Production Network Capacity Modeling for Strategic Network Planning

By

Philipp Simons

Ph.D., Materials Science and Engineering
MIT, Cambridge, MA, USA, 2021

M.Sc., Energy Science and Technology
ETH Zurich, Zurich, Switzerland, 2015

B.Sc., Physics
ETH Zurich, Zurich, Switzerland, 2013

Submitted to the MIT Sloan School of Management
in partial fulfillment of the requirements for the degree of

MASTER OF BUSINESS ADMINISTRATION

in conjunction with the Leaders for Global Operations program at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

March 2022

Signature of Author: ______________________________
MIT Sloan School of Management
March 11, 2022

Certified by: ______________________________
Sean P. Willems
Visiting Professor of Operations Management, MIT Sloan School of Management
Thesis Supervisor

Certified by: ______________________________
Michael J. Cima
David H. Koch (1962) Professor of Engineering and Professor of Materials Science and Engineering
Thesis Supervisor

Accepted by: ______________________________
Maura Herson
Assistant Dean, MBA Program
MIT Sloan School of Management

Production Network Capacity Modeling for Strategic Network Planning

By

Philipp Simons

Submitted to the MIT Sloan School of Management on March 11, 2022
in partial fulfillment of the requirements for the degree of
Master of Business Administration

Abstract

Strategic planning of manufacturing capacity requires data-based approaches to determine current and future constraints in a manufacturing network. While oftentimes, the desire to improve decision-making in strategic planning is strong among decision makers, and data on capacity generally exists in some form, there can be a lack of centrally coordinated efforts to harvest existing data as well as high degrees of inconsistency. In addition, modeling manufacturing capacity is an inherently complex problem due to varying modes of production, unclear units of measure, and complex global manufacturing networks.

In this thesis, a capacity model design is proposed for a global medical device manufacturer, and key aspects of the model functionality are demonstrated in a case study. At the core of the capacity model is a database structure using standardized data fields for capacity and demand data, including cycle times, shift structure, and space. The logic of the capacity model is developed, with the goal to capture supply chain complexities such as mixed model lines or various degree of automation. In short, the logic determines the required production time for the product portfolio under consideration, and assesses the available capacity by comparing this required production time with the total available time.

The logic is tested on a prototype product with a focus on mixed model lines. It is found that naming and product grouping inconsistencies require significant manual data manipulation, which – in combination with a lack of standardized, centrally available data – will form the biggest bottleneck in the implementation of the capacity model. Finally, an implementation roadmap is presented to offer guidance on converting the logic presented here into a functional model for decision makers in a supply chain strategy organization.

Thesis Supervisor: Sean P. Willems

Title: Visiting Professor of Operations Management, MIT Sloan School of Management

Thesis Supervisor: Michael J. Cima

Title: David H. Koch (1962) Professor of Engineering and Professor of Materials Science and Engineering

Acknowledgements

This thesis would not have been possible without the outstanding support from many who contributed to the success of this work, and my time at MIT more broadly.

First of all, I would like to thank Boston Scientific for providing me the opportunity to perform the work that led to this thesis in their organization. Thank you to the Ops Strategy team for welcoming me with open arms, and providing me with all the support I could have asked for: Thank you to my supervisor Marc and my mentor James, for all your guidance and support throughout this project. Thank you to Meqdad for helping me with data-related questions and for being a great buddy in the office. Thanks to Vinnie for helping me make sense of many intricacies and for sharing your experience and knowledge with me. Thank you to Dave for being a great mentor, critical thinker, and soccer fan. Thank you also to Jason for sharing your vast knowledge with me. A huge thank you also to Steve for being a great LGO sponsor, for supporting me and the other LGOs during our time at BSC, and for making this great collaboration possible. Thank you also to all the other people who provided me with data, answered my questions, and took time out of their busy schedules to mentor me. Finally thank you to Seth for being a great LGO advocate and for all your work in managing the partnership between LGO and BSC.

Thank you to my research advisors, Prof. Sean Willems and Prof. Michael Cima. Your input and guidance have shaped this project and how I think about operations as a whole. Thank you also to the LGO program for giving me the opportunity to join, in particular Thomas for offering this unique opportunity to me, to Patty for being an incredible support in resolving literally any issue that may come up, to Ted for helping me navigate LGO as a PhD, and to the rest of the LGO office for everything they have done to make this experience incredibly welcoming and rewarding. Finally, thank you to my PhD advisor Prof. Jennifer Rupp for supporting me in pursuing my MBA.

I would also like to acknowledge the incredible cohort of LGO students that have become my close friends over the last two years. It is the honor of a lifetime to study alongside you, and to learn from all of you. Thank you in particular to the Flying Kiwis, hands-down the best LGO co-team in history.

Thank you to my family: to my parents Axel and Antonette, and my brothers Jan and Peter, for keeping me grounded, and for supporting me in anything I do in life. None of this would be possible without you!

Finally, thank you to my better half Linn: Your support and love mean the world to me, and I cannot thank you enough for always being by my side.

Table of Contents

List of Figures

List of Tables

1. Introduction

1.1 Strategic Capacity Planning

The Encyclopedia of Production and Manufacturing Management defines [production] capacity as "the maximum possible output over a specified period of time. In cases where the output is non-homogeneous, capacity may be measured in terms of the available machine hours."[1] Production capacity therefore defines the amount of each product and their components a company is able to produce in a given unit of time. Production capacity is important for planning at various different time scales, as well as at various levels and functions in an operations organization. In terms of time scales, the shortest time scale at which capacity needs to be planned is in the short term, on the day-to-day scale, in order to plan actual hourly, daily, and weekly production volumes. At this scale, capacity can fluctuate based on staffing shortages, equipment downtime, or other short-term sources of capacity variation. On a medium term, at the order of a month to a year, companies need to plan their capacity in order to meet demand forecasts and execute tactical adjustments to the required capacity, such as increase in staffing or capacity utilization. On the long-term, from one to several years, companies need to plan the deployment of capital expenditure such as equipment

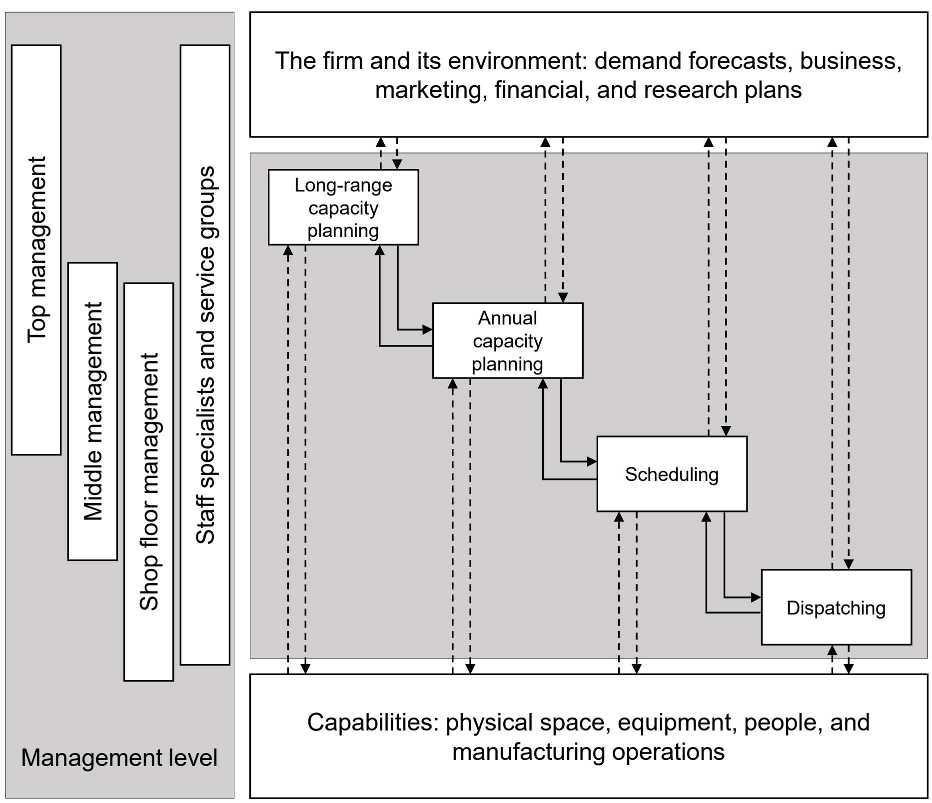

Figure 1-1 Capacity Planning Time Horizons in Operations (from Ref. [2])

or buildings, and make strategic decisions based on market access considerations, long-term market and company performance projections, acquisitions/divestitures, product portfolio changes, or other long-term strategic requirements. The levels and functions at a company which are stakeholders in the capacity planning process range from the local production staff and their front-line managers, to factory site leadership and corporate executive-level management. In addition, functions such as sourcing, supply chain planning, facilities management, or operational strategic planning all have a stake in the capacity planning process. **Figure 1-1**, which is adapted from a book by Hayes et al. (2005) depicts the capacity planning process both as a function of the time horizon and the level of managers involved at those horizons.[2]

Planning the manufacturing capacity of a company along the time horizons described above is part of the capacity strategy of a company, which in turn is a component of the operations or manufacturing strategy, a key discipline of corporate strategic planning.[3] To develop a capacity strategy, a company requires a series of input parameters and assumptions for the future, including the predicted growth and variability of demand, the construction cost and operational expenses of manufacturing plants, trends in manufacturing and product technology, expected behavior of competitors, and expected trends in markets and the supply chain.[4] Based on these predictions, the company can then decide how to strategically position itself, of which a capacity strategy is an important aspect. Developing a successful capacity strategy, and performing capacity planning along the different time horizons between the day-to-day and long-term scale, is a priority to any manufacturing company for a variety of reasons: Deployment of new capacity in the form of the construction of new manufacturing buildings constitutes a large investment which needs to be carefully evaluated. Companies aim to allocate their limited resources most effectively, and therefore, investment into new factories needs to occur when and where it is needed the most. Similarly, manufacturing resources need to be allocated effectively on a day-to-day basis. Furthermore, a good long-term strategy matters because lead times for new capacity is long: it can take three or more years in the medical device industry to construct and certify a manufacturing site. Assuming that a company aims to proactively deploy its capacity to meet market demand, it is therefore critical to have a successful strategy in place that ensures capacity is available when it is needed. Finally, manufacturing capacity is a key aspect in ensuring supply chain resiliency. Successful capacity planning allows for surge production resources to be available when there are demand spikes, while also ensuring that capacity utilization is level-loaded over time.

Complicating the need for an effective capacity strategy is the fact that production capacity is difficult to define and to measure.[4] However, at the same time it is vitally important for a company to accurately depict its current and future state of capacity in order to adequately lay out a capacity strategy. After all, "if you can't measure it, you can't improve it"[i]. Therefore, models of manufacturing capacity are of high interest

[i] Quote commonly credited to Peter Drucker.

for manufacturing companies, which aim to capture production capacity with as much accuracy and detail as possible.

Production capacity forms a multifaceted problem with three key areas of interest: manufacturing space, manufacturing equipment, and labor. Manufacturing space relates to the physical factory that houses the production. On the largest scale, manufacturing space is defined by the factory buildings with their square footage. However, defining manufacturing space requires additional levels of granularity of the individual rooms and even the physical footprint of a line or individual equipment. In addition, for medical device manufacturers, the manufacturing space dimension in the context of capacity also entails different types of space such as cleanroom space with different levels of atmosphere control, as well as whether space can be converted from a non-controlled to a controlled environment. Space is typically the capacity dimension which requires the longest times to deploy, and the most capital. In the most expensive and long-term case, additional space is deployed by building entire factories, which can easily take five or more years from the beginning of the planning process to the steady-state operation of the factory, and require investments of hundreds of millions of dollars. Manufacturing equipment refers to the tools and machines used to produce goods in a factory. Equipment is usually grouped in production lines, and can exhibit large degrees of variation in terms of complexity, degree of automation, size, throughput, process types, or requirements of ancillary resources. Ultimately, this translates to large variations in space requirements, lead times, and cost. In adequately capturing equipment as a dimension of manufacturing capacity, different degrees of granularity can again be considered. The highest degree of granularity for manufacturing equipment is defined by individual pieces of equipment, such as machines or workstations. Individual units of equipment are typically grouped into groups of machines, either by grouping identical or similar machines performing the same tasks, or by grouping equipment which perform multiple sequential processes, constituting a larger process step in the production. These groups of machines are part of a production line, which captures a series of production steps to produce some sort of intermediary or final product. Production lines can be grouped into larger units, such as business units, which in turn make up the entirety of a factory. However, there can be added complexities where groups of machines are part of multiple lines, for example because these machines constitute a process shared by many different products. As such, measuring equipment capacity can be challenging as the granularity of the considered grouping as well as the grouping logic of equipment can affect the way that capacity is measured. Equipment capacity is closely linked with the concept of the bottleneck, i.e., the production step or production steps which constitutes the rate-limiting step of a full production process. The production capacity of a manufacturing line is defined by the throughput of the bottleneck, and therefore, to capture capacity of a process, it is critical to measure the cycle time or throughput of manufacturing equipment at least to the granularity of the bottleneck. Finally, a third dimension of capacity lies in the labor force of a factory or production unit. Labor capacity relates to the degree of utilization of equipment and space: the shift structure and labor deployment defines the fraction of time where equipment is running versus standing still. Labor capacity therefore represents the amount of available time which can be used to convert a theoretical equipment throughput to an actual output of product. Furthermore, labor-related capacity aspects include staffing availability, training times, or short-term staffing fluctuations.

To measure capacity accurately and holistically, all three aspects of space, equipment, and labor need to be considered appropriately. Furthermore, they are naturally intertwined, as the availability of one of those aspects limits the availability of the other two (equipment requires physical space, as well as operators to control the equipment, space without equipment or labor is not productive, etc.). Finally, as described above, for each of these aspects, there is a range of different degrees of granularity, which can be considered in describing capacity, ranging from a network-wide or factory level view, all the way down to an individual production station, with the associated staffing and space footprint. Capturing this granularity is highly complex, but critical to allow for effective strategic capacity planning.

1.2 Company Background

Boston Scientific Corporation (also known simply as Boston Scientific, BSCI, or BSC) is a global medical device manufacturer with ~38,000 employees globally[5], headquartered in Marlborough, MA, USA. The company serves ~30 million patients every year, corresponding to a patient being treated with a Boston Scientific product once every second.[5] Competitors to Boston Scientific in the global medical device market include Johnson & Johnson, Inc., Abbot Laboratories, Medtronic, Stryker Corporation, Cook Medical, and Olympus Corporation. The organizational structure can be described as a hybrid divisional/matrix organization. The commercial business of Boston Scientific operates in seven divisions, namely Endoscopy, Urology & Pelvic Health, Interventional Cardiology, Peripheral Interventions, Cardiac Rhythm Management, Electrophysiology, and Neuromodulation. Moreover, Boston Scientific groups these divisions into three larger units: Endoscopy and Urology & Pelvic Health are grouped into *MedSurg*, Interventional Cardiology and Peripheral Interventions are grouped into *Cardiovascular*, and Cardiac Rhythm Management, Electrophysiology, and Neuromodulation form *Rhythm and Neuro*. The Endoscopy division produces and markets minimally invasive devices for gastrointestinal and pulmonary conditions, and the Urology & Pelvic Health division serves patients with solutions for urological, urogynecological and gynecological diseases. The divisions in the Cardiovascular unit offer minimally invasive therapies for heart, vascular, arterial, venous, and oncological diseases. Cardiac Rhythm Management and Electrophysiology offer treatments for irregular heart rhythms and heart failures, and Neuromodulation provides devices which address neurologic conditions such as Parkinson's or chronic pain. In addition to these commercial divisions, the company has several corporate functions which serve all divisions. These functions include the global supply chain organization, research and development, finance, quality, and business development. This research project was carried out in the global supply chain organization, with responsibility for the end-to-end supply chain of the company, ranging from materials sourcing, internal and external manufacturing, to sterilization and the distribution of finished products. The global supply chain organization consists of the production network, see below, distribution centers, sterilization, as well as global supply chain planning, sourcing, supply chain analytics, and supply chain strategy. Among these functions, global supply chain strategy is responsible for the strategic planning of all aspects of the supply chain, such as network optimization, site selection, operational due diligence of acquisitions, capability assessments, market access, or network expansions. The research for this thesis was performed within the global supply chain strategy team.

1.2.1 Manufacturing Network

Boston Scientific operates a global manufacturing network, with major production sites located in Minnesota, Puerto Rico, Indiana, Georgia, as well as Costa Rica, Ireland, Malaysia, and Brazil. In addition, the network contains several smaller facilities across the world. Several of these manufacturing sites also serve as divisional headquarters, research and development centers, process development hubs and locations for other corporate and divisional functions. The network of Boston Scientific further consists of a large sterilization site located in the Northeast United States, and two tier I distribution centers located in Massachusetts and the Netherlands.

The Global Supply Chain organization is led by the Senior Vice President of Global Supply Chain. Reporting to the SVP of Global Supply Chain are several Global Manufacturing Vice Presidents (GMVP), who form the link between corporate executive management and site leadership. Each GMVP is typically responsible for two to four manufacturing sites, and this distribution of sites between GMVPs is roughly by the commercial divisions that the sites predominantly serve. Therefore, GMVPs are also the senior supply chain leaders who form the liaison to senior leadership in the commercial divisions. Each site is managed by a Vice President who reports to a GMVP, referred to commonly as the site VP. The site VP has responsibility for all day-to-day management of the manufacturing site, including the financial performance of the site. Each production site is structured in business units, which represent a segmentation of the factory into smaller groups, typically by grouping similar production processes into a business unit.

1.2.2 Company Growth

At the time of research for this thesis, Boston Scientific is undergoing significant growth, with a projected revenue growth from approximately $11 billion in 2019 (the last full year before the global COVID-19 pandemic) to $16 billion in 2026.[9] This revenue growth corresponds to a significant growth in terms of units sold, which requires manufacturing capacity to enable the production of these units. **Table 1-1** lists the revenue of all divisions and the company as a whole for the years 2019, 2020, and 2021. The year 2019 corresponds to the last full year before the onset of the global COVID-19 pandemic, and as such forms a more reliable baseline for comparison of future company performance than the years 2020 or 2021. During those years, the COVID-19 pandemic caused a major recession, as well as waves of surging hospital occupancy that required many elective surgeries to be postponed. Since many of the products produced and marketed by Boston Scientific fall under the category of elective surgeries, these fluctuations led to strong fluctuations in demand of Boston Scientific products. In addition, the COVID-19 pandemic led to disruptions of many global supply chains, in particular of air freight due to cancellations in commercial airline traffic, and in the semiconductor industry. These disruptions were also felt by Boston Scientific and required the company to utilize its safety stock reserves in order to meet patients' needs. As can be seen from **Table 1-1**, there is strong growth within all divisions of the company, corresponding to the requirement to build more products. The overall growth forms the key motivation of the present study: in order to meet the projected growth targets, Boston Scientific needs to perform effective strategic planning of manufacturing capacity. Significant investment into manufacturing capacity will be required to provide the capability to sustain inherent product unit growth. In order to perform strategic planning of this growth trajectory in the operational landscape, it is critical to understand the current and future state of manufacturing capacity, and to proactively detect capacity constraints in order to alleviate them.

Division	Revenue 2019[6]	Revenue 2020[7]	Revenue 2021 [8]
Endoscopy	$1.9 B	$1.8 B	$2.1 B
Urology & Pelvic Health	$1.4 B	$1.3 B	$1.6 B
Interventional Cardiology	$2.8 B	$2.3 B	$3.0 B
Peripheral Interventions	$1.4 B	$1.6 B	$1.8 B
Cardiac Rhythm Management	$1.9 B	$1.7 B	$2.0 B
Electrophysiology	$0.3 B	0.3B	$0.4 B
Neuromodulation	$0.9 B	$0.8 B	$0.9 B
Total	$10.7 B	$9.9 B	$11.9 B

Table 1-1 Revenue of the commercial divisions of Boston Scientific in the fiscal years 2019, 2020, and 2021.

1.3 Previous Work

Boston Scientific has hosted several previous research projects in collaboration with the Leaders for Global Operations program at MIT, which performed research with respect to various aspect of the operations of the company.[10–16] Among these studies, the work by Awuondo[10] is closely connected with the present thesis. Awuondo carried out his research in the same organization within Boston Scientific as the present study, namely in the Global Supply Chain Strategy group, which was formerly known as the Global Operations Strategy group. In his work, Awuondo developed a Microsoft Excel-based planning tool for the manufacturing footprint capacity of the company network. In the tool, he includes important concepts relevant to operational strategic planning, such as scenario planning with adjustable parameters. Among these parameters are growth rates and improvement rates, and the input into the tool is the demand forecast, which is converted into future space requirements. The present study builds upon the learnings and methods from Awuondo, but aims to expand upon the presented concepts by developing a model framework which can capture supply chain capacity holistically, beyond the pure focus on the space component of capacity. Furthermore, the capacity model proposed here aims to use advanced data analytics and database tools, and proposes a business process that can be fully integrated with the existing planning and data processes in the supply chain organization.

1.4 Open Research Questions

Measuring and modeling supply chain capacity is a complex, interconnected problem common to many manufacturing companies of all sizes. As such, best practices of how to design, structure, implement, and sustain an effective supply chain capacity model is of relevance to supply chain strategic planners within and outside of Boston Scientific. A key question here lies in how to design a capacity model that captures as much of a supply chain's complexity with a minimally sufficient data input. Because supply chain data is often incomplete and not well-structured, it is critical to minimize the amount of data required to generate as much insight as possible. Furthermore, additional complexity in capacity modeling stems from the scale and complexity of the underlying supply chain. It is therefore an important question of how to effectively scale a capacity model to adequately represent a large, complex, global supply chain.

Answering the question(s) above requires an effective data structure. This data structure needs to satisfy a range of requirements:

First, it needs to be scalable such that it can be used to represent a complex supply chain with twelve manufacturing plants and tens of thousands of products. These products, and the underlying production processes, differ significantly in nature, and the data structure needs to be able to represent these differences, for example in the degree of automation or unit of measure of all products.

Second, the data structure needs to offer flexibility for future changes to requirements. For example, new acquisitions or the introduction of new products may lead to features of the supply chain that have previously not been present, or manufacturing processes may change due to technological shifts. Furthermore, methods of data gathering, analytics and digital manufacturing are all evolving rapidly, and the data structure needs to be ready to handle any such changes.

Third, the data structure needs to be able to represent connections between previously disconnected data sets. What this means is that there is a challenge of seemingly disconnected datasets originating from different parts of an organization. A key research question is how to establish a data structure and the corresponding business process to appropriately link data sets that are disconnected *a priori*.

Fourth, the data structure needs to be comprehensive, in that it needs to capture all attributes required for capacity modeling.

Fifth and finally, the data structure needs to function with various degrees of granularity of data. For example, data granularity may vary with regards to the specificity of production lines, i.e., whether data is available for a production line as a whole, or for individual work stations within the line. Moreover, demand planning may occur at varying levels of the product hierarchy, where in some cases, a range of similar products is grouped into one category of demand for the long-range plan, whereas in other cases, the demand forecast is available for individual products. The data structure needs to be able to adapt to these varying degrees of granularity.

From an organizational and cultural perspective, an important consideration is how to manage the change introduced by developing a supply chain capacity model. The successful implementation of an enterprise-wide capacity model poses a challenging problem, even if all technical aspects of the model were to be fully solved: It requires successful change management to implement business processes where the required capacity data is gathered and provided in the correct format, and maintained over time. This is best achieved if all stakeholders understand the benefit of the new approach, and are presented with clear advantages for their own function when the transition is successfully implemented. With regards to the core stakeholders, namely the operations strategy team driving the process of building a capacity tool, it is important to adapt existing methods of planning and decision making to the capabilities of the new tool. Finally, developing and implementing an enterprise-wide capacity tool and process is a transition that requires significant time and resources, likely at the order of one to two years for the implementation, with a team that contains project management, software development, analytics, data engineering, and operations management capabilities. Therefore, a question is how to successfully obtain support for such a project from senior leaders, and how to manage the end-to-end process of implementation.

2. Problem Statement

Boston Scientific has identified a need to improve the systems and processes of determining long-term capacity requirements and constraints of the production network. As previously described, the company is on a trajectory of significant growth, and consequentially, there is a need for increased production capacity in order to meet this growing demand. Furthermore, there is an increasing frequency of supply chain disruptions due to the COVID-19 pandemic, environmental events as a consequence of the effects of climate change, as well as a changing global political and trade landscape. In order to meet these challenges, there is a need to proactively model capacity constraints and requirements across the supply-side of the network on a 5-year horizon. With such a model, it will be possible to identify future capacity constraints, and to detect components of the manufacturing network which will require investment in the future to provide adequate capacity under the given boundary conditions. Furthermore, the increased uncertainty described above ensues a need for scenario planning, where different "what-if-scenarios" can be analyzed to understand the costs, benefits, and risks of certain supply chain investments. To achieve this, a robust and comprehensive definition of capacity is needed, that includes key aspects such as space, equipment capacity, and labor. The organization currently performs capacity planning in a distributed fashion, mostly at the level of the individual manufacturing sites or even smaller units within the sites, with limited centrally available data and lack of standardization in the types and structure of data. Therefore, a key challenge is to overcome this fragmentation of information across the organization, and to aggregate the existing fragmented information in a way that becomes useful for a variety of stakeholders such as operations strategic planning, supply chain planning, and manufacturing. This serves two purposes: first, to gain a more holistic understanding of current and future capacity and second, to understand which additional data should be gathered for further network optimization. This would ultimately allow for more effective strategic planning, and inform critical network-wide decisions such as investments or the placement of products at appropriate manufacturing sites.

3. Approach

Determining a company's manufacturing capacity can be performed taking either a top-down or a bottom-up approach. In a top-down approach, highly aggregated data such as company growth projections, total production capacity in units per year, and total production space are used to determine future capacity requirements and constraints in the system. For example, in this approach, one would assume a percentage growth rate in company revenue, corresponding to a percentage growth rate in units sold per year. This unit demand then corresponds to a space requirement, which would grow with the same growth rate as the unit demand, minus an improvement rate of space utilization. Now the total required space in this calculation can be compared to the total available or planned space, to understand the needs for capacity investments at the highest level. This top-down approach corresponds to the status quo of capacity planning at Boston Scientific, and offers the benefit of a fairly simple analysis, which can be expanded in granularity by considering regional or divisional growth projections, or site-specific improvement and utilization rates.

This method of extrapolation is vague and offers very limited insight into the specific constraints that will arise in the future. A second approach to perform capacity modeling is through a bottom-up approach, where the network capacity is measured with a high degree of granularity, for individual production lines or even stations, and then the overall network capacity is assessed by aggregating over the granular capacity of each station and for each product, in the form of assessing the rate-limiting step of each path. Specifically, it is possible in this approach to directly determine production units with future capacity constraints. The bottom-up approach requires a large amount of data from the manufacturing network, and relies on the effective handling of this large dataset to be successful. By performing calculations at a granular level, this approach offers the ability to identify specific constraints in the large network, which will provide decision makers the ability to respond proactively to these constraints. Importantly, it is straightforward to obtain aggregated results similar to the top-down approach from the bottom-up approach, but not vice versa: It is straightforward to aggregate a granular model, but very difficult to add granularity to an already aggregated model. Ultimately, manufacturing capacity is measured and executed at an individual station or an individual line, so granularity is critical in order to truly capture future constraints and analyze potential future scenarios without requiring assumptions or heuristics that disaggregate data.

The drawback of the bottom-up approach is that it may give the illusion of precision through its granularity. The accuracy of capacity calculations depends on the precision of demand forecasts of the future, which are inherently uncertain. Even if it is possible to measure the ability of a network to produce product with a high degree of precision, knowing if this capacity will be sufficient is impossible. It is therefore important to understand the limitations of even a highly granular model, which leads to the important feature of scenario planning, which helps decision makers understand how changes in forecasts or of model parameters will change the outcomes of the calculations. To perform scenario-planning, or “what-if-analyses”, input parameters and assumptions need to be varied and the resulting changes in model outputs are observed. Parameters that need to be varied for successful scenario planning of capacity are the demand forecasts, i.e., growth rates and upside potential, improvement rates of cycle times and yields, and predictions of future acquisitions and R&D activity.

In this thesis, the approach of developing a capacity model based on the bottom-up paradigm outlined before consists of five individual components. First, the current state of capacity-related data is assessed. This serves the purpose to identify the type of data available, the gaps in existing data, and the data structure and business processes currently in place to manage capacity-related data. The capacity model proposed here is based on assembling a central data base that can house a variety of data from the entire Boston Scientific organization, including demand data and site-specific capacity data. Importantly, for this thesis, only internal capacity data was considered, as considering supplier capacity adds further levels of complication in terms of data availability and standardization. This was deemed outside the scope of this research. Therefore, the second component is to design a database that has the appropriate fields to capture sufficient capacity data for a network-wide capacity model to be functional. In addition, the database design also includes selecting the appropriate architecture which allows for efficient data handling as well as performing analytics and data manipulations efficiently. Third, a capacity model is designed, in the form of establishing a modeling logic that can convert raw capacity data into meaningful outputs, such as reports on capacity constraints and space constraints. To verify the functionality of the model design, the fourth component of this work is to perform a case study on an individual product. This will allow to verify that the proposed type of data as well as the proposed data structure can sufficiently represent the production capacity of the selected product. Furthermore, the case study approach enables the detection of potential gaps in current logic.

It is important to keep in mind that a case study, by nature, only captures a fraction of the reality of the complex global supply chain. It is therefore prudent to perform a gap analysis of missing details which the selected product or a case study in general cannot capture. Finally, based on the learnings from the case study, the gap assessment, and the overall data structure and model design, a scale-up plan is developed. This scale-up entails scaling the proposed model and logic to the whole production network, including establishing business processes which serve to collect, manage, and maintain the data.

4. Capacity Model Design

4.1 Current State Assessment

As a first step toward a capacity model that can capture the entire manufacturing network of the company, it was necessary to assess the current state of capacity data and systems.

4.1.1 Work Content Graphs

Work content graphs are a type of document used by Boston Scientific to measure the labor and equipment resources required to perform a set of production steps. They are usually created for a line or a set of workstations, and contain time studies which measure the amount of manual labor and equipment time required to perform each production step. As such, work content graphs contain information on the cycle times of the production process which they represent, with a focus on the manual labor content of these production processes. The information contained in a work content graph may be for one or multiple products.

Importantly, work content graphs represent a theoretical upper bound of capacity of a given production unit. They are based on time studies of the manufacturing process under ideal conditions and in steady state, and therefore exclude real features of the production process that increase the effective cycle time, such as scheduled or unscheduled downtime, line ramp-up and ramp-down, breaks, or any other deviation from a perfect process execution. In addition, cycle times are typically unyielded, and therefore need to be corrected for by the yield of each production step.

Despite these short-comings, work content graphs constitute the most widely used and most highly standardized type of capacity data available within the manufacturing network of Boston Scientific. As such, they are a promising starting point in collecting network-wide capacity data. In order to better understand the possibility to utilize this type of data as a source for capacity data across the network, the sites were surveyed about their use of work content graphs. The survey found that the use of work content graphs varies significantly from site to site, see **Figure 4-1**. For example, 27% of sites responded that they maintain work content graphs for all products, whereas one site reported that it does not have work content graphs for any products. For the majority of sites, namely 64%, work content graphs exist for 40-99% of all products, i.e., sites maintain work content graphs mostly for key products, but there are typically some products for which no work content graph exists. Furthermore, the survey and informal interviews have shown that there are strong variations across the network with regards to how frequently work content graphs are updated. A common approach is to perform time studies of all production processes once per year as part of a structured business process of continuous improvement. In other cases, work content graphs are only updated when production processes are changed, and are otherwise left unchanged. Overall, this implies that the quality of information that can be extracted from work content graphs varies strongly across the organization.

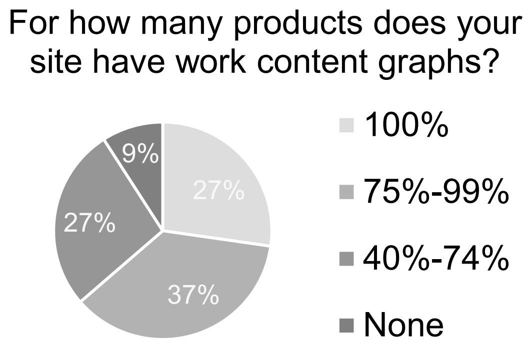

Figure 4-1. Survey result of site usage of work content graphs.

In addition, work content graphs only form one way with which sites manage their capacity. The role that work content graphs play in the day-to-day planning of capacity varies significantly, as does the way that work content graphs connect with the remaining tools for capacity planning at each site. For example, one site uses work content graphs as the source of truth for all of their capacity planning. At this site, the results from work content graphs are directly fed into a secondary capacity planning tool, which applies correction factors such the effective up-time of a line, the yield, and a general efficiency factor, to obtain the total capacity available to perform a certain production step. Due to the high importance that the work content graphs have at this site, they are generally updated frequently to ensure that data is current.

However, other sites do not use work content graphs with the same discipline as the mentioned example, creating an additional source of discrepancy between data quality between sites. And finally, even within sites, there is discrepancy in the use of work content graphs. Thus, a linkage is not necessarily straightforward between work content graphs of production processes that are, for example, sequential steps in producing the same product. In conclusion, work content graphs form an incomplete data set, that could not, by itself, serve to measure capacity across the network without improving the discipline of creating and maintaining work content graphs for all production processes across the network.

4.1.2 Site-specific Capacity Planning Tools

A key aspect to the processes relating to managing manufacturing capacity at Boston Scientific to-date is a decentralized approach to capacity planning. With very few exceptions, each site relies on their own local capacity planning tool. This approach has historical as well as organizational reasons. As the historical aspect, many sites in the Boston Scientific network have been integrated into the network over time through acquisitions. Each of these sites brought their own staff, tools, and processes, which included proven legacy methods how to navigate the capacity planning process. In some cases, a software that had been successfully used prior to an acquisition would continue to be used after the acquisition was completed. For the organizational aspect, sites operate with a relatively strong independence from each other and the central corporate management. Under this approach, sites are given the liberty to develop their own best practices in certain areas, which allows them to adapt to their specific needs.

Site capacity tools vary significantly with regards to the type of software, approaches to data structure, and the underlying definitions of capacity applied in these tools. For example, some tools are spreadsheets within Microsoft Excel, while other sites use dedicated stand-alone software solutions to manage their capacity. In terms of data structure, the tools show strong variations with regards to the data fields inside the tools, and the approaches used to calculate capacity. These variations make it challenging to define a standardized input format of capacity data, or to integrate locally housed capacity into a central database without significant manipulations to the format.

Moreover, similar to the use of work content graphs, there are large discrepancies with regards to the quality of data available, and the discipline with which data is collected and maintained. It was found that for some sites and products, data quality was very high, with thorough analysis of all steps of a production process. This was particularly the case for products that are of high interest to management because they exhibit fast growth, a high share of revenue, or require significant investments into capacity. However, for other products where interest was less, for example because their revenue stagnated or was comparatively low, the quality of capacity data available within local capacity tools was much less comprehensive. Other factors that affect the quality of capacity data are the cost and lead time of additional capacity.

In addition, it was found that with few exceptions, there was little formal connection between sites with regards to capacity planning. This means that in the case where multiple sites are involved in producing a finished product, for example by one site supplying components to a different site, there is little understanding by a downstream factory of the capacity of an upstream factory. In other words, while sites have a good understanding of the capacity within their own domain, they have a poor understanding of the capacity of their suppliers within the Boston Scientific network. This is problematic, as in this situation, sites plan their capacity as if they were a stand-alone system, whereas in reality, they are subject to constraints from other sites. This situation creates a lack of transparency into the true capacity of the supply chain.

Finally, even within the individual sites, there are variations to the degree of standardization among the different capacity tools. Some sites have implemented a capacity tool and best practices which are standardized across the whole site and adhered to. However, other sites show less discipline in how thoroughly locally standardized capacity tools are followed, or lack a standardized tool across the site.

The overall lack of standardization of capacity tools across the network, as well as of the underlying data, makes developing a network-wide capacity model impossible at the time this study is performed. However, generally, sufficient data to perform advanced, network-wide modeling of capacity exists within the company. It needs to be made useful through a process of loading existing data into a central database in a standardized format.

4.1.3 Control Tower

Boston Scientific runs a center of excellence which measures the day-to-day performance of the supply chain, internally called the Control Tower. The Control Tower is an example of the company moving toward a more data-driven way to operate, and to base decision-making on accurate, real-time data. The Control Tower generates data visualizations available to any member of the supply chain organization, which contain key performance metrics of the supply chain such as on-time delivery performance, inventory on hand, and a range of other business-critical metrics. The Control Tower data system is an example of a data-driven approach where information from separate sources is merged, such as from the company's ERP system, manufacturing execution system, production planning system, as well as unstructured data from other sources. The Control Tower's architecture is particularly interesting within the context of this thesis, since it has strong conceptual similarities with the required architecture of a network-wide capacity system. Data from a range of sources, and with different formats and structures, is automatically pulled on a daily basis. The data is then automatically cleaned and structured such that it can be stored in a relational database for further manipulation, analysis, and visualization.

As a recommendation, the existing infrastructure of the Control Tower should be leveraged when implementing and scaling up a network-wide capacity system, for the following reasons:

1) The database infrastructure of the Control Tower is similar to the required database architecture of a capacity model. Various capacity-related sources are already linked to the Control Tower, such as the supply chain planning system (see below), the bills of materials of all products, or the product

hierarchy of the ERP system. The Control Tower uses a cloud-based relational database system, which can form a blueprint for the capacity database.

2) Some of the data included in the Control Tower is identical to data that is required for capacity modeling. This data includes production planning data, which is synchronized from the supply chain planning system.
3) There is a large amount of existing know-how within the control tower team how to install an enterprise-wide database system with linkages into a multitude of software systems, as well as the handling of unstructured data. The team has experience on the required resources to implement sch as system, as well as the on how to perform front-end visualization and back-end data engineering. This experience can provide a significant advantage in designing and implementing the capacity database and model.

4.1.4 ERP System and Supply Chain Planning Software

Boston Scientific employs a company-wide enterprise resource planning (ERP) system to manage processes in human resources, finance, order processing, and procurement. In the context of the ERP system, products are sorted into a so-called product hierarchy with six levels, ranging from the division at the highest level of the hierarchy, groupings by so-called franchises and product families, down to the individual unique product numbers at the lowest level of the hierarchy. This product hierarchy is used in a range of planning processes in marketing, production, sourcing, and overall supply chain planning. The product hierarchy facilitates planning processes as planning at an individual unique product number is oftentimes not feasible. Products of a family or franchise often only differ by geographic localization such as labelling, or by other minor deviations which have a small impact on forecasting or medium-to-long-term planning. Therefore, planning processes of such products are simpler if they are carried at a franchise or product family level.

<u>SI&OP</u>

Boston Scientific uses a supply chain planning software system which is linked to the ERP system's data repositories, and contains a real-time image of orders, inventory levels, bills of materials, etc. In addition, the planning software contains demand forecasts for the next 18 to 24 months on a rolling basis. This raw data is then used by supply chain planning as well as production staff to carry out the Sales, Inventory & Operations Planning (SI&OP) process. During the SI&OP process, monthly demand forecasts are converted into production schedules, with the goal to balance inherent demand forecast fluctuations and thus create a level-loaded production schedule. This production schedule is referred to as the committed build plan, which is agreed upon by both the global supply chain organization as well as local site production engineers.

The SI&OP process is a tactical, forward-looking planning process with monthly cycles and a forecasting horizon of 18-24 months. Therefore, it inherently does not provide a complete view on longer term planning such as a five-year horizon of strategic planning. In terms of product granularity, the SI&OP process typically operates at a level of so-called representative parts, a form of product grouping that allows for demand planning across geographies etc. This poses the challenge of translating the representative part to actual units of production of the different underlying products, which is what is required for accurate capacity planning. Furthermore, the SI&OP process is focused on planning the supply of top-assembly products based on market demand. Sub-assemblies and components are planned manually as a consequence of the results of the SI&OP process for top-assemblies, and their planning is therefore highly reactive in nature. To streamline and accelerate the SI&OP process across the supply chain, it would be desirable to perform component and sub-assembly planning proactively, without having to wait for top-assembly planning to be complete. Implementing this process improvement is currently hindered by a lack of transparency into the supply chain, where upstream demand and capacity data is limited and only available as an output of finished-good-level planning. Therefore, the SI&OP process would benefit from the availability of more detailed, end-to-end supply chain data, in particular with regards to manufacturing capacity.

ERP system

The company is currently in the process of implementing a next-generation cloud-based ERP system, which serves as an opportunity to reshape the specifications of the ERP system based on changed needs, and to adapt business processes accordingly. In the context of this transformation, Boston Scientific has an opportunity to improve the data collection and maintenance of capacity-related data. For example, capacity data could become part of the data regularly collected from sites and stored within the data warehouse of the ERP system. This would have the benefit of converting the capturing of capacity data into a business process with strong support of the IT organization.

4.1.5 Long-Range Demand Forecast

Boston Scientific runs a process of long-range demand forecasting, which is driven by the Global Supply Chain center of excellence in collaboration with marketing teams in the commercial divisions. The purpose of the long-range demand forecast is to plan for future demand on a 5-year time horizon. The long-range plan carries out planning roughly at a product family or franchise level, and includes a degree of inconsistency with regards to the product hierarchy level at which this planning occurs. The data generally includes a nominal and an upside demand forecast, i.e., projections on the expected demand, as well as a percentage increase in demand which represents an upside potential beyond the nominal demand projections. This data is available on an annual basis (i.e., at the time of writing of this thesis, for the years 2022, 2023, 2024, 2025, and 2026).

Given that the long-range demand forecast includes nominal as well as upside demand forecasts, this data provides an opportunity for scenario planning and the analysis of what-if scenarios. Even with the raw data itself, the nominal and upside scenario are readily available as a basis for further what-if analysis, which would be based on adjusting the projected nominal and upside growth rates to understand how such changes modify associated capacity needs.

However, it is in the nature of forecasts that they include a large degree of uncertainty. Even with a comprehensive scenario planning approach, there is a significant chance that demand forecasts are entirely false and do not materialize. It is therefore important to also include unlikely extreme scenarios in the planning process, in order to allow for the development of contingency plans beyond the expected growth trajectory.

Overall, the long-range demand forecast is a useful dataset for the purpose of strategic capacity planning. While its accuracy is inherently associated with a large degree of uncertainty, and product grouping and naming is inconsistent within the dataset and with regards to other datasets, it forms a strong basis to understand future demand needs and can be easily stored in a central database for future analysis. The long-range plan therefore forms a key data source for the purpose of this project.

4.1.6 Overall Data Availability

Given the current state of data availability at the company, a network-wide capacity model is not feasible at this time. Besides the specific examples listed above, the following observations were made with regards to overall data availability.

<u>Lack of consistency of capacity data</u>

There is currently a lack of consistency among the capacity data that exists within the Boston Scientific network. As described above, it was found that each site uses a separate, locally created capacity calculation tool, which is commonly a spreadsheet in the software Excel. All tools are similar with respect to the high-level purpose and functionality: they aim to compare manufacturing capacity of individual business units or lines with the demand as reported from the global supply chain planning staff. However, these local tools differ significantly across the network, with variations along the structure of data, capacity definitions, and the business processes involved with creating, maintaining, and using the capacity tools. On the highest level, these capacity tools exhibit significant variations in the formatting and layout of spreadsheets, with differing structures of tables, or required columns of data between different tools. This means that even if these capacity tools contain similar information, aggregating this information, ideally in an automated fashion, across the whole network is challenging.

Furthermore, there exist differences in the processes of creating and utilizing capacity data among the sites. For example, some sites have thorough space planning processes, with detailed accounts for all use of space and regular updates to keep data recent. Other sites perform these actions ad-hoc as needed when major adjustments to the manufacturing footprint are performed. Furthermore, the format of data is fully inconsistent across sites, ranging from Excel spreadsheets to marked up images indicating space utilization. In addition, there are strong discrepancies with regards to data on cycle time, i.e., work content graphs, as described above.

These inconsistencies complicate the aggregation of capacity data across the network, and the lack of data for some products, processes, or sites, makes an overall network-wide collection of data impossible at the point of writing of this thesis.

Lack of consistency between different levels in the product hierarchy

There is a lack of consistency in planning demand and production capacity with regards to the level of the product hierarchy where planning is performed. The level where planning is performed ranges from franchises or product families down to individual SKUs. In particular, there is a lack of consistency between supply and demand planning: Oftentimes, the hierarchy level at which production capacity is planned does not align with the hierarchy level of the long-range demand forecast for the same product. In addition, the long-range demand plan is an ad-hoc process without a standardized format, commonly performed at a franchise or product family level, and sometimes the grouping of products even falls outside of the defined product hierarchy. In addition, the lack of consistency in the planning hierarchy level makes identifying products challenging due to inconsistent grouping. Individual SKUs have a unique product number, which could serve as a unique identifier for a product. However, the numbering system in the product hierarchy does not provide such unique identifiers to groups of products such as product families or franchises, nor is it possible to deduce a product group from a unique product number. This makes creating linkages between different levels of grouping impossible in many cases, unless data is manipulated manually based on domain knowledge of a data engineer. Specific examples of these are provided below, as well as in the case study in Chapter 5.

Use of decentralized systems

Additional challenges are caused by the lack of a centralized capacity planning system to date. Capacity planning occurs locally on an individual's computer, and not in a central tool with a central database. The lack of this central data system means that capacity data is not centrally accessible, which would simplify a central capacity planning tool. With the current state of decentralized data, in order to assemble a central capacity planning tool, it would be required to request capacity data from hundreds of different individuals for thousands of products, and then manually transfer this data into a central database. Such an effort requires significant resources and time, making it unrealistic at the current time.

Product nomenclature

There is a large discrepancy in the product names used within the organization, by individuals and within different datasets. Products may have three or more different names, such as a commercial brand name, a descriptor based on the medical indication for which it was developed, production names, and heritage names stemming from R&D working names or from prior to an acquisition. Furthermore, there may be additional variations to these nomenclatures due to the use of abbreviations, the use of colloquial product names, or inconsistent use of product sub-qualifiers such as for different sizes or product generations. These discrepancies are resolved only by human operators, as naming inconsistencies between different systems are not formally resolved in any kind of data system.

As an example of this highly complex and inconsistent nomenclature, the naming conventions in the Cardiac Rhythm Management and Neuromodulation divisions were analyzed. **Figure 4-2** displays the different naming conventions in the long-range demand forecast, and within the capacity planning tool used at a manufacturing site which will be called Site B for the remainder of this thesis. In addition, **Figure 4-2** contains the proposal for a grouping, which serves to resolve the discrepancies among the unclarity with product naming. As can be seen from the schematic, various products carry different names in the long-range plan and in the Site B capacity tool. For example, the product referred to in the long-range plan as Accolade is known as SB in Site B, where SB in turn is an abbreviation for Springboard. Such naming variations, while seemingly trivial, pose a stark challenge to the automation of capacity data analysis: both the long-range plan and the site capacity tool data sets refer to products by their respective name, and not some sort of unique identifier such as the unique product number. With the previous example of Accolade/Springboard/SB, there is thus no way for an automated system to know *a priori* that these products are in fact one and the same. The product naming discrepancies are further complicated by variation in product grouping between the different parts of the organization. This can be illustrated from the products of the Tachy franchise, and to an even graver extent, from the products in the Neuromodulation division. In the Tachy franchise, it can be seen that the long-range plan breaks out the NG3 and NG4 products, whereas the capacity tool at Site B only lists the NG3 product. In this specific case, this is caused by the development history and the underlying technical details of the product: The capacity tool which was inspected in-depth here stems from a business unit producing a certain component of the Tachy device. As it turns out, the two generations of Tachy productions referred to by the marketing organization as NG3 and NG4 share the identical version of this component, and thus, the manufacturing site refers to them as the same product, NG3. Again, there is *a priori* no reason to assume that NG4 and NG3 are the same product, and it was necessary to manually create this connection and group these products to accurately represent them in a capacity model. These product grouping discrepancies are even more prominent in the Neuromodulation division. Here, the commercial organization only delineates two types of devices, deep

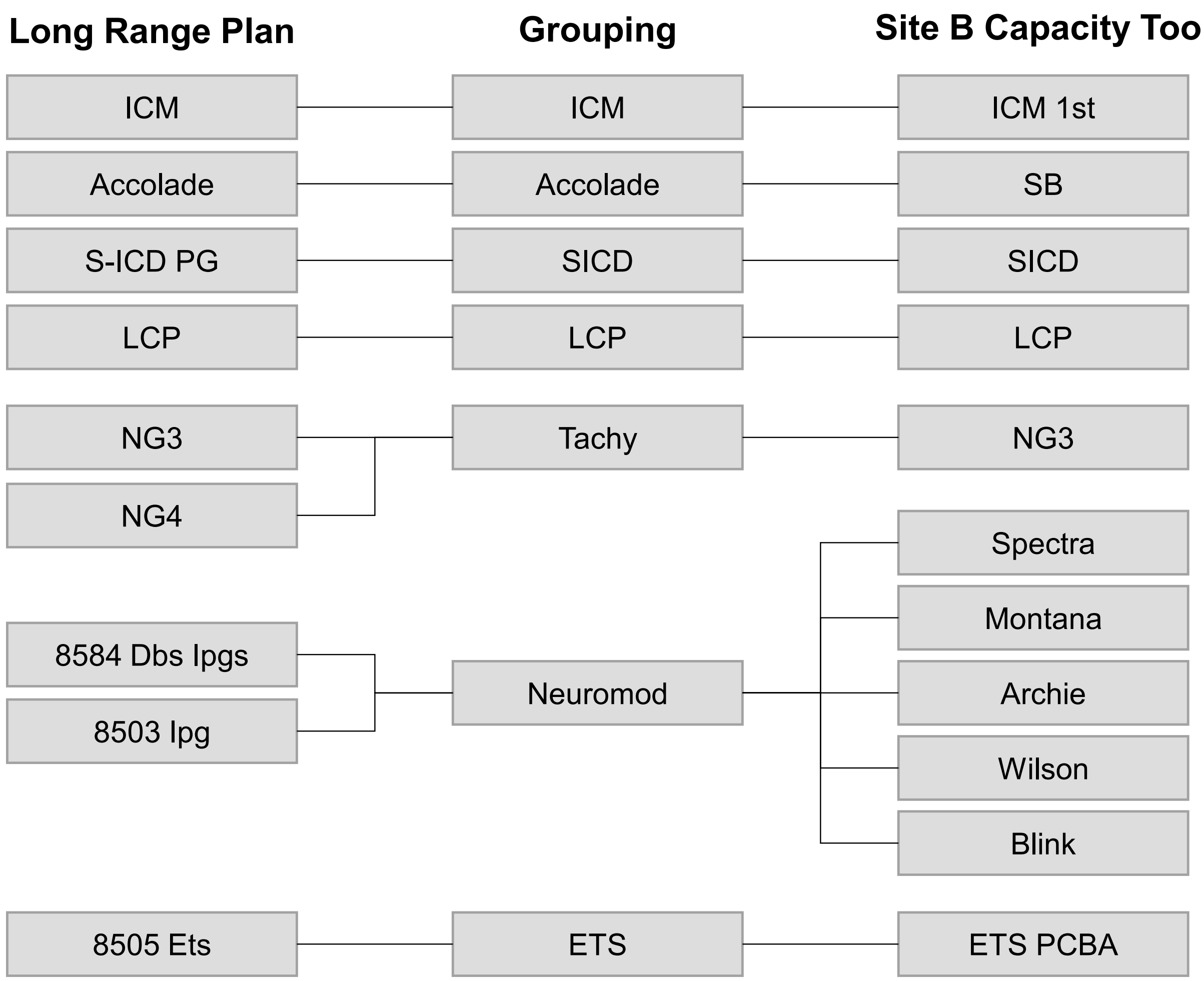

Figure 4-2. Example of inconsistent naming conventions and product grouping in the Cardia Rhythm Management and Neuromodulation divisions.

brain stimulators and spinal cord stimulators. Both are not referred to by these names in the long-range plan, and instead have the titles “8584 Dbs Ipgs” and “8503 Ipg”, respectively. In contrast, the manufacturing site at Site B performs planning based on five different product groups, namely Spectra, Montana, Archie, Wilson, and Blink. These products all exist in manifestations for deep brain stimulation and spinal cord stimulation, under the same Site B names. Instead, the different names refer to different generations of the medical devices, as well as different form factors and rechargeable/non-rechargeable configurations. They can therefore not be linked with the product types under which long-range planning is performed. In this remarkable example, the smallest possible grouping linking the production planning and demand forecasting datasets is the Neuromodulation division itself.

Again, the most significant feature of this naming inconsistency is that it is not *a priori* clear or discoverable without knowing about the linkages between different data sets. What this implies is that in order to successfully build a full-scale model of the entire supply chain of Boston Scientific, there will be the need for significant efforts in data cleaning requiring the input of subject matter experts familiar with the different products and their varying names across the organization. In Chapter 4.4, an approach to develop a dictionary and product grouping system to link various product names is proposed, which could be used to implement a systematic data structure of product names and their synonyms.

4.2 Unit of Measure

It is an inherently difficult problem to measure manufacturing capacity for a large and complex manufacturing network. The specific case of a medical device manufacturer with a diverse product portfolio adds to this complexity due to the large variety of types of products produced. In particular, in this complex system there are stark variations with regards to the definition of what defines a single unit of a product. For example, there are products which can simply be counted, i.e., the unit of measure is the number of individual copies of the same item. Beyond this most simple case, there are products, in particular raw materials and components, that measured in units of length, such as wires, braids of wire, polymer tubes. Originating from these components and sub-assemblies measured by length, there are products which differ only in their length, implying that a major difference is the amount of material taken from a continuous source such as coil. In a simple numerical example, one could imagine an assembly with three different lengths of 50 cm, 75 cm, or 100 cm. From a coil with raw material of length 1500 cm, one could therefore cut 30 pieces of 50 cm, 20 pieces of 75 cm or 15 pieces of 100 cm, or any combination thereof. Therefore, if the production system has a certain capacity to produce coils with 1500 cm of raw material, it is not clear, what the capacity of the network is to produce finished goods – the result depends on the required mix of lengths for the final product. Other potential units of measure for product quantity are by volume, i.e., liters or fluid ounces of a liquid, or by packages of individual items.

Additional challenges to measuring capacity originate from differing product sizes and quantities produced. Product and component sizes can range from the sub-millimeter range to large capital equipment with dimensions at the order of a meter. Simply counting the number of products produced cannot capture the difference in scale and complexity of producing an individual small component such as a piece of stamped metal with thousands of units produced per day, versus the assembly of a complex, large piece of capital equipment, where quantities are in the range of a few individual units per day.

Mixed model lines, which are discussed in further detail in the following section, add an additional dimension of complexity to the problem of accurately measuring capacity. Ultimately, to resolve the challenges associated with measuring capacity for different types of products and different units of measure for product, it is necessary to convert all production to the required production time. Production time presents a form of measurement that is comparable and clearly defined for all types of products. The details of this are described in Section 4.3 on mixed model lines, as the concept is closely linked to treating multiple products within a single line.

4.3 Mixed Model Line Modeling

Mixed model lines are production lines where the same production assets produce different types of products. In general, mixed model lines can take the shape of a line where the same tooling is used to produce different products, or where retooling occurs between batches of different products.

Allocating production capacity to the different products passing through a mixed model line is not possible if capacity is measured as units produced in a certain time period. This is because the product mix changes how much capacity is allocated to each product. As an example, a production line produces products A and B, and product A has an effective cycle time of 100 s, and product B has an effective cycle time of 20 s. If the factory runs one 8-hour-shift per day, then the line could produce 288 units of product A per day if it only produced product A, or 1440 units of product B per day if it only produced product B. However, since the line produces some mix of products A and B, the capacity in units per day for each of these products depends on how much of the other product is produced. It is therefore not possible to provide a definitive, generally valid number in units per day for products A and B. The solution to this problem lies in redefining how capacity is measured, namely in terms of the ability to produce what is demanded. For the same numerical example, assume that we have a daily demand for product A of 200 units, and for product B of 300 units, and we still have one 8-hour shift available to produce products A and B. The total time it would take to produce all 200 units of product A and all 300 units of product B is 26,000 s or 7.2 hrs. This required time to produce is less than the time we have available, and thus there is sufficient capacity to produce both products A and B. However, it is also obvious that the required production time is close to the available time, so if there is significant growth in demand for either product, there is the expectation that capacity will soon be limited. In that case, it would be necessary to add an additional shift to increase the available time, or to reduce the effective cycle time of product A or B or both e.g. through the addition of additional equipment or through process improvements. Note that inherently, there is not more capacity of product A or product B. If the overall available time were insufficient to produce enough of both products to meet demand, it is not clear a priori whether product A or B should be given priority. This is a decision to be made by management when constraints manifest.

The lesson from this example is that the appropriate way to measure capacity is by calculating the required time it takes to produce the demanded products, and to determine whether this amount of required time is above or below the available time to produce a product. If this analysis is performed in the appropriate time intervals (e.g. on a quarterly basis) based on demand forecasts, one can then identify current or future capacity constraints. It is also worth noting that management oftentimes prefers to be presented with summaries such as "we have the capacity to produce 10,000 units of product A and 50,000 units of product B per year" since this measurement is much easier to grasp than a pure time-based analysis. Based on the previous discussion, it is therefore very important to implement a cultural change where it becomes clear that the problem is not adequately captured by presenting a simple set of numbers in units per year for mixed model lines. Senior management needs to accept and drive the awareness that capacity is equivalent to the ability to produce

4.4 Database Fields

To develop a model of network manufacturing capacity, it is proposed to first implement a central database capturing capacity data from across the network in a standardized format. **Figure 4-3** provides an overview over the proposed database tables and fields, each containing relevant attributes required to calculate network capacity. The capacity database is proposed to have the following tables, which will be described in more detail below: Raw Capacity Data, Cycle Times and Available Times, Demand, Space Map, Space, Product Map, Product Synonyms, Product Groupings, and UPN List. The table fields described below represent the columns of each of these tables, and each row represents a data point within the table. It was found that this data represents the minimally sufficient set of data required to accurately capture capacity.

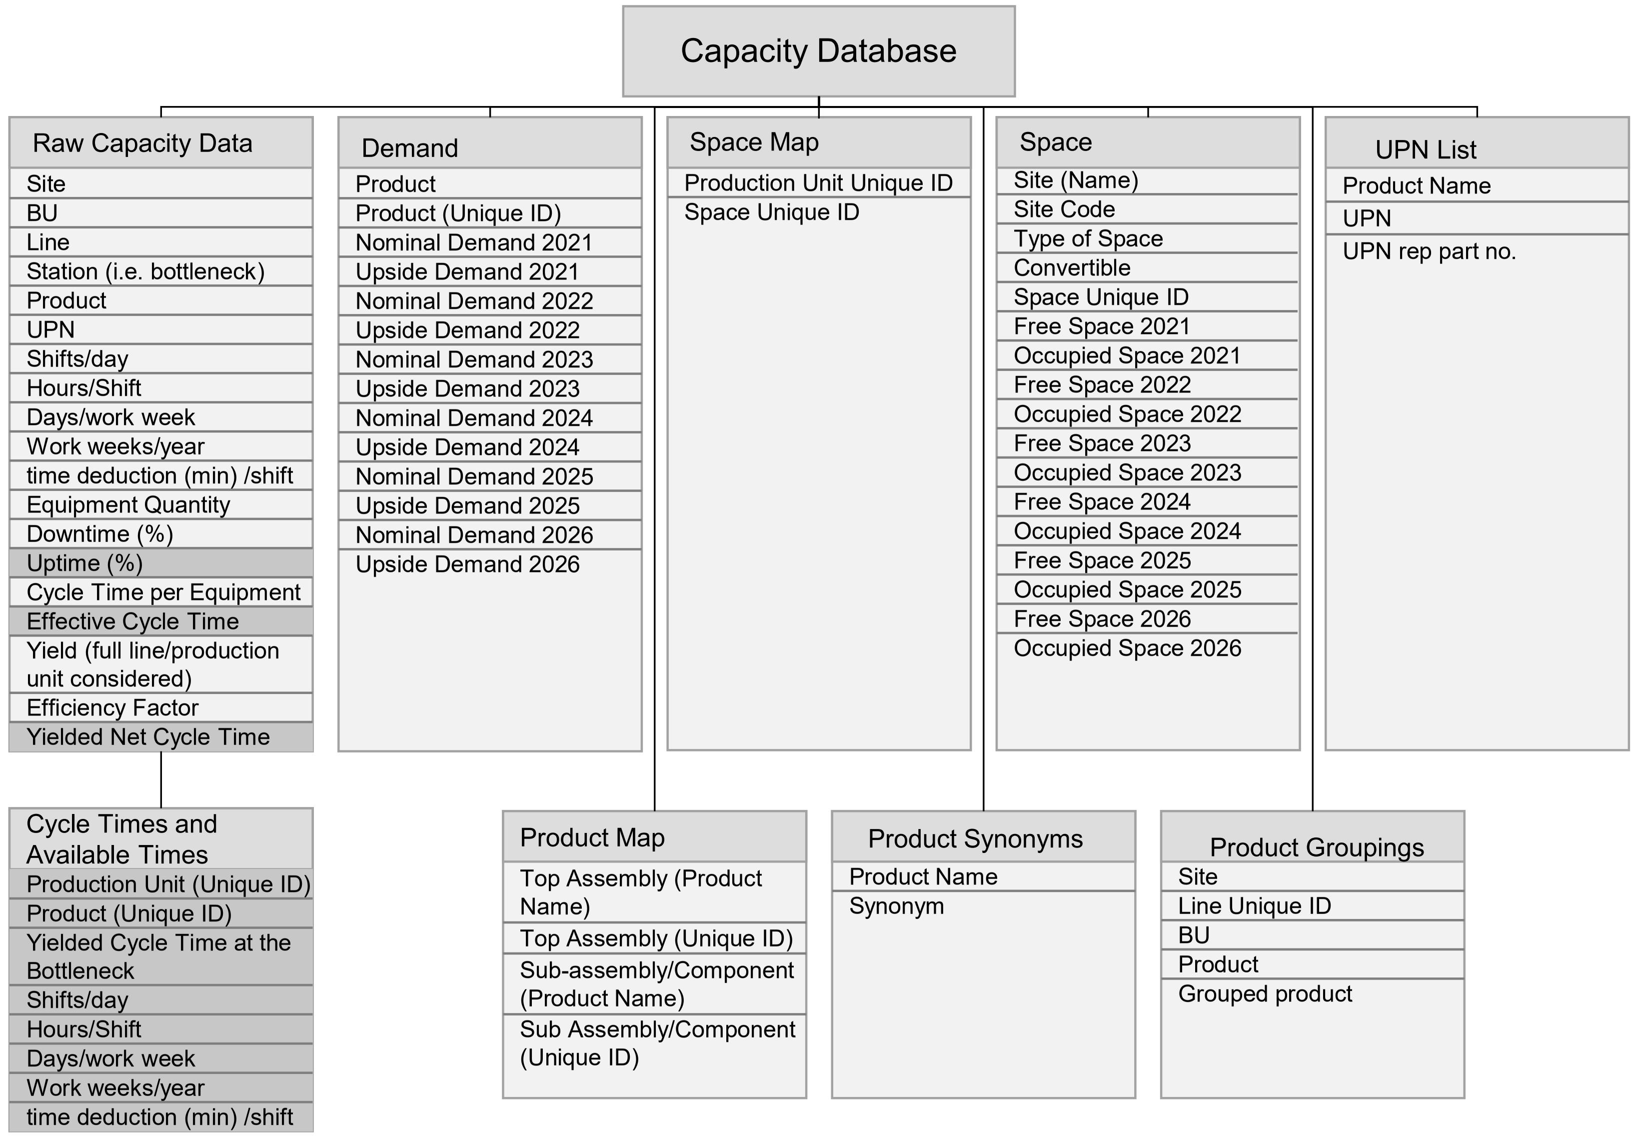

Figure 4-3. Overview over database fields for a central capacity database. Fields highlighted in dark-grey are calculated values based on other input values.

Raw Capacity Data

The table Raw Capacity Data contains attributes of production capacity for products and sites. An entry in this table consists of attributes identifying the site and production unit (table fields: Site, BU, Line, Station) and the product (table fields: Product, UPN). Capacity-related data includes the shift structure (table fields: shifts/day, hours/shift, days/work week, work weeks/year, time deduction/shift), and production-related attributes (table fields: equipment quantity, downtime in percent, uptime in percent, cycle time per equipment, effective cycle time, yield, efficiency factor, yielded net cycle time). Here, uptime is calculated based on the downtime, the effective cycle time is calculated based on the cycle time per equipment, the uptime, and the equipment quantity, i.e., the number of machines performing the same task, and the yielded net cycle time is calculated based on the effective cycle time, the yield, and the efficiency factor.

It is worth noting that the capacity logic should function in the two cases where only one cycle time is provided for an entire manufacturing line (i.e., the cycle time at the bottleneck), and where the cycle time of production units smaller than a full line are available.

Cycle Times and Available Times

This table consists of calculated values based on the input values of the Raw Capacity Data. In combination with the demand data, the table Cycle Times and Available Times represents the relevant quantities required to perform the analysis of required time and available time. Table fields include an identification of the production unit, the product, the yielded net cycle time, and the available time (shifts/day, hours/shift, days/work week, work weeks/year, time deduction/shift).

Demand

The Demand table contains data from the long-range demand forecast in a flattened and standardized form. The product is identified through its name and unique identifier, and for each year, the nominal and upside demand are represented in the table.

Space Map

The Space Map table serves to link production units with the space in which they are located. The production unit is represented by the production unit unique ID field, and the space is represented by the space unique ID field. By linking a production unit, which could be a line or a production station, with the space (i.e., room) in which it is located, analyses of space-related capacity becomes possible.

Space

The Space table captures a range of parameters relevant for capacity planning of production footprint. A set of attributes serves to uniquely identify a unit of space (table fields site (name), site code, space unique ID). Here, units of space can range from whole sites, to buildings, business units, and rooms. The space unique ID needs to be designated to each space, and should capture this hierarchy of space types. In addition, the table contains information about the type of space, and a Boolean value indicating whether the space is convertible. The rest of the table consists of the planned occupied and free space for the years considered, namely 2021 through 2026 for the purpose of this study. The total space of a unit is simply the sum of occupied and free space.

Product Map

The Product Map table serves to represent the product bill of materials in a flattened format. The attributes of this table are a top assembly name and unique ID, and a corresponding sub-assembly/component name and unique ID. Importantly, each sub-assembly or component for the same top assembly requires the entry of a new data point. Through this, a flattened data structure in the relational database is created.

Product Synonyms

The Product Synonyms table links a product name to a synonym of the same name. This serves as a dictionary that can help resolve the challenge of inconsistent product nomenclature across the organization. Again, the data is represented in a flattened format, i.e., for each synonym of the same product, a new row needs to be created.

Product Groupings

The Product Groupings table serves the purpose to link products to a product group. This table serves to formalize the observations made for products in the Neuromodulation division, where different parts of the company used different groupings to perform planning processes, see section 4.1.6, product nomenclature. The table contains fields to identify the line and site where a grouping occurs (table fields site, line unique ID, BU). This information is required because groupings may vary from site to site or even from line to line. The other fields identify the product and the group to which the product belongs.

UPN List

The proposed UPN List creates a connection between a potentially non-unique product name, and the unique product number. In addition, the representative part number is required as required as an attribute since it represents a more generally valid identifier than the UPN for planning purposes.

4.5 Capacity Calculation Logic

In the following section, a logical flow of capacity calculations is proposed, which serves to convert capacity data from various sources into outputs delineating capacity constraints in the manufacturing network, **Figure 4-4**. On the highest level, calculating network capacity requires a set of inputs from various sources, which are being fed into a set of calculations converting these raw input data into measurable capacity information. These calculations yield results which need to be aggregated and properly presented as outputs.

Inputs

A range of input data from different sources and in various formats is required in order to capture and calculate the capacity across the network of a complex manufacturing supply chain. Capacity data also necessarily requires the availability of demand data, in order to properly capture the product mix of mixed model lines, as described in section 4.3. The capacity data required for the proposed calculation logic includes the product hierarchy and information on bills of materials for all products. This information is critical to capture the flow of product through the supply chain, and to calculate supply and demand of sub-assemblies and components for a finished good. In addition, this data set includes information from the established product hierarchy, i.e. various levels of product grouping and aggregation. This becomes relevant to capture that different parts of the organization perform planning tasks at different levels of the product hierarchy. A second required data set is a complete list of production lines and individual stations on those lines, including the information of which products are being processed on each line. This data set constitutes a list of physical assets, and links them to the respective products. Importantly, for reasons which will be further elaborated in the case study section, it is desired to obtain a maximum level of granularity for this data: If available, production line data should be broken down to the level of individual production stations because for mixed model lines, the production bottle neck of each product may be a different station for each different product passing through the line. Further, capacity calculations require the cycle times of each product on each line, in the form of Design Units per Hour (DUPH) as determined by the work content graph. This data measures how much time is required, in steady state and under idealized conditions, to perform the bottle neck operation, and thus defines the output frequency of a production line. In order to

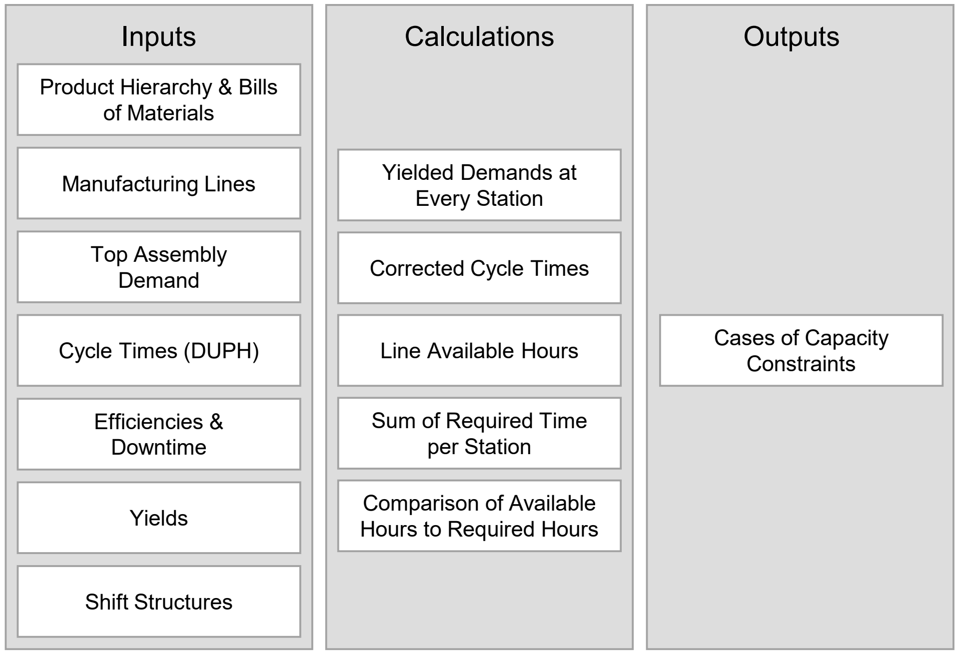

Figure 4-4. Overview over the capacity model calculation logic, feeding a range of inputs into a calculation logic, yielding outputs of capacity constraints detected in the manufacturing network.

determine the real production output of a line, efficiencies, downtime data, and production yields are required. With the help of this data, the theoretical cycle times or DUPH can be corrected to determine a net cycle time. With regards to data on the available time to produce product, shift structure data is required. Again, this data should be included with the highest possible level of granularity since shift structures vary even from line to line or business unit to business unit within production sites. Averaged or typical shift data could distort the calculated results, and cover up cases of capacity constraints that are specific to a specific line or station. Finally, as mentioned before, the capacity calculations require the input of demand data, which is typically provided in the form of the long-range demand forecast of top assembly products. Here, demand data exhibits variations at the level of the product hierarchy where planning is performed, and thus it is important to have the appropriate linkage to the product hierarchy to connect demand data with supply side data such as cycle times or production lines. In the model developed here, it is assumed that the data is available for the next five years. Importantly, this requires that for yield data, downtime, cycle times, etc. in future years, improvement rates have already been included in calculating this data as an input into the model. The input data fields were described in section 4.4, where the format of all tables and table entries is described to capture the data described here. The specific process of capturing such data is described in more detail in chapter 5, where specific case studies are presented with numerical examples.

Calculations

The input data needs to then be made useful by performing a sequence of capacity calculation steps, which convert raw input data into an output of detected capacity constraints. Further detail on the specific calculation steps is provided below. It is important to note that all calculations need to be performed for all future points in time of interest, for example on a quarterly or annual basis.

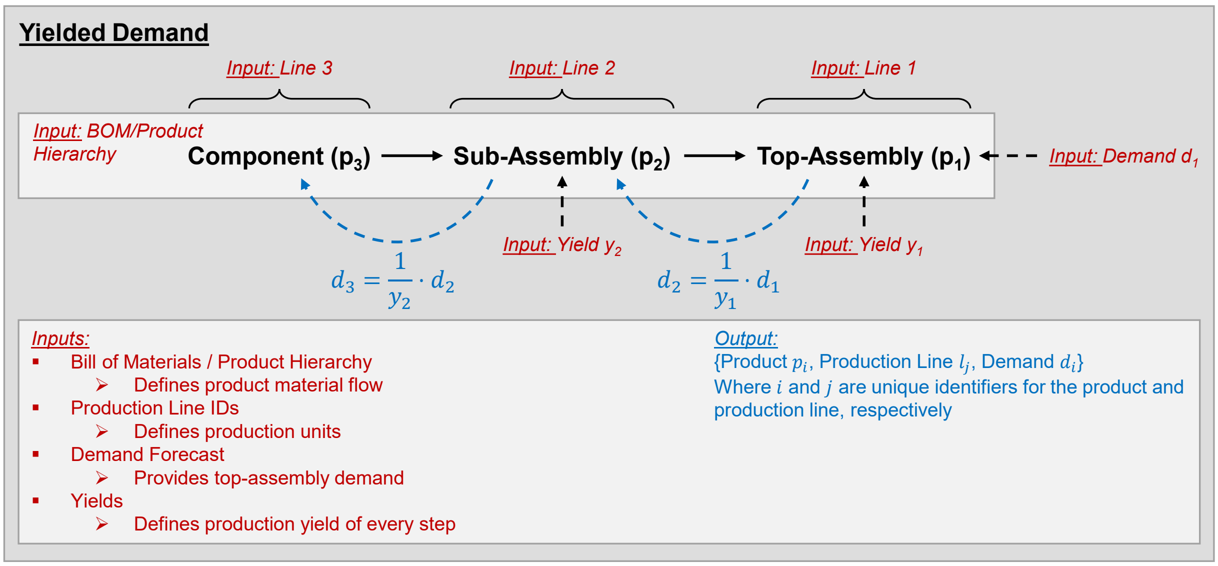

Figure 4-5. Calculation of the yielded demand of products upstream in the supply chain, based on the top assembly demand from demand forecasts, and on yields of each step along the manufacturing supply chain.

Demand data as provided from the long-range demand forecast inherently includes only the demand of finished goods, i.e., top assemblies. Since a complex supply chain such as the one of Boston Scientific represents a pull system, where sub-assemblies are serving as supply to the top assembly production, and components feed the sub-assembly production, each of these supply chain echelons needs to meet the demand of the layer which it supplies. This implies that the demand of sub-assemblies feeding into the top-assembly production line has to be corrected for the yield of the top-assembly production: The supplying layer of the supply chain has to provide sufficient input for the starts at the next layer, not just for the successful finishes.

Figure 4-5 schematically depicts the calculation of the yielded demand for each sub-assembly and component based on the yields of each production step and the demand of the top assembly. These calculations need to be performed for each top assembly, and for each sub-assembly and component within the supply chain of each product, respectively. At each layer, the demand of the supplying layer is given by

$$d_{l+1} = \frac{1}{y_l} \cdot d_l$$

Where d_{l+1} is the demand at the supplying layer, and y_l and dl are the yield and demand at the pulling layer, respectively. The calculation outputs a set of tuples of the form {Product p_i, Production line l_j, Demand d_i}. this means that for each product, which includes all top-assemblies, sub-assemblies and components, and for the corresponding production line, a demand value is calculated.

Corrected Cycle Times

$$Net\ Cycle\ Time = \frac{1}{yield} \cdot \frac{1}{up\ time} \cdot \frac{1}{efficiency} \cdot \frac{1}{DUPH}$$

Inputs:

- DUPH (Design Units per Hour)
  - Theoretical, optimal cycle time of a process
- Yield
  - Ratio (in %) of successful completes out of process starts
- Up Time
  - Ratio (in %) of time where equipment is running (1 – down time)
- Efficiency
  - Efficiency (in %) that corrects for other sources of reduced throughput

Output:

Net Cycle Time (in hours), calculated for each production step, and each product (including top-assemblies, sub-assemblies, and components)

Figure 4-6. Calculation of corrected cycle times, by correcting ideal cycle times, i.e. the inverse of Design Units per Hour (DUPH), with the yield, up time and efficiency of a process.

It was found that the most commonly available source of capacity data stems from work content graphs, which include measurement data on the theoretical, ideal cycle time of production steps in the form of design units per hour (DUPH). The cycle time of a production step is simply the inverse of the DUPH. Since this data represents an idealized state which does not account for imperfections and inefficiencies in the production process, it needs to be corrected for these imperfections, as schematically depicted in **Figure 4-6**. Here, the yield represents the fraction of successful finishes of a production step out of all starts. The up time is the fraction of time where the production process is actually running; most commonly this is equivalent to the equipment running. The up time in percent is 1 – down time, where the down time is the time fraction where the production process is stopped, for example due to equipment failure. Finally, the efficiency factor measures all other sources of inefficiency in the production process that could increase the cycle time relative to the ideal cycle time of the process. For example, this could include ramp-up and ramp-down phases after and before stoppage or breaks, or deviations from the ideal cycle time from newly trained workers.

Line Available Hours

$$\text{Available Hours (per year)} = (\text{hours per shift} - \text{hours of break per shift}) \cdot \#\text{shifts/day} \cdot (\#\text{work days/week} \cdot \#\text{work weeks/year} - \#\text{holidays})$$

$$\text{Available Hours (per week)} = \text{Available Hours (per year)} / 52$$

Inputs:

- Shift Structure:
  - Working hours per shift
  - Hours of break per shift
  - Shifts per day
  - Work days per week
  - Work weeks per year
  - Number of holidays (i.e. days of plant shut down) per year

Output:

Line available hours per year

Line available hours per week

Figure 4-7. Calculation of line available hours based on shift structure and other work time data.

Next, it is necessary to calculate the amount of time available at a factory for production. Importantly, there are two approaches to compute the available time: in the first approach, the current shift structure is used to obtain a total amount of time at which the production line is running. In the second approach, an ideal state of factory utilization is assumed, for example 70% utilization during steady-state production, with an additional 20% of utilization available for surge capacity and 10% available for maintenance. This corresponds roughly to a 24/5 shift structure, i.e. 24 hours of production during the 5 workdays of a week. For various reasons, many factories, or parts of factories, deviate from this ideal state, either by exceeding or by falling below a 70% utilization: While some production units are heavily utilized with utilization close to 100% (i.e., 24/7 operation), other production units only operate one or two shifts, corresponding to utilization rates well below 50%. Since measuring capacity ultimately reduces to comparing the amount of time required to produce the demanded products with the available time, it is important to understand the current available time based on the shift structure, corresponding to the first approach. This calculation is presented in **Figure 4-7** and explained here: The hours available per day are the length of a single shift, net of breaks multiplied with the number of shifts per day. Then, the hours of each day need to be converted to total available time per year, by multiplying with the number of work days per week, and the number of work weeks per year. Finally, the total number of available days needs to be corrected by any holidays where the manufacturing unit is shut down. It may be useful to convert this number of annual hours to the average number of hours per week, by dividing by 52. This allows for easy comparison with the described ideal state based on desired utilization factors.

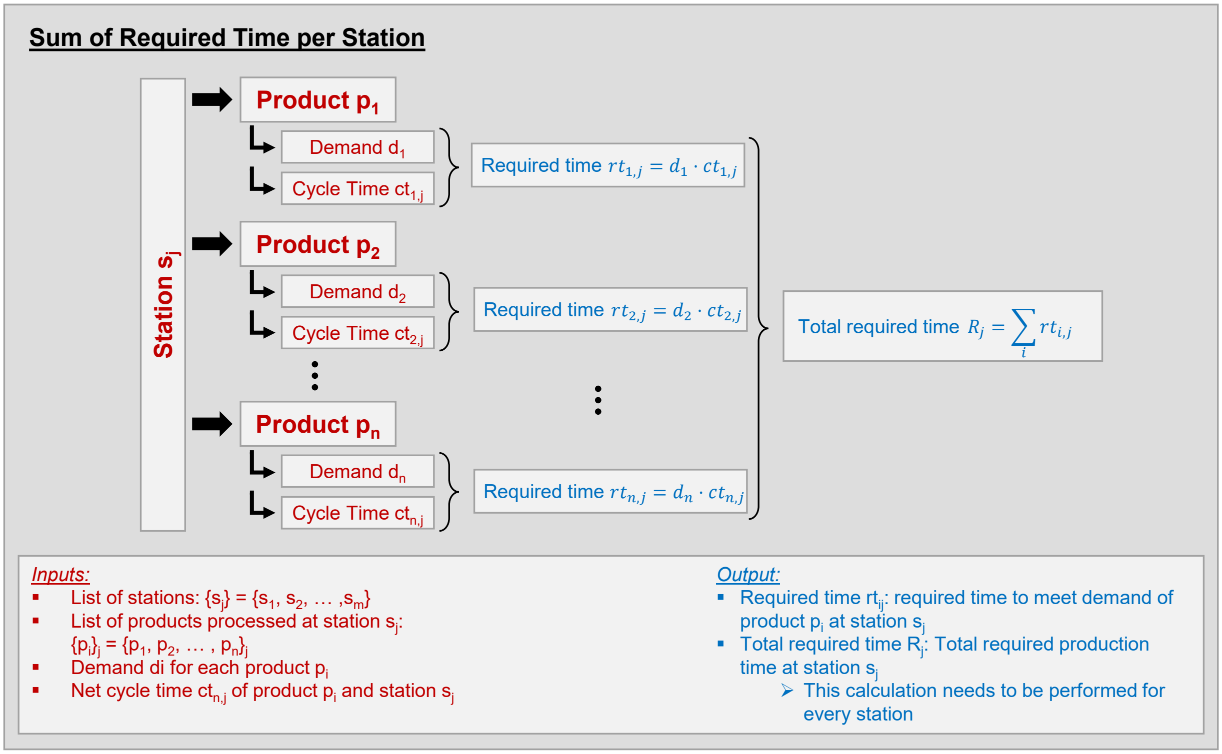

Figure 4-8. Calculation logic for the total sum of required time.

As described, the goal of the capacity model is to compare the available time to the required time at the most granular level of a production unit which the data can provide. In the most granular case, this occurs at an individual station level. Mixed model lines are a common feature for the production network of Boston Scientific, so the total required time at a station or line is the sum of required time for each individual product produced by the production unit, as depicted in **Figure 4-8**. For each product p_i processed by station s_j, the required time to produce that product is calculated by multiplying the previously calculated net cycle time with the yielded demand at this stage of the supply chain. This individual required time is then summed for all products to obtain the total required time at the station. Again, depending on the granularity of calculation (based on available data), this calculation can occur either at an individual station level, or at a line level, where only the bottleneck is considered. However, the bottleneck may not be the same for each product, so calculations performed at the individual station level will generate the most accurate results.

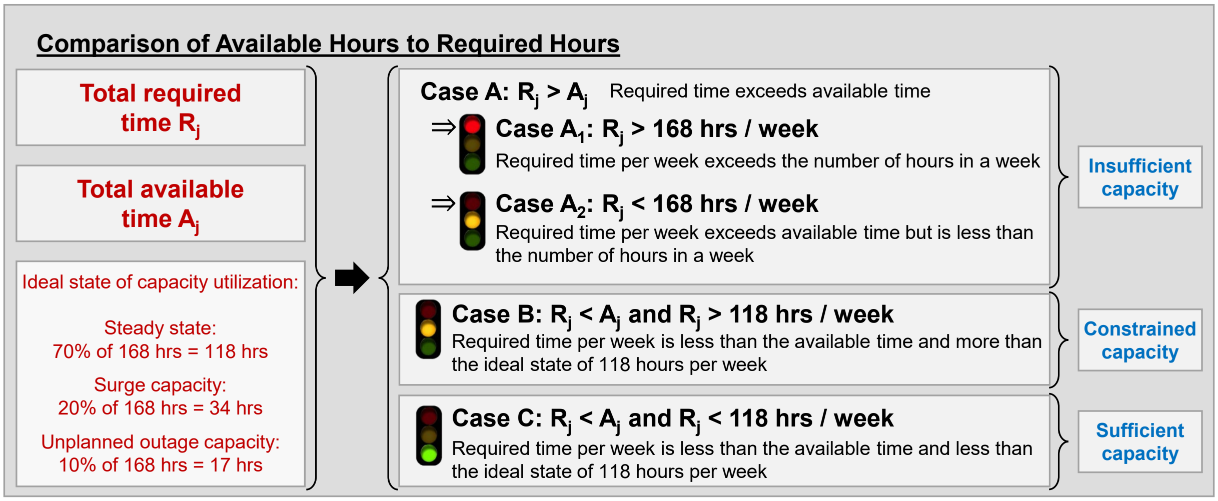

Figure 4-9. Comparison of available hours to required hours.

Finally, based on the previous calculations, it is now possible to actually assess the capacity of a production system such as a station or line. Assessing the capacity of a production system is equivalent to comparing the required time to the available time. In the proposed model, it is assessed whether the current capacity based on the actual shift structure is sufficient to produce product, and also whether the required time meets the ideal utilization goals of Boston Scientific, where shift utilization should be less than 70% to have available surge capacity and time available for maintenance. This logic is described in **Figure 4-9**, and broken down in three possible cases, A, B, and C. In case A, the required time R_j exceeds the available time A_j. If this is the case, the current capacity is not sufficient to meet demand. This scenario can in turn be broken down into two sub-cases. In case A_1, the required time exceeds 168 hours per week, i.e., it exceeds the total number of hours a week has, so it is physically impossible to produce all products required. This means that a hard capacity constraint exists, which requires some sort of investment to be alleviated. In the case A_2, the total required time is above the current available time, but less than 168 hours per week. In this case, the current shift structure does not provide sufficient capacity to meet demand, but it would be physically possible, through additional shifts or overtime, to meet demand with the currently available equipment and space. Both cases A_1 and A_2 correspond to insufficient capacity of the production system. In case B, the required hours fall below the available hours, but the required hours exceed 118 hours per week. In this case, the level of capacity utilization is higher than the desired ideal state of 70%, which means that there is limited capacity for surge. Therefore, case B corresponds to constrained capacity. Finally, in case C, the required time falls below both the available time and the desired state of 70%. In this case, there is sufficient capacity to meet demand with the current configuration of shift structure and equipment.

Outputs

The calculated results need to be aggregated and presented to decision makers, in order to make the useful and convert findings into actions which respond to detected capacity constraints. The proposed model determines capacity constraints for each production unit, using the cases outlined above, where capacity constraints are broken down into insufficient capacity and constrained capacity with different levels of urgency. These cases need to be aggregated and presented visually in order to extract information and inform decisions.

4.6 Space Logic

The model described in Chapter 4.6 outlines how capacity constraints are detected. As a next step, the model then needs to determine whether these capacity constraints can be alleviated by adding additional shifts, or if space additions are required. This schematically depicted in the flow diagram of **Figure 4-10**, which needs to be executed for all spaces/products, and for each planning cycle such as for every quarter or year for which the model should make predictions. At the start of this logic sits the output from the general capacity model, which produces cases of insufficient line capacity, i.e., where the amount of required production hours exceeds the available hours. By including information on the shift structure, the next step is to determine whether there is an ability to add an additional shift to increase the available hours. This can be achieved when a line is not yet operating at the ideal state of 70% utilization, i.e., a 24/5 shift structure (for simplicity, it is assumed that a 24/5 shift structure can be implemented at all factories in the network, ignoring potential labor-related issues that might arise). If a shift can be added, this should be done in order to meet capacity. If a shift cannot be added, this implies that additional space is required to increase capacity to the required level. At this point it is necessary to determine the additional space requirement to add capacity and alleviate the constraint. This should be done through a calculation using a *space scaling factor*, which yields the following equation:

$$a = c \cdot s \ \cdot (r - 1)$$

Where a is the additional space required, c the current space footprint of the line, s the scaling factor described below, and r the ratio of required hours to available hours.

The scaling factor is defined as follows, and based on the following rationale: doubling the output of a production unit does not necessarily require twice the space, as it may not be necessary to duplicate the full line. Instead, only parts of a line, namely the production steps with the lowest throughput, would need to be duplicated. So effectively, adding twice the production output may only require, say, 25 % more space,

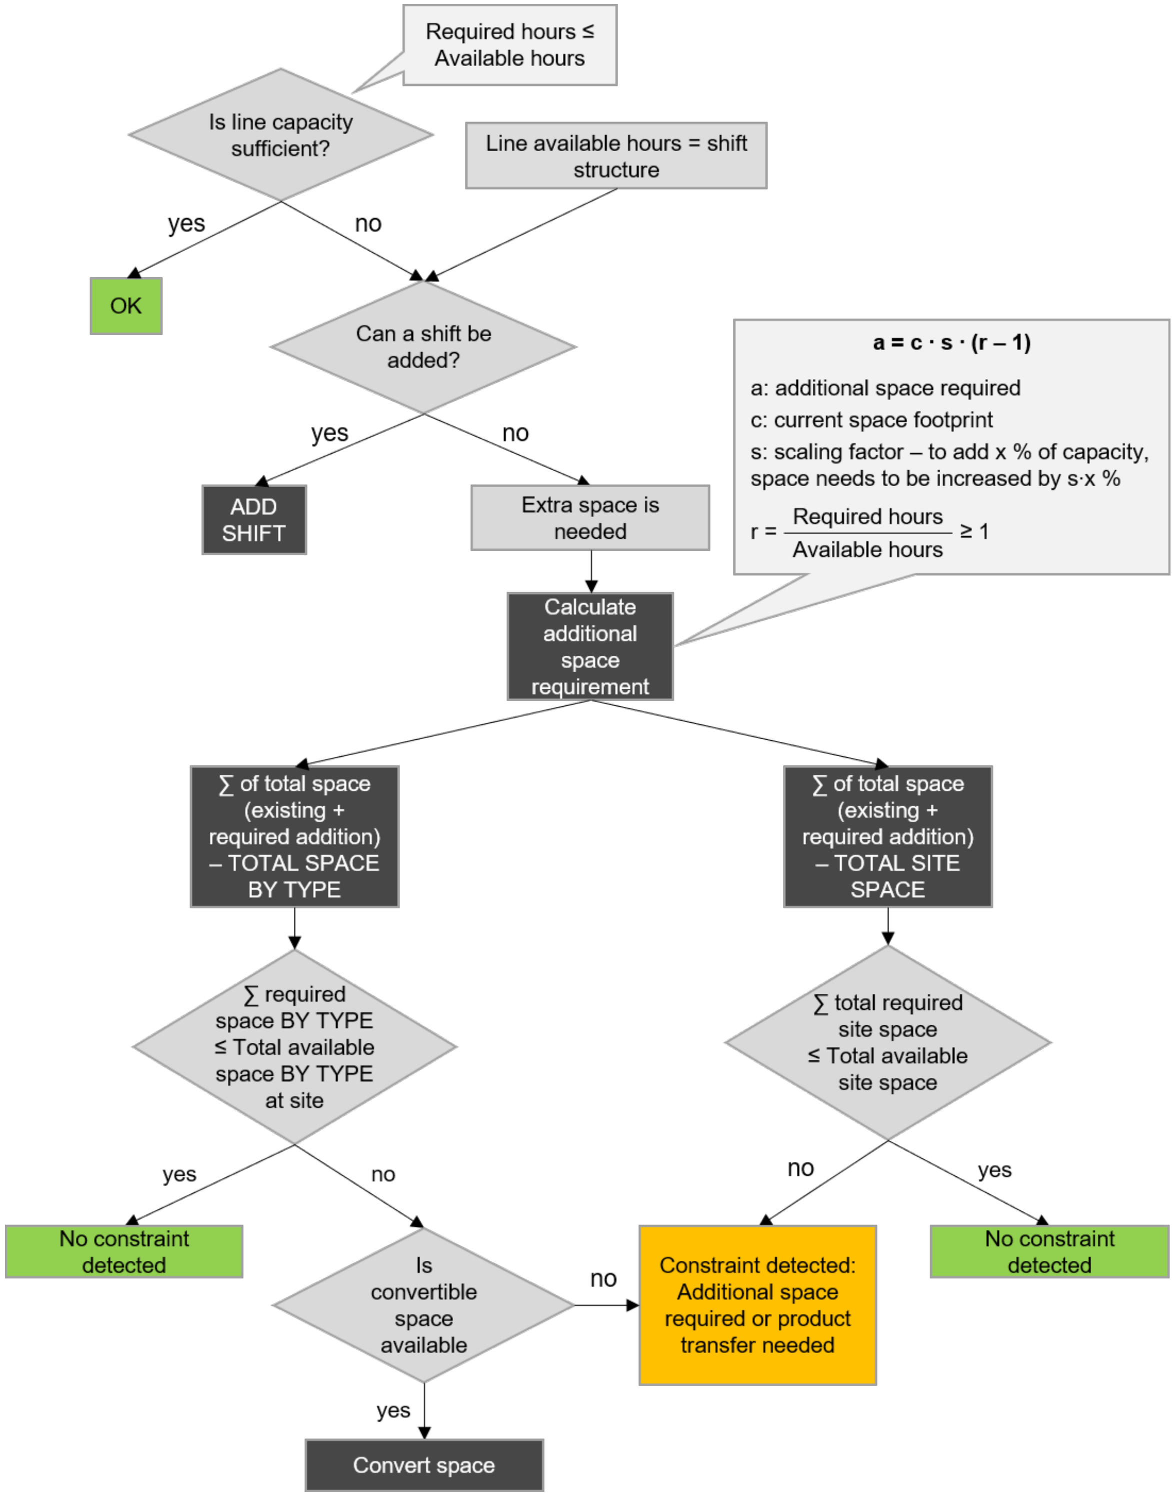

Figure 4-10. Flowchart of the logic to assess manufacturing space as a dimension of capacity, based on the outputs of the capacity model, i.e., the detection of capacity constraints. The decision tree determines if additional space is required or current footprint is sufficient to accommodate the additional capacity requirement.

in which case the scaling factor would be 25 %. Therefore, the scaling factor s is defined such that to add x % of capacity, space needs to be increased by $s \cdot x$ %, where $s < 1$.

Once the additional space is calculated, the next step is to verify whether the required additional space is available. Here, two criteria are relevant: First, whether the space required at a site exceeds the total available space. And second, whether the type of required space is available. The first criterion verifies whether space is available at all, and the second criterion ensures that the correct type of space, for example the appropriate level of cleanroom quality is available. If the total required additional space exceeds the available space, this indicates a capacity constraint, *i.e.,* the requirement to add additional space to the network. In the case of ensuring the availability of the correct type of space, if the correct space is not currently available, the next check is whether there is convertible space that can be appropriately converted to the required space type such as the right level of cleanroom control. If such space is not available, then again, a capacity constraint has been detected, requiring investment into additional space to alleviate the constraint.

As mentioned initially, this analysis needs to be performed for the whole network and for different points in time. Using this approach, future space constraints can be detected in the network, and ideally, if they are detected with sufficient lead time, appropriate measures to respond to future constraints can be taken before they arise. By doing so, capacity planning can be shifted from a reactive to a proactive mode, where capacity is being installed prior to an increase in demand manifesting itself. Given the high margins in the medical device industry, this is a favorable approach, as the opportunity cost of underage, where demand cannot be met, far exceeds the cost of excess capacity. In addition, there is an ethical obligation to have sufficient capacity for a medical device company, in order to be able to serve patients who require the life-saving products Boston Scientific offers.

5. Case Study

As laid out in the current state assessment above, the manufacturing network of Boston Scientific is a highly complex system, and to date, there a is gap of uniform capacity data that could be used to build a network-wide capacity model. To demonstrate the logic developed in section 4, a product was selected for case study of the performance of the model, which will be referred to as *Product A* in this thesis. In this case study, it is demonstrated how a capacity model could be implemented for a single product. These learnings can then be applied for the scale-up of the model beyond individual products.

To select an appropriate product for a case study, this product should satisfy a range of criteria that makes it a valuable example along the dimensions of relevance, complexity, and practicality:

<u>Relevance</u>: In order for a case study product to be relevant in the context of manufacturing capacity, it should contribute significantly to Boston Scientific's revenue today, as well as exhibit significant future growth of volume and revenue. If the product satisfies these criteria, this means that it is a priority to senior management. Furthermore, it implies that the product is representative of the future growth of the company, which is the key business driver requiring more systematic network capacity modeling.

<u>Complexity</u>: Products selected for a case study should capture as many aspects of supply chain complexity as possible, in order to be generalizable. These aspects of supply chain complexity include:

- Components, sub-assemblies, and top-assemblies pass through multiple factories within the network.
- At least some production occurs in mixed-model lines, and there is variation between batch and serial manufacturing.
- The production processes at different steps of the value chain vary in their degree of automation, from highly manual to fully automated.
- There are special handling requirements, such as cold-chain requirements, nuclear safety requirements, or short shelf lives.

<u>Practicality</u>: It should be practically feasible to build a capacity model for the selected product, meaning that sufficiently complete capacity data should exist for this product. Furthermore, this data needs to be accessible through both the relevant responsible team members and the data systems where the data is housed.

5.1 Case Study on Product A

5.1.1 Product Description

Product A is marketed by the Cardiac Rhythm Management division of Boston Scientific. It was selected for a case study in this thesis because it satisfies many of the criteria mentioned above. It represents a growing and relevant product in the Boston Scientific portfolio, and an example of a product in a rapidly growing segment of the healthcare market.[17] Furthermore, its supply chain contains many of the complexities that are characteristic for Boston Scientific and other global medical device manufacturers: multiple sites are involved in the production of the product, and many lines are mixed-model lines. The manufacturing systems of Product A and its components feature varying degrees of automation, and the manufacturing of lithium-containing batteries requires special handling including a dry-room atmosphere. Finally, the current-state assessment has shown that capacity data of the product is available, albeit in various formats and structures, which enables the creation of a capacity model for this product.

5.1.2 Process Flow Diagram of the Manufacture of Product A

Manufacturing of Product A occurs at two sites in the Boston Scientific network, which will be referred to as Site A and Site B in this thesis. **Figure 5-1** shows a high-level process flow diagram of the production steps required to produce Product A from individual components to the finished good. It can be seen that most components and sub-assemblies are produced at Site A, which is highly representative for the design of the Boston Scientific network: Site A serves mostly as an internal supplier of components to other factories. The components of Product A, components A, B, C, and D, are produced in different business units of Site A. Components C and D feed into sub-assemblies which are also produced at Site A, and component B feeds into a sub-assembly production process at Site B. All components and sub-assemblies are then supplied to the finished good production line at Site B. The top-assembly consists of four steps prior to sterilization and packaging, which are not considered in this thesis.

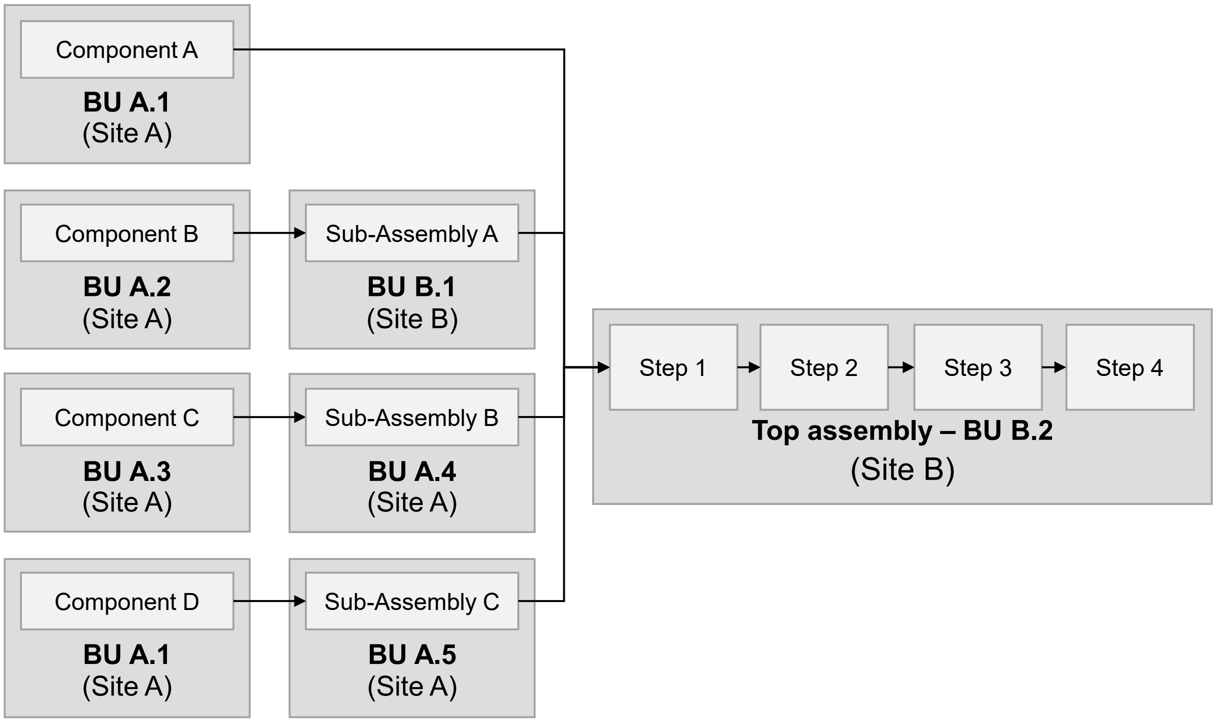

Figure 5-1. High-level process flow diagram of the manufacturing steps involved in the manufacturing of Product A from components to sub-assemblies and the top assembly.

5.1.3 Capacity Assessment

A prototype model for a mixed model line was assembled for the production of Sub-Assembly A at BU B.1 of Site B. For this line, it was found that sufficient data was available to perform stand-alone calculations, and additionally, gaps in data could be filled *via* manual inputs.

An important feature of this prototype is the ability to adjust certain input parameters for scenario planning. This is particularly important for parameters which are based on assumptions or forecasts of future developments, because these parameters are inherently uncertain. With the help of scenario planning, it is possible to determine potential breaking points in the future capacity evolution, e.g., by determining under which demand trajectory current capacity projections will be insufficient in the future. Importantly, in this first prototype, adjustable parameters were viewed as global, i.e., they apply to all products or processes. In future iterations, it should be considered if these parameters are chosen as global values, which would allow for simplicity for the operators, or as parameters which can be adjusted for individual lines, products, etc. In the view of the author, a good balance would be to maintain adjustable parameters at a factory level for capacity data, and at a product family level for demand forecasting data. Selecting this intermediate level of granularity, in combination with reasonable default values and the ability to adjust parameters at a global scale, will allow decision makers a reasonable degree of customization, without giving in to the illusion that there can be perfect predictions of the future.

While the prototype was a part of the capacity analysis of only Product A, the line investigated in detail constitutes a mixed model line. Due to the interdependency of capacity of all products passing through a mixed model line, it was required to perform a capacity analysis for all products passing through the same line. This made data capture significantly more complex than for an individual product, since the product mix corresponded to nine products in the demand forecast and eleven products in the manufacturing planning tools. However, by capturing this additional complexity, the model is more representative for the overall complexities in the supply chain, and generalized learnings can be extracted.

The following parameters were included as adjustable input parameters for the purpose of scenario planning:

Improvement rate – cycle time: This parameter sets the annual improvement rate, in percent, of the cycle times in the model. For example, an improvement rate of 5% would mean that year-over-year, the cycle time of all processing steps reduces by 5%.

Improvement rate – yield: The improvement rate of the yield defines by how many percent the yield improves annually. For example, an improvement rate of 5% of the yield would imply a year-over-year increase in yield by 5%.

Demand adjustment – nominal demand: This parameter serves to adjust the nominal demand as forecasted in the long-range demand forecast. Therefore, it is possible to investigate the capacity requirements based on higher or lower demand manifestations than nominally forecasted. The parameter is input in the form of a percentage deviation from the nominal forecast data.

Demand adjustment – upside demand: Similar to the demand adjustment for the nominal demand, this parameter lets a user of the model adjust the demand forecast for upside demand. Therefore, it is possible to perform scenario planning for stark deviations from the nominal forecasts, and observe potential breaking points in the capacity projections.

Beyond these adjustable parameters for scenario planning, the model consisted of three types of data: Input data from structured sources, manual entry data, and calculated data. Input data from structured sources represents data that was available from data sources such as site capacity models or the long-range demand forecast. It could be used without a large degree of manipulation, besides translating data structures from data sources to the data structure proposed in Chapter 4.4. Manual entry data is data that was not readily available from structured sources, but rather had to be collected through research, interviews, or explanations from subject matter experts. Calculated values are based on the first two types of data, and result from calculations and data manipulation. These values contain parts of the logic outline in Chapter 4.5.

The following data from structured sources was used as inputs into the model:

Cycle times and shift structure: The cycle times and shift structure were extracted from the capacity tool of the business unit B.1 at Site B. The data had to be reformatted to match the desired input format.

Demand: The demand data was extracted from long range forecast, as provided from the global supply chain center of excellence and the supply chain planning staff.

The following data had to be entered manually since no structured data was available:

Product synonyms: A list of product name synonyms was assembled which contained different names for the same products, containing synonyms for all involved products as used in the long-range demand forecast, capacity planning tool, as well as other sources of data. The list of synonyms can be found in **Table 5-1**. Here, we discuss all relevant product synonyms in the Cardiac Rhythm Management and Neuromodulation divisions, which were part of the considered mixed model line. As described in section 4.1.6, there are various types of product naming inconsistencies within the Boston Scientific supply chain and commercial businesses. For example, the types of synonyms range from spelled out abbreviations and the commercial names to a qualifier of the division, regional localizations, or indications of the product generation. In the model, all of these synonyms point to a *Product Master Name*, which serves as the linkage point between different names. It is worth reiterating that this connection between names and the list of synonyms was assembled manually, which implies that it is likely not exhaustive. Therefore, there is a likelihood for further complications in the future, when the model is scaled form a small set of products to the larger organization.

Product groups: The list of product groups, **Table 5-2**, contains the manually assembled grouping of products based on the different types of grouping in different data sources. Importantly, all products in a product grouping point towards the same master group name, which is given in the left column. The two product groups created for the Neuromodulation and Cardiac Rhythm Management divisions were the

Tachy group and the Neuro group. These groupings are consistent with the description in section 4.1.6 and **Figure 4-2**. For Tachy, it was found that two products in the long-range forecast named NG3 and NG4 correspond to only one product in the PCPA manufacturing line: The products represent to two generations of the same product series, and the built-in printed circuit board assembly has not changed between the two generations. Therefore, manufacturing plans production as only one product, which implies that a linkage between demand planning and manufacturing requires to group the products accordingly. In case of the neuromodulation division, it was found that the divisional planning of demand and the planning of manufacturing capacity occurs with entirely different types of groups, as outlined in section 4.1.6 and **Figure 4-2**. Here, the smallest possible grouping that can capture these differences is the division itself, and thus the grouping is termed Neuro, capturing all existing product names from both the long-range forecast and the local capacity planning tool.

Product Master Name	Synonym
ICM	Insertable Cardiac Monitor
ICM	LUX-Dx
ICM	CRM ICM
ICM	LUX-DX ICM US
ICM	LUX-DX ICM OUS
ICM	LUX-DX ICM 1.2 IDE US
ICM	ICM 1st
NG3	NG 3
NG3	Tachy
LCP	Leadless Cardiac Pacemaker
SB	Springboard
SB	Accolade
SB	Brady
SB	SB
SICD	Gen 2
SICD	Emblem
SICD	S-ICD PG
SICD	S-ICD
ETS	External Trial Stimulator
ETS	External Trial Stimulator
ETS	SCS ETS
ETS	ETS PCBA
ETS	ETS Flex

Table 5-1. List of product synonyms for the cardiac rhythm management and neuromodulation divisions, for products processed at BU B.1 of Site B.

Product	Grouped product
Tachy	NG4
Tachy	NG3
Neuro	DBS IPGs
Neuro	SCS IPGs
Neuro	Spectra
Neuro	Montana
Neuro	Archie
Neuro	Wilson
Neuro	Blink

Table 5-2. List of product groups for the cardiac rhythm management and neuromodulation divisions, for products processed at BU B.1 of Site B. Most importantly, neuromodulation products exhibit inconsistent grouping between the long-range forecast and the production site, making the division itself the smallest possible general grouping.

Based on the structured and manual input data, a range of calculations was performed, mirroring the capacity calculation logic outlined in Chapter 4.5. These calculations were performed sequentially, as some results served as inputs for subsequent calculations.

Joined groups: The joined groups represent the synthesis of product synonyms and product groups described above. They represent a "lowest common denominator" of product names and groups, linking all input data to the smallest possible group of products in order to connect different types of data collected with different product names/groups. This data table represents a flattened, table-version of the linkages laid out in **Figure 4-2**. Every input product name from the long-range demand forecast, and from the Site B capacity planning tool is first mapped to its product master name based on the synonyms table, and then to the appropriate product group, based on the product group table.

Cycle times by group: The table cycle times by group maps the raw cycle times as extracted form the Site B Capacity tool to the joined groups as previously defined. This is performed at the same level of granularity in terms of work centers as the raw data, which in the case of this prototype is by individual production station. It is worth noting that within a joined group, individual products do not necessarily have the same cycle times, and therefore, it was necessary to develop a definition of the cycle time for a group. For simplicity, cycle times were simply averaged in the form of the arithmetic mean, but other approaches are certainly viable and would represent the data more accurately. A grouping approach using a weighted average based on the actual product mix as represented in the demand would be more accurate, but create an additional loop in the calculation logic of the capacity model. Since in the case of the considered line, products grouped together are generally similar in nature, with similar cycle times, this additional loop was avoided to make the model simpler to understand and to implement at a larger scale.

Demand by group: The demand by group is the mapping of raw input demand to the groups established in the table *joined groups*. In the case of demand, the aggregation is simply a summation of the demand of individual products, and is therefore an accurate representation of the data, albeit with an obviously reduced degree of granularity.

Supply/demand match: The supply/demand match is the final step in the capacity logic and represents the core functionality of capacity modeling. In this calculation, the required and available times are compared and represented in a rudimentary graphical form, see **Figure 5-2**. The required times are calculated per station for each year, for both the nominal and upside demand forecast values: the cycle time of each product is multiplied by the annual demand to calculate the annual number of hours required to produce the product. For each station, these required hours are then summed up for all products, yielding the total required hours per station. The table furthermore includes the currently available hours based on the shift structure, as well as the total number of hours in a year as the theoretical upper limit of available time in a year. If the number of required hours is less than the currently available hours, the corresponding table field will be highlighted in green, indicating sufficient capacity. In the case that the required hours exceed the currently available hours, but are less than the total hours in a year, the corresponding table field is highlighted in yellow, indicating a capacity constraint under the current shift structure. Such a constraint could be alleviated by adding additional shifts, but without the requirement for further investment into equipment or space. Finally, in case the required hours exceed the total hours in a given year for the nominal or upside demand, the respective table field is highlighted in red. This corresponds to a lack in capacity which requires additional installation of equipment or space, as the constraint cannot be alleviated by an additional shift.

In addition to the nominal calculation results based on current cycle times and the long-range demand forecast, the supply/demand match table also includes a sensitivity analysis based on the adjustable parameters described above (improvement rates and demand adjustments). This second table, which is formatted in the same layout as the nominal results, shows how the capacity constraints shift in case of adjusted demand or when capacity-related manufacturing performance improves. As an example, it can be seen in **Figure 5-2** that for Station 1, under the given adjusted parameters (improvement rate: cycle time – 2%, improvement rate: yield – 8%, demand adjustment: nominal demand – 5%, demand adjustment: upside demand – 5%), capacity would be insufficient under nominal demand in year 3, whereas without adjustments, it would be insufficient only under upside demand in year 3. This indicates that there is a potential breaking point at Station 1 which could manifest quickly should demand rise quickly or improvement rates fall short of expectations.

Station	Current Available Hours	Max Hours in a Year	Nominal Required Hours Year 1	Upside Required Hours Year 1	Nominal Required Hours Year 2	Upside Required Hours Year 2	Nominal Required Hours Year 3	Upside Required Hours Year 3	Nominal Required Hours Year 4	Upside Required Hours Year 4	Nominal Required Hours Year 5	Upside Required Hours Year 5	Nominal Required Hours Year 6	Upside Required Hours Year 6
Station 1	3333	8760	8519	8767	8003	8284	8545	8959	9163	9714	9681	10142	10240	10737
Station 2	3333	8760	4125	4209	4044	4129	4104	4307	4233	4446	4284	4498	4338	4554
Station 3	3333	8760	1390	1423	1353	1387	1398	1467	1450	1529	1492	1566	1537	1613
Station 4	3333	8760	3632	3710	3540	3620	3605	3773	3701	3882	3756	3931	3814	3993
Station 5	3333	8760	5006	5122	4854	4975	4997	5238	5181	5454	5311	5568	5450	5715
Station 6	3333	8760	5249	5362	5097	5210	5200	5459	5443	5726	5538	5816	5638	5922
Station 7	3333	8760	1571	1604	1540	1574	1567	1643	1607	1688	1630	1709	1654	1735
Station 8	3333	8760	3436	3507	3394	3466	3453	3631	3532	3717	3584	3768	3639	3826

Scenario planning: adjusted values based on improvement rates and demand adjustments

Station 1	3333	8760	8945	9205	8403	8698	8972	9406	9621	10200	10165	10649	10752	11273
Station 2	3333	8760	4331	4420	4247	4335	4309	4523	4444	4669	4499	4723	4555	4782
Station 3	3333	8760	1460	1494	1420	1457	1468	1541	1523	1606	1567	1644	1613	1694
Station 4	3333	8760	3814	3896	3717	3801	3785	3961	3886	4076	3944	4128	4005	4192
Station 5	3333	8760	5257	5378	5097	5224	5247	5500	5441	5727	5577	5846	5723	6000
Station 6	3333	8760	5512	5630	5351	5470	5460	5732	5715	6013	5815	6106	5920	6218
Station 7	3333	8760	1650	1684	1617	1653	1646	1725	1687	1772	1712	1795	1737	1822
Station 8	3333	8760	3608	3683	3564	3639	3626	3812	3708	3903	3763	3957	3820	4017

Figure 5-2. Capacity calculation results per station and per year for nominal and upside demand (“supply/demand match”) for the provided input data (top table), and based on adjusted parameters for scenario planning (bottom table).

5.2 Generalized Learnings

The implementation of a prototype capacity assessment for Product A led to learnings which will be generalizable for the future scale-up of the capacity model. These learnings are described in the following section, and should be considered when planning for the implementation of the abstract model logic outlined in Chapter 4.

One of the major challenges associated with measuring capacity successfully, as described above, is the correct unit of measure of products, as well as the appropriate level of granularity of data included in the model. It has been described above in detail that the appropriate measure of capacity is in units of production time, and this was verified through the implementation outlined here (c.f., in particular, **Figure 5-2**). Furthermore, a goal here was to understand what constitutes a minimally sufficient level of data granularity for the raw capacity data, in particular of the level of product hierarchy and supply chain granularity required for successful modeling. It was found that capacity needs to be measured in individual production units, and data needs to be collected on a per station level. While this level of granularity is challenging to achieve and requires significant resources to implement, it will offer a degree of insights that cannot be captured by more aggregate, higher level data. In several cases observed for this work, it was found that for a mixed model line, the bottleneck of a line varies for different products passing through the line, and therefore, for an adequate level of model accuracy, it is necessary to capture sufficiently granular data that can delineate these varying bottlenecks. As a rule of thumb, capacity data should always capture at least all bottlenecks for all products passing through a line. Overall, simply put, it is always possible to aggregate granular data to obtain higher level insights, but is difficult to retroactively add granularity.

A second key learning from the case study is that all data inputs require to be in a standardized structure for a successful implementation of a larger-scale capacity model. While this observation appears trivial, it is critical for the success of such a model. It was found that a significant amount of time was spent with understanding raw capacity data as used by manufacturing sites, and by translating and manipulating such data to match the proposed attributes as defined in this thesis. Given the large variation that already exists in the quality and type of data within the Boston Scientific enterprise, it is critical to minimize complexity and variation wherever possible. Furthermore, from a practical implementation point of view, the local users of capacity tools usually have the largest amount of expertise in the tools' functionality and design. A feasible and efficient path to scaling the model would be for these subject matter experts to carry out the translation of their local capacity data into the centralized format.

In the current model, significant effort was required to determine adequate product groupings which can link different degrees of granularity. This is likely to only be exacerbated for an implementation at a larger scale, where instead of less than ten products, thousands of products need to be considered. This will likely require the most amount of manual data engineering by the implementation team with regards to the different tasks of the implementation process.

A specific observation from the case study presented above was that products within a product group can have different cycle times at the same station. Here, product group refers to the groupings that are required to link products in the different data sets, as described above. Because of these different cycle times within what becomes a single data point, it was required to somehow aggregate the cycle times of individual products to the effective cycle time of the overall group. For the purpose of the case study, the effective cycle time was simply the arithmetic mean (i.e., the average) of the individual products' cycle times. However, this is not necessarily accurate, and weighted average based on the true product mix would yield improved results. Given the multiple layers of uncertainty, and strong assumptions particularly in the demand forecast, it was found that this lack of accuracy should have minor implications on the overall findings of the model. Therefore, for simplicity, the arithmetic mean should provide an adequate level of precision without introducing additional computational complexities in the form of an additional feedback loop between demand and capacity data.

5.3 Gap Assessment

The previous case study by its nature offers only a partial view into the capacity of the highly complex network of Boston Scientific. While the product was selected to be representative of a broad set of products, the case study naturally cannot be equated with a full-scale capacity model. The intent of the case study is to serve as a basis for the future scale-up of the capacity model, and as such, it is important to assess the gaps of the case studies in order to understand the requirements which need to be met by a full-scale model.

The most immediate gap of the case study approach is that it only captures individual products, and not the complications that arise from the complex product portfolio of tens of thousands of products. From the case study of Product A, some of these complications can be observed in a simplified fashion. For example, the final assembly of Product A occurs on the same line as the entire product portfolio of cardiac rhythm management and neuromodulation devices. It was observed that for each of these devices, a different step in the production process represents the rate-limiting step, or bottleneck, of production. This means that there is no single bottle neck on this line, but rather, it depends on the product and mix of products being produced. In addition, for each of these products, there is a different supply chain with different constraints that feeds into the top assembly line. Because of these different constraints, a learning on Product A does not necessarily represent a learning that is applicable to other devices that pass through the same line.

In this context, there are additional complexities that arise from moving from an individual product to a set of products. By considering multiple products, especially the entire product portfolio, additional interactions between products have to be considered. Seemingly disconnected products may share common raw materials or sub-assemblies, or their sub-assemblies/components may share manufacturing resources. This all implies that capacity calculations upstream become increasingly complex.

The current case study captures a product that is produced in only two sites of the Boston Scientific network. Given the network's size and complexity with twelve major manufacturing sites, the case study therefore only captures a small fraction of this network. Other sites use different tools and processes than the ones harvested for the case study, and in particular, capacity related data such as from local capacity planning tools may contain different (and less) information than what was available for Product A.

In terms of manufacturing capabilities, the case study does not capture some of the intricate complexities of the network. In particular, special handling requirements such as for radioactive materials or in cold chains do not pose an issue for Product A, but are relevant for certain products in the Boston Scientific portfolio. All products can be measured in terms of counted units, which differs from components such as braided metal tubes or extruded polymer tubes. The product also differs strongly from large size, small volume products from the capital equipment business of Boston Scientific, which is becoming more and more relevant to the overall organization.

In the current prototype, the space logic described in Chapter 4.6 has not been implemented. This is primarily because the appropriate space data was not available at the level of granularity required to test the logic proposed here. At the same time, the space logic also requires highly granular space data for the entire site, which was outside the scope of the smaller scale prototype assembled here. Therefore, in the next iteration and scale-up process, a demonstration of the space capacity logic is required to test the functionality and determine its limitations.

6. Implementation Roadmap

In the following section, a roadmap for the scale-up and implementation of the capacity model and logic is proposed. This roadmap should serve as guidance to plan this challenging, cross-functional project, and give indications on the dimension which this project will take.

Overall, given the limited quality of data that currently exists, and the complexity of the Boston Scientific manufacturing network, it is to be expected that the implementation of a network-wide capacity model will require significant resources. Required capabilities include data engineering, analytics, visualization, supply chain planning, manufacturing/industrial engineering, and project management know-how.

6.1 Workshops to Develop Requirements

As a first step it will be important to asses the needs of all stakeholders involved in short-, medium-, and long-term capacity planning. To develop the requirements and specifications of the capacity tool, a series of workshops involving all these stakeholders should be held. At the workshops, stakeholders collaboratively develop a set of outputs and results which they would want to receive from a capacity model. Furthermore, they would determine which inputs are required to obtain these outputs, based on the data and logic presented in this thesis. These stakeholders would be able to asses which data of the proposed data structure is available, which data is missing, and how the data could be assembled for further development. Finally, the workshops would serve to align stakeholders and obtain their buy-in in the proposed model, by educating them how they would benefit from the success of this project.

6.2 Standardize Capacity Data Fields

As described in this thesis, it will be critical to standardize the database fields in order to enable a scalable solution that requires the smallest possible amount of iterative manual data engineering. Therefore, based on the data tables, fields, and attributes proposed here, as well as the outcomes from the stakeholder workshops described above, the capacity database fields need to be standardized and clearly defined prior to any implementation work. The benefit of investing into this design and specification step prior to implementation is that the definitions of database fields can be developed together with the business process that will be used to populate the data in the future. By considering which data is available, who would provide the data, and what format existing data has, the database fields will be more practical and faster to populate. The proposed database structure in this thesis should serve as a strong foundation for this step.

6.3 Build Data Architecture

After conceptualizing and defining the database fields and attributes, it will be necessary to implement the database architecture and infrastructure. This relates to installing server space with the appropriate relational database architecture, and to create the appropriate tables based on the definitions obtained previously. Furthermore, a process to populate these tables should be designed: This could be through implementing a type of entry form where data can be entered manually by an operator such as an industrial engineer or supply chain planner. Another feasible entry method could be the ability to upload tables which contain data in the desired, standardized format. For this purpose, user-friendly upload templates should be created, for example in the form of Excel spreadsheets.

6.4 Populate Database

After creating the raw data architecture, it will be required to populate the database with the desired data. This will be a resource-intense process that requires input from the entire manufacturing network, likely down to industrial engineers responsible for individual lines. It will be critical to ensure that all data inputs match the requirements and standardization. To achieve this, templates should be made available as described above, and significant effort should be spent on educating the network on the specifications and requirements of the data entry. By receiving high-quality data as input, less effort will be required for additional data manipulation, corrections, and for obtaining missing data. A particular challenge will be assembling a product name dictionary for the vast product portfolio of Boston Scientific. This was performed at a small scale in the case study presented in Chapter 5, but will require significant effort and resources to be populated at network-scale.

6.5 Implement Capacity Logic

As the core of the model, the actual capacity logic needs to be implemented. This implementation would execute the model logic designed in Chapter 4 of this thesis, and prototyped in Chapter 5, using the data in the database as designed and populated in the previous steps. The capacity logic will require a range of table manipulations to create the desired linkages between previously disconnected data. Moreover, the logic should include the ability to manipulate parameters for scenario planning and sensitivity analysis, as demonstrated for the demand forecast data and improvement rates in this thesis.

6.6 Build Minimum Viable Product

While all previous work to this stage was done on the back-end of the capacity model, the minimum viable product represents a first usable front-end solution. Therefore, in this development step, the focus lies on visualizing the outputs and results from the logic, and adequately aggregating results such that they are useful to decision makers. For example, high-level summaries of key capacity constraints should be assembled, and the most significant projected investments should be highlighted. Due to the granularity of the model and the underlying data, decision makers will then be able to dive deeper into these detected constraints, and investigate potential breaking points through sensitivity analysis. Developing the minimum viable product will require resources in the analytics and data visualization, combined with a strong understanding of the business needs for decision making in manufacturing capacity.

6.7 Iterate

After implementing the first full-scale minimum viable product, it will be necessary to iterate and improve the tool based on the learnings and results from this full-scale model. It is likely that flaws will be detected which should be addressed in an iterative fashion.

7. References

1 P. M. Swamidass, Ed., in *Encyclopedia of Production and Manufacturing Management*, Springer US, Boston, MA, 2006, pp. 74–74.

2 R. Hayes, G. Pisano, D. Upton and S. Wheelwright, *Operations, Strategy, and Technology Pursuing the Competive Edge*, John Wiley & Sons, Hoboken, New Jersey, 2005.

3 T. Hill, *Manufacturing Strategy Text and Cases*, McGraw-Hill Higher Education, Boston, MA, 3rd edn., 2000.

4 R. Hayes and S. Wheelwright, *Restoring our Competitive Edge*, John Wiley & Sons, Hoboken, New Jersey, 1984.

5 Boston Scientific, Boston Scientific - About Us, https://www.bostonscientific.com/en-US/about-us.html, (accessed 18 January 2022).

6 Boston Scientific, Boston Scientific - 2019 Annual Report, https://www.bostonscientific.com/content/dam/bostonscientific/corporate/annual-report/2019_Boston_Scientific_Annual_Report.pdf, (accessed 18 January 2022).

7 Boston Scientific, Boston Scientific - 2020 Annual Report, https://investors.bostonscientific.com/~/media/Files/B/Boston-Scientific-IR-V3/annual-reports-proxy-statements/2020-boston-scientific-annual-report.pdf, (accessed 18 January 2022).

8 B. Scientific, Q4 2021 Financial and Operating Highlights, https://investors.bostonscientific.com/~/media/Files/B/Boston-Scientific-IR-V3/Q4 2021 Financial Operational HighlightsFINAL.pdf, (accessed 9 March 2022).

9 Boston Scientific, Boston Scientific - 2021 Investor Day Presentation, https://investors.bostonscientific.com/~/media/Files/B/Boston-Scientific-IR-V3/documents/events/BSX-Investor-Day-2021-22Sep-final-final.pdf, (accessed 18 January 2022).

10 B. M. O. Awuondo, Massachusetts Institute of Technology, 2018.

11 A. R. Kalidindi, Massachusetts Institute of Technology, 2019.

12 B. Chang, Massachusetts Institute of Technology, 2014.

13 C. DiAndreth, Massachusetts Institute of Technology, 2019.

14 D. He, Massachnusetts Institute of Technology, 2020.

15 S. E. Laskowski, Massachusetts Institute of Technology, 2017.

16 S. Machinani, Massachusetts Institute of Technology, 2015.

17 Fortune Business Insights, IoT in Healthcare Market to Reach USD 176.82 Billion by 2026: Emerging Popularity of Smart Wearables and Remote Patient Monitoring to Add Impetus to Market, https://www.globenewswire.com/news-release/2020/02/20/1987603/0/en/IoT-in-Healthcare-Market-to-Reach-USD-176-82-Billion-by-2026-Emerging-Popularity-of-Smart-Wearables-and-Remote-Patient-Monitoring-to-Add-Impetus-to-Market.html, (accessed 20 January 2022).