Robust Advanced Modeling and Scheduling System (RAMS): From Space Exploration to Real-World Biomanufacturing

Michael Saint-Guillain (michael.saint@uclouvain.be)
Université Catholique de Louvain
Isabelle Sansdrap
Zentech
Michaël Brion
Takeda (Belgium)
Eric Pauly
Cellaïon

Case Report

Keywords: biomanufacturing, bioprocess engineering, industrial processes, scheduling, modeling

Posted Date: October 11th, 2022

DOI: https://doi.org/10.21203/rs.3.rs-2093105/v1

Robust Advanced Modeling and Scheduling System (RAMS): from Space Exploration to Real-World Biomanufacturing

Michael Saint-Guillain[1*], Isabelle Sansdrap[2†], Eric Pauly[3†] and Michael Brion[4†]

[1*]Université catholique de Louvain, Belgium.
[2]Zentech, Belgium.
[3]Cellaion, Belgium.
[4]Baxalta Belgium Manufacturing (Takeda), Belgium.

*Corresponding author(s). E-mail(s): michael.saint@uclouvain.be;
Contributing authors: i.sansdrap@zentech.be;
eric.pauly@cellaion.com; michael.brion@takeda.com;
†These authors contributed equally to this work.

Abstract

Advanced planning and scheduling (APS) systems aim at helping production and operation managers in organizing the manufacturing process of biotech products, finding near-optimal schedules meeting the demand, while taking operational resources and raw materials into account. We introduce an extension to APS, called *Robust Advanced Modeling and Scheduling* (RAMS), and present the first RAMS system: *Rombio*. Unlike existing APS systems, *Rombio* allows to (i) visually model the operational problem and context entirely (ii) generate and optimize schedules while taking uncertainty into account, and (iii) deal with a combination of various key performance indicators (KPIs). Probability theory enables us to cope with uncertainty, computing schedules that are robust to temporal deviations. Depending on the pursued KPIs, the resulting schedules optimize a combination of the following terms: the probability of satisfying the process constraints, the expected return/-efficiency/quality, and even the operators' wellness by minimizing its expected extra-hours. This introduces a new risk-aversion paradigm, replacing the well-known what-if and sensitivity analysis frameworks. Initially developed in collaboration with Nasa for space exploration, this versatile tool is applied to real-world biotechnology manufacturing in three different contexts: *diagnostics*, *medicines* and *stem cells*.

Keywords: biomanufacturing, bioprocess engineering, industrial processes, scheduling, modeling

1 Introduction

Project management is the cornerstone of efficiency in modern business. It is also one of the most challenging problems to solve, in particular when it comes to human project management, that is, human scheduling. Yet, most of the existing studies have only been done in the context of machine scheduling [1]. According to [2], project management, which also refers to human scheduling, accounts for about a third of the world's gross product.

In this paper, we introduce a visual tool, *Rombio*[1] (Robust Operations Management for Biotechs), and show how it has been successfully applied to three different case studies of real world human operations management, in three Belgian biotech companies: *Zentech*, *Takeda* and *Cellaïon*. Zentech is active in the manufacturing of screening and diagnostic kits. Takeda is a multinational pharmaceutical company, active in drug design and manufacturing. Cellaïon is active in stem cell manufacturing. Altogether, these three companies thus form a well diversified sample, in terms of product types and hence, manufacturing processes.

The system is novel in different aspects. First, it generates and optimizes schedules while taking uncertainty into account. Second, it allows the end-user to express the problem at stake directly in the graphical interface, using a visual modeling formalism. Third, it deals with a combination of various key performance indicators (KPIs), including the success probability of the optimized schedules. The technology embedded in *Rombio* is based and validated on extensive theoretical (as well as empirical) results in the domain of planning and scheduling under uncertainty (see Section 1.2).

Results

The results obtained from our three case studies convey three very important messages: **(a)** Even for very complicated and various different operational contexts, a common modeling framework exists, being user friendly, visual, and rigorous at the same time; **(b)** Even for real sized problems, computer-optimized solutions outperform the schedules hand-crafted by field experts in general, and serve as a strong basis for decision making, as the deciders can always adapt and reuse these depending on external factors; **(c)** Schedules obtained while taking uncertainty into account systematically outperform those obtained from deterministic assumptions in terms of reliability and

[1] Also known as *Romie* in the literature, when applied to space exploration.

expected KPIs, while preserving most of the solutions quality; the latter result remains valid even when provided very bad representation of the uncertainty.

1.1 Scheduling tools, APS and (R)AMS systems

As mentioned in our introduction, (planning and) scheduling problems constitute a very common challenge in the modern industry. As a consequence, several concepts and tools have been proposed. [3] present some of the most common ones (MS Project, Primavera, Artemis, *etc*).

There is a fundamental difference between software systems under the very large denomination of "*scheduling (or planning) tools*", and so-called "*advanced planning and scheduling*" (APS) tools [4]. Whereas classical tools somehow solely offer the ability to visualize and manipulate schedules, an APS comes with an optimization engine that automatically generates and improves schedules based on predefined generic constraints. In other words, the vast majority of non-APS tools, such as MS Project, are just complex Gantt chart visualization and edition systems. However, a Gantt chart implies having designed a schedule (*i.e.* ordered sequences of tasks scheduled in time and assigned to resources)! And the user is still left with that problem. Only an APS system, or a more elaborated one (such as RAMS we describe hereafter), is actually able to solve the underlying constrained combinatorial problem: the scheduling problem.

[5] already provides a comprehensive summary of the related frameworks and software systems, in space as well as in the industry. In particular, [6–8], among others, showed the limits of human self-scheduling in the context of space missions. The main differences between an APS and a (R)AMS comes from the limitations of APS systems, which can be summarized as follows. Any APS system either falls into *a)* being specifically designed for a particular application or operational context or *b)* not being able to generate *robust* schedules in light of uncertainty. This motivates a *(robust) advanced modeling and scheduling* (RAMS) system, and in particular the one presented in this paper being the first RAMS, having the following *technological innovations*:

1. A domain-independent graphical modeling interface, allowing the user to graphically draw the structure and constraints of its scheduling problem.
2. Thanks to recent results in probability theory, *Rombio*'s optimization engine allows the end-user to rapidly generate schedules that are reliable, robust to temporal deviations.

We often put the R in RAMS between parentheses, to highlight the fact that the robustness aspect of the problem is actually not the most central concept here. Unlike (R)AMS systems, existing APS tools do not allow the end-user to express complex requirements and constraints, and thus fail at meeting the resource and constraint structure involved in problems as complex as biomanufacturing. The only exceptions come from APS systems that are tailored for some problems, or in other words, in which these constraints are hard-coded.

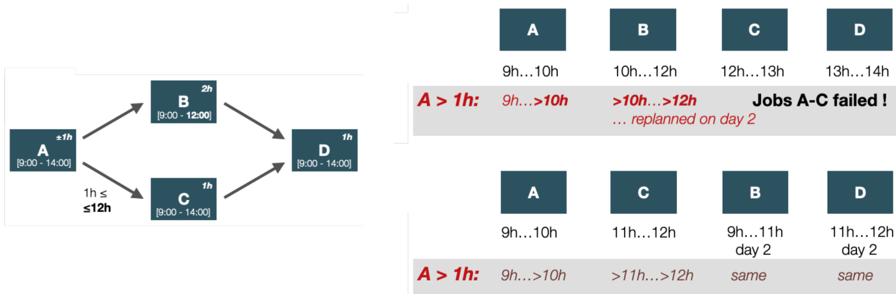

Fig. 1 Our illustrative example. There are only four tasks: A, B, C, D, to be scheduled on the same line. Each task has a processing time of 1 or 2 hours, and a time window spanning either the entire work day (9am to 2pm) or part of it (9am to 12am). Starting tasks B and C both require A to be completed, and D requires B as well as C to be completed. Task C must wait at least one hour after the end of A to begin, and should be completed during the same day.

Yet, such dedicated systems are terribly expensive. Most stakeholders can't afford it, especially when their constraints frequently evolve.

1.2 Theoretical Foundations

In stochastic contexts such as space missions or biomanufacturing, computing optimal schedules becomes significantly less attractive as problem data, such as the manipulation time of the modelled activities, are different from their predicted nominal values. This is what we refer to as *uncertainty*. In a constrained environment with shared resources and devices, when they arise such temporal deviations can propagate to the remaining operations, eventually leading to global infeasibility, that is, a project failure. Given a schedule, a central question is then the following: considering temporal uncertainty, what is the actual probability of success of the project?

Illustrative example

Consider the simple project depicted in Figure 1 (left). Suppose all tasks have to be scheduled on the resource, then one must necessarily begin with A and end with activity D. There are only two valid schedules, shown in Figure 1 (right). In fact, schedule (A, B, C, D) looks much more efficient, as all tasks are completed on the first day. On the contrary, schedule (A, C, B, D) requires an additional day. However, this is *only true on the paper*, when everything is predictable. If you account for (temporal) uncertainty, then the story is different. If operation A lasts for more than 1 hour, schedule (A, B, C, D) is *not valid anymore*: B will have to be resumed or rescheduled at day 2 (an additional day, that was not expected!). When starting C, we realize the worst: it actually requires to be processed the same day as A. Mission failed. Remark that if the true average processing time of A is 1 hour, then this scenario happens with at least 50% probability. On the contrary, (almost) whatever happens to A, under schedule (A, C, B, D) everything goes fine. This schedule is said to be *robust*. Its success probability is simply 1 minus the probability that A exceeds four hours (which we assume to be fairly unlikely).

Project management is hard

The problem scheduling a set of operations under constraints should be seen as a generalization of the well-known *NP*-complete *job-shop scheduling problem* [9], which has the reputation of being one of the most computationally demanding [10]. When taking uncertainty into account, the problem then becomes strongly *NP*-hard, an even harder family of problems. In a nutshell, *NP*-complete means that, no matter the available computational resources, the problem is conjectured as impossible to solve in practice, for realistic instance sizes, such as the number of activities and resources. In fact, whereas the problem depicted in Figure 1 admits only two solutions, in practice the number of possible schedules grows exponentially with the number of tasks and resources.

Previous researches

Based on the real case study of a Mars analogue mission in 2018, in [11] we proposed a first (incomplete) probabilistic formulation, as well as solution method, for the problem of scheduling a set of various human operated projects. In fact, the problem of scheduling a set of operations in a constrained context such as the *Mars Desert Research Station* (MDRS, Fig. 2) is not trivial, even in its classical deterministic version. We hence measured the gains and costs, on a priori mission planning, of robust schedules (optimized under uncertainty) compared to schedules optimized under classical deterministic assumptions.

In [12], the theoretical insights obtained from the former study were successfully extended to *probabilistic simple temporal networks* (PSTNs), a formalism able to mathematically describe operational problems in general, such as scheduling a space mission or a biomanufacturing campaign. In this paper written with the Jet Propulsion Lab (NASA), our probabilistic model is applied to the operation management of Mars 2020 planetary rover. We also formally define some of the most important theoretical concepts for describing schedule robustness to uncertainty, we introduce new ones, and give proofs for theoretical bounds. This contributed to filling the theoretical gap between specific mission planning and general operations management.

In [5] we proposed a very first version of the RAMS system, in which the user interface (in particular, the graphical modeling interface) was at a very preliminary stage, and showed how we successfully used the system to solve operational scheduling problems in three very different contexts: an analogue human space mission, a robots parameterization problem for cave exploration, and a first application to biotechnology manufacturing.

On novice self-scheduling

These previous studies mainly aimed at evaluating and demonstrating, both empirically and theoretically, the need and the advantages of using probabilistic assumptions (*i.e.* optimizing under uncertainty) in the context of operations

Fig. 2 Left: the Mars Desert Research Station in Utah. Right: extra-vehicular field operations.

management. In the context of space exploration, past missions (*e.g.* UCL to Mars 2018 [11]) have shown the importance of online reoptimization and, in particular, the need for the crew to autonomously adapt their science projects to unforeseen events. In [13] the scientific focus has been rather put on the user experience. More specifically, using techniques from human computer interaction (HCI), we measured how well a team of novice users succeeded (or not) at using our system to schedule, and reschedule online, their own activities. We obtained positive and encouraging results. Our RAMS system has demonstrated its applicability on a simulated mission in a Mars analogue habitat, during which a team of 8 analogue astronauts evaluated the *Rombio* system on their ability of self-schedule their own mission (both a priori and on-the-fly).

A new risk-aversion paradigm

What if analysis and *sensitivity analysis* are classical, well-known techniques for coping with uncertainty in operations management. In fact, [14] and [15] both argue for the importance, biomanufacturing, of simulation and the ability of performing what-if and/or sensitivity analysis in addition to optimization. What if analysis consists of optimizing several solutions (usually a few numbers), each solving a predefined scenario, such as best-case, average-case and worst-case scenarios. In [12] we formally prove[2] that the well known *what-if analysis* technique is fundamentally flawed, as it arbitrarily underestimates a schedule's risk. On the contrary, Rombio's optimization engine is proven to never underestimate it. Provided a schedule[3], *sensitivity analysis* approximates its average quality under uncertainty, how sensible (brittle) it is to stochastic variations. In that sens, the solutions computed by Rombio directly optimize their response to a sensitivity analysis. The proposed RAMS framework introduces a new paradigm, in which both what-if analysis as well as sensitivity analysis are turned deprecated.

---

[2] In [12], the degree of weak controllability (DWC) can be interpreted as a perfect what-if analysis, that is, when not only considering best, middle and worst cases, but all the possible scenarios. The demonstration is then reached at inequality (11), with the result $\mathrm{DDC}(N) \leq \mathrm{DWC}(N)$, where DDC is the true robustness of the schedule N.

[3] This however does not help at finding the right schedule in the first place!

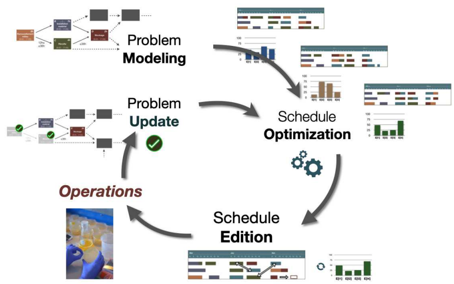

Fig. 3 Principal functionalities of *Rombio* tool.

1.3 Paper Contributions and Organization

In this paper, we present an innovative general tool, which can be configured to meet a large range of operational contexts, from space missions to various biomanufacturing domains, without domain-specific developments. This tool falls into the newly proposed category of Robust Advanced Modeling and Scheduling (RAMS) systems. This category aims at extending that of advanced planning and scheduling (APS), by enabling the end-users to model the operational problem at stake, and by computing schedules that are robust to uncertainties. The proposed RAMS system, *Rombio*, is described in Section 2.

Initially developed for scheduling the daily activities of a team of astronauts, in the context of a Mars analog mission, and later improved in collaboration with Nasa, and experimentally applied to Perseverance planetary rover's operations management under uncertainty, the technology has eventually been further extended to cope with operations management in biomanufacturing processes. Most of this paper is dedicated to the description made in Section 3 of the application case studies, conducted with three different Belgian companies involved in biomanufacturing: *diagnostics*, *medicines* and *stem cells*.

All three use cases are conducted using the same system. This is made possible thanks to both the uncertainty management and, above all, the graphical process modeling interface, which are at the core of the proposed new framework.

2 Overview of *Rombio*

Rombio is a web-based software (SaaS) system aiming at supporting the decision makers in their operations management and task scheduling. The key functionalities are depicted in Figure 3:

- User friendly, *visual modeling* of the problem at stake, in its own operational context: human and physical resources, operational constraints, key performance indicators (KPIs), execution uncertainties.

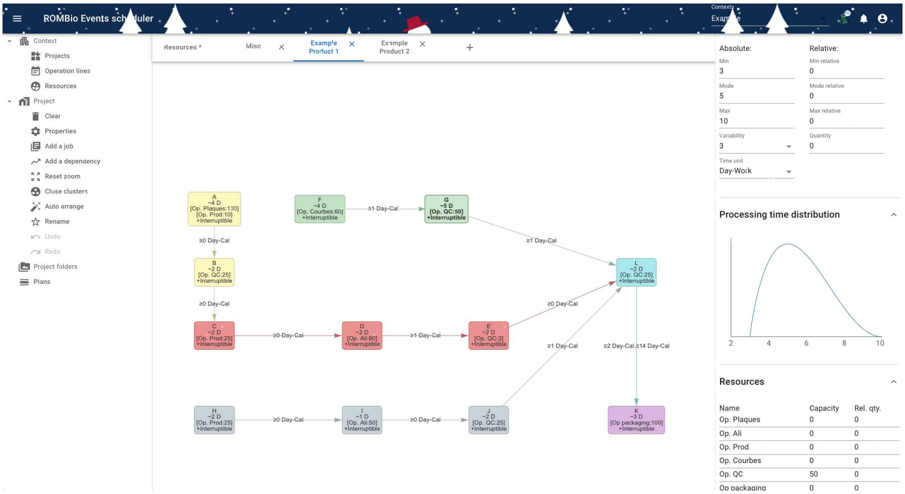

Fig. 4 The graphical modeling interface. An example of a model is shown, representing a manufacturing process involving 12 activities. Each activity is fully parameterizable: name, (random) processing time, resource usage, allowed types of operators/lines, *etc.*

- *Robust scheduling*: the optimization engine takes the time uncertainty on each task's duration into consideration, using modified-PERT distributions, yielding schedules with high probability of success.
- *KPI-guided scheduling*: The schedules are optimized while pursuing (a combination of) *various KPIs*, including success probability, expected cost, expected quality, and even operator wellness.

In a research domain in constant evolution, *Rombio* integrates state-of-the-art advances in robust scheduling under uncertainty [12]. Future versions will enable online monitoring of the operations, keeping the schedule and the underlying model consistent with the current state of the system, allowing the user to adapt and reoptimize future decisions based on past outcomes.

The Modeling Interface

Depicted in Figure 4, it provides to the user the ability to create, edit and manipulate process models. Operations can be added, deleted, moved around, etc. Arrows are used to denote constraints between pairs of activities. The classical basic constraint $A \rightarrow B$ is a dependency (or precedence) constraint, stating that activity B may only be scheduled after the completion of A, however more elaborated constraints can be added to the model, including stochastic delays between activities. In fact, both operations and constraints are fully parameterizable, as explained in Figure 4. Each random duration is characterized by 4 parameters: minimum duration, most probable duration (mode), maximum duration and variability ranging from 1 to 5. Together, these four parameters define a modified-PERT distribution, a probability distribution widely used in risk analysis [16], which has the advantage of enabling

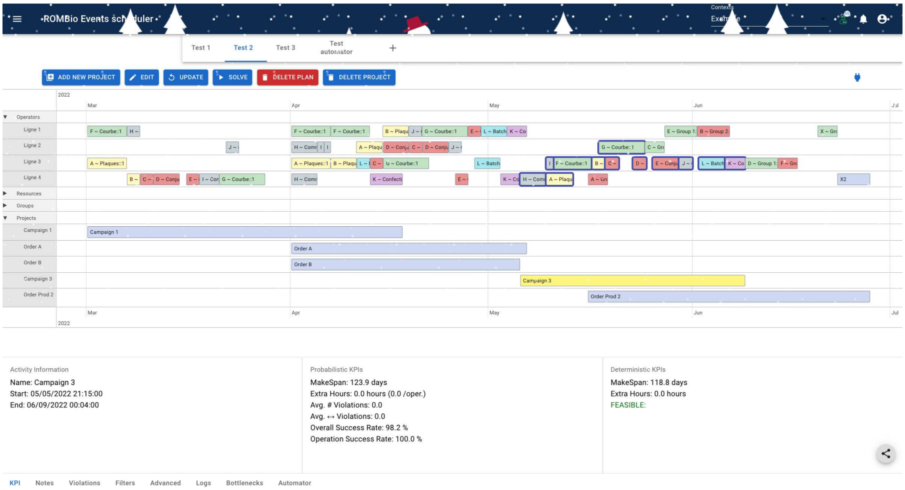

Fig. 5 The scheduling interface. Here for instance, 5 different manufacturing campaigns are scheduled; the first 4 come from the same model ("Example Product 1", Fig. 4), whereas the last project "*Order Prod 2*" has been inserted following a different model ("Example Product 2", not displayed in Fig. 4). The probabilistic (*resp.* deterministic) KPIs of the current schedule are listed at the centre (*resp.* right) of the bottom area, such as the probability of success (here: 98.2%).

asymmetric bounded probability distributions. More specific aspects of the modeling interface will be presented during the application cases of Section 3.

The Scheduling Interface

Depicted in Figure 5, it permits to visualize the existing schedules, but also create and optimize new ones, potentially based on different operational assumptions (such as resources, as extensively applied in Section 3.2) and/or models. Schedules can be duplicated, and the user may have an arbitrary number of different schedules, which allows for manual experiments. Starting from either an existing or an empty schedule, the user may modify it in 4 ways: add a project, remove a project, manual edition or optimize. Adding a project means inserting, in addition to already planned projects, all the activities corresponding to a particular model. In practice, a schedule will typically contain several occurrences of the same model: in fact, a model usually stands for a particular product, and several manufacturing campaigns (*i.e.* projects) may be planned for a given product. In fact, adding a new project to an existing schedule usually ends up in an infeasible solution, that is, a schedule in which some of the constraints, deadlines or resource usage are not respected. Whereas the interface allows for a manual edition of the schedule, the user may also ask the optimization engine (which runs on remote computation clusters) to "Solve" the problem, resolving the violations whenever they are, and eventually optimize the predefined KPIs. The displayed schedule is interactive: by clicking on any element on the chart, additional information is displayed, and possible further actions are proposed. The interface also shows both the deterministic and probabilistic KPI values, such as for instance, the overall success probability of the schedule. Finally, the user may also share a given schedule, by generating an URL that can be used by someone else to visualize the schedule in read-only mode (without accessing anything else), while enjoying the interactive functionalities. Other aspects of the scheduling interface will be presented during the application cases of Section 3.

3 Real-World Application Cases

This study presents a novel scheduling tool through three different case studies, involving three different biotech companies. Each case study happened sequentially, following and further validating different stages of the tool's development.

At the time of the first case study (Section 3.1), conducted in collaboration with *Zentech* company in 2019 and 2020, the graphical user interface (GUI) was not developed yet. The tool only consisted of a theoretical background, a modeling formalism and a versatile scheduling engine. The main research questions were then to: *(a)* empirically assess, on field, the ability of our collaborators to understand, reuse and further modify the proposed visual modeling formalism; and *(b)* determine the actual impact of uncertainty on their specific modelled manufacturing process, that is, the gain of schedules optimized in light of uncertainty over schedules obtained based on classical deterministic assumptions.

Section 3.2 describes the second case study, in collaboration with a biomanufacturing plant based in Belgium, and belonging to the multinational Japanese company *Takeda*. The study took place in 2021. *Rombio* is now given a brand new GUI prototype, allowing (for the first time) an end-user to model the problem using a user-friendly visual interface, run optimization processes and visualize optimized schedules. As the modeling framework and the technology were validated by the *Zentech* use case, the concerns were again two-folds: *(a)* the applicability of proposed technology to a completely different company and operational context; and *(b)* the practical utility of such technology in a provisional scheduling and forecasting context.

Finally, in Section 3.3 we eventually address one of the most important remaining questions: the ability of the end user to control and use our tool, through the provided GUI, including the understanding and use of the graphical modeling interface. In addition, this case study permitted us to confront our technology to another, significantly different operational context: the complex manufacturing process of living stem cells. This use case took place in 2021, with the Belgian biotech company *Cellaïon*.

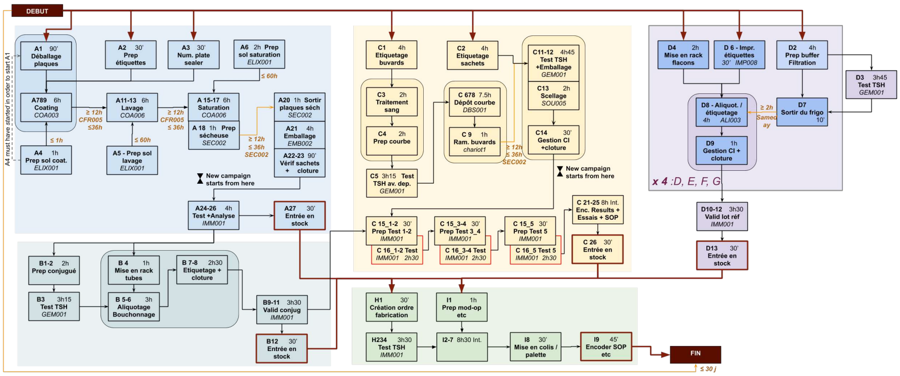

Fig. 6 A visual model that fully describes the tasks and constraints involved in one production campaign, for *Zentech*'s most popular product. Such a campaign involves around 85 activities, the most of these being subject to uncertain processing times. There are many dependencies between the activities, some with temporal constraints of the form *"activity B should be operated at most 36 hours after the actual completion time of A"*, which are naturally particularly problematic under temporal uncertainty.

3.1 Diagnostic Kits at Zentech: the modeling quest, and the price of robustness

Starting from the study conducted at the Mars Desert Research Station by [11], and as the day humans will live on Mars is still far ahead, we wanted to extend our tool as well as the underlying technology to tackle significantly different operational context, others than human space missions. The study was conducted with the Belgian company *Zentech*, specialized in the manufacturing of screening and diagnostic kits. They agreed to collaborate on the concrete project of modeling the scheduling problem involved in the manufacturing of one of their most popular products, and eventually solving this scheduling problem, at different scales. As observed and discussed further in this section, the impact of uncertainty tends to naturally increase with the problem size. From a development point of view, our collaboration permitted us to extend our scheduling formalism and technology and reach a higher stage of applicability, in the complex real world industrial context. Eventually, it also provided the very first occasion to confront the stakeholders, who were not planning experts, to our visual modeling formalism.

3.1.1 Modeling in the Industrial Real World Contexts

Figure 6 shows how their production problem was modelled, using the exact same *visual formalism* that is currently implemented in *Rombio*'s GUI. At the time of this use case however, the model was drafted by hand and further translated in a formal language, specific to our technology. Nowadays however, the GUI includes a modeling part, and the translation is automatic. The model depicted in Figure 6 only represents one single production campaign, involving around 85 tasks. The company would usually run up to three production

Fig. 7 A possible optimized solution to the problem described by the model depicted in Fig. 6 when three production campaigns are scheduled in parallel. The visualization includes the operators' task sequences as well as the resource (machines) usages. We see how the campaigns overlap at the bottom with the blue, green and purple bars: each denotes a campaign.

campaigns in parallel, whereas the operational human and physical resources remain fixed.

Figure 6 in fact shows the *graphical language* used in order to describe the activities and constraints involved in one manufacturing campaign of their product. It relies on simple diagrams that are easily understandable by our collaborators from the biotech company. In practice, the time needed by our collaborators, being novice schedulers, to master the proposed modeling formalism and reuse it revealed to be of five to ten one-hour meetings only.

3.1.2 Computing Robust Schedules

Experiments have been conducted by asking the Rombio system to generate and optimize schedules, for various problems involving from one to three production campaigns. Each time, 10 schedules were optimized under uncertainty, and 10 other schedules were optimized under classical deterministic assumptions. In fact, the artificial intelligence behind the optimization engine is actually quite likely to output different schedules each time at each run. An example of a schedule, involving three campaigns, is shown in Fig. 7.

Uncertainty on the Uncertainty

Unlike many other applications, project management (and especially human operations) suffers from the lack of historical observations. There is no "big data". As a consequence, the real probability distributions that describe our uncertainty (here, the activity durations) remain hidden, and difficult to estimate in practice. Our experiments hence take the uncertainty on this stochastic knowledge itself into account. Figure 8 illustrates how we proceeded. Three hypothetical couples of both estimator and real distributions are drawn. The estimator distribution is simply the distribution used to describe the duration of an activity. The estimator (blue) represents the current knowledge one has

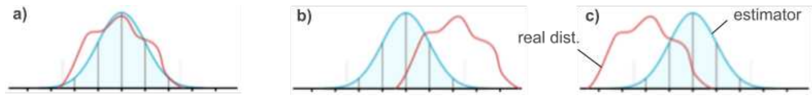

Fig. 8 Varying the quality of the stochastic knowledge, leading to three different experimental assumptions. Blue: estimated distribution, used at optimization stage. Red: real *hidden* distributions, revealed at execution stage. Assumptions: **a)** the estimators are of good quality, their means coincide with that of the real distributions; **b)** all activity processing times are globally underestimated; **c)** processing times are globally overestimated.

about the activity's uncertainty (red, hidden). A schedule is therefore optimized provided estimators only, whereas it is further evaluated by simulation on the real distributions. Note that case *c)* corresponds to a somewhat usual situation, in which the project manager tries to mitigate the risk by systematically overestimating the durations. We will soon see that one can do much better.

Optimizing Wellness: Stress Aversion

Classical KPIs, at least those considered in optimization, are usually cost-based (*e.g.* minimizing total production time). In human scheduling, stress aversion and wellness in general are key KPIs to consider as well, especially in the long term. The biotech company indicated that situations in which the employees are forced to do extra-time work, in order to stick with operational constraints and deadlines, necessarily result in an increased stress level. Therefore, we decided to consider the *expected total number of extra-hours* as one of our KPIs.

3.1.3 Deterministic VS probabilistic schedules

The average results of our experiments are given in Table 1. It is not surprising to see that the average *reliability of deterministic schedules significantly decreases as the size of the problem increases*. Here we only describe the context in which we make the (optimistic) assumption that all the durations have been globally over-estimated, because this is usually what people intend to do. When optimizing the efficiency (makespan) first (MS), moving from one to three campaigns decreases success rate from 17.1% to 2.5% for deterministic solutions, whereas probabilistic ones decrease from 98% to 96%.

A better robustness is globally reached when optimizing the extra-time first (EH), instead of makespan (MS). In fact, the less planned extra-time, the most likely the schedule is to be able to absorb unexpected deviations, and therefore the more flexible it is in the end.

Given three manufacturing campaigns and while minimizing the extra-time first (EH), deterministic solutions reveal an average of 11.3 extra hours (whereas only 1 hour was initially planned!). Probabilistic ones pass the simulations with only 3.8 extra hours on average, which is by the way very close to the amount of extra time initially planned (3.0). Therefore, the actual average amount of *unexpected extra-time can be divided by more than 3*, when optimizing under probabilistic assumptions.

		% Success			Extra-Hours				Makespan (days)			
		Exa.	Und.	Over	Plan	Exa.	Und.	Over	Plan	Exa.	Und.	Over
1C	MS	6.6	0.9	**17.1**	9.9	9.7	11.3	10.0	15.1	15.7	16.6	15.4
Det.	EH	60.9	37.5	40.2	0.0	4.6	5.9	4.3	17.2	19.8	20.7	18.5
1C	MS	98.4	83.7	**98.0**	1.4	4.1	6.1	3.1	17.2	17.2	17.3	17.2
Prob.	EH	99.9	99.6	100.0	1.0	1.9	2.7	1.0	26.9	27.1	27.4	27.0
3C	MS	0.5	0.0	**2.5**	16.6	29.1	34.6	26.7	38.1	41.9	42.5	**40.1**
Det.	EH	14.2	2.2	12.5	**1.0**	15.7	23.6	**11.3**	45.5	46.4	46.8	45.9
3C	MS	96.2	61.1	**96.0**	6.2	13.0	20.0	10.3	43.2	43.2	43.3	**43.2**
Prob.	EH	99.4	97.7	98.7	**3.0**	6.4	9.9	**3.8**	58.9	58.9	58.9	58.9

Table 1 Zentech case study: one (1C) and three (3C) manufacturing campaigns, optimized under either deterministic (Det.) or probabilistic assumptions (Prob.). Three different assumptions on a priori stochastic knowledge: exact mean, 10% under-estimations and 10% over-estimations. Results indicate the percentages of simulations (% Success) in which the optimized schedules respect all the problem constraints, when simulated using the "hidden" probability distributions. The *Plan* columns indicate the KPI values as predicted by the *a priori* schedule. Other columns show the average KPI values observed during simulations. *MS*: minimizing the cost KPI (makespan) first, extra-time second. *EH*: extra-time first, cost second.

The price of robustness

The difference of efficiency between probabilistic and deterministic solutions, when not taking the success rate into account, is what we call the price of robustness. Here for instance, when optimizing makespan first (MS), deterministic solutions to three campaigns are found with schedules of 40.1 days on average (no matter whether these schedules eventually succeed), whereas probabilistic solutions involve schedules lasting 43.2 days on average. In other words, the probabilistic approach suggests to *sacrifice 8% of the theoretical efficiency, in order to obtain schedules that are 38 times more reliable* in practice (2.5 to 96), and that even when all the processing times are globally over-estimated in the first place. Interestingly, a similar result was obtained in [11], in which the price of robustness was on 7% on average in the context of a space mission.

3.1.4 Conclusions and insights from the Zentech use case

This use case marked the very first proof of concept, in the industrial world, of the applicability of our technology. Empirical evidence showed that, without many historical observations and even provided a very bad a priori knowledge of the uncertainty, the optimized schedules significantly outperform those obtained using classical deterministic assumptions, while preserving most of the solution's quality. This is also true when all processing times are consciously overestimated. Moreover, we showed that using a probabilistic model enables us to optimize the operators' stress wellness, by minimizing their expected amount of extra-time. This *validates the concept of RAMS* over classical APS systems, namely by providing empirical evidences on the fact that **a)** a simple graphical formalism suffices to model a problem as complex as a real world biotech manufacturing process; and **b)** an optimization engine that takes uncertainty into account while generating solutions is clearly preferable to classical optimization techniques. Remark that these two conclusions were already obtained by [11] in the context of an human space mission, but not for biomanufacturing. In the next use cases, we will see that they also apply in other, slightly different, industrial contexts.

3.2 Pharmaceuticals Manufacturing at Takeda: capacity analysis and investments

The general purpose of this study is to optimize *Takeda Lessines* supply long-term plan (LRP), towards a higher manufacturing capacity while taking into consideration manufacturing constraints, as well as operation lines specificities and equipment capacities.

Takeda Lessines is a pharmaceutical company producing life-saving medicines for patients around the globe. The manufacturing plant, located in Lessines (Walloon Region, Belgium), operates 24h/7d and is dedicated to the purification, filling and packaging operations of plasma-derived products.

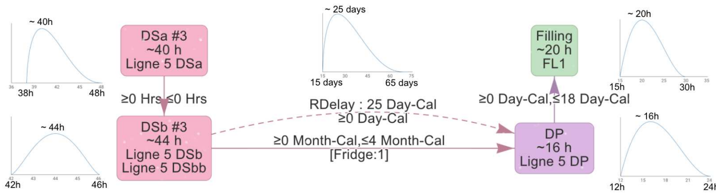

Fig. 9 The modeling of the GLA manufacturing process. All four operations have uncertain durations, represented by the probability PERT distributions, drawn from the user inputs: minimum, mode, maximum duration. Each operation is assigned to a specific operation line, and linked to its predecessor by temporal constraints. The model is structurally correct, yet the true operational and constraint values have been modified for confidentiality matters.

GLA is a sterile, ready-to-use, liquid preparation of purified human Alpha-1 Antitrypsin (AAT). GLA is indicated for chronic augmentation and maintenance therapy in adults with clinically evident emphysema due to severe congenital deficiency of AAT. The purification process consists of several chromatography and ultrafiltration systems to achieve a high-quality & highly purified intermediate, followed by filling operations using Restricted Access Barrier System (RABS) technology & final inspection and packaging with semi-automated lines.

3.2.1 Modeling and Scheduling

The model of the GLA manufacturing process, depicted in Figure 9, is quite simple: four activities (*DSa*, *DSb*, *DP*, *Filling*) linked by temporal constraints. Each of the tasks has several properties, in addition to its temporal uncertainty. For instance, the *DP* activity is carried out on a specific operating line *Ligne 5 DP*. Its processing time ranges from 12 to 24 hours, with a most probable duration (mode) of 16 hours. Furthermore, *DP* has two different temporal constraints with its predecessor *DSb*. The first one states that between *DSb* and *DP* the product uses one unit of space in the fridge, and that its maximum duration in the fridge is 4 months. The second temporal constraint is stochastic and represents the time required for releasing the unit of product, which may take from 15 days to 65 days, with a most probable duration (mode) of 25 days, leading to the asymmetric temporal distribution depicted in the top center of Figure 9.

The ***DS*** activities are modelled into two distinct subtasks *DSa* and *DSb*, in order to represent the fact that, whereas they both use the same production line *Ligne 5*, the first part (*DSa*) of the ***DS*** activity may not be processing in parallel of other activities on the same line, whereas the second part (*DSb*) may share the line (with *DSb* operations from different campaigns). A special temporal constraint between *DSa* and *DSb* states that there cannot be any waiting time between these two phases. Under this configuration, at most two distinct ***DS*** can be ran in parallel during the same week, with an offset that depends on *DSa*, as shown on Figure 10.

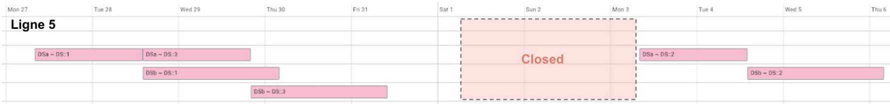

Fig. 10 A possible scheduling of three couples $DSa + DSb$, given that the operation line is closed from Saturday 6am to Monday 8am.

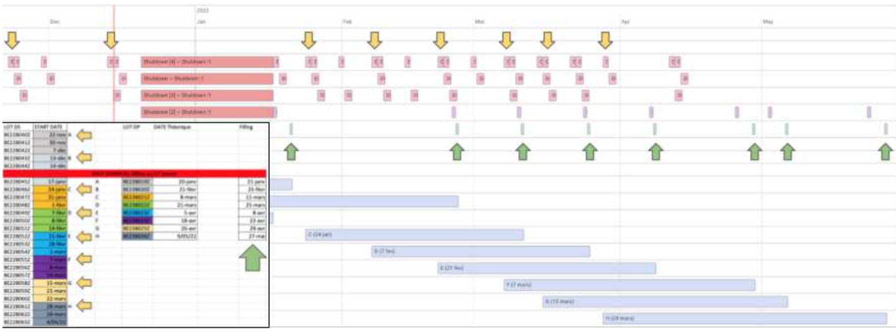

Fig. 11 An example of a schedule involving 8 manufacturing campaigns of *Takeda*'s GLA product, with a planned shutdown period (the four long red activities). Provided the starting date of each campaign (yellow arrows), the system accurately predicted the ending dates (green) expected by the company's field expert.

First schedules validating the model

In order to obtain first meaningful schedules, we have been provided a list of 8 manufacturing campaigns to be scheduled, together with the planned starting times (yellow arrows in Figure 11) of each of them. We were also given an interval of time during which all the production lines are closed (shutdown). Eventually, provided only the model presented in Figure 9 and the starting dates, the system was able to accurately predict the theoretical ending dates (green arrows) of each campaign, as planned by the field experts.

3.2.2 Robustness versus efficiency

In a second time, the same set of campaigns were considered again, but without imposing specific starting times. As a consequence, the optimization engine became free to schedule some campaigns sooner, eventually maximizing the overall efficiency. The produced schedules, such as shown in Figure 12, were in fact approximately 20% more efficient (25% if we don't take the shutdown into account) with total manufacturing time of one month instead of five. Unfortunately, these schedules were also very brittle, with a success probability of approx. 26%.

Detecting bottlenecks

Eventually, the system permitted us to detect the source of brittleness: the main bottleneck was due to the limited number of emplacements in the fridge. Because the manufacturing process uses part of these emplacements during

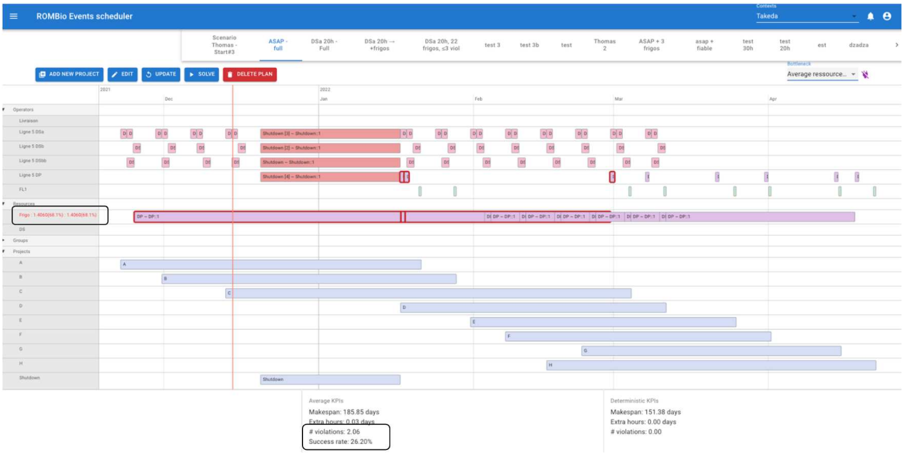

Fig. 12 By letting the system free to schedule all the campaigns, without imposing starting times, the optimized schedules are approximately 20% more efficient. Although being totally feasible on the paper, when taking uncertainty into account they however reveal very brittle. In particular, the system indicates that 68% of the potential issues are due to a conflict on a particular resource: the limited emplacements in the fridge.

the delay imposed by the release phase between *DS* and *DP*, delays in the other operations may propagate on the line and one may eventually run out of space in the fridge. In fact, *the system indicated that 68% of the potential problems were due to a conflict of the fridge resource*, rather than the violation of a temporal constraint.

Risk aversion

A possible solution would be to buy more space in the fridge; this will be addressed later. For the moment, we used the fact that *Rombio* allows us to parameterize a maximum level of risk, when computing the schedules. By imposing an acceptable level (according to field experts), the system produced optimized schedules while mitigating the risk to uncertainty, eventually integrating slack times in the planning to avoid conflicts on the fridge resource, as shown in Figure 13. While the new schedules were hence of reasonable robustness, it is interesting to note that their efficiency, in terms of total production time, revealed to be *equivalent to that of the initial schedules*, those optimized whereas the starting times of the campaigns were suggested by the company's field expert.

3.2.3 Improving the manufacturing process

One of *Takeda*'s concerns is the limited number of *DS* that can be launched in parallel in the same week. As shown on Figure 10, only two of them may cohabit during the same week. It appears however that, if the process (*i.e.* the model) could be adapted in order to decrease of 10 hours the time required by *DSa*, hence transferring this amount time to *DSb*, then a second DS could be

Fig. 13 When asking the system to orient the computation in order to obtain more reliable schedules (*i.e.* with higher success probability), the fridge emplacements bottleneck is compensated by introducing slack times (dotted circles).

started sooner, and eventually a third one on the same week. This is depicted in Figure 14.

Determining the appropriate investments

In order to fully exploit the potential of such adaptation of the model, which could significantly increase the efficiency (50%), it clearly appears that additional fridge emplacements are required. Indeed, we saw that even the current process was lacking some emplacements in order to be fully functional. In fact, in the current process, the system indicates that no more than 5 additional emplacements suffice to reach the maximum efficiency. Similarly, once the new process has been modelled, is it not difficult to simulate investments, by asking the system to optimize schedules provided a defined number of additional resources, such as the fridge emplacements. Eventually, repeating the operation with different numbers of emplacements *indicated that 12 additional ones are enough to reach the maximum manufacturing efficiency*. Interestingly, it appeared later one that this predicted number coincided precisely with the investment decided by the management.

Finding the right compromise between robustness versus efficiency

Recall that in Section 3.2.2, we applied a *risk aversion* strategy in order to reach an acceptable level of risk in the computed schedules. This is possible thanks to a feature of *Rombio*, which allows the user to define a maximum level of risk, a threshold that the system would try not to exceed while optimizing the solutions. The higher the threshold is, the more brittle will be the solutions, but also the more the engine will be able to focus on other KPIs such as efficiency by minimizing the total manufacturing time. Again, it is in fact easy to repeat the optimization process, while varying the predefined

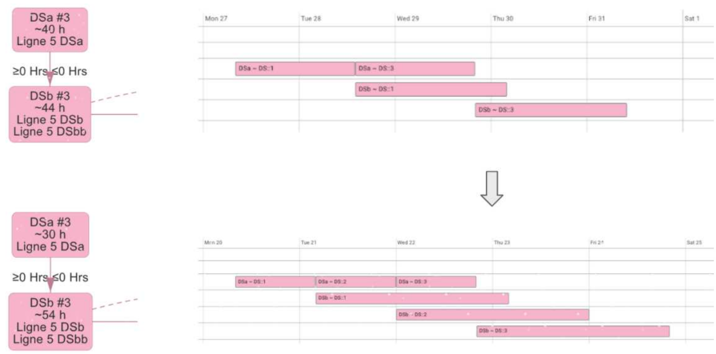

Fig. 14 By adapting the manufacturing process (the model) in order to transfer 10 hours of the *DSa* phase to *DSb*, we are eventually able to schedule three *DS* in parallel on the same week.

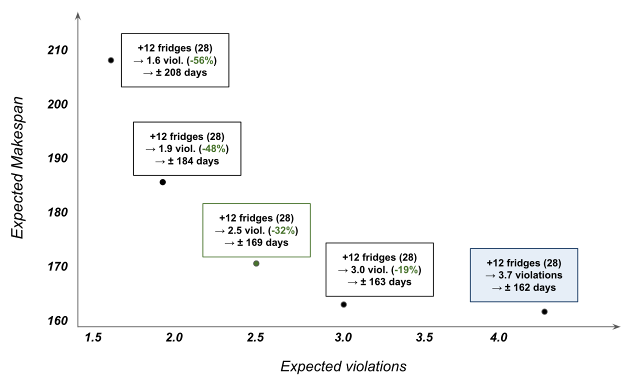

Fig. 15 A Pareto front obtained by plotting the different schedules, in terms of efficiency (*Expected Makespan*) and risk (*Expected violations*), obtained by optimizing while setting each time a different acceptable level of risk. Recall that all the solutions are feasible on the paper, but when considering the temporal uncertainty, they each come with different levels of risk. The levels of risk are expressed as an expected number of constraint violations (not only temporal, a conflict on a resource is also a violation). For instance, a conservative solution, with a risk of 1.6 would complete all the planned manufacturing campaigns in 208 days on average. On the other hand, the makespan could be ultimately decreased to 162 days by accepting a risk level of 3.7 violations on average. In the end, the manager decides by considering the expected cost of dealing with the constraint violations, compared to the savings provided by a shorter makespan.

threshold, hence producing different schedules, from the most dangerous to the most conservative. This is a common approach in bi-objective optimization, in particular when the objectives are contradictory. Here the two objectives are robustness and efficiency, and they are naturally contradictory. By plotting the results, we obtain in Figure 15 a so-called *Pareto front* (see *e.g.* [17]), a curve that describes all possible compromises between solution robustness and expected efficiency. For instance, we observe that *by sacrificing only 4% of the theoretical maximal efficiency* (169 days instead of 162), *the risk is decreased by 32%* (2.5 instead of 3.7).

3.2.4 Conclusions and insights from the Takeda use case

In this use case we successfully modelled and solved the problem of managing the operations along a real life production line, in a multinational pharmaceutical company. The previous use case of Section 3.1 was centered on the modeling complexity and the efficiency of the computed schedules. In this proof of concept in collaboration with Takeda, we saw that even provided very limited inputs (a few working times and a simple model), the system is able to accurately predict not only the actual provisional schedules, but also detect the major bottlenecks in terms of resources. In addition to the current state of the manufacturing process, the system also enabled us to consider an alternative, more efficient, process (*i.e.* a different model). The system eventually accurately predicted the investments, in terms of a particular resource, that were actually planned by the management in order to switch on the new process in the near future. Finally, by playing around with the risk threshold parameter of *Rombio*, we were able to compute a set of different solutions, from the most dangerous solution to the most conservative one. We hence obtained a so-called Pareto front, a curve that describes the set of *all possible compromises between robustness and expected efficiency.*

3.3 Stem Cells Production Cellaïon: scheduling with the living

Cellaïon develops a pipeline of second-generation products in systemic inflammatory space to repair and regenerate tissues and organs. The growing of stem cells represents a significantly different problem from that of manufacturing medicines or screening kits. In fact, working with the living imposes a higher level of uncertainty. The operations, conducted by human beings in sterile zones generate uncertainties of course, but the most of the uncertainties actually lie in the cell maturation times. In other words, the waiting times between the operations are more uncertain than the operations themselves.

Furthermore, in cell therapy one does not necessarily seek for raw efficiency. Naturally, the company is interested in being operationally able to maximize the number of cell batches that can be produced within, let's say, a year. Yet, producing even just one batch of stem cells at the time is so complex, that in practice, a marginal decrease in the natural variation costs could be much more profitable than producing a second batch. As a consequence, the company is more interested in the ability of forecasting the risks and natural variations, eventually finding how to reach better reliability, rather than designing a way of producing more.

We conducted this use case in collaboration with the Belgian company *Cellaïon* in 2021. In what follows, we present how the *Rombio* system has been used in order to address some fundamental questions the company faces, at a time when this company is completely rethinking and redesigning its activities.

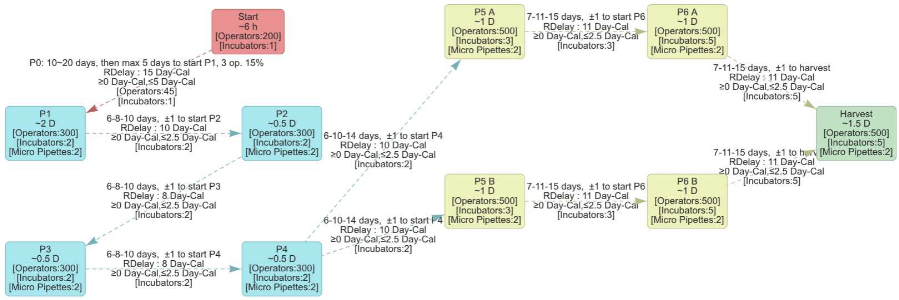

Fig. 16 The modeling of the activities, temporal constraints, and resources involved in the processing of one batch of stem cells. Because the maturation times of the cells are unpredictable by nature, all the constraints (dotted arrows) are stochastic, in the sense that the operators do not known in advance how long they will have to wait between two activities. Resources, such as incubators or human operators (expressed in % occupation), are exploited during these waiting times. Right after *P4*, the process is split in two parallel branches, each requiring a growing number of incubators. The model is structurally correct, yet the true operational and constraint values have been modified for confidentiality matters.

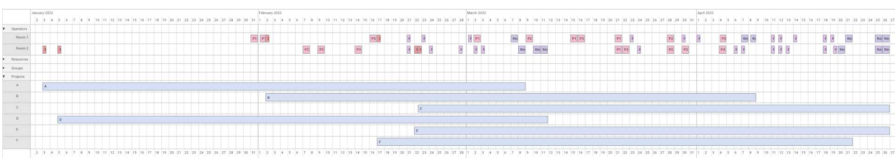

Fig. 17 An example of a schedule combining six stem cell production campaigns, spanning four months.

3.3.1 Modeling specificities

The modeling of a stem cell manufacturing process is depicted in Figure 16. The model has been constructed directly in the *Rombio* system, using the graphical modeling interface. In the depicted model, we see that all the temporal constraints are uncertain. In fact, each constraint represents a maturation time between the different phases of the process, which is stochastic. In fact, operators do not know in advance how long they need to wait between activities. Resources (both human and material) are exploited during the waiting times between the process steps. Actually, there are, however, strict constraints on the moments at which, for instance P_n task, can be operated depending on both the completion time of P_{n-1} and the time required to the cells to mature enough to be processed in P_n. If x days are required, then P_n must be operated between $x-1$ and $x+1$ days after P_{n-1}. Whereas this maturation time remains variable, yet the evolution of the cells can be monitored, which explains that the (a priori) unknown duration can actually be predicted (*i.e.* observed) a couple of days ahead.

1 zone 6 incub	1 batch	2 batches	3 batches	4 batches	5 batches	6 batches
5 op.	0.3 **3%**	1.1 **5%**	3.6 **12%**	n.a. (2)		
7 op.	0.3 **3%**	0.7 **3%**	2.3 **8%**	n.a.		
10 op.	0.3 **3%**	0.7 **3%**	2.1 **7%**	n.a. (1)		
12 op.	0.3 **3%**	0.7 **3%**	2.1 **7%**	n.a. (1)		
15 op.	0.3 **3%**	0.7 **3%**	2.1 **7%**	n.a. (1)		

1 zone 12 incb	1 batches	2 batches	3 batches	4 batches	5 batches	6 batches
5 op.	0.3 **3%**	1.1 **5%**	3.4 **11%**			
7 op.	0.3 **3%**	0.7 **3%**	2.6 **9%**			
10 op.	0.3 **3%**	0.7 **3%**	2.1 **7%**	n.a. (1)		
12 op.	0.3 **3%**	0.7 **3%**	2.1 **7%**	5.9 **15%**	n.a. (1)	
15 op.	0.3 **3%**	0.7 **3%**	2.1 **7%**	5.2 **13%**	n.a. (1)	

2 zones 6 incub	1 batch	2 batches	3 batches	4 batches	5 batches	6 batches
5 op.	0.3 **3%**	1.1 **5%**	3.6 **12%**	n.a. (1)		
7 op.	0.3 **3%**	0.7 **3%**	2.4 **8%**	n.a. (1)		
10 op.	0.3 **3%**	0.6 **3%**	1.6 **5%**	5.6 **14%**	n.a. (1)	
12 op.	0.3 **3%**	0.6 **3%**	1.6 **5%**	5.6 **14%**	n.a. (1)	
15 op.	0.3 **3%**	0.6 **3%**	1.6 **5%**	3.7 **9%**	n.a. (1)	

2 zones 12 incb	1 batch	2 batches	3 batches	4 batches	5 batches	6 batches
5 op.	0.3 **3%**	1.1 **5%**	3.6 **12%**	n.a. (1)		
7 op.	0.3 **3%**	0.7 **3%**	1.9 **6%**	6.8 **17%**	n.a. (3)	
10 op.	0.3 **3%**	0.4 **2%**	0.8 **3%**	1.9 **5%**	5.3 **11%**	n.a. (1)
12 op.	0.3 **3%**	0.1 **<1%**	0.6 **3%**	1.4 **3%**	3.9 **8%**	6.5 **11%**
15 op.	0.3 **3%**	0.1 **<1%**	0.6 **3%**	1.4 **3%**	3.1 **6%**	5.8 **10%**

Fig. 18 Capacity analysis study at *Cellaïon*: how the quality of the schedules evolves with the demand (number of cell batches) and the available resources. The values inside a table cell represent the quality of the best schedule found by the optimization engine. It is expressed in terms of average number of constraint violations, and percentage of problematic activities. A black cell indicates that no deterministic feasible schedule could be found.

3.3.2 Capacity and investments analysis

In the process of rethinking their activities, the company was interested in determining how the manufacturing conditions evolve with the number of stem cell batches that are processed in parallel. A batch of cells is processed in approximately 2.3 months. Provided the company's human and material resources, it should however be possible to produce concurrently, for instance, two or more batches in less than 4 months. While limiting the time horizon to 4 months, *Rombio* has been asked to find optimized schedules that combine from 1 and up to 6 production campaigns. Figure 17 shows an example of a schedule with 6 campaigns in parallel, organized over 4 months.

The company is however not accustomed to process six campaigns in parallel. In fact, the current resources do not allow it. Using our RAMS system *Rombio*, we were able to determine how the quality of the optimized schedules, in terms of reliability, evolves depending on the number of campaigns in parallel. Then, we repeated the computations while providing different resource configurations to the company (e.g. having 2 zones, 12 incubators, 10 operators, instead of 1 zone, 6 incubators, 5 operators). This eventually resulted in the four tables depicted in Figure 18, showing *how the production capacity evolves with both the imposed demand and the available operational resources.*

The results obtained in the capacity analysis study, in Figure 18, constitute key insights for the company to consider future investments. For instance, we see that solely increasing the number of human operators (top left table) is useless. Whereas it does not significantly improve the quality of the schedules, it does not allow to find deterministic feasible (*i.e.* fulfilling all the

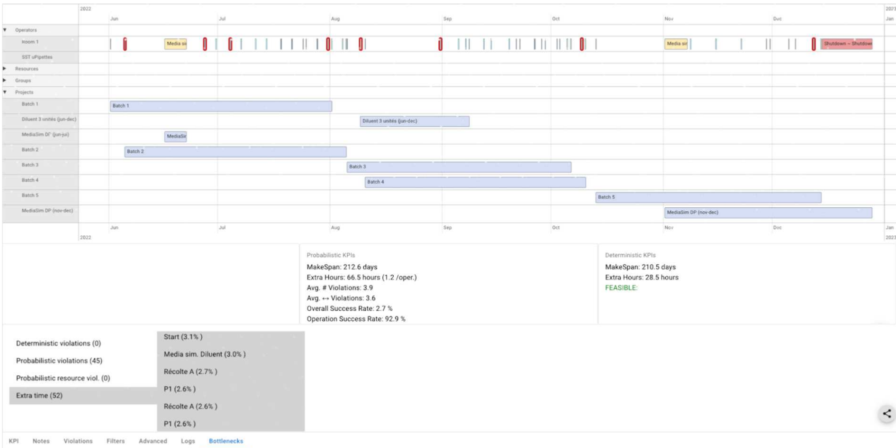

Fig. 19 An optimized long term schedule (7 months), involving five cell batches and three other productions. Here the *Rombio* interface spots the activities that are the most likely to generate extra-hours. Whereas the schedule involves 28.5 extra-hours in order to be feasible (*i.e.* zero violations on the paper), in practice when taking uncertainty into account this number is 66.5 extra-hours on average, with an expected 3.9 violations.

constraints under deterministic assumptions) schedules for 4 batches. Furthermore, increasing the number of incubators as well as operators do not provide a significant improvement, whereas some improvements are perceptible on 3 and 4 batches, when also moving from 1 to 2 zones. Eventually, we observe that only a combination of all investments (namely a second sterile zone, an increased number of incubators, and 12+ operators) permits to increase the production up to 5 or 6 parallel batches. But even with two parallel batches only, more reliable solutions can only be obtained by combining investments in all of these three resources (bottom right table of Figure 18).

3.3.3 Compromises in the long term

The company's expert provided an example of a forecast of the demand in stem cells. Five cell batches, as well as three independent productions, will have to be scheduled in 7 months. Given the resources of the company, still in terms of human operators, machines and production areas, the *Rombio* system has been solicited to solve this problem, producing optimized schedules.

An example of a schedule is shown in Figure 19. For the company, the principal quality criterion of a schedule, and therefore the first KPI to be optimized, is the expected number of constraint violations. It is however important to note that, in many cases, a constraint violation could be resolved by asking operators to work when they are not supposed to, that is, doing extra-hours. For example, an activity planned on a Friday may not necessarily be postponed to Monday, thus requiring a couple of operators to work on the weekend. In Figure 19, the system shows which are the activities that are the most likely to require extra-hours.

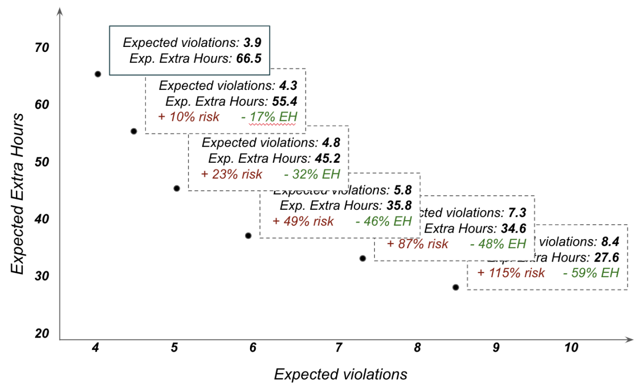

Fig. 20 Pareto front obtained by plotting the different schedules, according to their expected amount of extra-hours and their risk (expected violations). As for Fig 15, these compromises are obtained by optimizing with different risk limits.

Since extra-hours contribute to resolve constraint violations, increasing the allowed amount of it is likely to increase the success probability. Unfortunately, it also contributes negatively to the operators' wellness, inducing stress and, eventually, increases the amount of medical leaves. However, similarly increasing the acceptable level of risk results in a decrease of the expected amount of extra-hours. Similarly to the previous study of Section 3.2.3, we face a problem in which two contradictory objectives must be optimized. Many possible compromises actually exist, forming a so-called *Pareto front* (see *e.g.* [17]), a curve that describes all possible compromises between solution robustness and operators' wellness.

A Pareto front in the context of this study is shown in Figure 20. The minimum risk achieved by the optimization engine is 3.9, which would generate 66.5 extra-hours on average. Yet, by allowing the engine to consider slightly more risky schedules, the optimization process could focus on the second KPI, namely minimizing the expected extra-hours. For instance, there exists an alternative schedule, having 10% more risk, but reducing the extra-hours by 17%. Eventually, the adequate compromise will be selected by the company.

3.3.4 Conclusions and insights from the Cellaïon use case

In this use case, the principal operations in a stem cells manufacturing company have been modelled and scheduled, while considering both the process reliability and the operators' wellness. The reliability is again expressed in terms of probability of success, or expected number of constraint violations during the manufacturing process.

However, contrary to the two previous use cases (*Zentech*, *Takeda*), here Nature imposes its uncertainty between the actual operations, namely in the waiting times during cell maturation phases, in addition to the operations themselves. The resulting schedules are therefore highly brittle, and are likely to generate a huge amount of extra-hours. On the other hand, extra-hours are only manageable under a certain quantity. As *Rombio*'s engine is able to optimize while pursuing several KPIs, such as the risk *and* the expected extra-hours, the system demonstrated its ability of providing solutions for different levels of *compromises between risk and operators' wellness.*

4 Conclusions

In this paper, we introduce a new (manufacturing) operations management framework, called *(robust) advanced modeling and scheduling* (RAMS). Unlike classical tools, called APS (advanced planning and scheduling), with an AMS the user directly (and graphically) models the operational problem at stake, which allows for a wide range of potential applications. The R in RAMS stands for *robust*, as the system is not only able to generate optimized solutions (schedules), but these solutions are designed while taking uncertainty into account.

We described *Rombio*, a new software carrying these properties, hence being the first RAMS system. We apply *Rombio*'s technology to three different biomanufacturing use cases, in three different Belgian biotech companies, despite the different operational contexts: *diagnostics* (Zentech company), *medicines* (Takeda company) and *stem cells* (Cellaion company). These use cases show the benefits of using a probabilistic modeling approach, taking the time uncertainty of activities durations into account at optimization stage, which are clearly confirmed by the empirical average gains compared to classical (risk-aversion) approaches. Schedules obtained using our probabilistic optimization engine, at the expense of sacrificing only 8% of the theoretical efficiency, are 38 times more reliable on average than those obtained when not taking uncertainty into account. Use cases also show how an RAMS can be used to not only schedule operations, but also conduct capacity analysis studies, simulate investments and assess alternative manufacturing processes. Using our RAMS system, we were able to correctly predict the exact future investments already planned in a multinational pharmaceutical manufacturing company, in a few hours only. The system also permitted us to suggest new strategic (combinations of) investments, in all the three considered companies.

Future work includes the integration and assessment, in real conditions (biomanufacturing, space missions), of online management capabilities: online *monitoring* and *rescheduling.* Online monitoring aims at updating the schedule, as well as the underlying model, keeping them consistent according to past events and current state of the operations. This eventually allows for adapting and rescheduling future decisions, in light of past decisions and outcomes.

Author Declarations

Funding. The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.
Ethics approval/declarations. The authors have no relevant financial or non-financial interests to disclose.

Consent to participate. All authors contributed to the study conception and design. The first draft of the manuscript was written by Michael Saint-Guillain and reviewed by Isabelle Sansdrap, Michael Brion and Eric Pauly. All authors read and approved the final manuscript.
Consent for publication. All authors agree to publish this article in jour journal.
Availability of data and material. The datasets generated during and/or analysed during the current study are not publicly available due to confidentiality matters, but are available from the corresponding author on reasonable request and with permission of the companies involved in the study.
Code availability. The code and software access are not publicly available due to confidentiality matters, but are available from the corresponding author on reasonable request.

References

[1] Herroelen, W., Leus, R.: Project scheduling under uncertainty: Survey and research potentials. European journal of operational research **165**(2), 289–306 (2005)

[2] Turner, R.J., Huemann, M., Anbari, F.T., Bredillet, C.N.: Perspectives on Projects. Routledge, ??? (2010)

[3] Abramov, I., Poznakhirko, T., Sergeev, A.: The analysis of the functionality of modern systems, methods and scheduling tools. In: MATEC Web of Conferences, vol. 86, p. 04063 (2016). EDP Sciences

[4] Stadtler, H., Stadtler, H., Kilger, C., Kilger, C., Meyr, H., Meyr, H.: Supply Chain Management and Advanced Planning: Concepts, Models, Software, and Case Studies. Springer, ??? (2015)

[5] Saint-Guillain, M., Gibaszek, J., Vaquero, T., Chien, S.: Romie: A domain-independent tool for computer-aided robust operations management. Engineering Applications of Artificial Intelligence **111**, 104801 (2022). https://doi.org/10.1016/j.engappai.2022.104801

[6] Marquez, J.J., Hillenius, S., Kanefsky, B., Zheng, J., Deliz, I., Reagan, M.: Increasing crew autonomy for long duration exploration missions: self-scheduling. In: 2017 IEEE Aerospace Conference, pp. 1–10 (2017). IEEE

[7] Marquez, J.J., Hillenius, S., Healy, M., Silva-Martinez, J.: Lessons learned from international space station crew autonomous scheduling test (2019)

[8] Marquez, J.J., Edwards, T., Karasinski, J.A., Lee, C.N., Shyr, M.C., Miller, C.L., Brandt, S.L.: Human performance of novice schedulers for complex spaceflight operations timelines. Human Factors,

00187208211058913 (2021)

[9] Lenstra, J.K., Kan, A.R.: Computational complexity of discrete optimization problems. In: Annals of Discrete Mathematics vol. 4, pp. 121–140. Elsevier, ??? (1979)

[10] Applegate, D., Cook, W.: A computational study of the job-shop scheduling problem. ORSA Journal on computing **3**(2), 149–156 (1991)

[11] Saint-Guillain, M.: Robust operations management on mars. In: Proceedings of the International Conference on Automated Planning and Scheduling, vol. 29, pp. 368–376 (2019)

[12] Saint-Guillain, M., Vaquero, T., Chien, S., Agrawal, J., Abrahams, J.: Probabilistic temporal networks with ordinary distributions: Theory, robustness and expected utility. Journal of Artificial Intelligence Research **71**, 1091–1136 (2021)

[13] Saint-Guillain, M., Vanderdonckt, J., Burny, N., Pletser, V., Vaquero, T., Chien, S., Karl, A., Comein, A., Charmart, C., Wain, C., Casla, I.S., Jacobs, J., Manon, J., Meert, J., Drouet, S.: Enabling astronaut self-scheduling using a robust modelling and scheduling system (rams): a mars analog use case. In: Proceedings of the 16th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2022, ESA-ESTEC, Noordwijk, the Netherlands). (under publication), ??? (2022)

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

[15] Petrides, D., Carmichael, D., Siletti, C., Koulouris, A.: Biopharmaceutical process optimization with simulation and scheduling tools. Bioengineering **1**(4), 154–187 (2014)

[16] Kamburowski, J.: New validations of pert times. Omega **25**(3), 323–328 (1997)

[17] Ngatchou, P., Zarei, A., El-Sharkawi, A.: Pareto multi objective optimization. In: Proceedings of the 13th International Conference On, Intelligent Systems Application to Power Systems, pp. 84–91 (2005). IEEE