CHAPTER 3

Leveraging AI for Asset and Inventory Optimisation

Sid Shakya, Anne Liret and Gilbert Owusu

Introduction

Service organisations are typified by resources, both human and non-human. Human resources comprise the front and back office staff, and non-human resources include spares, network assets etc. Managing resources to meet customer demands is one of the key challenges in any large service organisations. It is well recognised that the proactive management of resources is one of the key contributors to the performance and profitability of service organisations (Shakya et al. 2013). Proactive resource management provides the framework to optimise the cost and quality of the products and services an organisation offers. It is one of many challenges that service organisations are faced with on a regular basis. There are, for example, specific challenges to be tackled in resource management, such as making decisions on many different types of resources that a company should maintain, and, more importantly, on managing the ways these different resources interact together to create products and service (Owusu & O'Brien 2013; Shakya et al. 2017).

A case in point is that fixing broadband at a customer's premises may involve a field technician, a vehicle, spare parts, a call centre operator and the network.

How to cite this book chapter:
Shakya, S., Liret, A., and Owusu, G. 2022. Leveraging AI for Asset and Inventory Optimisation. In: Wang, Y., and Pettit, S. (eds.) *Digital Supply Chain Transformation: Emerging Technologies for Sustainable Growth*. Pp. 39–60. Cardiff: Cardiff University Press. DOI: https://doi.org/10.18573/book8.c. Licence: CC-BY-NC-ND 4.0

These are all classified as resources. Each of these resources can have different capabilities (i.e. functions), capacity and/or a geographical requirement. For example, a specific field engineer may have broadband installation skills and plumbing skills, whereas another field engineer may have cable and fibre jointing skills. Similarly, a vehicle type could be a van with light equipment or it could be a truck that is able to carry large heavy equipment. This also applies to spare parts, call centre operators, or network services with various capabilities and different specifications, which may lead to many possible variations of service offerings. Managing these diverse resources to support the efficient delivery of products and services is a crucial and challenging problem.

In recent years, there has been a strong drive to leverage best practice in the supply chain management domain to service operations (Voudouris 2008). Recent research has shown that this increasingly involves the automation of planning processes (Owusu et al. 2013). The key objective of planning is to have the right resource available at the right time in the right place to fulfil customer demand. Advanced planning of resources helps firms to maximise utilisation and minimise waste, and by doing so it also helps firms to fulfil customer demands and maximise revenue, while minimising cost. One of the key prerequisites for successful planning is the optimal deployment of strategic resources to enable a frictionless delivery service.

In this chapter we deal with assets and inventories, one of the key resources that service organisations such as telecommunications and utility companies maintain. There are two dimensions – *strategic* and *operational* – to the deployment challenge. The strategic dimension focuses on deploying fixed assets to ensure that the organisation is set up for optimal performance. The operational dimension focuses on replenishing inventories (spares) for efficient delivery of services in alignment with service level agreements with customers (SLA). As with any other resource types, the timely availability of spare parts can have a positive impact on service quality. Therefore, it is very important to have the right spares in the right place and at the right time. Spare parts are normally kept in warehouses or distribution centres. It is therefore very important that the warehouses are also built in the right places, where they can provide the maximum value to the organisation while minimising the travel time and maximising the distribution coverage. This is a combinatorial optimisation problem. Real-world combinatorial optimisation problems involve a heterogeneous set of side constraints (i.e. rules). Modelling and maintaining such rules or constraints is non-trivial for complex problems. Operational requirements such as reuse and model configurability make AI (in particular heuristic search methods) a prime candidate for solving combinatorial optimisation problems for operational use.

The remainder of the chapter is structured as follows. The next section focuses on using AI to model and solve the challenge of strategically deploying assets. In the subsequent section, we focus on the replenishment of spares. We provide use cases to give insight into how we operationalised these models.

Strategic Deployment of Assets – the IoT and Inventory Management

Recent advancement in the IoT (internet of things) and connected technologies have had an impact on how warehouses are built and managed. Increasingly, warehouses are getting smaller and mobile in their nature, shifting from fixed structures, such as brick-and-mortar buildings, to mobile containers or even lockers. They are remotely monitored and operated, and the right to access is normally given on demand, accessible with a programmed device or unique passcode that is generated per visit. A fixed number of personnel (typically engineers and technicians) are assigned to these mobile warehouses based on their home locations and working locations, such that the distances they need to travel to get spares and get to the clients or repair sites are minimised. These mobile warehouses can be quickly deployed to different locations in a very short time. More importantly, they can be moved from one location to another and can be redeployed and reused, if required.

A typical case of redeployment in a telecom scenario would be that the demand for service at a certain location need to be shifted to another location because of the completion of a new housing project, which leads to the completion of the telecom infrastructure deployment in that area. In such a case, the mobile storage facilities near to the area can be relocated to a new area where a new project has started. Another example would be a situation where fewer repair jobs are required in certain areas and therefore spare parts are not required as frequently as in the past. This may be due to a change in technology or an upgrade to the telecom infrastructure. In such cases, the mobile warehouses can be moved to areas where they could be better utilised. Furthermore, in some cases, mobile warehouses are used by several lines of business (LoBs) simultaneously, both within the same organisation or contracted to an external organisation. There could be a complex SLA (service level agreement) with specifications such as who can use the mobile warehouses and how/when they can be used, typically operated through a booking system. A case for redeployment happens when a new LoB or new organisation is added in the service chain to use the mobile warehouses. In many cases, a new set of mobile warehouses would require a fresh deployment. This also occurs when an organisation leaves the service chain.

The location where these mobile warehouses can be deployed can also have constraints, such as the availability of a fenced perimeter, accessible by a large vehicle delivering spare parts, and a requirement for electricity. Therefore, certain locations are pre-determined as suitable hosting sites. The task then becomes finding the best location out of the suitable locations set to deploy or redeploy the mobile warehouses.

The problem is trivial if fewer mobile warehouses are involved with a small number of users and sites. The deployment decision could be made manually and accurately. However, this is not the case with most large service organisations. They can have hundreds of such mobile warehouses serving thousands of engineers with thousands of possible deployment locations to choose from. In such scenarios, a manual design of the deployment locations of mobile warehouses can be prohibitively time-consuming. More importantly, the design can be sub-optimal in terms of the distribution coverage and travel time required by engineers to acquire the spare parts.

A Use Case at BT

BT operates a large network infrastructure in the UK. It has over 22,000 field engineers maintaining over 5,000 exchanges, serving millions of customers and supporting many products and services. It uses a huge inventory and thousands of spare parts per day to repair or upgrade the network equipment, both at customer premises and at exchange buildings. The technicians travel to warehouses each morning to source the spare parts that they require to perform the tasks assigned to them for that day. Some parts could be specific to a task; others could be general spares such as cables and sockets. Keeping the correct number of spare parts in each warehouse is crucial for field operations. Furthermore, specialised spare parts sourced for a specific task should be delivered to the warehouse in a timely manner to complete the task in time and not to miss other impending appointments.

BT operates over 90 fixed warehouses and distribution centres across the UK where engineers can collect spare parts that they have ordered. Engineers travel routinely, sometimes more than once a day, to get the parts that they require. Fixed warehouse locations can cause some issues such as long travel times, particularly when the home location of the engineer is far from the warehouse, and more often when the site where the task has to be done is also far from the warehouse. In addition, those sites currently serve tasks at over 5,000 exchanges across the country. To increase efficiency, BT wants to increase the number of warehouses operational in the country from 90 to over 700 to minimise travel. Furthermore, those new warehouses will be mobile, capable of being quickly deployed and redeployed as and when needed. The new mobile warehouses would keep a set of small storage spaces, or lockers. Each of the lockers can keep parts required by a specific technician. Spare parts would be delivered to the locker as per the booking made, sometimes couriered for a rapid delivery. A technician would get fixed or one-time credentials to access the locker.

The lockers would be hosted at BT's exchange sites, which are capable of handling large delivery trucks, and with a fenced perimeter. However, finding the best 700 exchanges out of 5,000 possible exchanges to host the mobile warehouses is a difficult task, especially when they have to be redeployed every few months. Warehouses have to be deployed in such a way that:

1. The cumulative expected travel time for all technicians across the country is minimised.
2. The differences in distance from a warehouse to the served exchange sites should be minimal, to avoid long travel times to certain sites.

3. The differences in distance from a warehouse to the home locations of the assigned technicians should be minimised, to avoid cases where some technicians travel very short distances while others travel long distances.
4. The differences in the number of the tasks that warehouses are expected to serve per day should be minimal. This is to avoid a situation where some warehouses are extremely busy serving many tasks compared with other warehouses.
5. The differences in the numbers of technicians that the warehouses serve should be minimal. This is for the same reason for the case above.

Warehouse Deployment as an Optimisation Problem

Organisations are increasingly looking to AI techniques to solve large-scale industrial problems. They have been proven to provide good solutions for many real-world problems. Latest advancements in machine learning and deep data mining, and their successful application in some cutting-edge inventions such as self-driving cars, drones and augmented and virtual realities, has further highlighted the effectiveness of AI techniques, with organisations increasingly moving to adopt such techniques in their decision-making process.

The five conditions mentioned in the previous section can be considered the objectives of the warehouse deployment problem and the problem itself can be modelled as a combinatorial optimisation problem, where the goal is to find the 700 best locations out of 5,000 possible locations that satisfy the above objectives.

A manual approach to solving the same problem would involve a set of heuristics to find the 700 initial sites, evaluating them and manually exchanging sites to see if there is any improvement on the objectives. However, this would be very time-consuming, given the combinatorial nature of the problem. This is where AI techniques can be useful.

In the following sections it is briefly shown how the problem is modelled as an optimisation problem and how AI techniques are being used to solve this problem.

AI Approach to Solving the Problem

A simple greedy logic (GL) heuristic for solving the problem would be to choose *n* sites with a large number of tasks as the initial warehouse locations. Then create *n* cluster of sites by assigning all the remaining sites to the warehouse sites based on nearest distance. Evaluate the warehouse sites by calculating values for each of the five objectives and adding them together to get a weighted combined objective value. This objective value is sometimes known as solution fitness value. Then iteratively move to better deployment by changing a warehouse location to a neighbouring site one by one and accepting the new deployment if the combined objective is better.

The core idea of the heuristic here is that choosing the sites with large numbers of tasks as the initial warehouse location is likely to minimise the travel time. The motivation here is that the tasks for those initial warehouses will not require any travel. This, however, could conflict with other objectives. For example, minimising the travel time may not minimise the differences among the travel distances between sites. It may be that some sites require long travel times while others may require very short travel times, even though cumulative travel time may be minimal, resulting in the preference being given to one set of sites at the expense of others. In the optimisation literature, this type of solution is termed the local optimal solution.

A more sophisticated approach would be to use advanced search heuristics such as a genetic algorithm (GA) to optimise the warehouse allocation problem. GAs are a class of population-based evolutionary algorithms (Goldberg 1989) that find solutions for problems using the concept of natural selection and recombination to evolve a better solution (Larrañaga & Lozano 2002; Shakya & Santana 2012). One of the core strengths of GAs is the way they represent solutions (Goldberg 1989).

The obvious approach for this problem is to represent each deployment $x = \{x_1, x_2, \ldots x_n\}$ as a GA solution, where each x_i, which is the chosen exchange for deployment of warehouse *I*, is considered as a solution variable.

The next component of a GA is defined as the fitness, $f(x)$, of a given solution x, representing the quality of the solution. In this case, the combined objective value of five objectives for a solution x is used as the fitness value for solution x. The objective for the GA is then to find the best values for solution x such that fitness $f(x)$ is minimised. This optimisation task can be expressed as below (eq. 1):

$$\min_{x = \{x_1, x_2, \ldots x_n\}} f(x) = \alpha Trv(x) + \beta \Delta Tsk(x) + \gamma \Delta Tch(x) \tag{1}$$

Here, $Trv(x)$ represents the first three objectives, related to travel. represents objective 4, related to difference in task, and $\Delta Tch(x)$ represents objective 5, related to difference in served technicians. Given a solution x, it is trivial to calculate values for each of these three terminologies. The parameters α, β and γ are the weights applied to the respective terms. This provides explicit control over which of the objectives to be prioritised as per the requirements of the design in different scenarios.

Unlike traditional methods, where an algorithm works directly with the problem definition, GAs try to evolve the solution by working with the population of solutions as a whole and selecting the better solution while letting worse solutions die off over multiple generations, so as to increase the density of good solutions in the population and recombining them to create even better solutions.

The difference between GAs and the greedy logic (GL) is that GL works with a single solution and tries to find the best locations by assuming that assigning the sites with high volumes of tasks as warehouses are likely to minimise the travel time, thus not properly considering other objectives. In contrast, GAs continue to evolve to find the best locations that balance all the objectives against one another and explore a wider part of the search space by working with multiple solutions at once.

Business Impact

The AI logic described in the previous section was incorporated into a tool called Intuitu strategic planner (Figures 3.1 and 3.2), which is being used by BT to plan the mobile warehouse deployment. In Figure 3.2, each central node

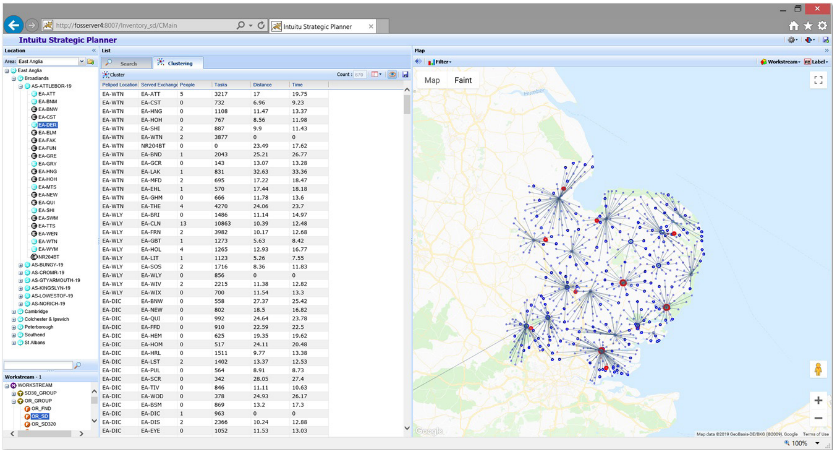

Figure 3.1: Intuitu Strategic Planner tool.

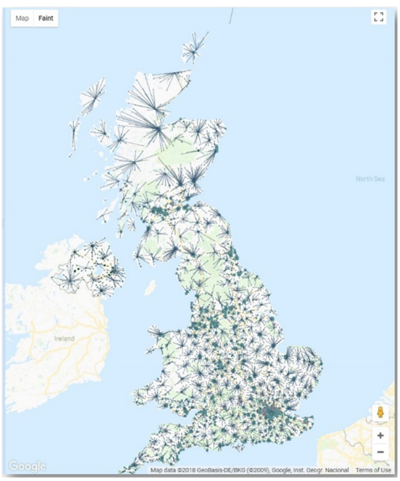

Figure 3.2: A design by Intuitu for 600 locations.

represents a proposed warehouse location and each edge pointing to the central node represents a site that is served by the central node. Our approach has made noticeable improvements for the business in terms of the cost savings due to reduced deployment time and reduced travel time for technicians to source spare parts and to perform tasks assigned to them with increased SLAs.

The combined value of the benefits by the project over a five-year period is estimated to be millions of pounds in cost savings due to a reduction in deployment travel time for technicians. Other benefits include improving the quality of the decision-making process by enabling what-if scenario modelling.

A Use Case of Operational Replenishment of Inventories and Assets

In the previous section we saw a use case of the strategic dimension to inventory and asset management, which focused on the strategic deployment of warehouses to ensure that the organisation is set up for optimal performance. In this section we present a second use case. The focus is on the operational aspect of inventory management. As mentioned before, the operational dimension to inventory management is about replenishing inventories (spares) for the efficient delivery of services in alignment with the agreed service level agreed with customers (SLA).

A typical operational journey starts with a demand linked to a fault at a client site impacting one or multiple parts (pieces of asset). This demand is associated with a service assurance that constrains the faulty part to being repaired within a pre-agreed time, which could vary from two weeks down to less than one day. Once the fault is identified and the availability of the relevant spare part has been confirmed, a job is created and assigned to a field technician for further survey and fixing operation. When a part is picked up by the engineer out of a warehouse stock, there is an update process over the supply chain, to ensure any future task allocation decision takes into account the remaining volume of accurate spares. Figure 3.3 outlines the typical flow of operations in an organisation responsible for asset-constrained service of maintenance.

It is well recognised that it is possible to apply AI reasoning to recommend a proactive decision for the transfer of spares to the warehouses according to configurable asset management policies, while optimising the overall cost such as storage of asset and shipping taxes (Desport et al. 2016). The use case here was to build a plan of asset transfer between warehouses over a given number of days and compare the plan against different asset replenishment policies and demand profile. Such decisions are typically made to address the asset replenishment problem. When the same asset is common across the business or can be used to serve different clients, the asset management includes:

1. an estimate of how many and where/for which client the asset will need to be replaced;

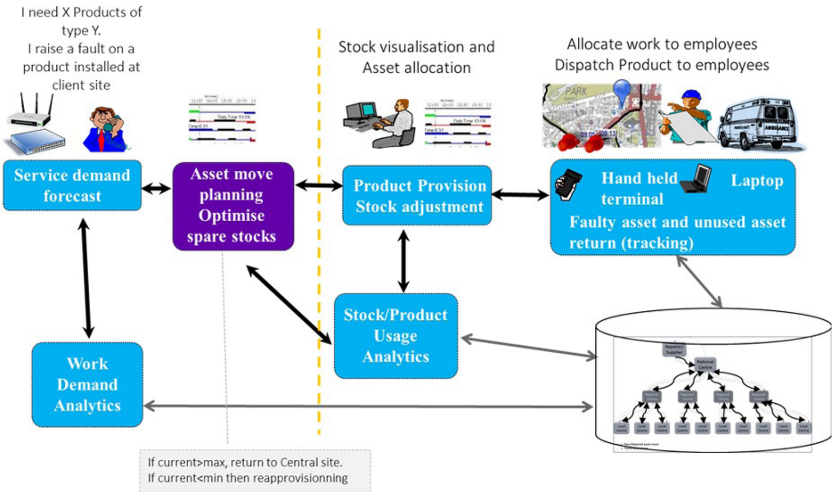

Figure 3.3: AI-supported asset-constrained service flow.

2. an evaluation of the optimum volume of spare stock of such pooled assets that allows the customer needs to be served under the agreed commitments;
3. what policy for the replenishment of those assets is recommended according to the criticality of customer contract or the asset cost.

In the telecommunications, facilities and energy sectors, maintenance service providers have to address the large-scale real-world variants of asset management across multiple customers and over a distribution network of widespread stock warehouses and repair centres. A typical distribution network in e-logistics is composed of a repairer or supplier node, a main national centre, regional centres dependent on the national centre, and a set of local centres dependent on the regional centres. The topology of the network defines the allowance and cost of asset transfer between two nodes of the network (i.e. two centres). When a new or repaired asset comes into the chain, the process moves it from supplier node down to one of the suitable centres. When a faulty asset is returned to the repair centre, the process moves it from local centres to the repairer.

A Typical Use Case

As a case in point, a maintenance agreement of premium service assurance can require a full fix within four to eight hours, which means all the following steps have to be completed within that time range:

1. assigning a field technician,
2. collecting the spare part at the warehouse,

3. travelling to the client site, and
4. replacing the faulty part.

It is well known that the field force scheduling aspect of the problem only can be operationally automated and enhanced using a vehicle routing problem model and heuristic search solution approach (Liret 2008). Nevertheless, the problem variant with asset collection at warehouse and completion in very short response time requirement is a challenge to performing optimally. As a matter of fact, field engineers could not have time to collect spares, travel and perform the job unless a spare is present in the nearest warehouse. More precisely, to solve this problem, a proactive planning solution is needed to plan a minimum number of spares ahead so that, operationally, there is sufficient relevant spares in place at the nearest warehouse to the client site. The volume of spares to plan is inherently a function of an acceptable business risk, usually estimated by services assurance management teams and supply chain operations teams.

There is thus a proactive replenishment planning problem to solve in order to transfer the right amount of the right spare asset in alignment with the service assurance. To solve that problem in a sustainable manner, one needs a hybrid approach towards the asset movement planning in order to optimally meet the service assurance across as many client sites as possible, while recommending the risk of shortages for the warehouse and service contracts. When a fault profile is predicted (based on client, asset features and geography), the approach would extend to targeting any required replenishment action (asset transfer) ahead of the fault estimated date. The assumption in this section is that assets can be used across multiple clients and countries. Thus, the question is how to optimise the storage of these assets for seamless reuse.

Whereas the strategic aspect of mobile warehouse deployment supports the right positioning of warehouse of assets for service assurance, the operational aspect of asset transfer plans supports the decisions over which product and which spare volume should be replenished (i.e. transferred) at which warehouse (mobile or fixed). In both cases, IoT technologies are key enablers of a zero-touch approach across the service chain.

Asset Move for Automated Replenishment Supported by the IoT

As mentioned in previous use case, the IoT is changing the way businesses maintain equipment, write service agreements and set customer expectations in the process – i.e. exploring a new approach to maintenance that is driven by insights, instead of errors. Today, most companies offer scheduled maintenance as part of an equipment service contract. The IoT now enables a shift from just consuming data from connected devices to gaining visibility into the current state of equipment and using that information to deliver a different type of field service while operating the inventory and asset network in a more proactive manner. We can now use data from sensors that indicate an asset's health to act based on the probability of a fault occurrence. It has been suggested that the IoT will shortly enable businesses to shift from recurring preventative maintenance plans to the proactive monitoring of devices and predictive maintenance (Pintov & Brandeleer 2019). Indeed, IoT platforms can host artificial intelligence (AI) components that monitor trends and predict which installed asset is likely to fail. AI and data science provide the ability to process massive amounts of information (given the right dataset), which help to inform the need for increasing the volume of spares in particular warehouse at certain date.

This kind of reasoning will reduce the likelihood of maintenance delays while improving customer satisfaction. This approach, however, poses a number of challenges. For client service assurance management, we want to ensure the SLA can always be met on existing maintenance contracts, with a recommended risk of reaching a shortage in spares in the event that a new client site or contract is evaluated.

There are number of questions that have to be answered:

1. According to a risk function, do we have enough spare to cover short SLA contracts? Are they at the right locations? If not, how much time will it take to get the right coverage?
2. Do we have enough equipment in stock at the right locations to cover contract requirements? If not, can we recommend actions (plan of asset move, invest in new asset) to meet the short SLA service assurance?
3. Do we have stock in surplus, i.e. assets that are not used nor related to any potential contract?
4. Knowing a demand profile, what is the most suitable asset decision that de-risks the service (i.e. minimises the volume of unmet demand or penalties)?
5. What acceptable risk rate can we afford with a given stock and client sites scenario without investing? What is the best risk rate to apply for a given client, product or geographical area, knowing the reported faults, proactive asset moves, and risk rate used in the past?

The high-level problem could be outlined as in Figure 3.4. The functional component is notified by a number of inputs such as a change in the customer sites and warehouses –which could be represented by an address and a capacity of installed or stored parts, and a distribution network linking these sites according to some policy of transfer. Moreover, to be able to analyse the state of equipment and estimate their fault likelihood, a certain level of accuracy in the volume of spares (different status) in addition to a certain agility in updating the real-time data is expected. Another key input is the model for penalties and client priorities.

The output of a typical asset replenishment process would be a plan over a number of days recommending a set of transfers of assets from the warehouse to other locations closer to the client site, while minimising the overall cost

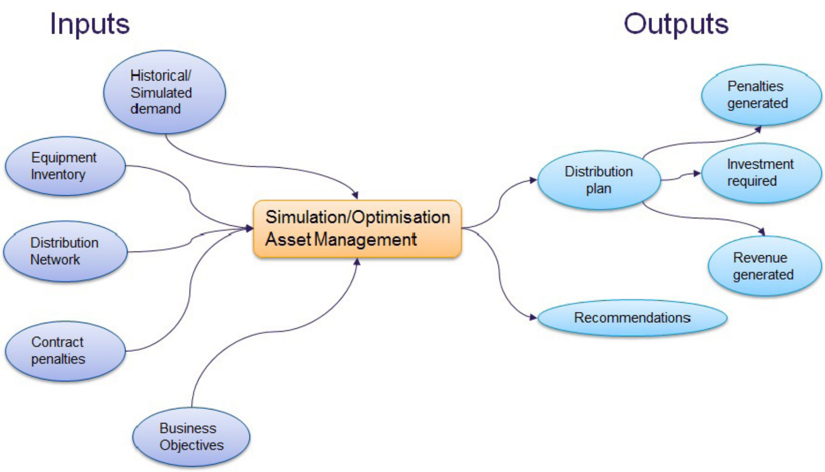

Figure 3.4: Problem statement.

(storage, transfer, penalty fees) and assessing the asset products that are estimated to be at risk (since there are installed occurrences not covered by a spare one), and some recommendations for investing in critical assets.

Replenishment Optimisation Problem

Desport et al. (2016; 2017; 2019) proposed a heuristic search-based model optimising the planning of asset volumes at each location of a list of sites, in closed-loop chains, taking into account a pre-known demand, unit costs for transferring part from one location to another one, and unit costs for storing the assets. This model addresses the pure asset move problem. Asset move planning covers only one part of the replenishment problem (the reusable one in a closed loop). However, to apply this approach to real-word, client-wise costs, service assurance risk estimates and some equivalence knowledge reasoning between reusable products are required. Introducing these features into the asset move planning problem allows us to assess the benefits of an augmented AI-based automated supply chain decision, which has traditionally been made by human and most frequently driven with a siloed view of each client's product needs. We use AI to gain flexibility in decision-making and recommendation against uncertain and dynamic demand trends in assets (fault prediction and asset provisioning due to contract).

We propose a hybrid AI simulation approach with the aim of de-risking asset decisions for customer service. We model the problem as an extension of Desport et al. (2017):

1. In addition to the demand profile, a *minimum spare stock amount* per client contract is defined according to a service assurance policy and an *acceptable risk rule*. The risk value is defined as a rule for a given product,

and minimum stock level as a function of risk value, the spare for the product in each warehouse, and the SLA of contract on the same product. For instance, a simple risk rule is: '1 spare is required for 10 installed parts at client site', '2 spares for 11 to 20 installed parts', etc. If two client sites are mapped to a city node (e.g. PARIS), one client having 32 CLK routers installed and another client 14 CLK routers installed, then the minimum stock of CLK product for node PARIS will be 1+ (46 modulo 10), i.e. five spare pieces of CLK product.

2. The client's priority is defined as a cost of penalty if a service is not covered as per the minimum stock constraint defined in (1); for any potentially missing spare, a risk is estimated, and a penalty cost defined.
    1. Cost incurred when a delay in service occurs: this penalty can be, for instance, a function of contractual fixed fees for each day of delay, and of a variable component function of the unit penalty per product and day of delay, and of the number of assets estimated at risk because not covered by a spare at nearest warehouse node. Both components can be weighted by a product-wise or customer-wise factor.
    2. Cost of healthy/faulty storage, function of (asset, node): each asset stored in a warehouse will incur a cost, which is configurable per product and warehouse node.
    3. Cost of not reaching the minimum stock for any tuple (node, product): a cost will be incurred on a daily basis if the latter is not satisfied.
    4. Cost of shipping a number of products during a transfer function of (origin, destination, asset): any valid move will incur a transfer cost (which could be zero, for instance when bringing an asset back to the main repair centre).
    5. Cost of repairing an asset or sending it to the manufacturer, function of (asset).
3. Valued topology of distribution network: each warehouse (depot) and customer site is abstracted as a node of an oriented network. Depots can be of various superficies and storage can cost vary depending on the region. In the problem model, tuning the storage cost allows us to iteratively identify the suitable policy of storage. For instance, if assets from the same product are installed on different clients based in different cities all in the same region, there is a choice between storing spares at the regional warehouse or distributing the spares in smaller volumes in each city local warehouse. The level of service assurance as well as the transfer cost influences this decision.
4. Handling faulty parts: in a supply chain, faulty assets are transferred back to the repair centre either internally to provider, or directly to the manufacturer with the target that after a period of time a healthy asset can be reinjected into the chain (closed-loop supply chain). With the support of the IoT, we can imagine an automated approach that allows triggering the transfer back and the reinjection of a given product in a given warehouse.

AI Approach to Solving the Problem

The problem can be modelled as a constrained optimisation problem (Hooker 2012; Taleizadeh, Niaki & Aryanezhad 2010), where decision variables represent possible actions, a set of constraints allows business rules to be represented, and a list of costs items is to be minimised. Basing the solving algorithm on AI techniques known as meta-heuristics, an iterative process is considering suitable moves (partial valuing of the solution) and performing the best moves in the search space. For each valid move (satisfying the constraints), the impact on the costs is computed and then, depending on the meta-heuristic strategy chosen, the move will be accepted or rejected. Moves in this problem can be any of the following:

1. bringing in new asset using available capital expenditure (CAPEX),
2. moving assets between different storage warehouses, and
3. repairing assets and reinjecting them into the network.

The strategy chosen in this use case is based on a heuristic search applying a best improvement neighbourhood selection: each feasible move represents the transfer of a volume of asset from one node to another node in the distribution network. Each transfer incurs a cost and an update of the volume at each impacted node. The transfer is validated by the heuristic search only if it leads to a cost improvement that reduces the overall penalty cost across all clients, while limiting the storage and transfer cost. Thus, the volume of transfer will be bundled (grouped) to avoid extra transfer cost. Further, the asset will not be moved if an equivalent product is already in place in sufficient volume. The latter approach requires the modelling of an equivalent model between products and its incorporation into the risk rule and minimum stock evaluation. With regard to constraints, a constraints network restricts the list of valid asset moves and actions considered by the algorithm; from the topology of the distribution network, a set of constraints defines valid transfers (oriented) between centres along with their duration, and status of assets (healthy or faulty) and type of product (heavy or not). This is important to reflect organisational policy (faulty parts are to reach a repair centre, for instance), as well as allowing sufficient agility in the adjustment of the topology without invalidating the replenishment optimisation method.

Business Impact

The output of such optimisation and AI-driven tool includes:

1. asset moves plan 1–7 days (all involved in a four-hour SLA) (plus estimated reduction for financial risk as impact) to cover service assurance with existing spares in supply chain, as outlined in Figure 3.5;

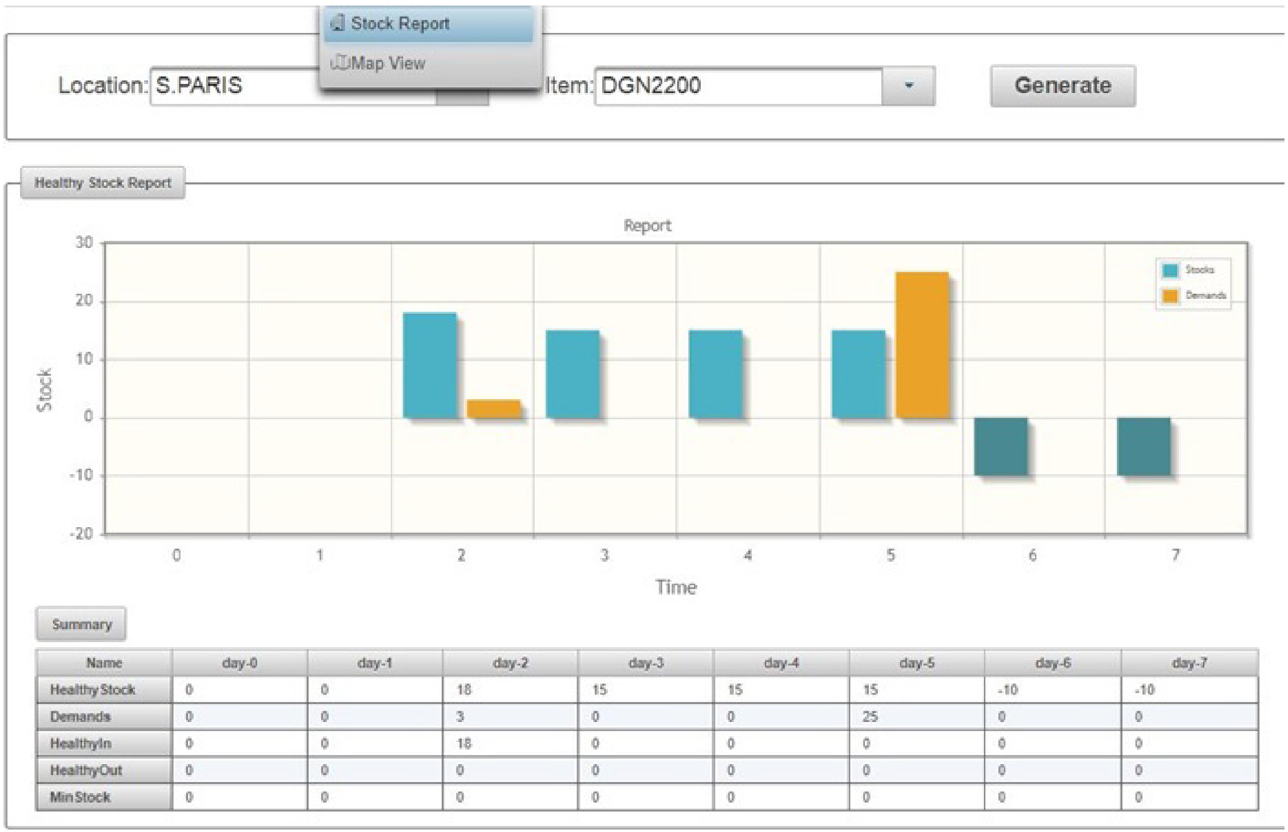

Name	day-0	day-1	day-2	day-3	day-4	day-5	day-6	day-7
Healthy Stock	0	0	18	15	15	15	-10	-10
Demands	0	0	3	0	0	25	0	0
HealthyIn	0	0	18	0	0	0	0	0
HealthyOut	0	0	0	0	0	0	0	0
Min Stock	0	0	0	0	0	0	0	0

Figure 3.5: Example of report dashboard of recommended asset transfer plan.

Node	Product Id	Nb Equipe under Contract	minimum stock	Stock Before	Stock After
S_PARIS					
S.PARIS	PWR-2700-AC/4	2	45	0	6
S.PARIS	CLK-7600	1	30	0	11
S.PARIS	7600-MSFC4	1	31	1	10
S.PARIS	WS-F6700-CFC	1	30	3	30
S.PARIS	WS-X6724-SFP	1	23	0	23
S.PARIS	PWR-MX104-AC-S	2	22	0	0
S.PARIS	FAN-MOD-4HS	1	22	5	10
S.PARIS	CISCO7604	1	19	2	5
S.PARIS	CWDM-SFP-1490	2	17	3	3
S.PARIS	RE-MX-104	1	11	0	0
S.PARIS	MX104-MX5-AC	1	11	1	1
S.PARIS	7600-PFC3C-10GE	1	10	0	0
S.PARIS	MIC-3D-20GE-SFP-E	1	10	1	1
S.PARIS	MEM-XCEF720-256M	1	11	2	2
S.PARIS	RSP720-3C-10GE	1	8	8	3
S.PARIS	FANTRAY-MX104-S	1	11	3	3
S.PARIS	X2-10GB-LR	2	6	0	6
S.PARIS	C9300-48T	4	1	0	0

Figure 3.6: Result of risk on products and asset replenishment recommendation.

Table 3.1: Example of service de-risking impact following asset move plan deployment.

ODE ID	Product	Min-Stock	Stock of spares before	Stock of spares after	unmet 4 hours SLA before optimisation	unmet 4 hours SLA after optimisation
WH.RENNES	CLK-7600	1	2	1	1	0
WH.STRASBG	CLK-7600	1	0	1	−1	0
WH.PARIS	CLK-7600	6	0	6	−6	0
WH.ORLEANS	CLK-7600	1	0	1	−1	0
WH.LILLE	CLK-7600	1	0	1	−1	0

2. recommendation of equipment/contract at risk (not covered), as in Figure 3.6;
3. recommendation of a site or warehouse nodes, contract, or asset where an action is needed such as purchase, resell or move;
4. recommendation on stock not primarily used (could inform about the feasibility of engaging a new contract within an existing mutualised pot of spares);
5. review of stock and client sites alignment through a geographical view (map), identifying shortage, risk and surplus before and after optimisation (Table 3.1).

Figure 3.5 outlines the qualitative and quantitative impacts that asset move automated planning and processing could provide. Impact can be observed at various levels:

1. *Healthy Stock*: evolution of healthy stock throughout time horizon (negative value reflects non-met demand);
2. *Demand*: number of assets requested on that particular day;
3. *Healthy In*: number of healthy assets received on that day;
4. *Healthy Out*: number of healthy assets lost (due to demand or external move);
5. *MinStock*: minimum stock required for that asset on that day.

Figure 3.5 illustrates a plan over seven days from the perspective of a local warehouse in one large city. On day 2, the PARIS warehouse has received 18 new DGN2200 assets and has a demand for three of those assets. On day 3, we can see that the demand was reduced from the stock, leaving 15 assets left. Starting day 2 with a stock of 20 spares, from day 6, as a result of two peaks of demand, the warehouse is missing 10 spares to serve the total demand over this week's period. This example typically shows the kind of risky situation that could happen when replenishment and provisioning is not planned against a proper service demand profile.

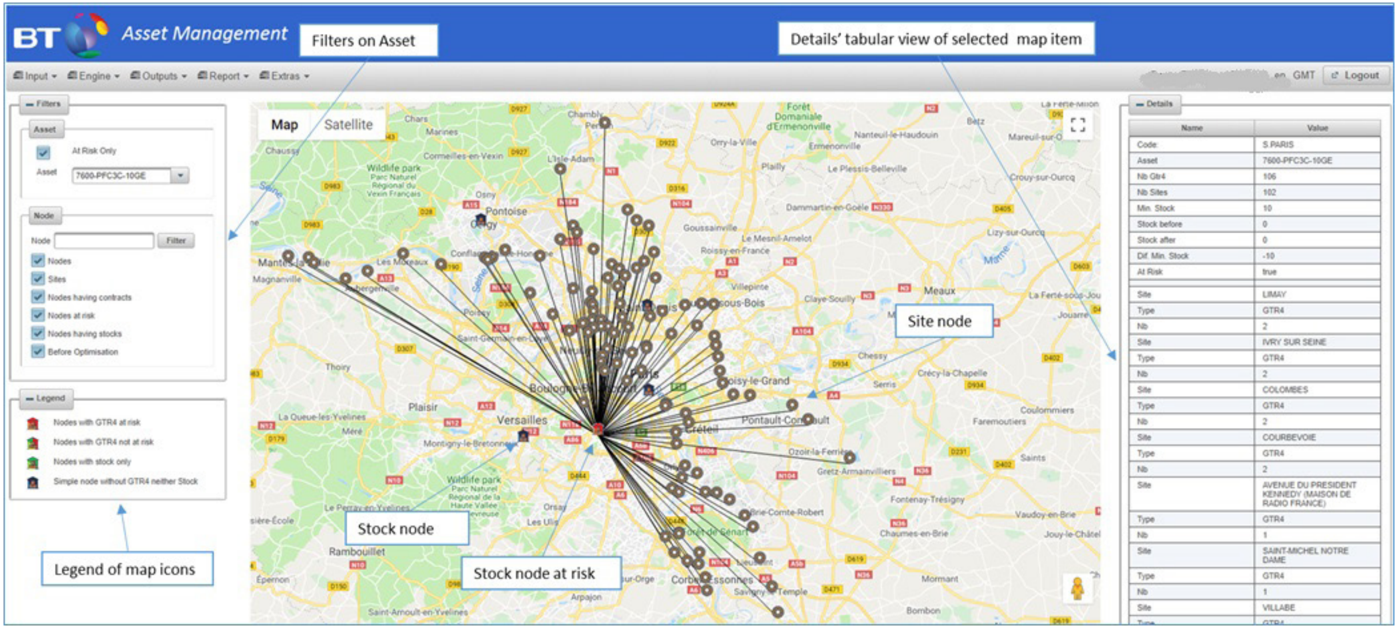

Figure 3.7: Map from one depot for asset 7600-PFC3C-10GE.

On the other hand, when asset move plan has been followed, as in the example of Table 3.1 below, the risk was initially highlighted for product CLK-7600 – the volume of uncovered installed asset before optimisation – which is represented by the value of '*unmet 4 hours SLA before*'. This value is negative, which means that some spares are missing and the risk of the installed asset is not entirely covered by the service. The risk is then reduced to the point of having no contract at risk after the deployment of the optimised asset move plan: the value of '*unmet 4 hours SLA after optimisation*' is 0. The local warehouse RENNES was storing one extra spare compared to the minimum volume needed, which was moved to warehouse PARIS. The other spares were provisioned by the main parts centre (not visible in the table).

Figure 3.6 shows the overall report after optimisation. In the local warehouse PARIS, the asset move optimisation led to a full coverage of service assurance (*MinStock*) for WS-F6700-CFC, X2-10GB-LR, and attempted to reduce the risk for other products CLA-7600 and CISCO7604 but could not fulfil the SLA-wise *MinStock* objective. In this case, the machine recommends investing in or re-evaluating the risk rule.

Analysing the impact for all the products installed and on spare in stock, the overall results show that, before optimisation, the status was:

1. Volume of uncovered MinStock level: 201 spares assets were missing or misplaced.

After optimisation, the status on a real-world use case became:

1. Reduction by 43% of the penalty cost.
2. Volume of uncovered MinStock level: 121 spares required and still missing, so 60% reduction in the volume of misplaced items.

A map view as shown in Figure 3.7, displaying stock nodes (warehouses) and customer sites, allows the visualisation of nodes that have stock and are at risk for a specific asset. It can also be used to rapidly assess the positive impact of the asset optimisation process. In total, 106 contracts of 4-hour SLAs are present and spread over 102 sites. This node is considered at risk with a deficit of 10 items to reach an appropriate level of coverage for those contracts.

The impact of the generated asset moves plan is evaluated as a reduction of the cost potential caused by lack of spare parts and consequently maintenance service failure. This reduction of cost can be derived from the outcome of asset move optimisation planning when the heuristic is guided by the risk rate associated to each asset product. This needs to be balanced against the cost of transferring assets (shipping, packaging) – to avoid moving too frequently and to facilitate a grouped transfer of assets.

Conclusion

Optimising the deployment of assets and the replenishment of spares is key to the successful performance of service organisations. Central to this viewpoint is the ability to model scenarios in dynamic environments. Changes in operations tend to impact the way organisational resources such as assets and spares are utilised to deliver services – a combinatorial and optimisation problem. AI techniques are known for their efficiency in finding good solutions in polynomial time using heuristic search methods and their modelling capabilities to express constraints declaratively. In this chapter, we have presented AI-based approaches for asset management in a service chain. We have described two use cases along the two dimensions of maintaining organisational resources – strategic and operational. We proposed a decision-making approach where AI-constraint propagation and heuristics search techniques are used to proactively recommend the levers for optimising the overall service assurance level and the performance of inventory-dependent services. This enables asset and spare managers to optimise the deployment of mobile warehouses and the replenishment of spares with the objectives of:

1. automating strategic coverage of inventory sites and operational asset decision allocation to warehouses while minimising CAPEX investments;
2. having a list of potential locations for new warehouses and identifying which places are the best located for them, according to a demand profile;
3. recommending warehouse topology changes (removing, resizing) without damaging the asset maintenance service assurance of the organisation;
4. optimising CAPEX while reducing penalties and maximising revenue;
5. optimising minimum and maximum stocks per site/equipment/customer;
6. assessing whether the existing stock in depots is sufficient to cover a given provisioning and fault demand for a given asset;
7. providing the optimised asset move plan given a policy; and
8. estimating what CAPEX investment is needed to meet demand in the event of no feasible plan being found.

Some of the key lessons learnt during the rollout of the capabilities included engaging end users to validate (1) key requirements, and (2) the outputs of the models. We used an agile approach coupled with rapid prototyping in this regard. This approach enabled us to deploy the capabilities right the first time. The capabilities have led to significant operational benefits and better service outcomes for our operational teams. The feedback from the business was excellent, highlighting the flexibility that AI techniques have to offer in asset management in a service organisation.

References

BT. (2019). BT Group plc annual report 2019. Retrieved from: https://btplc.com/Sharesandperformance/Annualreportandreview/2019summary/assets/documents/BT_annual_report_2019.pdf.

Desport, P., Lardeux, F., Lesaint, D., Liret, A., Di Cairano-Gilfedder, C. & Owusu, G. (2016). Model and combinatorial optimization methods for tactical planning in closed-loop supply chains. ICTAI-IEEE 2016 Conference.

Desport, P., Lardeux, F., Lesaint, D., Di Cairano-Gilfedder, C., Liret, A. & Owusu, G. (2017). A combinatorial optimisation approach for closed-loop supply chain inventory planning with deterministic demand. *European Journal of Industrial Engineering, 11*(3), 303–327.

Desport, P., Lardeux, F., Lesaint, D., Di Cairano-Gilfedder, C., Liret, A. & Owusu, G. (2019). A sequence-based metaheuristic for tactical distribution planning in closed-loop supply chains. 9th IFAC Conference on Manufacturing Modeling, Management and Control (MIM 2019), 28–29 August 2019, Berlin, Germany.

Goldberg, D. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, Boston, MA.

Hooker, J. N. (2012). *Integrated methods for optimization*, 2nd edition. Springer, New York and London.

Larrañaga, P. & Lozano, J. A. (2002). *Estimation of distribution algorithms: A new tool for evolutionary computation*. Kluwer Academic Publishers, Boston, MA.

Liret, A. (2008). Work allocation and scheduling. In C. Voudouris, D. Lesaint and G. Owusu (eds), *Service chain management: Technology innovation for the service business,* (pp. 139–152). Springer, Berlin, Heidelberg.

Owusu, G., O'Brien, P., McCall, J. & Doherty, N. F. (2013). *Transforming field and service operations*. Springer, Berlin, Heidelberg.

Owusu, G. & O'Brien, P. (2013). Transforming field and service operations with automation. In G. Owusu, P. O'Brien, J. McCall & N. Doherty (eds), *Transforming field and service operations with automation: Methodologies for successful technology-driven business transformation,* (pp. 15–28). Springer-Verlag, Berlin, Heidelberg.

Pintov, U. & Brandeleer, G. (2019). IoT and the future of field service. Field Technologies Online, April 2019.

Shakya, S. & Santana, R. (2012). *Markov networks in evolutionary computation.* Springer, Berlin, Heidelberg.

Shakya, S., Kassem, S., Mohamed, A., Hagras, H. & Owusu, G. (2013). Enhancing field service operations via fuzzy automation of tactical supply plan. In *Transforming field and service operations,* (pp. 101–114). Springer, Berlin, Heidelberg.

Shakya, S., Lee, B. S., Di Cairano-Gilfedder, C. & Owusu, G. (2017). Spares parts optimization for legacy telecom networks using a permutation-based evolutionary algorithm. 2017 IEEE Congress on Evolutionary Computation (CEC), 1742–1748.

Taleizadeh, A. A., Niaki, S. & Aryanezhad, M. (2010). Optimising multi-product multi-chance-constraint inventory control system with stochastic period lengths and total discount under fuzzy purchasing price and holding costs. *International Journal of Systems Science, 41*(10), 1187–1200.

Voudouris, C. (2008). Defining and understanding service chain management. In C. Voudouris, D. Lesaint and G. Owusu (eds), *Service chain management: Technology innovation for the service business,* (pp. 1–17). Springer, Berlin, Heidelberg.