar National Research Foundation for a period of three years, starting from January 2015.

References:

[1] M.A. Cohen, N. Agrawal, V. Agrawal: "Winning in the aftermarket", Harvard Business Review, 84(5), p.129, 2006

[2] C. C. Sherbrooke: "Optimal inventory modeling of systems: multi-echelon techniques", Springer, 2004.

[3] S. R. Ghahroodi et al.: "Integrated Field Service Engineers and Spare parts Planning in Maintenance Logistics",. under review..

Please contact:

Ahmad Al Hanbali, Sajjad Rahimi Ghahroodi, Henk Zijm
University of Twente, The Netherlands
a.alhanbali@utwente.nl,
s.rahimighahroodi@utwente.nl,
w.h.m.zijm@utwente.nl

Utilising the Uniqueness of Operation Days to better Fulfil Customer Requirements

by Sara Gestrelius

No two days are exactly the same on the Swedish railways. Despite this, most trains are granted only one train path that they are supposed to use every day of operation. Restricting each train to a single train path wastes infrastructure capacity and prevents train operators from getting the capacity they require. In a project funded by the Swedish Transport Administration, SICS Swedish ICT used optimization to plan for each operation day individually. The results show a major improvement in customer requirement fulfilment.

Train operators apply for track capacity in September each year, and it is the Transport Administration's job to combine the trains in the applications into a yearly timetable. In order to make the timetable problem manageable for manual planning, only one train path is generally constructed for each train. All conflicts that this train faces with some regularity should be resolved in this one train path. However, the traffic pattern is different on different days, and to make this one train path conflict-free for all days the planner is forced to include extra time both for conflicts that occur, for example, solely on Mondays, and also for conflicts that occur solely on Wednesdays, even if including extra time only once would have been enough. This wastes capacity and results in the train path having unnecessary stops and time supplements on the day of operation.

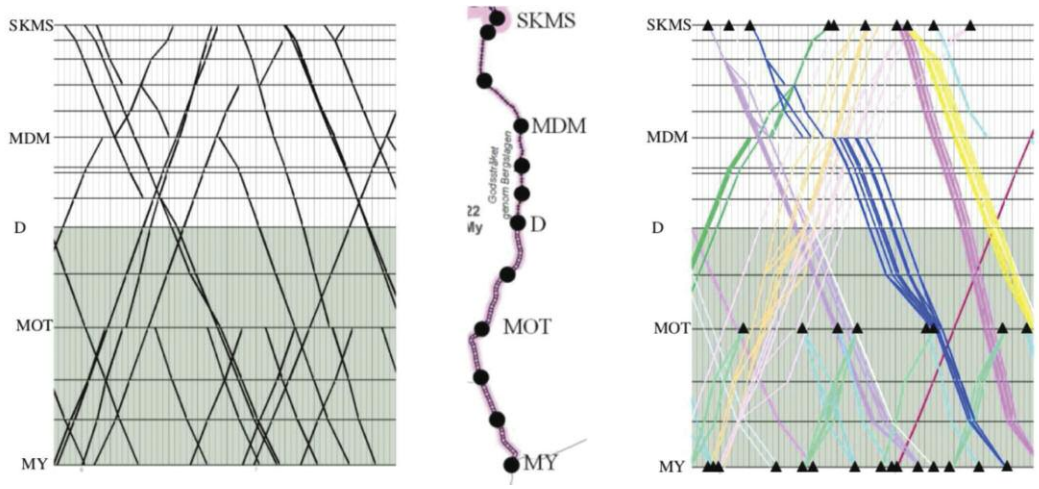

Figure 1: The line between Skymossen (SKMS) and Mjölby (MY) that was used as a case study. The outcome of the current process is shown to the left, and the new proposed process to the right. The delivery commitment times are marked with black and include the entire train path for the current process and a subset of times (marked with triangles) for the proposed process.

SICS Swedish ICT has been doing research on train timetable problems and processes for several years, and one of the core ideas is that there are only certain points in the railway network where the exact timings of a train have to be fixed, namely points where there is some commercial activity. We call these points "delivery points" and the time promise a "delivery commitment". In the current process the exact timings of the entire train path are included in the delivery commitment. By committing only to certain times rather than to entire train paths the Transport Administration retains the freedom to re-plan the operation to some extent, and may also use different operation plans (i.e. train paths) for different operation days. This new way of only committing to certain times is currently being implemented in Sweden. In Figure 1 the times included in the delivery commitment for the two different processes are marked in black.

In a project started in 2014 and due to end in 2016 we have been investigating how the delivery commitments generated in the yearly timetable process are affected by planning for each individual operation day rather than an example day with all traffic. Further, mathematical models that can be used in a planning support tool handling each day individually are being developed and tested. More precisely, the timetable problem was modelled as a Mixed Integer Program and rolling horizon planning was used to generate a yearlong timetable. As stated above, the delivery commitments generated have to be upheld every day of operation. This means that the when all individual days of a train have been planned for, the latest arrival time and the earliest departure time at delivery points can be put in the delivery commitment. Also, if a short run time is required the runtime from the worst-case day can be promised.

The method was tested on the line between Skymossen and Mjölby in Sweden [1]. The test case includes both single and double track sections and a mixture of passenger and freight trains. Further, it is assumed that some train operators value getting the exact times they've applied for (true for most passenger trains), while others have indicated that a short running time is more important (true for some freight operators). The application times were sampled to be close to the times in the finalized 2014 timetable.

The results show that the average running time was improved for 20 of the 67 trains where a short running time was preferred, and the worst running time was improved for eight trains. When it comes to timeliness the number of application times that could be granted exactly as applied increased from 193 to 260 out of a total of 296. That is, the number of times the customers could get the arrival or departure times they had applied for increased from 65% to 88%, indicating that there is great potential in planning for each individual day.

During the last year of the project, the aim is to develop a method for handling larger geographical areas, and also testing various objective functions and possibly iterative solution methods to improve the delivery commitment generation even more. For example, it would be interesting to explore methods that inhibit the possibility of one or a few operation days getting exceptionally bad train paths and thereby resulting in the train getting poor delivery commitments.

Link:
https://www.sics.se/projects/framtidens-leveranstagplaneprocess

Reference:
[1] S. Gestrelius, M. Bohlin, and M. Aronsson, "On the uniqueness of operation days and delivery commitment generation for train timetables", IAROR ICROR RailTokyo2015, 2015.

Please contact:
Sara Gestrelius, Martin Aronsson, Markus Bohlin, SICS Swedish ICT
sarag@sics.se, martin@sics.se, markus.bohlin@sics.se

Planning Complex Supply Networks Facing High Variability

by Ulrich Schimpel and Stefan Wörner, IBM Research

Interwoven production lines may be complex, with variable yield and production times, various sub-components competing for processing capacities, and fixed batch sizes. Furthermore, inventory costs need to be minimized and fluctuating customer demands need to be satisfied 98 % of the time. Such complex production lines need to be optimized using a combination of techniques. We describe an approach using a simplified mathematical model that allows for sensitivity analyses, followed by a discrete event simulation to adequately represent the complex business environment.

In supply chain management, there exists a fundamental principle: A better visibility of your supply network processes increases their velocity and reduces their variability [1]. This principle is of increasing importance to successfully operate supply chains in today's fast-paced world. Imagine you run a manufacturing network with assembly locations world-wide that produces hundreds of products, each requiring dozens of components - also referred to as "bill of material" (BOM) (Figure 1). The target is to fulfil customer demand 98% of the time despite a significant fluctuation in demand, production times, lead times and yield. Achieving this target becomes especially tricky in an environment of multi-purpose resources that are being shared among a set of different products and thus the variability of one product affects a multitude of other products.

The challenge

The described scenario is a daily reality for many companies. One option is to eliminate most variability and waste by well-known concepts such as just-in-time (JIT). However, this is not feasible in environments with a strong inherent variability of individual processes or very long lead times. Also, the business complexity usually inhibits a precise holistic formulation and finding a globally optimal solution. Over the past decade, the authors have been developing feasible approaches for exactly such situations in the semi-conductor industry [2].

A usual primary objective is the satisfaction of delivery times and quantities of all products. A common secondary objective is the minimization of inventory of items required for production lines. Important constraints are "capacity groups" of products that share the available time-variable capacity. It might be necessary to build products ahead of time to avoid a bottleneck capacity in the future. This decision is complicated by existing uncertainty. An opposing constraint on pre-building is the "inventory budget" that must be met at the end of each quarter for accounting reasons. This usually results in splitting or delaying purchases, such that the inventory is likely to hit the accounting books shortly after the next quarter begins. Other aspects to be considered are business rules, lot-sizes, and quantity-dependent stochastic yields. Both of the latter tend to magnify the variability throughout the supply network, starting from the customers. This is known as the "bull-whip effect"

The solution approach

Lacking exact solutions for such an environment, the authors apply a combination of two different techniques. First, a simplified mathematical model is solved which allows for sensitivity analyses. Second, a discrete event simulation evaluates the result from the simplified model in a close-to-reality business context.

The analysis starts with capturing the historical data including all its variability and errors. After cleaning the data, percentiles are derived for the stochastic lead times, and several possible demand scenarios are generated. These demand scenarios can easily be illus-

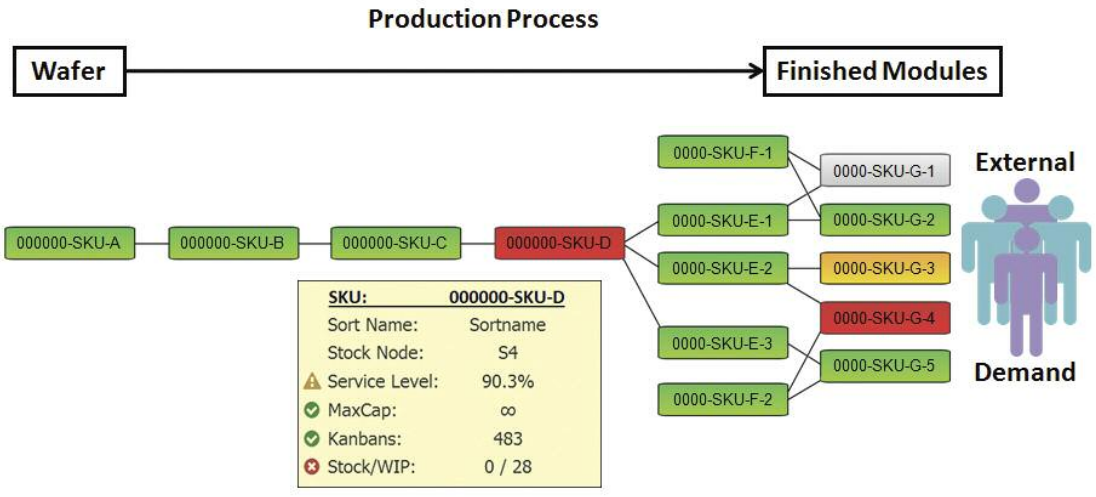

Figure 1: A typical bill of material with colour-coding in the semi-conductor industry for a set of finished goods (right) being produced from a single wafer (left).

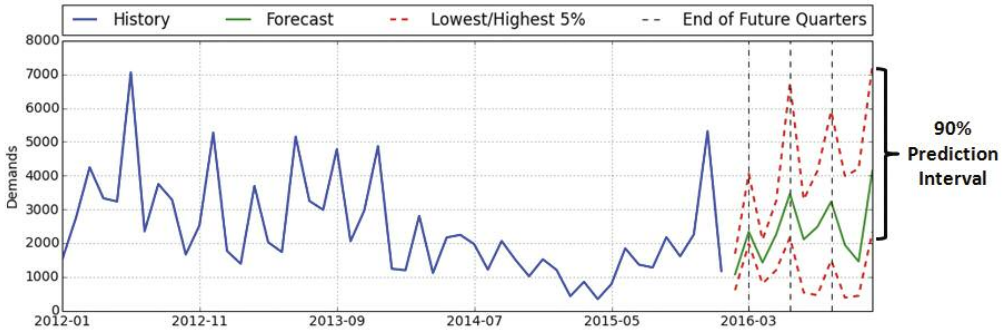

Figure 2: Range of expected demand scenarios originating from stochastic customer behaviour (average and extreme upper / lower predictions).

trated by their average and extreme instances for evaluation and reporting purposes as shown in Figure 2.

All data is fed into the mathematical model, which is processed by optimization software to obtain the optimal solution for this simplified setup. It does not help to obtain a fixed schedule, since the stochastic reality will invalidate such a static solution instantly. It is essential to determine an executable policy that is able to react to the concrete situation and that is robust in achieving good results even in the presence of variability. Applying those policies in combination with a periodic re-run of the entire process proves to be a powerful mechanism for volatile environments. The optimization also allows sensitivity analyses to be performed, which identify the bottlenecks that are most prohibitive for obtaining a better result - i.e. a higher service level or lower inventory - in the different capacity groups and over time. This information is valuable since even minor adjustments often lead to significant improvements but are very hard to determine owing to complicated network effects. In Figure 3, the red horizontal limits in periods 1, 2, 5, and 6 indicate critical bottlenecks in contrast to the green limits.

The simulation uses the obtained policies and runs hundreds of scenarios with different realizations of all stochastic parameters. The result of each scenario incorporates the full complexity of the business environment. This includes specific sequences of processing products at each point in time, dependencies regarding the availability of components, the maximal limit of produced units per day or the maximal "work in progress" (WIP) for a specific product. The different scenario results are aggregated to obtain a plausible range of what the company can expect in the future for each performance indicator of interest. It is straightforward to indicate potential future problems and their correlation on the BOM-tree via colour-coding (Figure 1). This strongly enhances the visibility across the entire supply network, facilitates a targeted problem resolution, and prevents more severe disruptions and variability. Of course, automated alerts can be sent to mobile devices to trigger an exception process.

In summary, the key ingredient for success for these complex projects is the right mixture of innovation and stability — by combining "cutting-edge" and "practice-proven" elements within models, algorithms, software, infrastructure and processes. The "intelligent glue" between these elements originates from both a deep expertise in this area and considerable creativity.

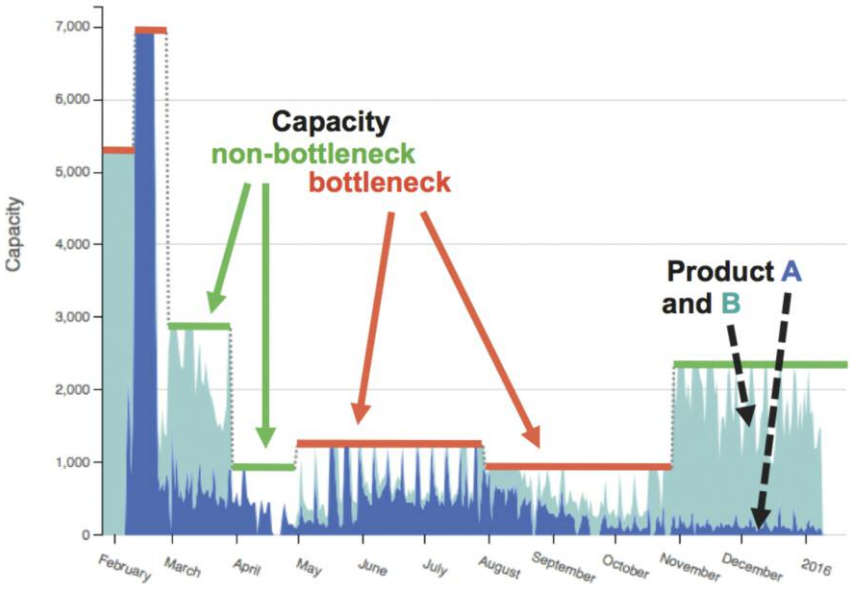

Figure 3: Capacity and utilisation chart over time (red limit: critical, service-level affecting bottleneck).

Link:
http://www.research.ibm.com/labs/zurich/business_optimization/

References:
[1] APICS: APICS Dictionary. 14th edition. Chicago, 2013.
[Ut11] Jim Utsler: Taming the supply chain. Interview with M. Parente et al., IBM Systems Magazine, July 2011, 12-15, 2011.

Please contact:
Ulrich Schimpel
IBM Research Zurich, Switzerland
uls@zurich.ibm.com

European Research and Innovation

Business Process Execution Analysis through Coverage-based Monitoring

by Antonello Calabrò, Francesca Lonetti, Eda Marchetti, ISTI-CNR

Nowadays, Business Process Model and Notation (BPMN) represents the de facto standard language for creating a description of processes and then developing executable frameworks for the overall planning and management of these processes. We describe a methodology for the development of execution adequacy criteria able to identify the main entities of the business process that are covered during execution, and to issue a warning if some entities are not covered.

The monitoring of business process execution is crucial not only for business process management and validation but also for performance analysis and optimization. When assessing the thoroughness of a business process execution it is important to identify entities that have not been observed for some time and to be able to check whether something is going wrong. Existing work in this field focuses on monitoring and analyzing the factors that influence the performance of business processes. Key performance indicators (KPIs), including time-based and cost-based parameters, are defined together with their target values based on the business goals.

The aim of our approach is to use monitoring facilities in order to measure the adequacy of the business process execution, identifying the entities that should be covered (activities, connection objects, swimlanes, etc.), assessing what percentage has been covered, and issuing a warning if some entities are not covered. These coverage-based business process execution adequacy measures make it possible, on the one hand, to detect unexpected behavior or security flaws in the business process execution and, on the other, to improve business process planning and performance optimization. This idea has been inspired by coverage-based test adequacy, which has been extensively studied in software testing, e.g. referring to the coverage of entities in the program control-flow or data-flow, and nowadays constitutes a fundamental instrument for test suite evaluation.

The main components of the monitoring framework that measures the business process execution adequacy are:

- BPMN Path Explorer: This component explores and saves all the possible entities (Activity Entity, Sequence Flow Entity, Path Entity) reachable on a BPMN.

- Complex Event Processor (CEP): This is the rule engine that analyzes the events generated by the business process execution and correlates them to detect patterns in order to monitor the business process execution adequacy.

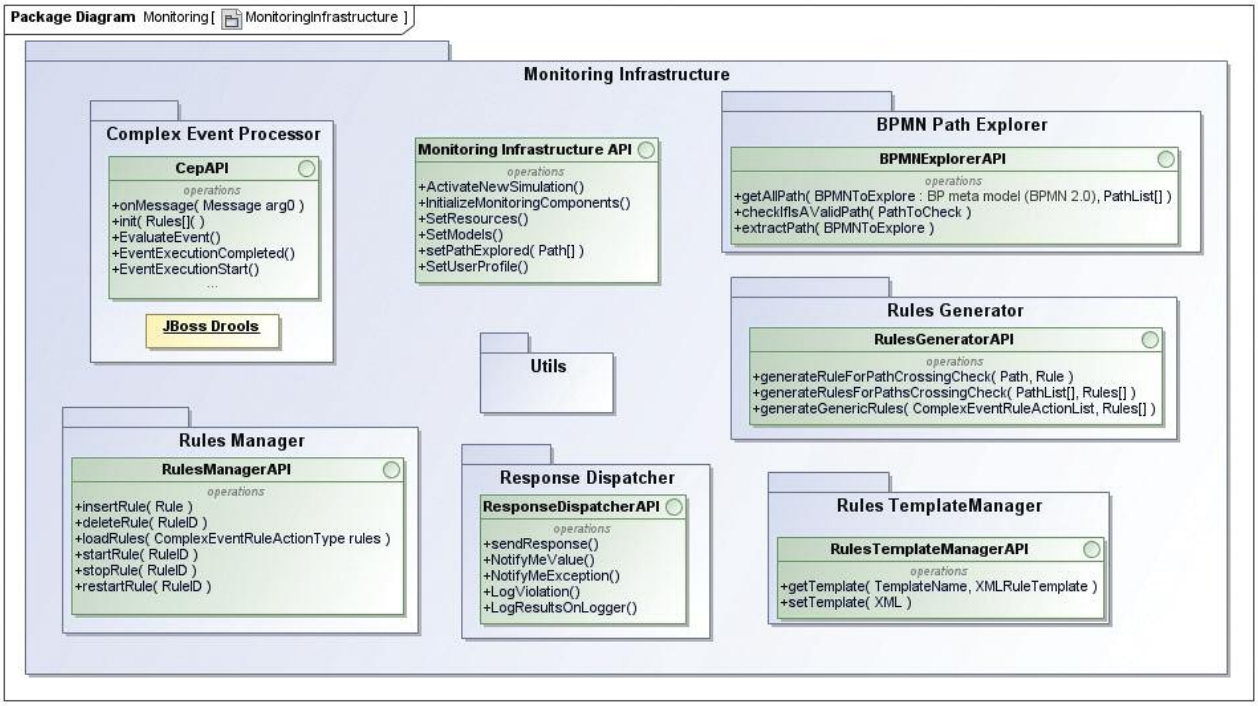

Figure 1: Monitoring Framework.

- Rules Generator: This component generates the rules that monitor the business process execution. It uses the templates stored in the Rules Template Manager. These rules are generated according to the specific adequacy criterion to be assessed and the entities to be covered. For each entity, the rule generator generates one corresponding rule for the CEP. A generic rule consists of two main parts: in the first part the events to be matched (the entities to be covered) are specified; the second part includes the events/actions to be notified after the rules evaluation.

- Rules Template Manager: This component is an archive of predetermined rule templates that are instantiated by the Rules Generator. A rule template is a rule skeleton, the specification of which has to be completed by instantiating a set of template-dependent placeholders. The instantiation will refer to appropriate values inferred from the specific adequacy criterion to be assessed.

- Rules Manager: The complex event detection process depends directly on the operations performed by the Rules Manager component, which is responsible for loading and unloading sets of rules into the complex event processor, and firing it when necessary.

- Response Dispatcher: The Response Dispatcher is a registry that records the business process execution adequacy monitoring requests. When it receives notice of a rule firing or pattern completion from the CEP, it stores coverage information. It then produces statistics on the overall percentage of the covered entities and sends warning messages regarding those entities that are not covered to the consumer/requester of the business process adequacy evaluation.

A first assessment of the approach proposed was made on a case study developed within the Learn PAd European project [1]. It was demonstrated that the coverage measurements provided were useful for assessing the adequacy of a business process based learning session.

Reference:

[1] Bertolino, Antonia, Antonello Calabró, Francesca Lonetti, and Eda Marchetti. "Towards Business Process Execution Adequacy Criteria." In Software Quality. The Future of Systems-and Software Development, pp. 37-48. Springer International Publishing, 2016.

Please contact:

Antonello Calabró, ISTI-CNR, Italy
antonello.calabro@isti.cnr.it

Quality of Experience-assessment of WebRTC Based Video Communication

by Doreid Ammar, Katrien De Moor and Poul Heegaard, NTNU

Web real-time communication has enabled hassle-free, no installation, in-browser applications such as Google hangout and appear.in. Multi-party video conferencing has now finally been made easy. But how can we provide acceptable quality of experience in such an interactive service? Our research aims to gain insight into what matters, and how to assess, design, and manage the services accordingly.

Web real-time communication (WebRTC) has become popular in recent years, with numerous free of charge applications, such as appear.in and Google hangouts, used in both private and professional contexts. Such applications enable real-time (audiovisual) communication in the browser with multiple parties, without need for plug-ins or other requirements. As a result, such applications are very easy to use and for some applications (e.g. appear.in [L1]), it is not even necessary to create a user account. Users can simply connect and access WebRTC-based applications from a wide range of devices (e.g. smartphone, tablet, laptop), using a web-browser that supports WebRTC (e.g. Chrome, Firefox) or a dedicated App.

However, WebRTC-based multi-party video communication is highly interactive and Web-RTC based applications can be used from multiple devices and in very different contexts. Moreover, the implications for users and their quality of experience (QoE), i.e., their degree of delight or annoyance when using the application [1] are not fully understood. Users may for example experience different types and gradations of quality impairments (e.g. video freezes, bad or no audio) during a conversation, which may make smooth communication nearly impossible. It is vitally important that application and service providers address such issues, in order to prevent users from getting so annoyed that they stop using the application and switch to a competitor. Preventing user frustration and providing the best possible experience requires thorough insights into the various technical and non-technical factors that may influence users' QoE.

In the scope of the Telenor-NTNU research collaboration project "Quality of Experience and Robustness in Telecommunications Networks", which started in 2015, we aim to:

- identify the most relevant influence factors, with a primary focus on technical (quality of service-related) factors, but also including non-technical (contextual, human level) influence factors;
- investigate in which ways and to what extent factors influence users' QoE and corresponding user behaviour; and
- from an understanding of these relationships and thresholds, provide input to the development of QoE-aware adaptation strategies to reduce/avoid user annoyance, and foster user delight.

The WebRTC based application called appear.in, which enables video communication for up to eight parties and which can be accessed using browsers that support WebRTC, is used as a concrete use case.

A multi-method approach is used to gain more insight in the relevant influence factors and investigate the relationship between performance-related parameters and users' QoE. We have conducted a literature study, analyzed historical data, and used exploratory focus groups with appear.in users. In our ongoing work, we perform real-time logging and analysis of performance and session-related statistics using a research version of appear.in. For this purpose, we use both Google Chrome's WebRTC internal functionality and the WebRTC Analytics interface getstats.io. Different types of data are being collected and analyzed, including subjective, explicit user feedback (collected at the end of a session), implicit and behavioural user feedback (collected during a session), and objective performance-related statistics. So far we have run a series of tests with two parties according to different test scenarios and network conditions (see Figure 1). We collected real-time session statistics by means of Google Chrome's WebRTC-internals tool. Although the Chrome statistics have a number of limitations, we have found that they are useful for QoE research. The results high-

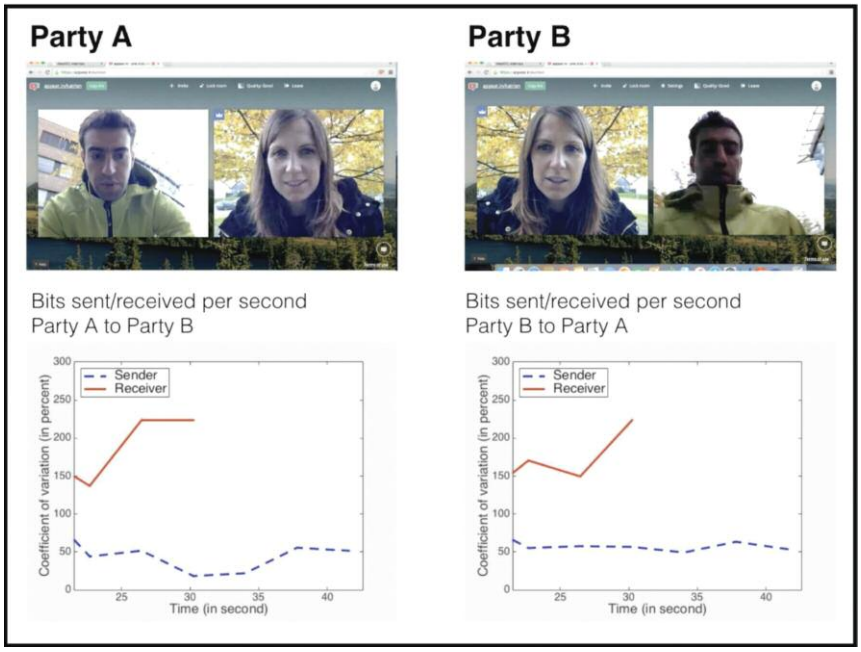

Figure 1: A two-party video communication using appear.in.

light the relevance of some of the investigated performance statistics for detecting potential QoE issues.

We are now planning to conduct both controlled laboratory and longitudinal "living lab" empirical studies involving real users. For this we are developing an experimental test platform, will log performance stats in real-time, collect explicit and implicit user feedback in test users' natural environments, and perform in-depth data analytics and data visualization. Of particular interest is how users´ tolerance levels towards different types of QoE impairments and users' expectations evolve over time and in a real-life context, as a complement to the 'instantaneous' QoE assessment in the lab setting.

We believe that our research activities will help to formulate concrete recommendations on how to meet and exceed users' QoE requirements and to generate research-based knowledge on how to foster user delight and avoid annoyance in the context of WebRTC-based video communication.

The Telenor-NTNU research collaboration project "Quality of Experience and Robustness in Telecommunications Networks" involves a multidisciplinary team, including current ERCIM fellow Doreid Ammar and former ERCIM fellow Katrien De Moor, and is a collaboration between the following European research institutions: The Department of Telematics at Norwegian University of Science and Technology – NTNU (Trondheim, Norway), Telenor research (Trondheim, Norway), represented by Dr. Min Xie and the Department of Communication Systems at Blekinge Institute of Technology – BTH (Karlskrona, Sweden), represented by Prof. dr. Markus Fiedler.

Link:
[L1] Appear.in: https://appear.in

Reference:
[1] A. Raake, S. Egger: "Quality and Quality of Experience," in Quality of Experience: Advanced Concepts, Applications and Methods, in: S. Möller and A. Raake (Eds.), Springer, 2014.

Please contact:
Doreid Ammar
ERCIM fellow at NTNU, Norway
Tel: +47 73 59 43 21
doreid.ammar@item.ntnu.no

D2V – Understanding the Dynamics of Evolving Data: A Case Study in the Life Sciences

by Kostas Stefanidis, Giorgos Flouris, Ioannis Chrysakis and Yannis Roussakis, ICS-FORTH

D2V, a research prototype for analysing the dynamics of Linked Open Data, has been used to study the evolution of biomedical datasets, such as the Experimental Factor Ontology (EFO) and the Gene Ontology (GO).

Datasets are continuously evolving over time as our knowledge increases. Biomedical datasets in particular have undergone rapid changes in recent years, making it difficult for engineers and scientists to follow the evolution and keep up with recent developments.

To address the problem of managing evolving datasets and understanding their evolution, we have developed D2V [L1], a research prototype designed for studying the dynamics of web data and the associated Linked Open Data (LOD). Specifically, D2V is able to detect and analyse changes and the evolution history of LOD datasets, thereby allowing remote users of a dataset to identify changes, even if they have no access to the actual change process. Interestingly, it also empowers users to perform sophisticated analysis on the evolution data, so as to understand how datasets (or parts of them) evolve, and how this evolution is related to the data itself. For instance, one may be interested in specific types of evolution, e.g., classes becoming obsolete, or in the evolution of a specific entity (e.g., a specific disease or genome). Our tool aims to become a critical addition to the arsenal of data analysts and scientists for dynamicity analysis in biomedical or other datasets.

For example, consider Alice, a specialist in the field of genetic diseases, who is interested in the connections between human genome and diseases, and frequently annotates the Disease Ontology [L2] with her research findings. As ontology terms change, often becoming obsolete, i.e., replaced by other terms, Alice needs an easy way to identify the changes related to the object of her study and understand how these changes affect her research. To help Alice, D2V allows her to manage, detect, and view changes in a variety of ways.

In particular, D2V handles two types of changes, with the aim of making changes intuitive and human-understandable. The first is simple changes, which are fine-grained changes defined at design time and provide formal guarantees on the soundness and completeness of the detection process. The second type is complex changes, which are custom-built and defined at run-time by the user to satisfy application-specific needs; for example, complex changes may be used to report coarse-grained changes, changes that are important for the specific application or user, changes with special semantics,

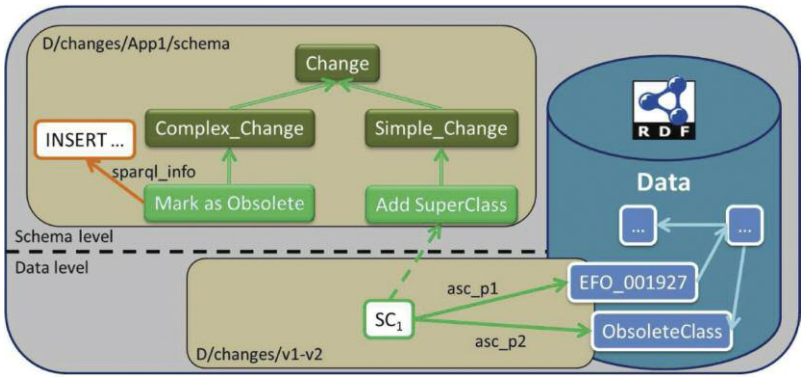

Figure 1: The ontology of changes.

We designed and developed D2V at FORTH-ICS, an ERCIM institute, in collaboration with researchers from ATHENA and EMBL-EBI, in 2014-2015. We plan to enhance our software with additional features, and incorporate additional biomedical datasets.

The development of D2V was funded by the EU FP7 projects DIACHRON and IdeaGarden. Further details can be found at [1], [2].

or changes that should not happen at all (their detection being like an alert for an abnormal situation).

To detect changes, we rely on the execution of appropriately defined SPARQL queries, a W3C standard for querying LOD. The answers to these queries determine the detected changes, which are represented as instances of an ontology of changes (Figure 1). This allows the connection of the detected changes with the actual data using standard LOD principles, blending the data with the evolution history and supporting navigation among versions, cross-snapshot and historic queries. The analysis of the evolution history is based on SPARQL.

The retrieval of the detected changes from the ontology allows them to be presented to the user, through interactive interfaces and visualization paradigms (see Figure 2 for a visualization of EFO [L3] evolution). We provide different views of the evolution history for different types of analyses: the user can see the evolution history of a given URI (term-centric view – e.g., return all changes associated with "adrenal gland disease"), of the dataset as a whole (dataset-centric view – e.g., view all changes in the Disease Ontology), or of specific versions (version-centric view – e.g., view all changes between a given pair of versions); or the user may be interested in a change-centric view, where the instantiations of a given change are reported (e.g., return all classes that were made obsolete); the user can filter the different results to a fixed set of changes, change types, or versions; or visualize evolution along a series of consecutive versions, or for an arbitrary pair only.

Links:
[L1] http://www.ics.forth.gr/isl/D2VSystem/
[L2] http://disease-ontology.org/
[L3] http://www.ebi.ac.uk/efo/

References:
[1] Y. Roussakis, et al.: "A Flexible Framework for Understanding the Dynamics of Evolving RDF Datasets", Best Student Paper Award, ISWC-15, Research Track.
[2] Y. Roussakis, et al.: "D2V: A Tool for Defining, Detecting and Visualizing Changes on the Data Web", ISWC-15, Demonstrations Track.

Please contact:
Kostas Stefanidis, ICS-FORTH, Greece
Tel: +30 2810 391635
kstef@ics.forth.gr

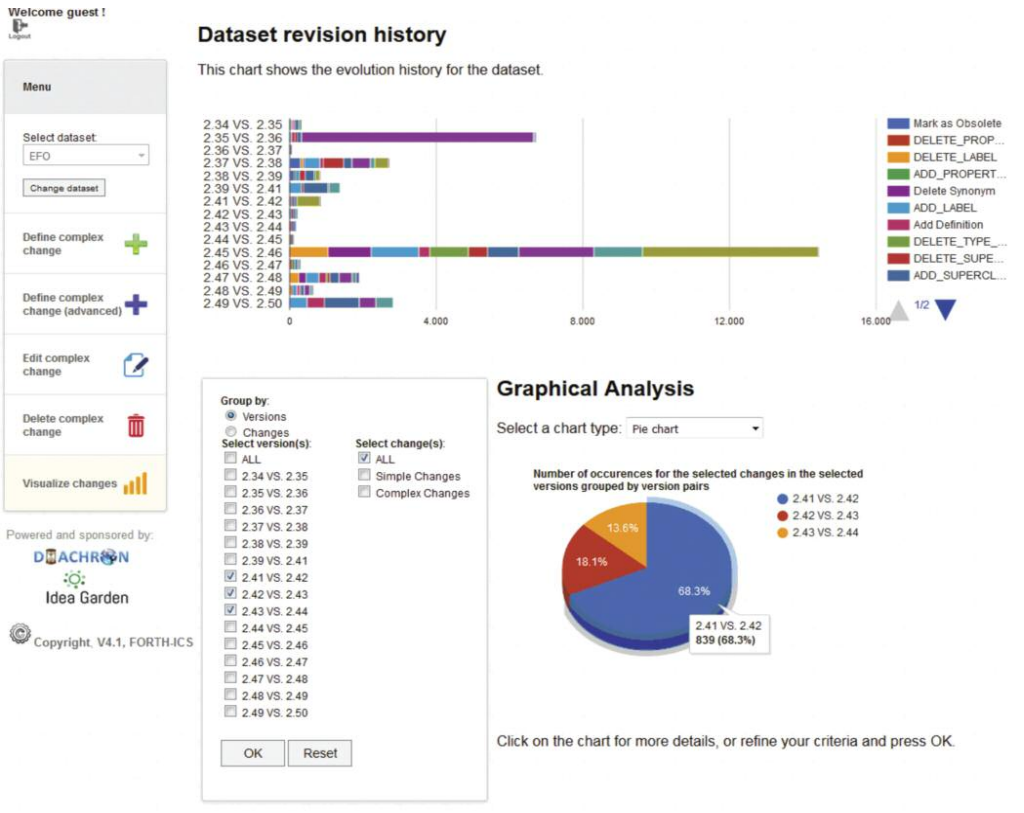

Figure 2: Dataset-centric and version-centric view.

Detection of Data Leaks in Collaborative Data Driven Research

by Peter Kieseberg, Edgar Weippl, SBA Research, and Sebastian Schrittwieser, TARGET

Collaborative data driven research and development is one of the big issues in BigData-Analysis. Still, several obstacles regarding the preservation of ownership over the shared data need to be overcome to enable the techniques to unfold their full potential. This especially concerns the detection of partners that leak valuable data to the outside world.

In the current age of big data, machine learning and semantic systems are playing increasingly important roles in data analysis. In the medical sector in particular, large volumes of data are stored in well-protected databases at various institutions, such as hospitals and research facilities. But more traditional industries, such as production environments and factories, have also increased the incorporation of sensor data for optimization purposes or in order to make new features and services available. The industry 4.0 paradigm is often seen as a trendsetter for the next decade, with different speeds of adoption by different industries.

Scientific analysis of data sets could often be improved tremendously by cooperating with other institutions like research facilities or suppliers, which either have access to additional data (e.g. complementary information on the production process of raw material) or are researching new analytical methods or services. Especially when developing algorithms in the area of machine learning, the possibility of tailoring the algorithm to the specific problems, including investigating the best definition for the parameters, can have a vast impact on the overall performance. Thus, many industrial development and research programs could profit from exchanging relevant data between the participants.

While data protection is a well-researched area, see for example [1], the problem remains that even when using proper anonymization, the leakage of data sets can lead to severe problems. The reasons for this range from financial concerns, e.g. when the data sets in question are offered for sale, to strategic reasons, e.g. the anonymized data provides information about the interests and/or capabilities of the institution or industry, to the desire to keep control over the distribution of the data. Additionally, the threat of collusion attacks that can be used to subvert anonymization arises when differently anonymized data sets are provided to multiple data analysts. Fingerprinting techniques can be used in order to mitigate the risk of unwanted data leakage by making it possible to attribute each piece of distributed data to the respective data recipient.

While some approaches exist for fingerprinting structured data as is typically stored in tables, these usually rely on the availability of a substantial portion of the leaked data to identify the leak. In the DEXHELPP-project [L1], we thus conduct research on a combined approach for anonymization and fingerprinting in a single step, as both needs are apparent in many data driven environments: The data set for each recipient is anonymized slightly differently using k-anonymity based on generalization and providing roughly the same remaining data quality, thus resulting in data sets with different anonymization strategies but close to equal value (see Figure 1).

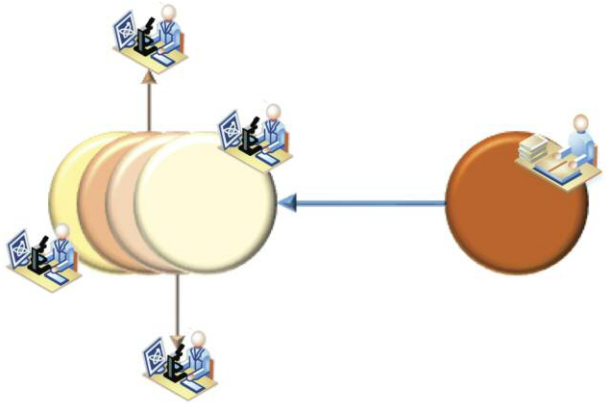

Figure 1: Data owner provides different versions to recipients.

Since generalization achieves k-anonymity by applying generalization to each individual record in the data set, only a single record is needed to detect and identify the leaking data recipient (see Figure 2):

(A) A data record is encountered in the wild.
(B) The generalization strategies for the field attributes are identified and the resulting identification pattern is generated.
(C) The identification pattern is matched against a list holding all data recipients and the patterns of their respectively received data sets, thus allowing identification of the leaking recipient.

One of the main problems of this approach lies in the danger of colluding attackers and resulting inference attacks. Thus, the generalization patterns must be selected in such a way that any arbitrary combination of these data sets neither results in a reduction of the anonymization level k, nor in the ability to hide the identities of the colluding attackers (see [2]). Furthermore, the theoretical approach based on utilizing generalization and k-anonymity will be extended to incorpo-

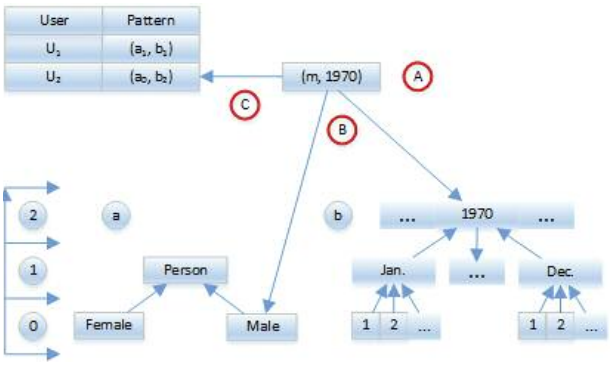

Figure 2: Detection of the data leak.

rate more advanced concepts for anonymization, as well as more diverse mechanisms for achieving anonymity. The main reason for this research step lies in the various other prerequisites for data anonymization, especially in the medical sector, that are not fulfilled by basic k-anonymity, e.g. considering dependencies between data records or statistical attacks.

During the DEXHELPP-project, the devised algorithms will not only be tested with synthetic data, but also with real medical records. This is especially important, as the actual performance of this fingerprinting approach largely depends on the ability to cluster the data into equivalency classes that form a partition of the whole data set. This again relies very much on the actual distribution of the data inside the data set, especially when refraining from removing outliers. Principal performance identifiers are not only related to KPIs that are typically identified with performance, but mainly with attributes like the number of possible fingerprints in a given data set, the distance between two fingerprinted sets, as well as the achievable remaining quality.

In conclusion, our work on fingerprinting data sets in the DEXHELPP project will result in practical fingerprints for collaborative data driven research and development environments. The resulting best fingerprinting strategy will be implemented inside the DEXHELPP test server environment, which was designed to securely perform collaborative medical research on data provided by the Austrian healthcare providers whilst preserving privacy.

Link:
[L1] http://dexhelpp.at/?q=en

References:
[1] B. Fung, et al.: "Privacy-preserving data publishing: A survey of recent developments", ACM Computing Surveys (CSUR) 42(4), 14 (2010).
[2] P. Kieseberg, et al.: "An algorithm for collusion-resistant anonymization and Fingerprinting of sensitive microdata", Electronic Markets - The International Journal on Networked Business, 2014.

Please contact:
Peter Kieseberg, SBA Research, Vienna, Austria
pkieseberg@sba-research.org

HOBBIT: Holistic Benchmarking of Big Linked Data

by AxelCyrille Ngonga Ngomo, InfAI, Alejandra García Rojas, ONTOS, and Irini Fundulaki, ICS-FORTH

The HOBBIT project aims at abolishing the barriers in the adoption and deployment of Big Linked Data. To this end, HOBBIT will provide open benchmarking reports that allow to assess the fitness of existing solutions for their purposes. These benchmarks will be based on data that reflects reality and measures industry relevant Key Performance Indicators (KPIs) with comparable results using standardized hardware.

Linked Data has grown rapidly over the last ten years [L1]. Organizations are increasingly interested in using solutions based on Linked Data. However, choosing the right solution for their needs remains a difficult task. HOBBIT's rationale is to support organizations that aim to use Linked Data technologies at all scales (including Big Data) in the choice of appropriate solutions. To this end, HOBBIT [L2] will provide benchmarks for all the industry relevant phases of the Linked Data lifecycle [1] according to the approach shown in Figure 1.

In particular, the H2020 Hobbit EU project will create benchmarks for the following stages:

1. Generation and Acquisition: Benchmarks pertaining to the transformation of unstructured, semi-structured and structured data into RDF.
2. Analysis and Processing: Benchmarks pertaining to the use of Linked Data to perform complex tasks such as supervised machine learning.
3. Storage and Curation: Benchmarks pertaining to the storage, versioning and querying of RDF data stored in corresponding solutions.
4. Visualisation and Services: Application centric benchmarks pertaining to queries used by software solutions which rely on large amounts of Linked Data.

The HOBBIT project will provide innovative benchmarks based on the following premises:

- Realistic benchmarks: Benchmarks are commonly generated with synthetic data that reflect a single and specific domain. HOBBIT aims at creating mimicking algorithms to generate synthetic data from different domains.
- Universal benchmarking platform: We will develop a generic platform that will be able to execute large scale benchmarks across the Linked Data lifecycle. The platform will provide reference implementations of the KPIs as well as dereferenceable results and automatic feedback to tools developers..
- Industry relevant Key Performance Indicators (KPIs): In addition to the classical KPIs developed over the last decades, HOBBIT will collect relevant KPIs from industry to make the assessment of technologies based on the industrial needs possible.

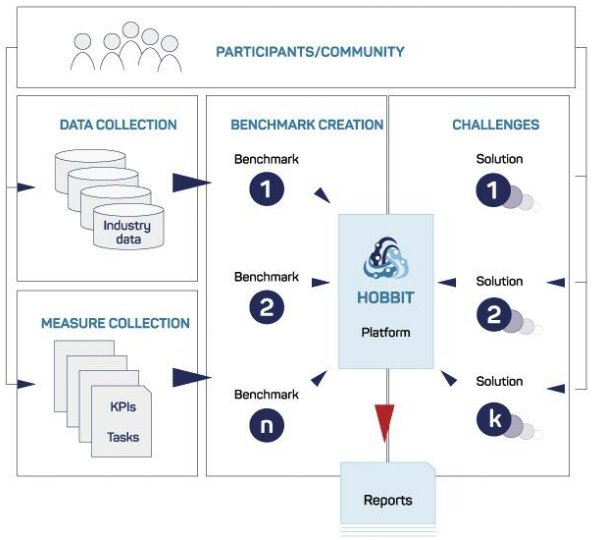

Figure 1: HOBBIT Ecosystem.

Comparability of results: Within its challenges, HOBBIT will provide the possibility of deploying benchmarks on hardware provided by the project but also in the cloud. Therewith, we will ensure that the results achieved by the frameworks benchmarked by HOBBIT are comparable.

Independent HOBBIT institution: HOBBIT will be grow to become a biasfree organization that will conduct regular benchmarks and provide European industry with up-to-date performance results. We target to initiate the organization in 2017. To achieve the goals of HOBBIT, we aim at bringing together three main categories of actors: solution providers, technology users and scientific community. In the first stage of the project we will target all three categories of actors that will provide us with (1) supplementary KPIs and (2) datasets from their domain for the mimicking algorithms used to create the benchmarks. Once the first version of the different benchmarks is available, we will challenge solution providers and the scientific community to evaluate their solutions.

These dereferenceable benchmark results will be provided in an open format [2]. The members of the HOBBIT community [L3] will be provided with the early results and early insights pertaining to the benchmarks and results of the project. HOBBIT's consortium is composed of 10 organizations: InfAI (Germany, coordinator), Fraunhofer IAIS, (Germany), Foundation for Research and Technology Hellas (Greece), National Center for Scientific Research "Demokritos" (Greece), iMinds VZW (Belgium), USU Software AG (Germany), Ontos AG (Switzerland), OpenLink Software (UK), AGT Group GmbH (Germany) and TomTom Polska (Poland).

Links:
[L1] http://linkeddatacatalog.dws.informatik.uni-mannheim.de/state/
[L2] http://project-hobbit.eu/, https://twitter.com/hobbit_project
[L3] http://projecthobbit.eu/getinvolved/

References:
[1] S. Auer, et al.: "Managing the lifecycle of linked data with the LOD2 stack", in The Semantic Web–ISWC 2012 (pp. 116), Springer.
[2] R. Usbeck, et al.: "GERBIL: general entity annotator benchmarking framework", Proc. of the World Wide Web Conference 2015.

Please contact:
Axel Cyrille NgongaNgomo,
University Leipzig, Germany
ngonga@informatik.uni-leipzig.de

Alejandra Garcia Rojas
Ontos AG, Switzerland
alejandra.garciarojas@ontos.com

Irini Fundulaki
ICS-FORTH, Greece
fundul@ics.forth.gr

ERCIM "Alain Bensoussan" Fellowship Programme

ERCIM offers fellowships for PhD holders from all over the world. Topics cover most disciplines in Computer Science, Information Technology, and Applied Mathematics. Fellowships are of 12-month duration, spent in one ERCIM member institute. Fellowships are proposed according to the needs of the member institutes and the available funding.

Application deadlines: 30 April and 30 September.

More information: http://fellowship.ercim.eu/

Smart Solutions for the CNR Campus in Pisa

by Erina Ferro, ISTI-CNR

The research area of the Italian National Research Council (CNR) in Pisa is a 130,000 m^2 village where 3000 people live and move daily. It can thus be considered as a laboratory where smart solutions to problems that can be encountered in any town or city can be tested.

The Smart Badge registers someone's presence when her or his own badge has been lost or forgotten, or is not usable as the person is temporarily working in another structure. The smart badge app, installed on a smart phone, offers three main functionalities: on-site check-in/check-out, remote registration, a list of recent entries/exits. In order to implement the remote check-in/check-out, the only hardware subsystem required is a GPS system, available on most smartphones: the user can see his own position and confirm the registration. For the local check-in/check-out, the app displays a dynamically generated QR-code with a user identifier, a code for the operation, and a unique one-time password that must match the corresponding password generated on the reader.

The main features of our low-cost Smart Parking application are: a multi-technology platform, interoperability, flexibility, utility, and low-bandwidth requirements for data transmission. Three monitoring technologies have ben integrated: four commercial ground sensors, installed in four parking lots and used as ground truth, nine newly developed embedded cameras, and nine newly developed smart cameras; each type of camera monitoring about 20 lots simultaneously. In both cameras, data are elaborated on board, and the binary information (empty/busy) regarding the monitored spaces is transmitted to the cloud dedicated to the smart area infrastructure. The smart cameras are constituted by a Raspberry Pi, equipped with a camera module; an algorithm has been developed that exploits some visual processing techniques to analyze the captured images, and a neural network trained using a deep learning approach to determine the occupancy status of the parking spaces.

The embedded cameras have a main board that manages both vision and networking tasks; other components are provided by the power supply system that manages charging and applies optimal energy savings policies. A dedicated vision logic has been designed and implemented.

The Smart Shared Mobility (SSM) allows users to share urban and suburban trips, customized wrt their preferences, requirements, and ride sharing habits. SSM includes a web application, accessible from the Smart Area web site, and the mobile app GoTogether. The solution provides highly personalized suggestions, both in response to explicit user requests, and proactively, selecting a set of most appropriate solutions based on the user's history in the app usage and on reminder features. The solution is completely integrated with the Smart Area Cloud, through which it can access data of the other services (such as the smart parking) to provide additional features, and can also push its own data to the cloud, in order to enrich the entire system.

The Smart Campus Group of the CNR Institutes in Pisa: ISTI, IIT and IFC.

The Smart Building solution not only monitors the energy consumption of each single office, but also implements energy saving policies by recognising the presence of people in the room, and enables policies to detect and react to inappropriate intrusions in the monitored offices. To achieve these goals, a wireless sensor network (WSN) has been developed, where each node is equipped with many different types of transducers. The ZigBee technology is employed; each node of the WSN is connected to a Zigbee sink to which an IPv4 address is associated. Each in-building sensor is characterized by a physical location. The solution develops a pervasive infrastructure represented by the middleware used to build location-based services and is used to integrate different kinds of sensors thanks to its modular nature.

The Smart Navigation application (for smart phone and tablet) consists of an interactive, HTML5-based map of the CNR area in Pisa, and an indoor positioning service. The map is used to see the user's own position within or without the buildings on the campus, to explore the floors of the buildings, to search for rooms or people, or to obtain directions to a specified destination. The interactive map is a web application consisting in a rich interactive map of the CNR campus, available to users with modern desktop or mobile browsers. No external proprietary services, such as Google Maps or similar, are used at runtime by the web application. The user indoor positioning facility constitutes a large research field; the solution we are implementing is based on Wi-Fi fingerprinting. The idea is to exploit the ability of the Wi-Fi receiver of a smartphone, carried by the user, to scan the existing Wi-Fi networks and, for each one, to measure the received signal strength (RSS) and to compare it with a database of measurements performed in known locations.

All data generated by the above mentioned applications converge in the Smart Area infrastructure, which is constituted by an OpenStack-based "Infrastructure as a Service" cloud architecture. In the cloud, meteorological data derived from the eight meteorological stations of CNR located throughout Italy is also collected.

Link: http://www.smart-applications.area.pi.cnr.it

Please contact:
Erina Ferro, ISTI-CNR, Italy
erina.ferro@isti.cnr.it

Call for Papers and Participation

FMICS-AVoCS 2016: International Workshop on Formal Methods for Industrial Critical Systems and Automated Verification of Critical Systems

CNR, Pisa, Italy, 26-29 September 2016

In 2016, FMICS and AVoCS join their forces to hold a workshop combining their themes of formal methods and automated verification. For FMICS, this will be the 21st, for AVoCS the 16th edition.

The aim of the yearly workshop of the ERCIM working group FMICS is to provide a forum for researchers who are interested in the development and application of formal methods in industry. FMICS brings together scientists and engineers who are active in the area of formal methods and interested in exchanging their experiences in the industrial usage of these methods. The FMICS workshop series also strives to promote research and development for the improvement of formal methods and tools for industrial applications.

The aim of the AVoCS workshop series is to contribute to the interaction and exchange of ideas among members of the international research community on tools and techniques for the verification of critical systems. It covers all aspects of automated verification, including model checking, theorem proving, SAT/SMT constraint solving, abstract interpretation and refinement pertaining to various types of critical systems which need to meet stringent dependability requirements (safety-critical, business-critical, performance-critical, etc.).

Contributions

Full paper submissions must describe the authors' original research and results and should clearly demonstrate relevance to industrial application. Case study papers should identify lessons learned, validate theoretical results (like scalability of methods) or provide specific motivation for further research and development. Submissions should not exceed 15 pages LNCS style. All submissions will be reviewed by the Program Committee who will select papers based on the novelty, soundness and applicability of the ideas and results.

FMICS-AVoCS also encourages the submissions of research ideas in order to stimulate discussions at the workshop. Reports on ongoing work or surveys on work published elsewhere are welcome. The Program Committee will select research ideas on the basis of submitted abstracts according to their significance and general interest. Abstracts should not exceed 3 pages and they will not be part of the published proceedings.

Important Dates

- 25 April: submission of full papers
- 19 June: notification for full papers
- 10 July: final versions of full papers
- 10 August: submission of research ideas
- 17 August: notification of research ideas
- 26-29 September: FMICS-AVoCS workshop

The workshop proceedings will be published by Springer in their LNCS series, while authors of the best full papers will be invited to submit extended versions to a special issue of Springer's International Journal on Software Tools for Technology Transfer.

More information:
http://fmics-avocs.isti.cnr.it/

Call for Nominations

Minerva Informatics Equality Award

The Minerva Informatics Equality Award, presented by Informatics Europe, recognizes best practices in Departments or Faculties of European universities and research labs that encourage and support the careers of women in Informatics research and education.

On a three-year cycle, the award will focus each year on a different stage of the career pipeline:

- Developing the careers of. female faculty
- Supporting the transition for PhD and postdoctoral researchers into faculty positions
- Encouraging female students to enroll in Computer Science/Informatics programmes and retaining them.

The 2016 Award is the first of this annual Award and is devoted to gender equality initiatives and policies to develop the careers of female faculty. It seeks to celebrate successful initiatives that have had a measurable impact on the careers of women within the institution.

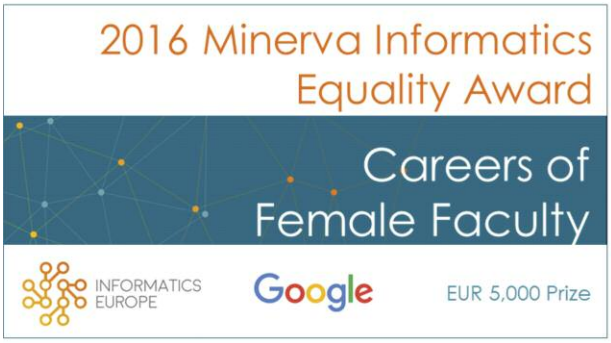

The 2016 Minerva Informatics Equality award is sponsored by Google and will be presented during the 2016 European Computer Science Summit (ECSS 2016) in Budapest, Hungary, October 2016.

More information:
http://kwz.me/UK

SHIFT2RAIL - European Railway Research of 2015-2024

SHIFT2RAIL is a joint technical initiative that will handle all railway research within Horizon 2020. The founding members are Trafikverket, Network Rail, Bombardier, Alstom, Siemens, Ansaldo STS, CAF and Thales, together with the European Commission. The total budget for the program during 2015-2024 is about 800 million €. Shift2Rail has five main innovation programs that can briefly be described as development of vehicles, signal systems, infrastructure, information systems and freight traffic. The overall ambitious targets are to double the capacity of the European rail system, and increase its reliability and service quality by 50%, all while halving the lifecycle costs.

As a part of the Swedish rail research program "Capacity in the railway traffic system" (in Swedish abbreviated "KAJT"), SICS Swedish ICT will be a part of Shift2Rail, acting as a partner to the Swedish infrastructure manager Trafikverket. The first research projects within Shift2Rail are expected to start in autumn of 2016. SICS Swedish ICT will mostly be active in areas related to resource optimization, like yard- and timetable planning, areas in which SICS has vast previous experience.

"We have noticed a great interest in the upcoming Shift2Rail-projects, both from researchers around Europe and from the railway industry. We are looking forward to developing the close relationship to Trafikverket and other research organizations in this field during the next years and we are proud to be a part of this", says Martin Joborn, senior researcher at SICS Swedish ICT and coordinator of the research program KAJT.

More information: http://www.shift2rail.org

The ExaNeSt Project - Fitting Ten Million Computers into a Single Supercomputer

The next generation of supercomputers must be capable of a billion billion calculations per second. These are referred to as Exascale computers and with this ability to undertake such volume of calculations, they will transform our understanding of the world through advanced simulation and problem solving. It will take ten million processors working together to achieve Exascale; so how can this be achieved?

A step towards the Exascale vision is being made by a EU funded ExaNeSt project, which is building its first straw man prototype this year. The ExaNeSt consortium consists of twelve partners, each of which has expertise in a core technology needed for innovation to reach Exascale. ExaNeSt takes the sensible, integrated approach of co-designing the hardware and software, enabling the prototype to run real-life evaluations, facilitating its scalability and maturity into this decade and beyond.

Being able to move, process and manage unprecedented volumes of data would allow greater insight into many areas of our lives including climate change, cosmology, drug design, energy safety, national security, material science, medicine and countless other scientific and engineering disciplines. Current technology is faced with many technical limitations in reaching an Exascale architecture. Key barriers are energy and cooling demands, compact packaging, permanent storage, interconnection, resilience and not least application behavior. ExaNeSt addresses these using energy-efficient ARM cores, quiet and power-efficient liquid cooling, non-volatile (e.g. flash) memories integrated into the processor fabric, and the development of innovative, fast interconnects that avoid congestion.

More information: http://www.exanest.eu

ERCIM Membership

After having successfully grown to become one of the most recognized ICT Societies in Europe, ERCIM has opened membership to multiple member institutes per country. By joining ERCIM, your research institution or university can directly participate in ERCIM's activities and contribute to the ERCIM members' common objectives to play a leading role in Information and Communication Technology in Europe:

- Building a Europe-wide, open network of centres of excellence in ICT and Applied Mathematics;
- Excelling in research and acting as a bridge for ICT applications;
- Being internationally recognised both as a major representative organisation in its field and as a portal giving access to all relevant ICT research groups in Europe;
- Liaising with other international organisations in its field;
- Promoting cooperation in research, technology transfer, innovation and training.

About ERCIM

ERCIM – the European Research Consortium for Informatics and Mathematics – aims to foster collaborative work within the European research community and to increase cooperation with European industry. Founded in 1998, ERCIM currently includes 21 leading research establishments from 18 European countries. Encompassing over 10 000 researchers and engineers, ERCIM is able to undertake consultancy, development and educational projects on any subject related to its field of activity.

ERCIM members are centres of excellence across Europe. ERCIM is internationally recognized as a major representative organization in its field. ERCIM provides access to all major Information Communication Technology research groups in Europe and has established an extensive program in the fields of science, strategy, human capital and outreach. ERCIM publishes ERCIM News, a quarterly high quality magazine and delivers annually the Cor Baayen Award to outstanding young researchers in computer science or applied mathematics. ERCIM also hosts the European branch of the World Wide Web Consortium (W3C).

"Through a long history of successful research collaborations in projects and working groups and a highly-selective mobility programme, ERCIM has managed to become the premier network of ICT research institutions in Europe. ERCIM has a consistent presence in European Community funded research programmes conducting and promoting high-end research with European and global impact. It has a strong position in advising at the research policy level and contributes significantly to the shaping of EC framework programmes. ERCIM provides a unique pool of research resources within Europe fostering both the career development of young researchers and the synergies among established groups. Membership is a privilege."

Dimitris Plexousakis, ICS-FORTH

Benefits of Membership

Institutions, as members of ERCIM AISBL, benefit from:

- International recognition as a leading centre for ICT R&D. ERCIM, a European-wide network of centres of excellence in ICT, is internationally recognised as a major representative organisation in its field;
- More influence on European and national government R&D strategy in ICT. ERCIM members team up to speak with a common voice and produce strategic reports to shape the European research agenda;
- Privileged access to standardisation bodies, such as the W3C which is hosted by ERCIM as to other bodies with which ERCIM has also established strategic cooperation. These include ETSI, the European Mathematical Society and Informatics Europe;
- Invitations to join projects of strategic importance;
- Establishing personal contacts among executives of leading European research institutes during the bi-annual ERCIM meetings;
- Invitations to join committees and boards developing ICT strategy nationally and internationally;
- Excellent networking possibilities with more than 10.000 high-quality research colleagues across Europe. ERCIM's mobility activities, such as the fellowship programme, leverages scientific cooperation and excellence;
- Professional development of staff including international recognition;
- Publicity through the ERCIM website and ERCIM News, the widely read quarterly magazine.

How to Become a Member

- Prospective members must be outstanding research institutions (including universities) within their country;
- Applicants shall address a request accompanied by short description to the ERCIM Office. The description must contain:
  - Name and address of the institute;
  - Short description of the institute's activities;
  - Staff (full time equivalent) relevant to ERCIM's fields of activity;
  - Number of European projects currently involved in;
  - Name of the representative and the alternate.
- Membership applications will be reviewed by an internal board and may include an on-site visit;
- The decision on admission of new members is made by the General Assembly of the Association, in accordance with the procedure defined in the Bylaws (http://kwz.me/U7), and notified in writing by the Secretary to the applicant;
- Admission becomes effective upon payment of the appropriate membership fee in each year of membership;
- Membership is renewable as long as the criteria for excellence in research and an active participation in the ERCIM community, cooperating for excellence, are met;

Interested in joining ERCIM? Please contact:

ERCIM Office
2004 route des Lucioles, BP 93
06902 Sophia Antipolis Cedex
Tel: +33 4 92 38 50 10
Email: contact@ercim.eu
http://www.ercim.eu

ERCIM
European Research Consortium for Informatics and Mathematics

ERCIM – the European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development, in information technology and applied mathematics. Its member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

ERCIM is the European Host of the World Wide Web Consortium.

Consiglio Nazionale delle Ricerche
Area della Ricerca CNR di Pisa
Via G. Moruzzi 1, 56124 Pisa, Italy
http://www.iit.cnr.it/

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering, N 7491 Trondheim, Norway
http://www.ntnu.no/

Czech Research Consortium
for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
http://www.utia.cas.cz/CRCIM/home.html

SBA Research gGmbH
Favoritenstraße 16, 1040 Wien
http://www.sba-research.org/

Centrum Wiskunde & Informatica
Science Park 123,
NL-1098 XG Amsterdam, The Netherlands
http://www.cwi.nl/

SICS Swedish ICT
Box 1263,
SE-164 29 Kista, Sweden
http://www.sics.se/

Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
http://www.fnr.lu/

Spanish Research Consortium for Informatics and Mathematics
D3301, Facultad de Informática, Universidad Politécnica de Madrid
28660 Boadilla del Monte, Madrid, Spain,
http://www.sparcim.es/

FWO
Egmontstraat 5
B-1000 Brussels, Belgium
http://www.fwo.be/

F.R.S.-FNRS
rue d'Egmont 5
B-1000 Brussels, Belgium
http://www.fnrs.be/

Magyar Tudományos Akadémia
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
http://www.sztaki.hu/

Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
http://www.ics.forth.gr/

University of Cyprus
P.O. Box 20537
1678 Nicosia, Cyprus
http://www.cs.ucy.ac.cy/

Fraunhofer ICT Group
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany
http://www.iuk.fraunhofer.de/

University of Southampton
University Road
Southampton SO17 1BJ, United Kingdom
http://www.southampton.ac.uk/

INESC
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, nº 378,
4200-465 Porto, Portugal

Universty of Warsaw
Faculty of Mathematics, Informatics and Mechanics
Banacha 2, 02-097 Warsaw, Poland
http://www.mimuw.edu.pl/

Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
http://www.inria.fr/

Universty of Wroclaw
Institute of Computer Science
Joliot-Curie 15, 50–383 Wroclaw, Poland
http://www.ii.uni.wroc.pl/

I.S.I. – Industrial Systems Institute
Patras Science Park building
Platani, Patras, Greece, GR-26504
http://www.isi.gr/

VTT Technical Research Centre of Finland Ltd
PO Box 1000
FIN-02044 VTT, Finland
http://www.vttresearch.com