IFS Planning and Scheduling Optimisation

Executive Summary

At the heart of Total Cost of Ownership

Planning & scheduling optimization at scale

Optimization can deliver outstanding business improvement, but: "The devil is in the detail"

One of the biggest challenges in planning and scheduling optimization is the fundamental need for a deeper understanding of the mathematical approach taken – mathematical optimization strategies run the risk (as a subject) of being pushed to specialists and given limited focus at the business executive level.

This summary is to encourage you as a reader to reject this push to experts and understand the deeper fundamentals and how it will impact your business – this subject truly matters, and the mathematics used to solve business problems in this area are not all the same. Fundamentally, we strongly suggest you "pull back the curtain" on this subject and the associated marketing messages and understand the impact of different optimization techniques. We have field evidence of vendors in this space "cheating the maths" to get better results and then failing to deliver on promises. Team IFS want to expose you to the reasons this problem is challenging, and then what can be done to get real world improvements.

100-person years of applied PhD science

Optimization is a several decade old field of research, and one that is seeing a renaissance with the buzz word filled advent of "Artificial intelligence" (AI). We've written this short document, specifically for IFS customers looking at the subject, to avoid the risk of generic answers from software vendors. Such answers do not address the heart of the questions with optimisation - no suppliers should get away with simplistic answering in this important area due to the huge impact is can have on business performance.

However, to avoid a high-level answer, understanding the mathematical requirements for optimal problem solving - producing the best solution to a combination of weighted needs of people, time, contracts, skills, inventory, flexibility, profit margin, etc - at scale in the specific environment your business operates in, is at the very heart of the challenge.

To appreciate the depth of this challenge and what it really takes to solve the relevant optimisation problems, we feel that all levels in the decision-making process deserve the opportunity to challenge IFS, and any other suppliers, to avoid "smoke and mirrors".

To facilitate this, please allow yourself the opportunity to dive a little deeper with the material below. We also invite you to meet with the IFS research team, that have been working on this specific problem for the last two decades, to explain our approach in detail.

As an example, most engineers are familiar with the "travelling salesperson problem" for working out optimal routes for a field-based engineer. What is also well known among Operational Research scientists is that solving such multi-hop problems, where skill sets, optimization of resource, distance travelled, job criticality, SLA penalties, (among potentially many other data sets) is a multi-factorial mathematical problem – and seriously hard.

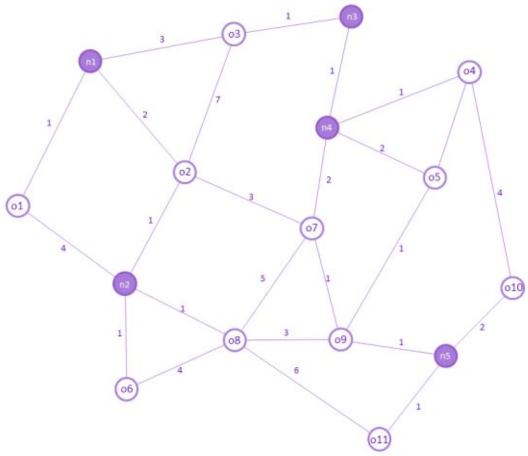

Large versions of these problems are practically impossible for humans and even with powerful computational hardware, they usually require many hours to solve, or the solutions are so tightly constrained that they do not offer the ROI that companies will need to be world class. IFS has invested many tens of millions, working with a dedicated team of research scientists for many years to craft and combine world-class optimization research, mathematical best practices and specialist solving engines needed to deliver against these challenges <u>when at real-world scale</u> (not in the lab)

Field Service planning and optimization needs more than one approach to improve the Total Cost of Ownership and deliver the best business returns for your company.

IFS Planning & Scheduling Optimization (PSO) has its roots in complex numerical analysis. It was developed in the 1990s as a sophisticated mapping and scheduling application for use by the emergency services in the U.K. The system helped save lives and made efficient use of their resources while meeting legal SLAs measured in minutes. This is important to know, because at the heart of all optimization approaches is a choice, either discrete (batch) or continuous solving. Please understand that we do not say “real-time” here, deliberately. Both these mathematical methods can have hardware thrown at them and be “fast” or near real-time (within limits) – but the physical approach to the mathematics is different with continuous vs discrete solving.

DYNAMIC CONTINUOUS SCHEDULING

One of the major components in IFS PSO is the Dynamic Scheduling Engine (DSE). The DSE is a general-purpose, continuous optimization engine, which deals with goals and constraints, and seeks to find the best solution to a problem. It **runs at all times** on the mathematics problem it is presented with. It is not a discrete, batch or “start again” type engine.

Optimize the enterprise – one size does not fit all problem spaces

Our team spent a decade creating this engine, behind the scenes, that automatically classifies the mathematical problem to solve and matches with to one or more of 35 different optimization algorithms: some are well known in the industry (for example, simulated annealing, genetic and pheromone-based algorithms) but also proprietary algorithms developed by IFS and held as closely guarded commercial secrets.

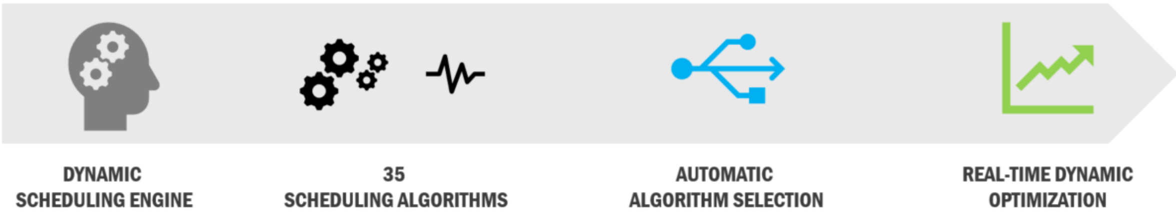

Based on the characteristics of the problem the engine is presented with, the DSE automatically decides which algorithm(s) is/are best suited. On average, a combination of between two and six algorithms are expected to solve most problems, since we've learned that combinations generally work better than a single algorithm alone.

Having chosen the algorithms, the DSE uses meta-heuristics (data that guides) to dynamically adapt each algorithm's parameters to the business goals which have been specified. This calibration is continually evaluated so that if the nature of the business problem changes over time (like if an engineer goes off sick, or an urgent job arrives), the system will automatically adapt and self-learn without the need for any human intervention or re-implementation effort. We believe this to be unique in the industry.

Why does this matter? – This is the HEART of the matter

Most solution approaches in this area revolve around a few well-known optimization techniques, all of which suffer from a problem called local maxima – where the engine gets "stuck", due to time constraints for the window within which you must come up with a solution. For example, in a batch scheduling process where the optimization runs every 10 minutes, the algorithms are typically expected to arrive at an answer within 2 minutes from a standing start. Imagine asking your dispatch team to only work for 2 out of every 10 minutes?

This is "good enough" at a smaller scale and for problems that do not have such an impact on the Total Cost of Ownership. This is typically why batched scheduling optimization engines split your team into lots of smaller areas by enforcing work zones or regional geographical boundaries, creating inefficiencies and resource planning headaches. Optimisation specialists sometimes call this inferior method "fancy calendaring", in recognition of its limitation of benefits.

However, most companies are solving problems for hundreds, or even thousands, of engineers across multiple locations with multiple data sets (things like skills, SLAs, etc) – but also solving within tight financial requirements – all in a dynamically changing environment, ALL the time!

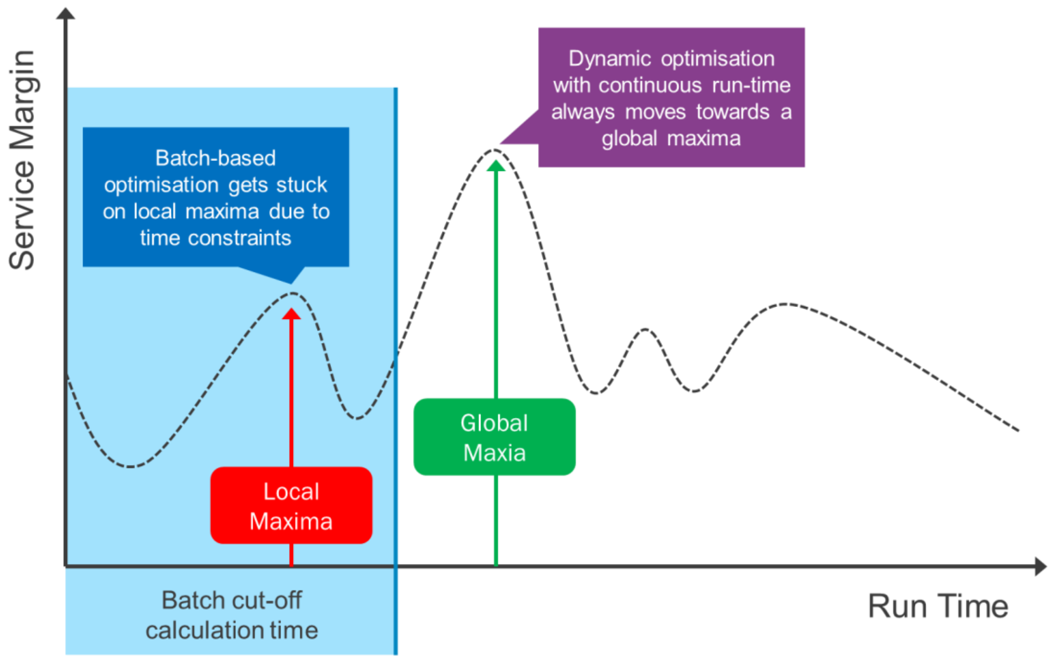

The IFS approach is confidential, but broadly speaking, while the optimization algorithms are running, mathematical constructs called hyper-heuristics are fine-tuning the parameters continuously, to reflect the difference between starting a schedule from scratch and adjusting one which is already fairly good (a local maxima).

> *"We have recorded savings across all our contracts which means Cubic has been able to expand the business by taking on several new contracts, as well as some substantial variations to existing contracts, without having to increase the number of engineers and resource controllers. We are so pleased with the impact of the scheduling system that it is, along with other new technologies we have introduced, being used as a differentiator when bidding for new business."*
>
> MIKE GOSLING
> Manager of Management Information Systems

So, the DSE is automatically adapting to its environment at three different levels – to our knowledge, no one else has spent the time, money and talent to do this at scale in these complex problem areas – IFS rivals are always restricted, they just avoid telling you, or they often hide the fact by either using a reduced data set, or falsely relaxing the SLA requirements when they demonstrate to you.

This is what gives us both the flexibility and scale at the level you require to stand out.

- Algorithm choice (pick the right tool for the job because one size doesn't fit all problems)
- Algorithm parameter tuning (tweak all the time because this is the real-world)
- Hyper-heuristics (guide and learn to save time and get results)
- Continuous solving (always run a live solve that you build on – don't start again every time)

On top of this, all the adaptation is automatically learned, we often find creative solutions to customer problems very quickly.

Multi-time Horizon Planning (MTHP) and Optimisation

As a final requirement in this space to maximise Return On Investment, the same set of unique software solvers can be used for scheduling and planning across multiple time domains – see the graphic below:

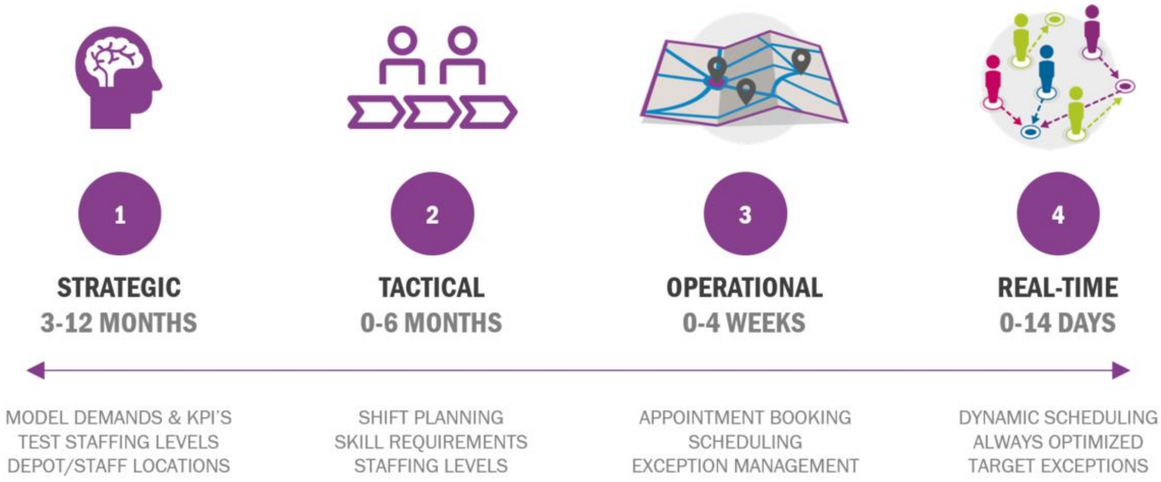

To really appreciate the scale of the challenge put forward by most companies, the Chinese game of Go has 100 legal moves for every position, and the number of possible games of 400 moves or less is on the order of 10800 – meaning that there are about 10720 possible games for every atom in the known universe – and we've only recently created a computer program that has beaten a human at the game.

2 People, 4 Activities
- **120 solutions**

2 People, 5 Activities
- **720 solutions**

4 People, 10 Activities
- **1,037,836,800 solutions**

5 People, 30 Activities
- **12,301,367,000,000,000,000,000,000,000,000,000,000 solutions**

Now scale this mathematics problem up to a working year at a typical company with tens of thousands of jobs and thousands of engineers with minute by minute change potentials – all to SLAs while maximizing return on investment – the computational load is mind boggling.

Therefore, this subject is utterly critical to the solution put forward to your business.

Optimisation Plus Machine Learning

Not content to just use the power of our algorithms to define the best schedule we also seek to learn and provide more accurate data inputs into the optimisations problem, thus increasing its predictive accuracy. Currently this is done in two ways. First by the introduction of an Automated Intelligent Travel Profiles add-on, which provides more accurate future travels times for conurbations based upon millions of actual journeys taken. Secondly, the actual duration times of types of work can be learned and these replace the default work duration times. These learned data inputs, based on actual historical values, enables the optimisation algorithms to provide an even more accurate future plan.

What does this look like in the real world?

CUBIC TRANSPORTATION SYSTEMS

Cubic Transportation Systems is a leading provider of integrated systems and services for transportation and traffic management, providing ticketing and fare collection services for enterprise customers. As part of its IFS Planning & Scheduling Optimization solution, the company has been using the IFS Dynamic Scheduling Engine since 2009.

When Cubic first started using the DSE, they noticed that the system was making different scheduling decisions, compared to how it operated previously. For instance, there may be three repair jobs, all at the same customer location, and previously an engineer would have travelled to this location and done all three jobs. But sometimes, the DSE would schedule the engineer to do just the urgent jobs at that location and then travel elsewhere to another urgent job, leaving the lower-priority jobs to be done later either by the same engineer or a different one – turns out it was a better solution. For them as the business cost of missing SLAs outweighed the travel cost of returning to that location.

Another example that Cubic noticed is that, when an engineer started work, they would often travel straight past a job that needed doing in order to get to other jobs which were more important or where a group of jobs was clustered together or were more appropriate for the engineer. They questioned this initially until it became clear that the DSE was optimizing the whole schedule and that an engineer would return to do the original job later, thus maximise Return On Investment.

The DSE was not told to do this. It was given the constraints – skills needed, SLAs (target completion times), travel times – and it chose how to solve the problem optimally with the application of Machine Learning described previously. By focusing on the right constraints, and solving AT ALL TIMES, it was able to improve the SLA hit rate while maintaining an optimal Return On Investment against all the constraints in place – even as jobs and appointments shuffled during a working day.

Concluding thoughts:

The technical team here at IFS hope this has given insight into the fact that optimization in software is not all created equally, and one size does not fit all. In addition, doing this at scale is an exceptional challenge and still an active field of engineering and academic research, one we have pioneered at IFS for many years already. We understand that this can be a rather technical and potentially dry subject, but the team here utterly believe that in your business, this technology has the potential to deliver something unique and world class if we collaborate to understand the application of this science.

We invite your team to dive deeper into the heart of this problem and speak to our team around this subject.

For general information about IFS Applications, visit [ifs.com](ifs.com)