Supply Chain Digital Twin

A Case Study in a Pharmaceutical Company

João Afonso Ménagé Santos
joaoasantos@tecnico.ulisboa.pt
Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
November 2019
Abstract

***Abstract*—The integrated supply chain is a network where all the business areas are dependent on each other. While they are extremely powerful structures, the logistics behind these interconnected networks of people, products, machines and information are highly complicated, and solutions for their optimization are increasingly required. The technological advancements seen in recent years, have allowed for better optimization of its processes, and data-driven solutions are being extensively adopted, due to their accurate results. The concept of the digital twin, when applied to the internal supply chain, has the possibility of aiding in its management, along with its capability of increasing awareness to the company's stakeholders and decision-makers, while allowing the deployment of accurate simulation models. This work presents a digital twin of a pharmaceutical internal supply chain, along with a simulation-based rough cut capacity planning tool, capable of giving estimates of the required monthly capacity for the different areas of the organization on the long-term. The work was a case study performed at a pharmaceutical company. The developed digital twin offers a graphical user interface, with several views into the past and present tasks performed and the evolution of the key performance indicators. The simulation tool is also included in the user interface, giving the possibility to decision-makers of creating their own scenarios and performing the simulation.**

***Index Terms*—Digital Twin, Internal Supply Chain, Operations Research, Simulation-based Rough Cut Capacity Planning, Pharmaceutical Supply Chain**

I. Introduction

CURRENT trends in logistics have placed the integrated supply chain (SC) as a process of unmatched capability in improving the overall performance of the organizations SCs. The integrated SC considers all the phases of production and support as dependent, resulting in a complex network of people, machines, infrastructures and products. The planning and logistics behind these structures have undeniable impact but are extremely complicated. Strategies for optimization and improvement of these activities are being increasingly investigated.

An industry which has adopted the integrated SC is the pharmaceutical industry. This industry has faced a shift on its characteristics, *e.g.* the liberalization of the global markets (exposing companies to global competition) or the widespread appearance of generics [11]. Additionally, the number of drugs in the companies portfolios has been steadily increasing, on account of the yearly number of approved drugs by the responsible entities being also rising [9]. Both these reasons have prompted a shift in the industry: *big pharma* companies are nowadays more focused on research & development, leaving the manufacturing of the drugs to contract (development and) manufacturing organizations, C(D)MOs. These organizations operate by manufacturing products (and industrial processes in case of CDMOs) by contract to other pharmaceutical companies. Their large, diverse and ever mutating product portfolios have created massive and complex SCs, which increasingly require optimization for a smooth operation.

The 21st century has seen the rapid development of modern technologies and computational power has become much cheaper. Concepts such as the Internet of Things (IoT), enabling the collection of data at distinct levels within an organization, have brought the required tools to optimize the integrated SC. Forecasting models, much more accurate than ever before, can now be developed, using collected data to obtain models based on demonstrated performance.

The objective of this work is the creation of a Digital Twin (DT) of the internal SC of a pharmaceutical CDMO, with the goal of increasing awareness to stakeholders and decision-makers, by delivering information regarding past and present tasks and performance indicators. Additionally, the DT incorporates a forecasting tool, capable of performing simulation-based rough cut capacity planning (RCCP).

II. Concepts

A. Digital Twin

A DT can be defined as a *dynamic virtual representation of a physical object or system, using real-time data to enable understanding, learning and reasoning* [1]. Although its definition varies from source to source, the basic idea consists on a digital representation of an asset (being it tangible [entity] or intangible [system]), which uses IoT to receive meaningful real-time data and reaches conclusions based on both the developed model, how it has performed in the past and how it is performing at the present. The basic schematic of a DT is shown in Figure 1.

DTs are frequently confused with both monitoring tools and simulation models. In reality, DTs bring together both concepts, effectively delivering a visualization tool with improved simulation models [7]. DTs differ from simulation models in the sense that they receive real-time data to generate better predictions. Regularly, simulation models have complete descriptions of the object or system in study, but often lack its historical performance and almost always lack their current state. By having both,

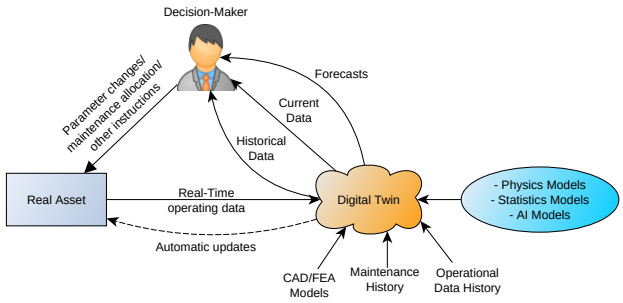

Fig. 1: Basic schematic of how a DT works. Note that the *automatic updates* transition may be optional.

the generated simulations can be verified and improved by simulating on past data and the predictions that the model can create are based on current states, which will deliver more data-driven and accurate responses. DTs also supersede monitoring tools in the sense that all the data that these tools possess and display is also available by DTs. Additionally, DTs have access to forecast data, created by its accurate simulation models.

B. Supply Chain

Mentzer et al. [8] define SCs as *"a set of entities (organizations or individuals) directly involved in the upstream and downstream flows of products, services, finances, and/or information from a source to a customer"*. The definition of internal SC can be seen as the internal processes and agents of the SC, with the source and customer ends of the SC being adapted into a network of internal operations. Pharmaceutical CMOs typically feature several areas, 4 of which will be further studied:

- **Manufacturing (M):** area responsible for the manufacturing of the products. Its operations can be divided into two categories: production and batch production record (BPR) review (after production).
- **Quality control (QC):** area that verifies the quality of the produced goods. Its operation can be divided into 3 categories: in-process control (IPC) – analysis during production; release (R) – analysis after production; release review (RV) – review of the analysis, after the release.
- **Quality assurance (QA):** area that ensures the quality of the produced goods and paperwork. Its operations are divided into 2 categories: BPR release – after the BPR review process by the manufacturing teams; release of the QC analytical package – after the QC RV stage.
- **Warehouse (WH):** area that directly influences the production through two operations: dispensing of raw materials and final product (FP) storage.

The internal supply chain operations of manufacturing a pharmaceutical product, entailing all of its components (mainly production, QC, QA and warehouse operations), is represented on figure 2, showing a business process model and notation (BPMN) graph of the manufacturing process.

C. Capacity

On a manufacturing organization, available capacity is measured for a given production plant, area or workcenter and for a specific range of time. It corresponds to the total available time (in the considered period) multiplied by the number of resources related to the selected scope. This can be expressed both in $worker \cdot hours$ or $machine \cdot hours$. Capacity utilization is then a measure of how intensively a resource is being used at a given time, with relation to its available capacity. A similar concept is the effort (or utilized capacity) that a certain activity requires. It can be defined as the amount of work that is required to complete it.

The sales and operations planning (S&OP) cycle is comprised of the distinct stages in a corporate plan that are sequential and always repeating and that feature different objectives at different stages of the cycle. Generally, 4 time fences can be distinguished in the S&OP cycle, as shown in table I, along with the typical times for each fence.

Horizon	Short-Term	Medium-Term	Long-Term	Strategic
Time	1-8 Weeks	1-3 Months	1-24 Months	3-5+ Years

TABLE I: Time horizons for the different S&OP cycle stages

These time horizons are of extreme importance since they define the different planning stages. Additionally, the aim within each time horizon changes greatly, from an operational and task-oriented point of view in the short to medium-term, to a planning and capacity focus in the long-term, to a strategic and managerial perspective in the strategic horizon.

D. Rough Cut Capacity Planning

Capacity planning can be defined as the process of determining and evaluating the amount of capacity required for future manufacturing operations. This capacity can often be in terms of labor, machinery, warehouse space or supplier capabilities. The RCCP step comes as the capacity plan at the tactical level, which regards the master production schedule, the plan made by the company at the long-term, regarding production, inventory and workforce. RCCP is a capacity planning tool made at the long-term, used to adjust the required and available capacity and to change the master schedule or available capacity, if necessary. Using the results from an RCCP, the master schedule can be modified in order to solve capacity inconsistencies by moving scheduled dates, increasing/decreasing scheduled production quantities or subcontracting additional workforce, for example.

E. State-of-the-art & Proposed Solution

Very few applications of DTs on the SC or intangible assets were found. Kritzinger et al. [4] provide a literature review of scientific articles on DTs, but these appear to be mostly DTs of tangible assets. Commercially available RCCP tools are frequent but no occurrence of simulation-based RCCP tools were found.

Applied to the pharmaceutical industry, Costigliola et al. [2] developed a discrete event simulation model of a QC

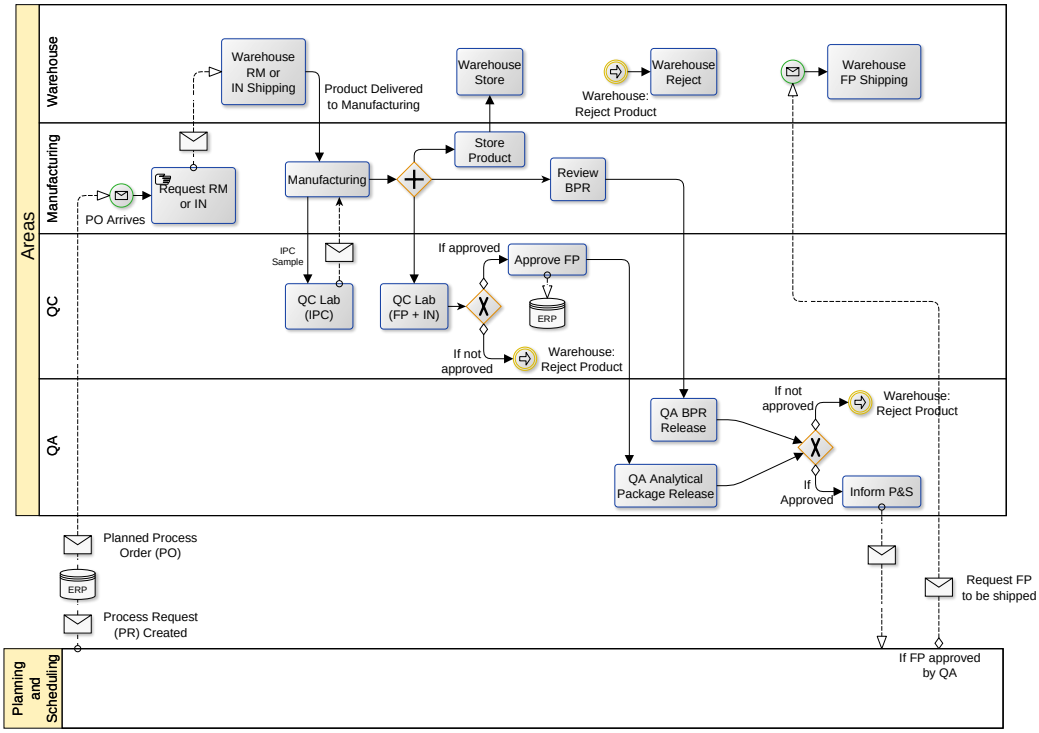

Fig. 2: BPMN graph of the internal supply chain considered

laboratory, intended to be used as a decision-support tool for scheduling and planning (operational level); Lopes et al. [6] presented a benchmarking platform to estimate the performance of new QC facilities (strategic level); Papavasileiou et al. [10] developed a Monte Carlo simulation approach to task scheduling and cycle duration calculation.

The presented solution is a DT to support the internal SC of the pharmaceutical organization in study. The DT largely follows the basic concept of DT provided by Ivanov et al. [3], where the authors describe DTs as a *combination of simulation, optimization and data analytics*. Similarly to the works by Costigliola et al. [2] and Lopes et al. [6], the tool is intended to act as a decision-support tool, but not applied exclusively to QC laboratories, and at the long-term timeframe. It provides a framework for visualizing operational information regarding the plant's activities, schedule operations and historical data and provide accessibility to the simulation tool. The simulation-based RCCP tool calculates the approximate monthly utilized capacity for the different areas, based on the concept of demonstrated performance. Furthermore, an approach similar to the one used by Papavasileiou et al. [10] was employed, in order to verify interferences between tasks using assets considered to be bottleneck, in the sense that they are constrained to only be used by one project at a time.

III. Implementation

A. Data Extraction

Data was extracted from the company's enterprise resource planner. This data was divided into 4 categories:

- **Campaigns durations:** for each campaign manufactured in the production plant, information regarding the project, manufacturing duration and quality release duration were extracted, along with the batch size and other additional information that could be of use in future versions of this tool. The extracted data and its relation to the reality can be graphically observed in the timeline from figure 3.

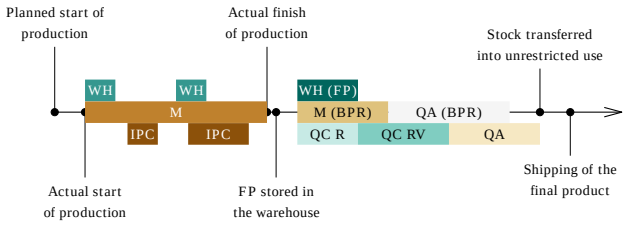

Fig. 3: Extraction data points and their relation to the real processes. The bars represent the real starts and ends of the different stages, while the vertical lines represent the dates that can be extracted.

The quality release (QR) stage corresponds then to all the processes after the manufacturing stage. This is done because there could not be extracted any start and end of the processes individually, and the complete set of tasks is considered as one.

- **Planned Orders:** for the extraction process, these are considered orders that have been confirmed and are within the range of 3 months and 2 years, corresponding to the long-term timeframe. The extracted information regarded the project, planned start date, planned deadline and batch size.
- **Required Resources:** *worker-hours* and *equipment-hours* associated with each project. This information is required to calculate the utilized capacity by each project, towards each area. The extracted data contemplated for each project all the tasks in the project's sequence of operations, which included information

regarding the tasks area, effort and duration. The efforts were then aggregated by area and project.

- **Available Capacities:** data regarding each area and each month's available capacity were extracted, to enable the calculation of the ratio of capacity utilization (utilized capacity divided by available capacity).

B. Distributions Fitting

To model the variability inherent to the processes, the collected data regarding manufacturing and QR durations was studied, with the objective of fitting a theoretical probability distribution function (PDF) to the observed durations, in such a way that allowed their values to be randomly sampled by the simulation tool. Note that for confidentiality reasons, the time units are not particularized, expressed instead in $[TU]$.

The basic steps for choosing theoretical PDFs and fitting them to the observed data were:

1) **Choosing the set of contender PDFs:** since the extracted data did not have sufficient granularity for continuous distributions, only discrete PDFs were considered. Law [5] proposes the use of the Poisson, binomial and negative binomial distributions. These PDFs were the only ones considered for the fitting process. Note that these PDFs belong to the exponential distributions family, which do not take negative values, having a lower bound at 0. Since the observed processes duration frequently also had a lower bound (inherent to chemical processes), an additional parameter was added to the distributions, an offset to the x input value, corresponding to the observed lower bound. The 3 contender PDFs are formulated in equations 1.

Binomial

$$p(x) = \frac{t! \cdot p^{(x-x_0)}(1-p)^{t-(x-x_0)}}{(x-x_0)!(t-(x-x_0))!} \tag{1a}$$

Negative Binomial

$$p(x) = \frac{(s+(x-x_0)-1)!}{(x-x_0)!(s-1)!} \cdot p^s(1-p)^{(x-x_0)} \tag{1b}$$

Poisson

$$p(x) = \frac{e^{-\lambda}\lambda^{(x-x_0)}}{(x-x_0)!} \tag{1c}$$

2) **Preprocessing the data:** after the data was extracted, a few basic steps were applied in order to remove incorrect data, such as eliminating observations with negative durations. Given the similar nature in processes, it can be assumed that all the projects duration follow approximately the same bahavior; to this end, only projects with more than 30 observations were accepted to discover the most representative theoretical PDF.
3) **Removing outliers:** the datasets in study are not significantly large, with a maximum of less than 500 observations. This enabled the removal of outliers using *Tukey's Fences*. To remove the outliers using this concept, points outside the range defined by $[Q_1 - k(Q_3 - Q_1), Q_3 + k(Q_3 - Q_1)]$ are removed. Q_1 refers to the first quartile while Q_3 to the third quartile (note that Q_2 is the median). A value of $k = 1.5$ is proposed for outlier removal. Applying this outlier removal algorithm showed satisfactory results, with outliers correctly (and conservatively) removed and manifestations of variability correctly kept.
4) **Statistical study:** observing the main statistical properties of the distributions can provide insights into what theoretical PDF to choose *a priori*. According to Law [5, p. 322], as a general rule, PDFs can be chosen based on their Lexis ratio, τ: $\tau = 1 \equiv$ Poisson, $\tau < 1 \equiv$ binomial, $\tau > 1 \equiv$ negative binomial. Considering that the results show 12/20 occurrences of $\tau > 1$, it can be roughly predicted that the most adequate PDF would be the negative binomial. An additional empirical method for determining the PDF based on the distributions' statistical properties is through the Cullen and Frey graph. This method considers the kurtosis and the squared of the skewness to make a prediction about the most appropriate PDF, in the manner depicted in figure 4. As can be seen from the figure, the negative binomial PDF appears to be the most appropriate one, according to the existing data.

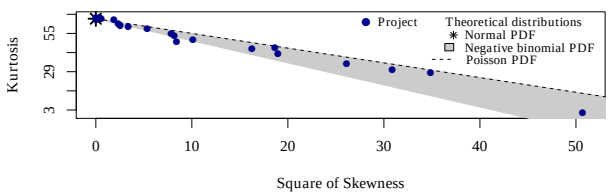

Fig. 4: Cullen and Frey graph for the projects in study

5) **Selecting the Goodness-of-fit (GoF) test:** GoF tests are tests made between a theoretical PDF and the real observations, with the objective of evaluating how well the theoretical PDF is fitted to the data, according to its own metrics. The Chi-squared (CS) test was chosen for two reasons: (1) the CS is a GoF test which evaluates the fit between the PDFs not on a cumulative way; many other GoF tests do it in such a way, which generally leads to better results in the tail of the distribution, but since there is no need for good precision on the tail, it is less advisable; (2) the CS test is by far the most used and reviewed method, being widely recognized and approved, what is often preferred. The CS GoF test is formulated as shown in equation 2.

$$\chi^2 = \sum_{i=1}^{n} \frac{(O_i - Np_i)^2}{Np_i} \tag{2}$$

$$\begin{cases} O_i \equiv \text{number of observations for TU } i \\ N \equiv \text{total number of observations} \\ p_i \equiv \text{theoretical probability of TU } i \\ n \equiv \text{number of cells} \end{cases}$$

6) **Fitting the theoretical PDFs to the data:** to fit the theoretical PDFs to the existing data, an optimization process is performed, with the objective of minimizing the CS value. The mathematical formulation is as shown in equation 3. Note that it assumes p_{i_x} with

x as being one of the 3 PDFs as shown in equation 1, and the optimization parameters $\{s, t, \lambda, p\}$ according to the chosen PDF.

$$\begin{aligned} \min_{t,s,\lambda,p} \quad & \chi^2 = \sum_{i=1}^{n} \frac{(O_i - Np_{i_x})^2}{Np_{i_x}} \\ \text{s.t.} \quad & t, s, \lambda > 0 \\ & p \in [0, 1] \end{aligned} \tag{3}$$

The results from the fitting process are condensed in table II.

PDF	CS Value	Time $[ms]$
Poisson	324.2	1.1
Binomial	476.3	15.5
Negative Binomial	**46.3**	**10.3**

TABLE II: Results of the fitting process, in terms of CS value and time taken for the optimization.

7) **Selecting the most representative PDF:** given the results of the fitting process and the empirical deductions made from the statistical properties of the distributions, it is possible to conclude that the Negative Binomial PDF is the most appropriate PDF for the processes being modeled.

C. Simulation-based RCCP Implementation

The first consideration made for the construction of the RCCP algorithm was the adoption of an upper bound for the obtained PDFs, point at which these are truncated. Since a lower bound already exists, the adoption of an upper bound removes the possibility of obtaining far-out values, which could badly influence the results of the simulation. The approach to do so is: define an upper bound (in percentage); obtain the duration value which corresponds to the selected upper bound; sample values from the distribution – if the obtained value is larger than the upper bound, re-sample until otherwise.

The simulation-based RCCP tool was based on Monte Carlo simulation. This class of computational algorithms relies on random sampling of values in order to find a pattern or tendency, and theoretically, is able to solve any problem with probabilistic interpretation. In the case at hand, it was seen that the manufacturing and QR durations had probabilistic characteristics that could be measured, which are propagated to the area's efforts. The mathematical formulation of the Monte Carlo method is done through the problem's probability space, which can be seen in expression 4 ($N \equiv$ *number of orders*).

$$\begin{cases} \boldsymbol{\Omega} = \{D_{M_1}, D_{QR_1}, D_{M_2}, D_{QR_2}, \cdots \\ \qquad \cdots, D_{M_N}, D_{QR_N}\}, \Omega \in \mathbb{N} \\ \boldsymbol{\mathcal{F}} = 2^{\Omega} \\ \boldsymbol{P(x)} = \prod_{i=1}^{N} P(D_{M_i}, s_{M_i}, p_{M_i}, x0_{M_i}) \cdot \\ \qquad P(D_{QR_i}, s_{QR_i}, p_{QR_i}, x0_{QR_i})_j \end{cases} \tag{4}$$

with $P(x, s, p, x0) \equiv$ *Negative Binomial PDF*

Regarding the probability space of each PDF, each value of D_{Mi} or D_{QRi} corresponds to a duration in terms of manufacturing or QR, corresponding to project i, taking a value from $\{D_{min_i}, \cdots, D_{max_i}\}|_M$. Note that the maximum value derives from the truncation that was previously mentioned.

For a simulation as the one whose results are shown in section IV-D, a total of 547 orders were considered, resulting in a sample space with around 10^{1392} scenarios. Calculating the monthly capacities for every scenario and observing the most probable utilized capacity for every month and every area would be computationally unfeasible. The use of Monte Carlo simulation is justifiable in such a situation: by randomly sampling values (according to their probability of occurrence), not all scenarios can be simulated, but a convergence can be found, which would eventually lead to the most probable scenario.

The implemented algorithms are based in either earliest due date (EDD) or latest start date (LSD). The basic difference between both methods is what date is considered when sampling values: EDD method fixes the start dates and then adds the manufacturing duration plus the QR duration to arrive at the *earliest due date*; on the other hand, the LSD method fixes the deadline (with a given time buffer) and by sampling manufacturing and QR durations arrives at the latest start date that can be considered for the project to finish on time. Generally, the EDD method is preferred. Regarding the implemented algorithms, these can be simply subdivided into 5 sections:

- **Sampling:** considering the set of planned orders (with information regarding project, start date and deadline), each campaign's manufacturing and QR durations are sampled from their respective PDFs. The values are re-sampled until being contained within the confidence level set. Furthermore, the start and end dates of each process (manufacturing and QR) are also calculated.
- **Monthly utilized capacity:** considering that the RCCP operates on the long-term, vectors for each area are created with a length of $2[y] \cdot 365[d/y] \approx 700[d]$, comprised of all null values. Then, for each planned order, the corresponding daily capacity for each area is calculated (based on the sampled values and on the efforts from the recipes). After all the planned orders are considered, the capacities are aggregated per month.
- **Bottleneck asset (BA) utilization:** given a set of assets considered to be BAs, the tasks on each BA regarding every project are obtained, by being scaled according to their duration on the recipes and on the sampled manufacturing duration.
- **Aggregation of simulations:** both the monthly utilized capacity and the BA utilization have to be aggregated, since a series of i independent results exist, with i being the number of Monte Carlo iterations chosen. The basic central tendency and variation metrics used are the median and interquartile range (IQR) since these metrics are more robust to skewed distributions than the mean and standard deviation. For the monthly utilized capacities, the aggregated

results are the median of the monthly capacities, with a deviation of plus or minus 1 IQR. Regarding the BA utilization, the process depends on the type of simulation chosen, EDD or LSD. If EDD is chosen, the start date and duration of each process are used, with the start being its median plus or minus its IQR, and the end being the median plus or minus IQR of the sum of start and duration. For LSD, the end and duration are used, with the end being its median plus or minus its IQR and the start being the median of the end minus the duration and the variation the sum of the IQRs.

One last step is performed, which is detecting the interferences between BAs utilization. The process is done after the aggregation and is done BA-wise, with different BAs being independent. The algorithm detects if there are any interference between two tasks of the same BA. If so, the tasks are considered to have *interferences*. Additionally, tasks that have interferences but only on their extended length (considering the minimum IQR of the start and the maximum IQR of the end), are considered to have *possible interferences*. All the other tasks have *no interference*.

- *(Optional)* **Optimization:** the optimization process is optional and has the objective of removing any interference between the tasks. It can be performed by the median, correcting the interferences between BAs themselves or by a factor of the IQR, either 0.25, 0.5, 1 or 2. This way the method basically operates by considering the tasks as an extended version of each, by the chosen factor. After the optimization, the resulting monthly utilized capacities are also calculated. Consider $\mathbf{BA} = \{BA_1, \cdots, BA_m\} = \{BA_i\},\ i \in [1, m]$ as the set of BAs, chosen by the user. i corresponds to the index of the BA, and m is the total number of BAs. Each BA features a series of activities from different projects; considering the BA with index i (BA_i), $\mathbf{P_i} = \{P_{1_i}, P_{2_i}, \cdots, P_{o_i}\} = \{P_{j_i}\},\ j_i \in [1, o_i]$. Here, j_i corresponds to the index of each activity (for BA_i), with the number of tasks being o_i.

$$\begin{aligned} \min_{n} \quad & n = \sum [P_{j_i} \cap P_{k_i} \neq \emptyset] \\ \text{s.t.} \quad & i \in [1, m] \\ & j_i, k_i \in [1, o_i] \\ & j_i \neq k_i \\ & n \geq 0 \end{aligned} \tag{5}$$

IV. Results

A. Convergence

Evaluating the convergence of the results is a fundamental step, since it can give clear insights into whether or not the simulation is in fact trying to reach a representative solution. In theory, the more simulations ran, the greater the confidence of an average value being the most representative. For a specific scenario, the simulation was ran for a series of different number of iterations, ranging from 10 to 50000. An example of the obtained results is shown in figure 5, where the evolution of the histogram of monthly utilization for a randomly selected month and for the manufacturing area is shown, along with each distribution's mean, standard deviation, median and IQR.

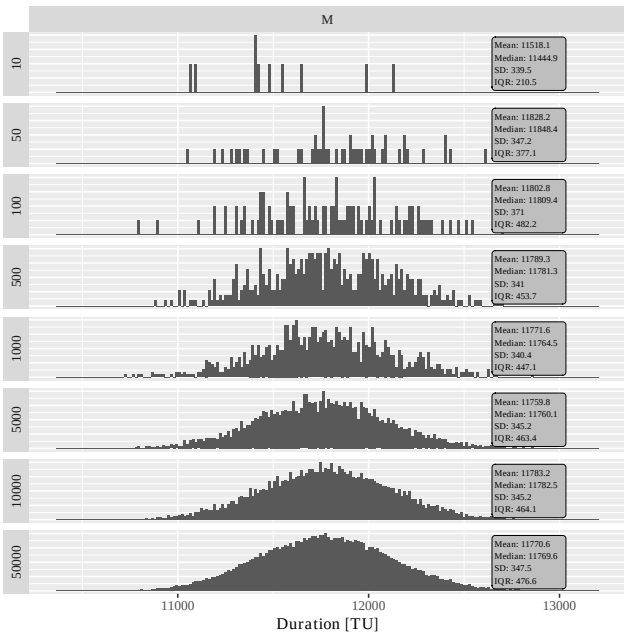

Fig. 5: Evolution of a month's utilized capacity probability distribution for manufacturing.

As can be seen from the figure, the resulting distribution tends to converge to a normal distribution shape, with the mean and median only slightly changing. Note that while this example shows the evolution towards a Gaussian distribution, it does not always happen, depending on the month and on the area. However, what seemed to be applicable to every distribution was the fact that their mean and median always converged, fluctuating only slightly. In fact, and as can be seen in figure 6, the median converges on all the areas, even though on some it does no vary greatly.

Considering that the medians of utilized capacities for each area at 50000 iterations are certainly the best estimates of the monthly utilized capacities (named as the *50000-result*), the relative error between each iteration's median and the *50000-result* was calculated, and the area-wise total is presented in table III. This information shows how accurate each simulation at different number of iterations is from the *50000-result* result.

Iterations	Total Relative Error
10	14.6%
50	6.0%
100	8.5%
500	0.7%
1000	0.3%
5000	0.5%
10000	0.8%

TABLE III: Total relative error between the areas at each number of iterations and the *50000-result*.

As can be seen from the results in table III, the relative error becomes neglectable from 500 iterations on.

Regarding the BAs, the convergence was studied in a slightly different manner: the number of interferences,

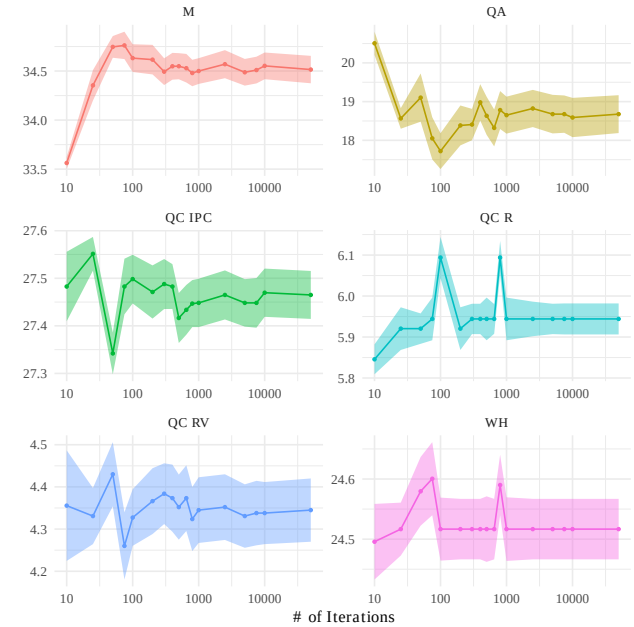

Fig. 6: Evolution of the median of a month's utilized capacity for all the areas. The shaded region corresponds to the median plus or minus $1/10 \cdot IQR$.

possible interferences and no interference was recorded per simulation and its evolution with the number of iterations was obtained. The result is shown in figure 7.

Fig. 7: Evolution of the number of BAs' interferences *per* number of iterations.

The graph shown in figure 7 shows an interesting bahavior. Initially, the number of interferences starts to decrease while the number of no interferences remains constant and the number of possible interferences increases. At around 500 iterations, this convergence stops and there is no more fluctuation in the number of interferences. This bahavior can be explained by asset utilizations that become better defined, reducing the number of interferences; these former interferences become possible interferences which explain the rise in such category. After stagnating, increasing the number of iterations will reduce the variance which leads to a reduction in possible interferences, resulting in an increase in the number of assets without interferences.

B. *Efficiency*

Efficiency of the algorithms implemented is an important aspect because it determines how long the user would have to wait for the simulation to be ran. Ideally, simulations with many iterations would be preferred for their improved results; however, the time taken by these would be too big, rendering the simulations unfeasible, since this is a tool supposed to aid in scenario exploration, where multiple simulations would be required. While code efficiency was taken into account and functions were chosen according to how efficient they were at performing their required tasks, an additional approach was used to reduce the time taken by the simulation: the adoption of parallel computation. The results of time taken by the implemented functions of parallel *versus* non-parallel computing, along the number of iterations ran are shown in figure 8.

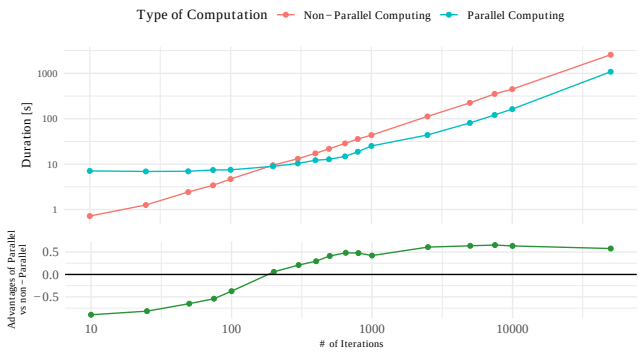

Fig. 8: Graphs of efficiency comparison: (1) time taken by each type of simulation *versus* the number of iterations; (2) *per* iterations of the simulation, the relative advantage of parallel versus non-parallel computing.

As can be seen from the graphs, for a small number of iterations the non-parallel computing is much faster, completing the simulation in around $1/10$ of the time taken by the parallel computing algorithm. This happens because the parallel computing algorithm requires the creation of a cluster of the CPU cores. This process has a certain fixed duration, but is only done once. This justifies the bahavior of the parallel computing evolution: it remains almost constant when there are not many iterations and the total time is virtually the set-up time of the cluster; the time starts to decrease comparatively to non-parallel computing at larger iterations, when the time taken by the simulation itself starts to be more significant. The advantages of parallel computing in terms of efficiency start to become apparent at 200 iterations; considering that according to the convergence analysis 500 iterations proved to be the simulation with the smallest number of iterations that had already converged, the use of parallel computing is justified.

C. *Validation*

The process of validating the model is extremely important to verify whether or not the simulations obtained are trustworthy. The basic procedure for doing so starts by selecting a scenario for the simulation that has already happened, in order to run the simulation and then compare the obtained results with the actual results. The data used for the validation regarded a period of 9 months at the pharmaceutical company in study. The data included planned start date and project for each campaign and for its validation the monthly utilized capacities in each area. Note that these capacities were not available and were therefore calculated according to the rules described in section III-C from the real processes' durations. Additionally, the first 3 months are not considered since they

are comprised of orders within the short and medium terms, and the simulation tool only considers the long-term orders. The simulation type was earliest due date, for 500 iterations. The results per area are shown in figure 9.

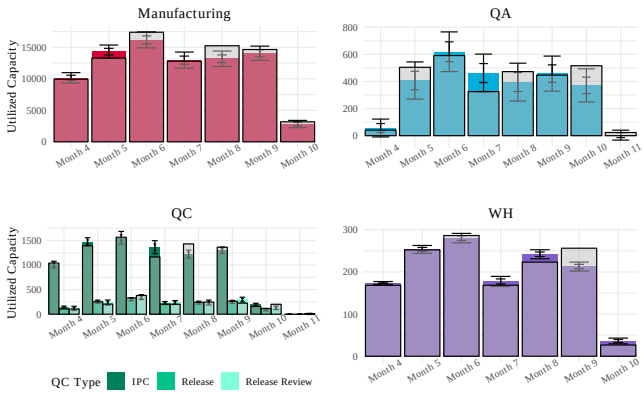

Fig. 9: Capacities validation graph. The graphs of the four areas are shown with the QC graph being divided into the 3 sections. The shaded regions over each bar correspond to the real capacities, while the color bars correspond to the simulated capacities. The error bars are distinguished between 1 IQR and 2 IQRs: the broader-width error bars correspond to 2 IQRs, while the narrower-width bar corresponds to 1 IQR. Note that QA and QC RV have utilized capacities in month 11 because these processes are always the last processes and therefore tend to occupy months ahead of the manufacturing.

As can be seen from the graph from figure 9, the simulated capacities are generally within acceptable levels of the utilized capacity. In fact, table IV shows the count and percentage of occurrences of the real capacities being within 1 or 2 IQRs of the median simulated value.

IQRs		M	QA	QC IPC	QC R	QC RV	WH
1	**Count**	2	3	2	3	3	1
	%	28.6%	37.5%	28.6%	42.9%	37.5%	14.3%
2	**Count**	6	7	4	6	6	3
	%	85.7%	87.5%	57.1%	85.7%	75%	42.9%

TABLE IV: Occurrences of monthly capacities being within the 1 or 2 IQR range

The information presented in table IV shows that there is a considerable amount of real capacities contained within 1 IQR of the simulated ones (31.5%) and even larger when considering 2 IQRs of the simulated utilized capacities (72.3%). Having in mind the rough nature of this simulation tool, it can be concluded that the results are very promising. Note that the BAs interferences were not validated since they are not intended to be precise but rather to try and solve possible clashes between the asset's occupation.

D. Forecasting

To test the forecasting capabilities of the simulation-based rough cut capacity planning tool, a specific scenario was ran, corresponding to a series of planned orders and orders on short/medium term (current orders), with their respective project and planned start date. While the current orders cannot be modified in any way, it is important to consider them in the simulation since they affect the utilized capacity in the following months. A total of 299 planned orders and 248 current orders were considered, for 500 iterations. Furthermore, the simulation type chosen was the earliest due date. The resulting monthly capacity utilization graphs for each area are shown in figure 10.

Regarding the BAs utilization, the Gantt chart represented in figure 11 shows a set of assets considered unique and unchangeable, with the clashes between tasks depicted.

An optimization step is then possible. The user can select the desired scope of the optimization, whether by the median or by 0.25, 0.5, 1 or 2 IQRs. The objective of this process is detecting and removing clashes between tasks by their median value or by a factor of their IQR and calculating the resulting utilized capacity. Considering the example previously described and an optimization by 0.5 IQRs, the resulting Gantt chart is as shown in figure 12.

In terms of monthly utilized capacity, it could be seen that while before being optimized, the results tended to be more concentrated on the initial months, greatly decreasing to near zero values afterwards, after being optimized the utilized capacities tended to be more distributed. This makes complete sense since the tasks are also more distributed, leading a greater distribution of the projects and consequently to capacities less concentrated on the initial months.

E. User Interface

Besides being a platform for delivering the simulation tool presented in section III-C, a graphical user interface (UI) was built with the objective of delivering clear, intuitive and customizable ways of showing information regarding tasks and operations happening at the production plant, as well as each area's key performance indicators. This information is presented regarding not only the present but also historical information, allowing the users to observe past tendencies and patterns, enabling their decisions to be more data-driven.

The user interface was developed using the *R* programming language and its capabilities for front-end application development, through the *shiny* package. Regarding visualization, the developed tool features 5 tabs:

- **Overview:** shows the current essential key performance indicators for the selected area, along with a map of the production plant showing either the location of the building (if a single building is selected), the location of the buildings from an area (if an area is selected), or the global map, indicating for every building the areas there operating (if the global view is selected).
- **(activities) by Building:** featuring a map of the production plant, the users can select the building according to their needs. By doing so, the map is zoomed in the building and a table shows the tasks in production, produced in the past or scheduled to be produced in the near future, depending on the date range selected. If no building is chosen, all the tasks are shown for the selected dates.
- **(activities) by Project:** users can choose the desired project displaying a schematic network of the relationships between its intermediate and final products,

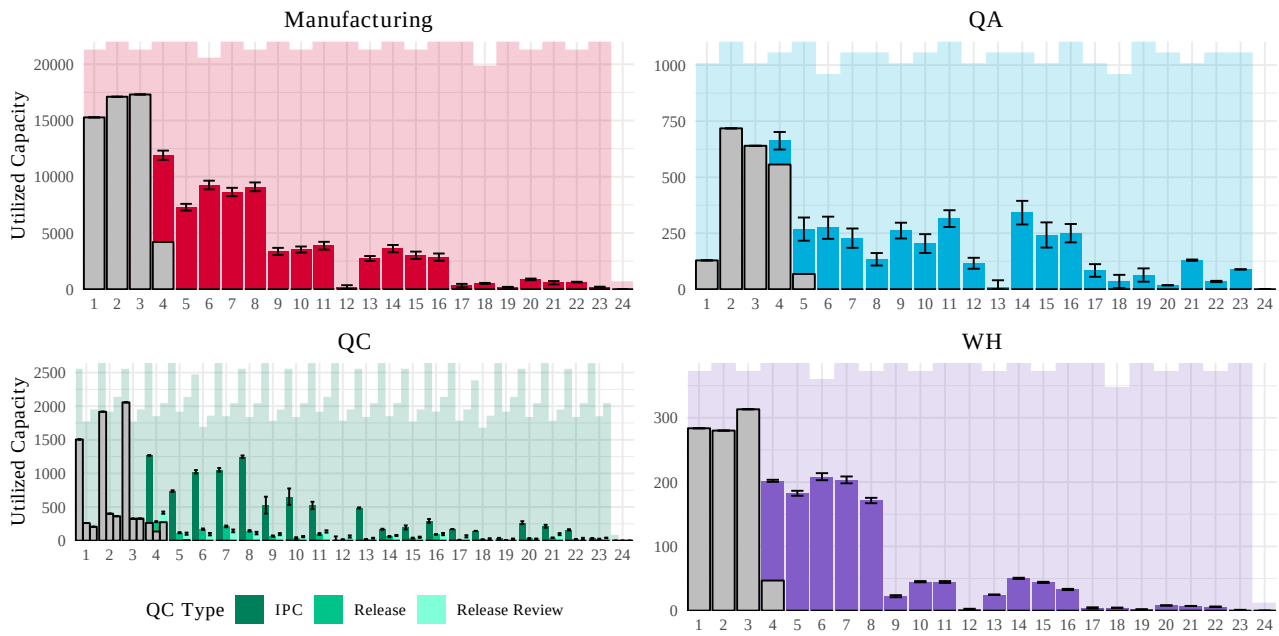

Fig. 10: Forecasted capacity evolution per month and area. The full color bars correspond to the actual simulated capacity for the month and area in question, with an error bar indicating $\pm$ 1 IQR. The grey bars correspond to the capacities from the current orders. The shaded background area corresponds to the limit capacity of each area.

Fig. 11: Gantt chart of the BA's utilization. Tasks color-coded as grey signify that they are current orders, which can no longer be modified.

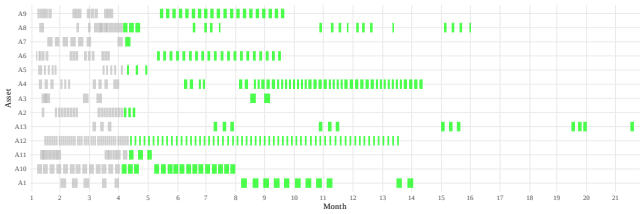

Fig. 12: Gantt chart of the BA's utilization after optimization. Grey tasks $\equiv$ *(un-optimized)* current orders; green tasks $\equiv$ planned orders *(optimized or not)*.

and the map of the production plant, displaying the buildings related in any way with the production of any of the products. All the tasks related to the project are also displayed, according to the selected date range.

- **KPIs:** tab showing the current and the evolution of the key performance indicators (KPIs). The area or building can be selected (global view is also available) showing the current 6 most representative KPIs at the desired timescale (weekly, monthly or yearly). Additionally, each KPI can be chosen for the user to view its historical development, on the selected timescale.
- **Schedule:** a schedule of the operations is shown by its production line. The activities are shown according to a specific date range, and the results may be filtered by building or by project.

Furthermore, the UI provided the platform for the user to connect with the simulation tool, through 2 tabs:

- **Project Database:** a comprehensive database of each project, allowing the user to view their historical information, a Gantt of its processes, the histogram of its manufacturing and QR durations, its bill-of-materials and its usual adherence to start date. This can be helpful for stakeholders in charge of using the simulation tool, giving them clearer insights into how campaigns of a certain project tend to occur.
- **RCCP:** tab that offers the simulation tool itself. Besides delivering the simulation results and allowing the optimization as explained in the previous sections, the tab allows the user to view the existing planned orders, edit their start dates and deadlines and create their own user-defined orders. Additionally, previously saved simulation scenarios can be loaded (and saved), the effects of specific campaigns can be seen on the capacity and asset utilization graphs and a file containing the results can be exported for offline view.

An example of a screenshot of the developed UI is as shown in figure 13, where the (Projects) by Building tab is shown.

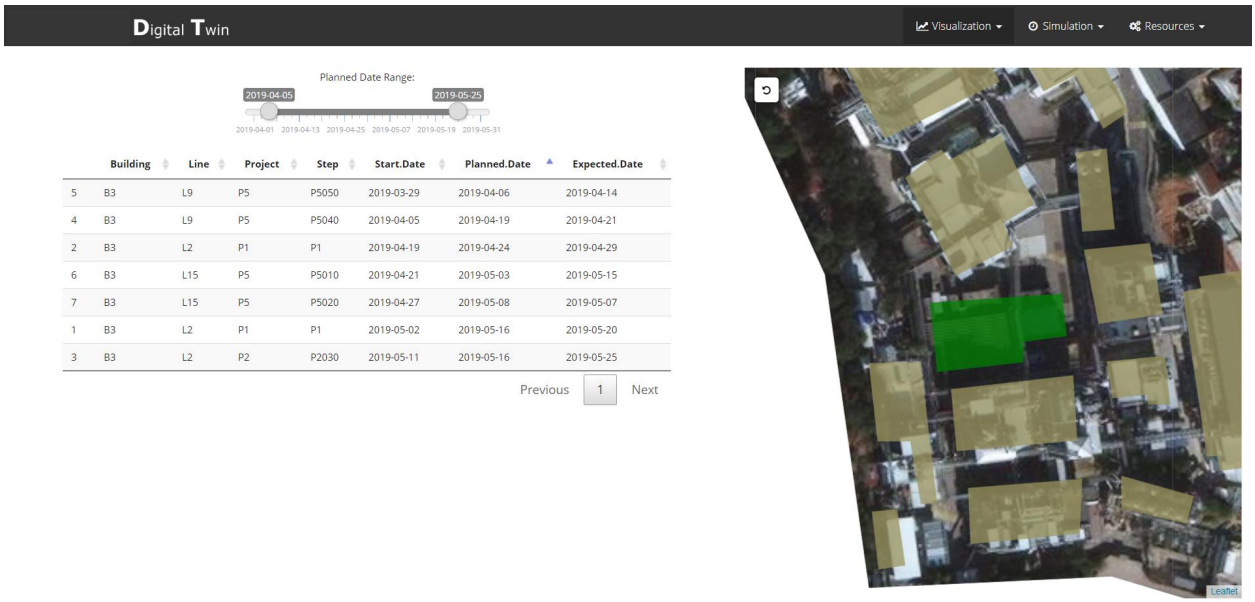

Fig. 13: (Projects) by Building tab, from the developed UI.

V. Conclusions

The developed tool was successful in the two objectives that were proposed: visualization and simulation. The visualization component of the tool was able to deliver intuitive views into the operations happening at the production plant at a given time; several methods of delivering these views were implemented in order for the users to view the information in the way that best fits to their needs. The simulation tool created was able to successfully generate accurate forecasts (in the scope of a rough-cut tool), regarding the necessary monthly capacity, or the monthly percentage of the maximum capacity needed. The simulations ran were validated and showed promising results for a rough-cut tool, with the possibility of improving the base capacity utilization estimations made the ERP, since demonstrated performance values were considered instead of the values strictly derived from the recipe.

Acknowledgements

The author would like to thank the pharmaceutical organization where this project was done, *Hovione Farmaciência S.A.*, the Institute of Mechanical Engineering of the *Instituto Superior Técnico*, *IDMEC – IST*, and his supervisors, both at company side and at the university side.

References

[1] R. N. Bolton, J. R. McColl-Kennedy, L. Cheung, A. Gallan, C. Orsingher, L. Witell, and M. Zaki. Customer experience challenges: bringing together digital, physical and social realms. *Journal of Service Management*, 29(5):776–808, 2018. doi:10.1108/JOSM-04-2018-0113.

[2] A. Costigliola, F. A. Ataíde, S. M. Vieira, and J. M. Sousa. Simulation model of a quality control laboratory in pharmaceutical industry. *IFAC-PapersOnLine*, 50(1):9014–9019, 2017.

[3] D. Ivanov, A. Dolgui, A. Das, and B. Sokolov. Digital supply chain twins: Managing the ripple effect, resilience, and disruption risks by data-driven optimization, simulation, and visibility. In *Handbook of Ripple Effects in the Supply Chain*, pages 309–332. Springer, 2019.

[4] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn. Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11):1016–1022, 2018. doi:10.1016/j.ifacol.2018.08.474.

[5] A. M. Law. *Simulation Modeling and Analysis*. McGraw Hill Education, international fifth edition, 2015. ISBN 978-1-259-25438-3.

[6] M. R. Lopes, A. Costigliola, R. M. Pinto, S. M. Vieira, and J. M. Sousa. Novel governance model for planning in pharmaceutical quality control laboratories. *IFAC-PapersOnLine*, 51(11):484–489, 2018.

[7] A. M. Madni, C. C. Madni, and S. D. Lucero. Leveraging digital twin technology in model-based systems engineering. *Systems*, 7(1):7, 2019.

[8] J. T. Mentzer, W. DeWitt, J. S. Keebler, S. Min, N. W. Nix, C. D. Smith, and Z. G. Zacharia. Defining supply chain management. *Journal of Business logistics*, 22(2):1–25, 2001. doi:10.1002/j.2158-1592.2001.tb00001.x.

[9] A. Mullard. 2018 FDA drug approvals, 2019. doi:10.1038/d41573-019-00014-x.

[10] V. Papavasileiou, A. Koulouris, C. Siletti, and D. Petrides. Optimize manufacturing of pharmaceutical products with process simulation and production scheduling tools. *Chemical Engineering Research and Design*, 85(7):1086–1097, 2007. doi:10.1205/cherd06240.

[11] N. Shah. Pharmaceutical supply chains: key issues and strategies for optimisation. *Computers & chemical engineering*, 28(6-7):929–941, 2004. doi:10.1016/j.compchemeng.2003.09.022.