insights

"USING OUR WEIGHT & BALANCE SOFTWARE, AIRLINES GAIN VALUABLE TIME FOR CARGO SORTATION BETWEEN INBOUND AND OUTBOUND FLIGHTS"

Karel Tavernier,
B. Rekencentra

GOING BEYOND RULES OF THUMB

How algorithm-based optimization is helping businesses make better decisions and become more efficient.

EXECUTIVE SUMMARY

"WHAT USED TO TAKE TWO MONTHS, NOW TAKES JUST THREE SECONDS. WE CAN SOLVE LARGER, MORE COMPLEX PROBLEMS FASTER THAN WE COULD SOLVE MUCH SIMPLER ONES 20 YEARS AGO."

Michael Trick,
Carnegie Mellon University

GOING BEYOND RULES OF THUMB

Every business process has the potential to become more efficient. And yet, without reliable automated support, decision-makers continue to follow 'rules of thumb': broad principles based on past experience that, while helpful, cannot produce truly optimal results. In the past few years, however, massive increases in computing power and the development of smarter algorithms that are capable of thinking faster have turned true operational optimization into reality. With the potential to cut operating costs such as production time and energy consumption by as much as 30 percent, algorithm-based optimization is ready for prime time—and poised to change the way we work.

THE INTEGRATION CHALLENGE

Operational optimization offers enterprises of all sizes and from just any sector the opportunity to significantly improve the way business-related decisions are made—from developing the ideal playing schedules for professional sports leagues to getting ships through the locks at the Port of Antwerp faster and more efficiently.

However, convincing companies to integrate optimization algorithms into their decision support systems and production process means bridging the gap between research-based best practices and real-world implementation.

COLLECTING AND CLEANING THE DATA

For many companies, both generating and capturing real-time data with the exactitude required by optimization algorithms won't be easy. While in some cases, the data are unavailable, in others, the format or quality of the data is insufficient. Algorithms will need the implicit knowledge that human planners use to 'read between the lines' and interpret incomplete or inconsistent data.

TRANSLATING THE RESULTS

Companies will need patience and make the necessary resources available to help their personnel fully understand the rationale behind the solutions.

Since the results of an optimization procedure are not always obvious, an element of translation is often required: if the human planners responsible for their implementation fail to understand the results, they may seriously doubt the algorithms' quality and fall back on their old rules of thumb.

Consider, for example, the task of loading cargo into a plane. If the amount of cargo exceeds the available cargo space, leaving the least valuable containers behind might seem ideal for maximizing the value of each trip. However, in reality, the optimal configuration presented by load and balancing software could yield a different selection to account for the cargo's impact on fuel consumption, loading and unloading time, etc.

THE WAY FORWARD

In the past decade, iMinds' researchers have worked hard to make optimization algorithms faster, smarter and more robust. Moving forward will necessarily imply additional partnerships with industry, academic and public sector partners to overcome the barriers to integration and to ensure that these algorithms can be adopted by organizations of all sizes.

Therefore, iMinds is looking to develop modular, affordable optimization approaches that can be easily implemented even by smaller companies. On top, iMinds will also work with Flemish software developers to establish the requirements for easy-to-use middleware that will seamlessly interface with the optimization algorithms.

FOR MORE INFORMATION

Does your organization have an optimization problem it's looking to solve? Contact
Greet Vanden Berghe
(iMinds - ITEC - KU Leuven)
greet.vandenberghe@cs.kuleuven.be

GOING BEYOND RULES OF THUMB

> "OPERATIONAL OPTIMIZATION AIMS TO WRING OPTIMAL PERFORMANCE OUT OF ALL KINDS OF SYSTEMS."

Every process has the potential to become more efficient than it is. Yet even the most gifted human mind has only a limited capacity for figuring out how to achieve such efficiency—especially when variables multiply and scenarios become exponentially complex.

Operational optimization aims to transcend the bonds of human brains and wring optimal performance—or as close to optimal as possible—out of all kinds of systems, from air freight transportation to the scheduling of sports tournaments.

Referred to in some circles as "the science of better", operational optimization is hardly a new concept. It's only in the last few years, however, that algorithmic breakthroughs and huge leaps forward in computing power have been able to propel operational optimization far beyond the capacity of human planners.

For the past decade, iMinds' researchers have been working to identify and address the essential requirements for making optimization algorithms faster, smarter and more robust. Through their efforts, algorithm-based optimization can now yield gains of as much as 30 percent on key performance indicators such as completion time for production orders, energy consumption and operating costs.

These kinds of massive improvements are available to organizations of all sizes—but to be realized, ongoing research and industry collaboration must bridge the gap that currently lies between research-based best practices and real-world implementation.

WHY RULES OF THUMB ARE NO LONGER ENOUGH

Optimization problems are everywhere. Yet when most companies tackle these challenges, they give the task to human planners who have little in the way of automated decision-making support. To solve optimization problems manually, these planners often fall back on rules of thumb: broadly applicable principles based on past experience or results. These types of rules are helpful, but cannot produce truly optimal results, especially not when applied to today's highly complex scheduling, rostering, routing and packing queries.

Specifically, rules of thumb cannot account for the vast range of inter-related factors at play in modern production and service operations. As a result, human planners tend to focus on optimizing only one piece of the puzzle, or they arbitrarily divide the whole into smaller sub-problems to make them more manageable—actions that actually prohibit the identification of a truly optimal solution.

Despite their limitations, rules of thumb continue to dominate decision-making—in part because few organizations have any real understanding of algorithm-based operational optimization. Corporate leaders are unaware of the efficiencies that could be gained, and while frontline workers may have a gut feeling that their company's work assignments and schedules could be improved, they have no way to know if incremental enhancements are actually making things better.

A further obstacle to overcome is the fact that quite often what 'looks' optimal to the human mind—often for aesthetic or logical reasons—is not actually the optimal solution in practice.

IT MAY LOOK WRONG, BUT IT'S OPTIMAL

For a concrete example, consider the task of loading cargo onto a plane. It may seem self-evident that the optimal arrangement would be whatever fits the most containers into the aircraft's hold to maximize the value of each trip. Yet other variables have to be considered, too. The plane must be

>>

“TODAY’S BUSINESS CHALLENGES PUSH OPERATIONAL OPTIMIZATION WELL BEYOND THE LIMITS OF HUMAN CAPABILITY.”

GLOBAL LEADERS IN OPERATIONAL OPTIMIZATION

iMinds’ researchers are among the world’s best in operational research and optimization. They have won international competitions such as the 2014 VeRoLog Solver Challenge, the 2014 EU FET-Open Inductive Constraint Programming Challenge, the 2011 Cross-Domain Heuristic Search Competition (CHeSC) and the 2010 META Eternity II contest.

They’ve also lent their expertise to co-organize a number of major events, including the 2013 Multidisciplinary International Scheduling Conference (MISTA) in Ghent, Belgium.

properly or it won’t fly. It can’t be so overloaded that all profit made on its cargo is burned up in excess fuel consumption. And the time required to unload and load at stops along the way also needs to be factored into the overall cost/efficiency equation.

When all these factors are considered, the optimal solution may not ‘look’ like the right solution at all. The most cost-effective configuration, for example, might involve taking off with some empty space still inside the plane, even if that defies conventional thinking.

This tension between intuitive notions of ‘right’ versus algorithmically calculated optimizations is visually illustrated in the following puzzle. If one were asked to fit as many identical rectangular boxes into a rectangular area, a conventional solution might look like the first option in the figure found at page 7. Twenty-one boxes, all neatly and logically arranged. But an algorithm isn’t worried about how ‘pretty’ the solution looks—it just wants to find the optimal outcome for the parameters it’s been given. The result? A total of 23 boxes in the same space.

While it’s likely a person could reach the same solution given enough time, this is a relatively easy task with a limited number of variables and constraints. The challenges facing today’s businesses are exponentially more complex, pushing the task of operational optimization well beyond the limits of human capability. People simply cannot solve these problems manually in any reasonable amount of time.

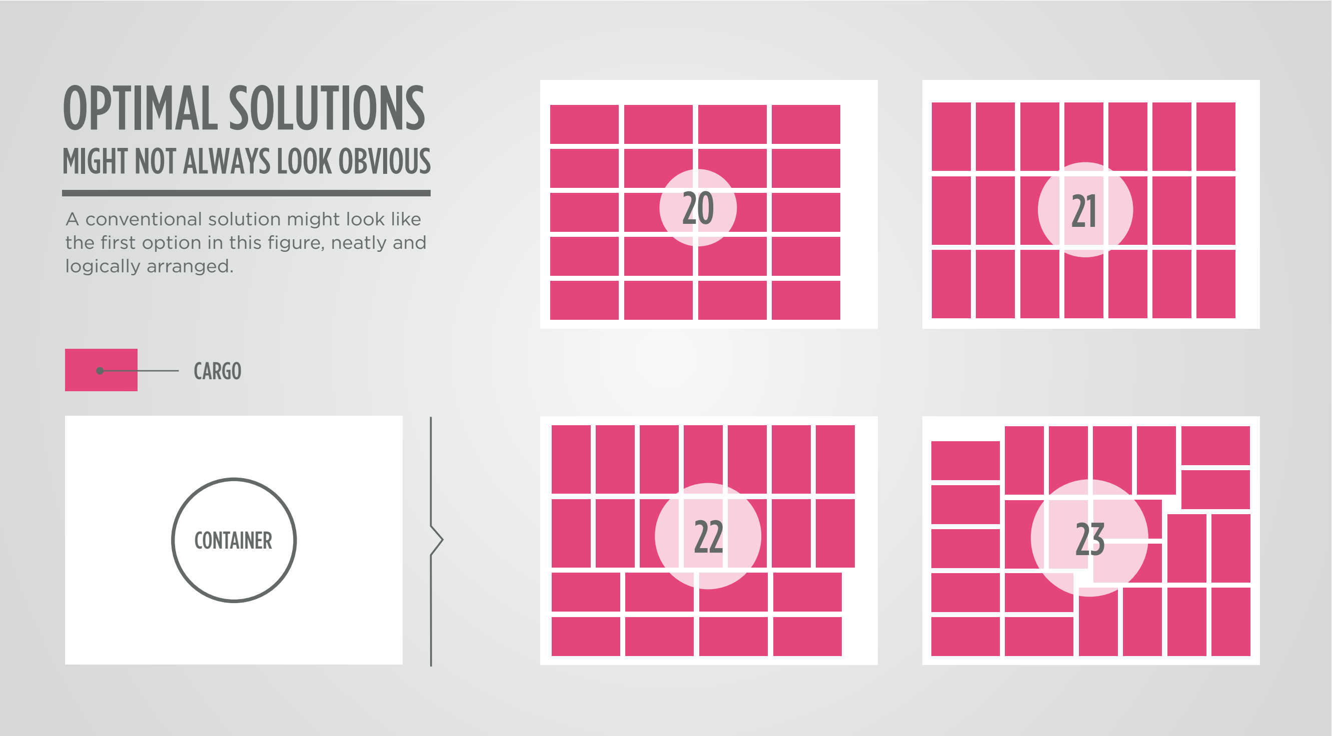

THE REAL-WORLD APPLICATIONS OF OPERATIONAL OPTIMIZATION

The potential applications of algorithm-based operational optimization are far-reaching, offering large and small enterprises from virtually any economic sector the opportunity to significantly improve the speed and efficiency of their business-related decisions.

The Port of Antwerp provides a good example. As the second largest port in Europe, it's essential that barges and container shuttles are able to get their goods in and out of port as quickly as possible. Because it is a tidal port, however, only so many ships can travel upriver through its series of locks at one time. Lockmasters have to decide which ships will be taken together and how they will be arranged within the locks—and which will have to wait.

With little time to make their decisions, lockmasters rely on rules of thumb based on ship size and cargo: vessels of 'X' size can fit in a lock with ships of 'Y' size; those carrying certain types of goods might be given higher priority; etc. But with optimization software developed by iMinds, lockmasters at the Port of Antwerp will soon be able to see the optimal lock configuration and schedule with just one click—greatly reducing the amount of time ships spend waiting to get into port.

Cutting leather for use in cars and clothes provides a smaller-scale example. It can take a person hours to figure out how to get the maximum number of pieces out of an irregular-shaped hide—and he or she will still probably end up with a sub-optimal solution. Algorithms being developed by iMinds, however, can determine a good-quality cutting pattern in a matter of seconds—saving not only time but also money by ensuring less leather is discarded as waste.

Regardless of business size or sector, perhaps the most significant benefit of optimization algorithms is that they can handle so many more variables than a human planner. Consider the difficult task of building a work schedule for nurses in a large hospital. Creating a good-quality schedule that >>

takes into account each nurse's skills, qualifications, availability and personal preferences in terms of when and where they'd like to work is practically impossible for a human. Advanced algorithms like those iMinds created to power Saga Consulting Group's Health Care Personnel Suite, however, have no problem making such computations.

By doing so, these algorithms can greatly alleviate the pressure placed on human planners. The people responsible for generating a hospital's monthly work schedule—in many cases, the head nurses—might spend up to 70 percent of their time on that task. These are often senior personnel with a great deal of organizational knowledge. By giving them support on computationally hard tasks, they can spend more time doing the jobs they're actually trained to do.

A similar scenario plays out in factories around the world, which often struggle to determine the best time to turn their machines on and off. It isn't just a matter of uptime: the ideal solution will also take into account factors such as energy consumption, order volume, total cost of ownership and more.

IMPROVING OPERATIONAL EFFICIENCY IN HOSPITALS

The nurses' work schedule isn't the only thing that needs to be optimized in a hospital. The flow of patients must also be managed so operational inefficiencies (such as bed shortages) are minimized and patients receive care in a timely fashion. iMinds is building decision-support algorithms that will improve admissions planning related to the inflow of elective patients, helping hospitals make better use of their key resources to avoid excessive workload in the operating theatre and ensure a balanced patient load across wards.

ALGORITHMS: THE HEART OF OPERATIONAL OPTIMIZATION

The efficiencies made possible by operational optimization are all driven by advanced algorithms that turn a company's data into actionable solutions. While these algorithms have existed for a number of years, the organizations that committed significant amounts of money to optimize their operations 10 to 15 years ago were likely disappointed with the results: the computational power, modeling and algorithm development skills were simply not at the level required to enable fast and optimal decision-making support.

A lot has changed since then.

Better off-the-shelf hardware and sheer computing 'horsepower' have certainly helped—but it's the development of algorithms that can 'think' in faster and smarter ways that is making true operational optimization a reality. A production planning model that would have taken 82 years to solve using the algorithms available in 1988 can now be solved in one minute—an improvement by a factor of 43 million.

> "A PRODUCTION PLANNING MODEL THAT WOULD HAVE TAKEN 82 YEARS TO SOLVE USING THE ALGORITHMS AVAILABLE IN 1988 CAN NOW BE SOLVED IN ONE MINUTE."

The algorithms used a decade ago were often based on linear mathematical models that clearly described what was possible—and what was not. Although they were capable of solving specific types of problems, they were unable to scale, converge or adapt if the problem changed—greatly limiting their real-world applications. Heuristics (simple rules that help accelerate and automate the decision-making process) were also common. While they proved to be flexible and easy to implement, they often traded precision and accuracy in favor of speed.

Recent advances, however, have seen the potential of mathematical optimization approaches increase exponentially—and iMinds' researchers are on the front lines. Leveraging their skills and experience in computer science and engineering, they are combining the best of both mathematical and heuristic models—a method known as 'matheuristics'—to build incredibly powerful and flexible algorithms.

Matheuristics use advanced mathematical approaches to develop optimal solutions for the most important parts of a problem, and then apply heuristics to quickly generate starting solutions or mathematical bounds on the quality of the solutions for the rest of the problem. The future of algorithm development will rely on this type of hybridization, borrowing from a variety of domains (machine learning, mathematical optimization, etc.) to develop robust solutions for industrial problems.

For example, optimization algorithms are now being hybridized with the latest research on artificial intelligence. Today's algorithms are adaptive to different problems—with some capable of making adaptations on the fly in the middle of a process—and can overcome issues resulting from missing or faulty information.

Other key developments have been seen in the area of decomposition—that is, the ability to solve a problem by merging the solutions of smaller, related problems. A courier business, for example, needs to design its routes so trucks can be loaded as fully as possible and reach the greatest number of customers in the least amount of time. Yet the load can't be determined without knowing the route—and

>>

DATA-FLOW INTERCEPTION

While there is no single ideal architecture for collecting the real-time data needed by optimization algorithms, there are a number of common practices and patterns. Data-flow interception, for example, allows companies to introduce 'probes' that transparently intercept and manipulate data as it is sent from one computer to another. The data can then be encrypted and sent to the software or application that features the optimization algorithms.

HEURISTICS THEN AND NOW

Heuristics can be used to quickly produce a solution that is 'good enough' for a given problem. While early heuristics tended to sacrifice optimality for speed, recent advances are helping to make heuristics more accurate and reliable—meaning they can deliver better solutions than traditional methods without giving in on speed.

- *Metaheuristics: Providing an additional layer of decision-making support, metaheuristics observe the overall process (rather than the problem) at a high level to decide when and how to apply a lower-level procedure or heuristic.*
- *Learning heuristics: Also known as 'genetic algorithms', these heuristics take the best parts of one solution and apply them to similar problems. They're capable of 'learning' something about the structure of a problem so that general algorithms can be applied to very specific situations.*

the route can't be plotted without knowing if all the packages to be picked up will fit inside the truck. Decomposing the problem to solve the two pieces independently, and then 'recomposing' them to identify the tradeoffs between them, will help yield better results.

SOLVING THE INTEGRATION CHALLENGE

If the latest technology and research are creating algorithms that are faster and smarter than ever before, what is preventing companies from implementing them into their everyday decision support systems?

It's a matter of integration—or more precisely overcoming the difficulties associated with integration. Blending optimization algorithms and software with legacy production processes can be complex, costly and require many businesses to reconsider how they approach data collection, privacy and security. Companies also need to commit time and resources to ensure optimization tool users understand the rationale behind the more efficient (yet sometimes counterintuitive-seeming) solutions being generated.

It's not surprising, then, that it is mostly large manufacturers and institutional settings that are considering adopting optimization software today. The entry barrier seems too high for most other types of businesses to consider, despite the long-term gains that will eventually offset the initial cost.

After all, many small and medium enterprises (SMEs) could realize significant improvements in overall efficiency and productivity as their schedules are less complex. That's why iMinds is looking at ways to develop modular, affordable optimization solutions that can be easily implemented by smaller businesses. Working in close collaboration with Flemish software developers, iMinds is establishing the requirements for easy-to-use middleware that can seamlessly interface with the optimization algorithms—and fill in the missing piece of the integration ecosystem.

COLLECTING AND CLEANING THE DATA

Optimization algorithms require high-quality, real-time data. For many companies, generating and capturing the right data at the right time—with the exact level of granularity required by the algorithms—may not currently be easy. In some cases, the data simply isn't available: nobody in the company thought to capture it before. Here, serious discussions will be required to determine what data needs collecting and how it will be gathered. This could require the installation of new ERP or SAP systems, or the retrofitting of older machines to ensure the right probes and sensors are in place to capture the relevant data needed for optimization.

In other cases, the data is available but in a format or quality not sufficient for use with the algorithms. While a human planner might be able to 'read between the lines' and interpret it, taking into account all kinds of explicit and implicit knowledge, this so-called 'dirty' data is too incomplete or inconsistent for a computer to process.

Theoretically, algorithms should also be able to read between the lines and 'clean' the data as it enters the system so that only the most relevant and revealing information is taken into consideration—but only if companies can delineate and define for them all the factors that influence their decision-making processes.

DRAWING ON INTERNATIONAL EXPERTISE

The United States is particularly advanced when it comes to implementing operational optimization approaches. To bring these best practices home to Belgium, iMinds has partnered with leading U.S. researcher Michael Trick, Professor of Operations Research at Carnegie Mellon University. Thanks to today's smarter, more powerful algorithms, Trick can generate the playing schedule for Major League Baseball 50 million times faster than he could 20 years ago.

>>

TRANSLATING THE RESULTS

Perhaps the most significant barrier to integrating optimization algorithms has nothing to do with technology at all. Rather, it is a matter of convincing people that the results are indeed better than whatever could be achieved before. Because the results produced by an optimization algorithm don't always 'look right', an element of translation is often required: human planners responsible for implementing the solutions need to understand their quality. Otherwise, they may second guess the algorithms' output and fall back on their old rules of thumb.

It must also be made clear that the algorithm can never be held responsible for the final outcomes—that must fall to the person using the optimization software. For this reason, the solutions produced by the algorithms must be extremely intuitive and, above all, acceptable to the end user so they feel confident and comfortable in assuming responsibility for the end result.

THE WAY FORWARD

Evolving from rigid linear programming to today's faster and smarter hybrid algorithms, operational optimization has finally progressed to the point where it is delivering real, tangible benefits to businesses across Flanders and around the world. Moving forward, additional partnerships with industry, academic and public sector partners will need to be formed to overcome the barriers to integration and ensure optimization algorithms can be adopted by organizations of all sizes—especially SMEs, which make up an integral part of local economies.

In this regard, iMinds will continue to raise awareness of the benefits of operational optimization. In some cases, this will mean demonstrating the difference between today's optimization algorithms and those from more than a decade ago. Other companies are simply reluctant to adopt new processes and will need to be reassured that optimization algorithms can be introduced without having a negative impact on their existing systems.

Efficiency gains as high as 30 percent are achievable through algorithm-based optimization. It's about time we started to look past our thumbs—and trust the latest research and technology to take some of the stress and complexity out of the way we plan our day-to-day business operations.

FUTURE FOCUS

What kinds of challenges do iMinds researchers want to take on next? Some areas where operational optimization could be applied include space allocation (how to make more efficient use of space within office buildings and other settings); chip routing and packing (where to place components on computer chips to minimize surface distance and heat); diamond cutting (to get the biggest, most valuable gems out of rough diamonds); and the optimization of elevator operations in large office towers.

“WITH EFFICIENCY GAINS AS HIGH AS 30 PERCENT, IT’S ABOUT TIME WE STARTED TO LOOK PAST OUR THUMBS.”

OPERATIONAL OPTIMIZATION

CASE STUDY

OPTIMIZATION TAKES FLIGHT WITH WEIGHT AND BALANCE SOFTWARE FOR AIRLINES

The weight and location of every container has a significant impact on a cargo plane's center of gravity: if not properly balanced, the plane won't fly. Yet for all the complexity associated with loading aircraft, many planners still work with pencil and paper. By doing so, they're missing a massive opportunity to optimize payloads and reduce fuel consumption.

Software and solution provider B. Rekencentra is changing that with its System Automated Balance and Load Engineering (SABLE) software. Powered by an algorithm developed by the research group iMinds - ITEC - KU Leuven, SABLE can decrease the time it takes to balance a fully loaded 747 from three hours to five minutes. B. Rekencentra's Technical Director ***KAREL TAVERNIER*** *explains how SABLE works—and why it's been adopted by 35 airlines around the world, including every airline flying for global shipping company DHL.*

Q: What are the disadvantages of manually calculating an aircraft's weight and balance?

Karel Tavernier: To begin with, it's a very time-consuming process. This is partly because there is a lot of trial-and-error involved—if at the end of your calculations you see you are outside the acceptable limits, you have to throw away your sheet of paper and start all over. And with every airline having its own load and trim sheet for each type of aircraft it flies, load planning also requires a high level of skill and extensive, recurring training.

Finally, because the task is so complex and the amount of time available is short, most planners are happy once they get the center of gravity right—they can't afford to think about optimizing any other parameters. So most aircraft are never really used to their full potential.

Q: How does your software make things easier for load planners?

Karel Tavernier: With manual load sheets, planners have to go back and forth between many different documents to transcribe information about the cargo and calculate the plane's weight. SABLE speeds up this process by interfacing automatically with about 10 different systems to get the information it needs: the check-in

>>

> "SINCE ITS INTRODUCTION, THE SABLE SYSTEM HAS PROVEN TO BE STABLE, RELIABLE, USER-FRIENDLY AND FLEXIBLE, ENABLING DHL TO MAXIMIZE THE USE OF SPACE AVAILABLE ON ITS AIRCRAFT RIGHT UP TO THE LAST MINUTE."

Neale Millett
Head of Global Airside
and Standards, DHL Aviation

system to get passenger figures, the warehouse or baggage system to get cargo figures, the flight dispatch system and even the aircraft itself to get fuel data. Because we use the exact limits described in the manufacturer's weight and balance manual, our calculations can go right to the upper limits knowing that safety will not be compromised.

Q: Does it look the same as the load sheets planners are used to?

Karel Tavernier: SABLE actually presents a graphical view of the aircraft's decks, so what you see is what you get. You just click and drag containers to the desired position and all the actual checks are done in background: maximum position weight, linear loading limits, cumulative loading limits per deck and even adjacency rules for dangerous goods, which are especially difficult to control manually. The end result is the printed load sheet: the official, legal document the captain has to sign before the plane is allowed to fly. If any violation is detected, SABLE will not allow that load sheet to be generated.

Q: What are the big-picture benefits for airlines?

Karel Tavernier: Because less time is spent doing the planning, airlines gain valuable time for cargo sortation between inbound and outbound flights. Beyond that, SABLE enables optimization on three levels. First, the software can determine the maximum payload that can be carried by each plane, which increases revenue for the airline. By minimizing the lateral imbalance of the aircraft, pilots experience better flying conditions. And by moving the center of gravity as far aft as possible, fuel consumption can be significantly reduced.

In fact, we found that by moving the center of gravity just 0.1 percent more aft, a passenger airline that burns 1.5 billion kilograms of fuel per year can reduce its annual fuel consumption by one million kilograms—saving approximately a million dollars. The more aft you go, the more you save. And that's actually a very conservative estimate as our study involved only passenger aircraft, where cabin layouts limit what you can do to

HOW IMINDS APPROACHED THE RESEARCH CHALLENGE

"Weight and balance is subject to non-linear safety constraints that constantly change during a flight—and therefore restrict the applicability of optimal mathematical solvers. Instead of including the non-linear constraints into our model, or fitting a piecewise linear function to the non-linear constraints, we assumed that the constraints should be satisfied at specific critical flight conditions (e.g., zero fuel weight, take-off weight, landing weight). After computation, we check the assumption's validity by confirming the feasibility of the aircraft's envelope on the computed solution."

Greet Vanden Berghe
iMinds - ITEC - KU Leuven

move the center of gravity. With cargo planes, you have both the main and lower decks to play with, allowing for greater optimization.

Q: What role did the iMinds research group ITEC play in SABLE's development? And how has the software changed since it was first released in 2002?

Karel Tavernier: When I first met with the iMinds research team at KU Leuven, we had already developed the weight and balance calculator. But its optimizing function, which relied heavily on existing rules of thumb, did not perform well—in fact, it would often propose solutions that could not pass final checks because limits were being violated. So after explaining the problem and what I wanted to achieve, they designed and developed the mathematical algorithm for a new optimizer module. The biggest challenge for iMinds was the sheer scope of the project: collecting and analyzing the data for each aircraft type, then building those into the algorithm so the center of gravity could be optimized according to various targets and flight conditions.

After the initial development there was a handover of the source code. Since then, we've been making small modifications on our own, adding new interfaces and functional improvements—many requested by our users—to make SABLE faster, more user-friendly and easier to maintain.

Q: Where does SABLE go from here?

Karel Tavernier: With a focus on automating more of the operational procedures, our next steps will come in two areas. The first is commercial trim, where we want to help airlines position cargo more effectively so they don't have to offload and reload freight when landing at every new destination. The second pertains to the cargo offer, where we want to give warehouses the necessary data to build custom-sized pallets based on an aircraft's remaining freight capacity and available container positions. While this is still a very new idea, it will allow the plane's payload and center of gravity to be optimized even further.

We are currently working with iMinds on these and other projects. Because of their excellent work, we currently have the best weight and balance optimizer in the world—and it is continuing to evolve in very cool ways.

ABOUT B.REKENCENTRA

Founded in 1983 and headquartered in Ranst, Belgium, B. Rekencentra NV is a software provider specializing in aircraft weight and balance, container terminal and aircraft communication solutions. The company also offers consulting, system integration, administration and training services and support.

OPERATIONAL OPTIMIZATION

CASE STUDY

UNLOCKING EFFICIENCY IN THE PORT OF ANTWERP

Europe has nearly 41,000 kilometers of waterways—at the heart of which the Port of Antwerp provides a vital link between the world's biggest maritime routes and Europe's major centers of production. More than 200 container shuttles arrive and depart every week from here, the continent's second busiest port, making the smooth handling of traffic vitally important. Yet lockmasters still rely on manual methods to determine which ships can be sent together through each of the port's six locks.

The nature of Antwerp's lock operations is about to dramatically change, however, thanks to sophisticated optimization algorithms that were developed by iMinds - ITEC - KU Leuven and will be brought to market through a start-up company led by iMinds researcher ***JANNES VERSTICHEL****. By simplifying and accelerating the decision-making process, the Port will be able to get ships through its locks and into harbor faster—and in doing so, attract more business from around the world.*

Q: What's the process today for planning lockages at the Port of Antwerp?

Jannes Verstichel: Because barges can call in as little as 10 minutes before arriving at port, lockmasters stick by the 'first come, first serve' rule—there's just not enough time for anything else. But even that system doesn't always work. There's increasing pressure to plan lock operations around financial considerations: giving priority to ships carrying perishable goods, for example. At the same time, ocean-going ships always have priority over smaller barges because, in most cases, they can pass through the locks only during high tide. So if a large vessel arrives at a certain time, the barges have to be planned around that—with some going in with the bigger ship and others told they'll have to wait for a different lockage.

Q: How will your software make things easier for lockmasters?

Jannes Verstichel: Right now they're using a drag-and-drop computer interface to arrange smaller rectangles (the ships) within a bigger rectangle (the lock). It's like a puzzle with constantly changing pieces—and it can take up to 10 minutes to find a solution that's acceptable. We're adding an 'optimization' button to the interface they already use: with just one click the puzzle is solved, arriving at the same optimal configuration the lockmasters would reach if they had the time to do so. The software currently used by the Port of Antwerp is like a really nice car that's missing its engine—you need to push it manually if you want to go anywhere. We're building that engine.

Q: How will the Port of Antwerp benefit from your software?

Jannes Verstichel: We're giving lockmasters instant decision-making support. As soon as a barge calls in,

>>

ABOUT THE PORT OF ANTWERP

Europe's second busiest port—and one of the top 15 ports in the world—the Port of Antwerp in Belgium handles 186 million tonnes of international maritime freight each year. It's also the world's largest port in area, spanning more than 13,000 hectares—about 20,000 football fields—across which are 30 docks, 21 bridges and six locks of various sizes. A seventh lock is currently under construction; once completed it will be the largest lock in the world.

they can evaluate when it can be scheduled into the lock, which will help decrease barge wait times. Of course, the Port does not receive any direct financial gain from this type of optimization—the barges still pay a fee whether they wait five minutes or five hours. But if the locks are operated more efficiently and ships know they can get their goods in and out of port faster, that becomes a major selling point to attract new traffic. By working on lock operations at this high level of practicality, our software will provide the Port of Antwerp (and Flanders) with a very strong competitive differentiator in the years to come.

Q: And what about the barge operators? How will they be affected?

Jannes Verstichel: The moment a barge calls in to announce its arrival, the lockmaster can give an approximate time the vessel should arrive at the lock. If a barge operator knows it will be about three hours before he can get into the lock, he can lower his speed to reduce fuel consumption. Alternatively, if the lockmaster says he can get in sooner if he arrives in an hour—but will have to wait three hours if he doesn't—the barge operator can speed up to catch that earlier lockage.

> "OUR SOFTWARE WILL PROVIDE THE PORT OF ANTWERP WITH A VERY STRONG COMPETITIVE DIFFERENTIATOR IN THE YEARS TO COME."

Q: Do you expect any challenges when it comes time to fully implement the software?

Jannes Verstichel: No, I don't think so. As our decision-support system generates very intuitive solutions, there should be little to no learning curve at all for the lockmasters. Of course, from a technical point of view, there will be some important factors to address: in particular, the fact that the algorithm is only as good as the data it receives. Making sure the software can get the right type of data for every ship under every circumstance—the kind of information lockmasters already know instinctively—is something we're still working on, but we're definitely going in the right direction.

Q: What are the next steps once your software moves beyond the prototype phase?

Jannes Verstichel: Our finished commercial software will be up and running by 2016. At that point, we'll start looking at a much more technically challenging problem: the synchronization of lock operations between the left and right banks of the Port to reduce wait times for ships transferring from one side to the other. We also plan on incorporating tugboat operations into our algorithms. Large ocean-going vessels must be towed by tugboats through the locks—but there are only a limited number of tugboats to go around. To deliver a truly optimal solution, our software should take into account how many tugboats are available so ships aren't stuck waiting for tugboats to arrive.

Our aim is take the 'green wave' approach used in automobile traffic planning—where you try to maximize the number of green lights along a person's journey—and apply it to inland waterways. Barge operators should have to wait only when they arrive at that first lock; once they're transferred, they shouldn't have to wait again until they reach their destination.

Q: Do you have any plans to expand beyond Flanders?

Jannes Verstichel: Although we've been working closely with the Port of Antwerp to develop our decision-support system, the principles behind it—the optimization routines—are actually quite generic. This means it could be applied easily to any lock system of any size anywhere in the world, not just in ports. In fact, we're already exploring opportunities in the Netherlands to develop prototypes and implement versions of our software there. And also the Panama Canal is facing similar challenges related to the complexity and diversity of its ship traffic.

And we will not limit ourselves to just lock operations. We want to look at the entire supply chain: not only how the ship arrives at the port but also how it unloads its cargo and then how that cargo is sent to the customer's warehouse. Ultimately, we intend to expand to truck transportation and terminal operations, optimizing the way goods are transferred between ships, trucks and container yards.

ABOUT JANNES VERSTICHEL

Jannes Verstichel will lead the start-up company that will implement the optimization software at the Port of Antwerp and sell it to other interested parties around the world. In 2013, he earned a PhD in Engineering (Computer Science) at iMinds - ITEC - KU Leuven for his work on the lock scheduling problem. He is currently a post-doctoral researcher at iMinds - ITEC, where his research is funded by a grant from the Agency for Innovation by Science and Technology (IWT Vlaanderen).

SOLVING BIG PROBLEMS FASTER

With three decades of experience in operational research, ***MICHAEL TRICK*** *from Carnegie Mellon University is world-renowned for his insights and expertise in scheduling and routing optimization. His company, Sports Scheduling Group, has developed the playing schedule for Major League Baseball for the past 10 years. Recognizing the unique perspective he brings to the field—and the multiple breakthroughs his team at Carnegie Mellon University has made over the years—we asked Michael to share his thoughts on the past, present and future of operational research.*

Q: What are the major trends you're seeing in operational research?

Michael Trick: Operational research used to be limited to scheduling machines in a factory or working through a production system. Today it's used to allocate marketing budgets, optimize financial portfolios and guide supply chain decisions, such as where to hold inventory and how much to hold. As a result, we're now seeing more and more companies relying on optimization to help them make increasingly sophisticated decisions.

We've also gotten better at solving problems where certain data are unknown, allowing for probabilities that could come into play. When I'm building the schedule for Major League Baseball, for example, there's no way to know which games will be rained out. However, I can create schedules that are more adaptable to the possibility of rainouts: for instance, by giving teams common off-days that make it easier to reschedule postponed games. It's about building an underlying model that can accommodate the full range of what's possible.

Q: How has the field changed over the years?

Michael Trick: I'll give an example to make the point. In 1995, I created

THE IMINDS CONNECTION

Like any other discipline, further advancements in operational research are only possible through collaboration. To better incorporate the great work being done overseas into its own research, iMinds - ITEC - KU Leuven sent a PhD student to learn from Michael and his team at Carnegie Mellon University. The student's one-year stay in the United States wrapped up in August 2014.

iMinds also shares Michael's passion for setting academic benchmarks for sports scheduling problems. In fact, iMinds - ITEC - KU Leuven currently holds the best results for the 'travelling umpire' problem, which asks researchers to assign umpires to a tournament of games while minimizing travel and ensuring each umpire works at each venue at least once (but avoids working the games of the same team too often in short succession).

>>

> "OPERATIONAL RESEARCH IS NO LONGER ABSTRACT CALCULUS: WE'RE SOLVING REAL DAY-TO-DAY BUSINESS CHALLENGES."

a model to answer one question: what is the best possible Major League Baseball schedule for the month of April? It took two months of computation to solve. Two major things have happened since: today's computers are about 10,000 times faster than the one I was using back then, and the underlying optimization software is about 5,000 times faster. Multiplying those two numbers, I'm now 50 million times faster at solving that problem. What used to take two months now takes just three seconds. This presents a fundamental shift in the capabilities of optimization. We can solve larger, more complex problems faster than we could solve much simpler ones 20 years ago.

Q: What makes a problem difficult—or less difficult—to solve in an optimal way?

Michael Trick: That gets at the heart of what operational research is all about. There are computationally hard problems and computationally easy problems—but you can't just look at a problem and know which class it falls into.

Many scheduling and routing problems are inherently difficult to solve, such as the famous 'traveling salesman' problem. This is where you need to figure out how to visit 100 cities so that you minimize both the time spent and distance traveled. After the first city is chosen you have 99 possible destinations to choose from, then 98, and so on. With so many possibilities, the problem becomes very challenging. What operational research tries to do is separate what's easy from what's hard, and then find optimal or near-optimal solutions to those problems in a reasonable amount of time.

Today we have theorems and techniques that allow us to attack the easy problems and guarantee optimal solutions, and others that attack the hard problems to get near-optimal solutions. With Major League Baseball, there are 2,430 games in a season—making it extremely difficult to build a truly optimal schedule. For these larger problems, it's all about finding good, usable, practical solutions. That's the essence of operational research.

Q: What drew you to the field?

Michael Trick: In operational research you have a very concrete opportunity to apply mathematics in a way that helps people make better decisions at a very practical level. When you see the types of problems operational research can solve, mathematics suddenly becomes a heck of a lot more interesting. It's no longer abstract calculus: you're solving very real, day-to-day business challenges.

Q: Business challenges of all kinds, it would seem—including sporting events. How might optimization research apply to something like the Olympics or the World Cup?

Michael Trick: With the Olympics, you need to schedule multiple events over dozens of venues. Each individual venue is relatively easy to schedule. Things get more interesting when you start considering the bigger picture: how do you schedule all of these events so they can be televised in an effective way? You don't want to show the finals of a major swimming event at the same time as a high-profile track event. Not only do you need to schedule the events so that everything gets completed, but you also need to make sure there is a narrative line running through them that produces a wonderful television experience. How do you optimize for that?

The World Cup is actually rather easy to schedule. There aren't many teams and not a lot of games: in fact, you could probably do it by hand. The Belgian Pro League would be much more challenging. Every weekend there are eight games—and figuring out how to schedule the 16 clubs so they don't, for example, play three home games in a row, is a hard thing to do. Once you get up to that size, it's difficult for humans to do it at any quality.

But to be honest, the hardest scheduling I've ever done is for my kids' weekend soccer league. The number of players, the parental considerations… it's far more difficult than any of the college sports I've worked on!

Q: What's next for operational research?

Michael Trick: There's a lot of data out there, and organizations need to know what to do with it. A lot of companies right now are investing in data collection, but gathering it is useless on its own. Data has value only when it drives decisions. The decisions are what count. This is why I've been looking at business analytics, which combines data mining and operational research.

Data mining yields simple decisions: if you send a certain type of customer a catalog, how likely is it that he or she will order

>>

something? By adding operational research to that data, you can create more complicated decision structures. For example, once you have an idea as to which customers might buy your products, you can try to optimize the allocation of your marketing budget relative to the models you've established. It's about adding an extra layer to the data, and it can give companies a competitive advantage.

Q: What are the biggest challenges currently facing operational research—and what will be the key to solving them?

Michael Trick: We still have a lot of work to do to truly integrate data mining with operational research. We're just barely beginning to understand how these complicated models of human behavior integrate with allocation, routing and other techniques. And as a field, we need to do more to convince companies that if all they are doing is collecting data, they're wasting their time and money. I see company after company talking about the data they collect, but they have no idea how to translate that information into better decisions.

Part of the problem is that very few people understand both data mining and operational research. Creating easy-to-use software that includes both sets of tools will definitely help in that regard, in helping businesses really unlock the value of the data they're collecting. The way the two fields are taught in our schools also needs to be changed; right now, people who learn one technique are often never exposed to the other.

Finally, we'll definitely need more research like the kind being done by the team at iMinds. They're working on a lot of interesting things that will help companies of all sizes solve difficult problems faster and more efficiently—and I can't wait to see what happens next.

ABOUT MICHAEL TRICK

Michael Trick is a Professor of Operations Research and the Senior Associate Dean of Education at the Tepper School of Business at Carnegie Mellon University. A Fellow of the Institute for Operations Research and the Management Sciences (INFORMS) and founding editor of INFORMS Online, he has solved supply chain and scheduling issues for organizations like Major League Baseball, the United States Postal Service, Motorola and Sony. You can read more of his thoughts on operational research at: http://mat.tepper.cmu.edu/blog.

“COMPANIES HAVE NO IDEA HOW TO TRANSLATE THE DATA THEY COLLECT INTO BETTER DECISIONS.”

GREET VANDEN BERGHE

Head of CODeS, Ghent, part of iMinds - ITEC - KU Leuven.
Head of the Computer Science Technology Cluster.

FOR MORE INFORMATION

about iMinds' expertise in the field of Operational Optimization, please contact **Greet Vanden Berghe**
greet.vandenberghe@cs.kuleuven.be, +32 9 265 87 03

iMinds editorial team: Sven De Cleyn, Koen De Vos, Thomas Kallstenius, Els Van Bruystegem, Wim Van Daele, Stefan Vermeulen

Copy: Ascribe Communications
Design: Coming-Soon.be
Photography: Lieven Dirckx, Nils Blanckaert

Additional content will be published on www.iminds.be/insights.

LIST OF SUBJECT-MATTER EXPERTS WHO CONTRIBUTED TO THIS PAPER (in alphabetical order):

TOM DHAENE
(iMinds - IBCN - Ghent University)
Professor

SAM MICHIELS
(iMinds - DistriNet - KU Leuven)
Industrial Research Manager

ANN NOWÉ
(VUB Artificial Intelligence Lab)
Professor

KAREL TAVERNIER
(B. Rekencentra)
Technical Director

MICHAEL TRICK
(Tepper School of Business, Carnegie Mellon University)
Senior Associate Dean, Faculty and Research and Professor in Operations Research

GREET VANDEN BERGHE
(iMinds - ITEC - KU Leuven)
Professor

JANNES VERSTICHEL
(iMinds - ITEC - KU Leuven)
Post-Doctoral Researcher

TONY WAUTERS
(iMinds - ITEC - KU Leuven)
Post-Doctoral Researcher