Portfolio-wide Optimization of Pharmaceutical R&D Activities Using Mathematical Programming

Hua Wang[1], Jon Dieringer[2], Steve Guntz[2], Shankarraman Vaidyaraman[2], Shekhar Viswanath[2], Nikolaos H. Lappas[1], Sal Garcia-Munoz[2], Chrysanthos E. Gounaris[1*]

1. Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15229

2. Small Molecule Design & Development, Eli Lilly and Company, Indianapolis, IN 46225

Abstract

The R&D management in any major research pharmaceutical company is constantly faced with the need to make complicated activity scheduling and resource allocation decisions, as they carry out scientific work to develop new therapeutic products. This paper describes how we develop a decision support tool that allows practitioners to determine portfolio-wide optimal schedules in a systematic, quantitative, and largely automated fashion. Our tool is based on a novel mixed-integer linear optimization model that extends archetypal multi-mode resource-constrained project scheduling models in order to accommodate multiple rich features that are pertinent to the Chemistry, Manufacturing, and Controls (CMC) activities carried out within the pharmaceutical R&D setting. The tool addresses this problem at the operational level, determining schedules that are optimal in light of chosen business objectives under activity sequencing, resource availability, and deadline constraints. Applying the tool on current workload data demonstrates its tractability for practical adoption. We further illustrate how, by utilizing the tool under different input instances, one may conduct various tactical analyses to assess the system's ability to cope with sudden changes or react to shifting management priorities.

`Key words`: pharmaceutical R&D; CMC (chemistry, manufacturing, and controls); project scheduling; mixed-integer linear optimization

*Corresponding author: gounaris@cmu.edu

Introduction

Eli Lilly and Company (Lilly) is a global health care leader that unites caring with discovery to create medicines that make life better for people around the world. These medicines are manufactured in 8 different countries and sold in 120 countries worldwide.[1] The Small Molecule Design and Development (SMDD) organization within Lilly Research Labs is an innovation focused organization that designs drug products to meet patient needs, develops commercial drug substance and drug product processes, and supplies medicines for clinical trials for small and medium (peptide) molecules. SMDD activities span the spectrum of development from pre-clinical studies all the way to commercialization.

The arrangement of research and development (R&D) activities and allocation of human resources have become more complex and challenging over the years because of portfolio expansion and increasing competition. Yet, to-date, no structured and standardized approach has been deployed to making the operational decisions across the full portfolio of products within the SMDD organization. For the past few years, operational decisions for the drug development process, such as time span and resource allocation for development activities, are determined by experienced personnel, who bring their expert skills to different parts of the R&D process. Additionally, frequent changes to portfolio projects, driven by clinical and toxicological data read-outs, cause constant resource trade-off decisions on a case-by-case basis. A portfolio-wide resource optimization tool is therefore sorely needed. The existence of variability in judgmental decisions due to human bias can lead to time delays and incorrect resource allocations. Moreover, each product team carries out the research plans pursuing respective interests, which as a result may cause a conflicting portfolio-wide research plan given the budget and resource limitations as well as individual drug development timelines that need to be met.

The traditional approach calls for iterating the research plans to balance the needs of each team until a practical portfolio-wide schedule is achieved. Such reactive manual input often leads to suboptimal decisions, or even the inadvertent perturbation of decisions that were otherwise acceptable or even near-optimal. Meanwhile, the portfolio planning costs keep increasing, not to mention the burden in the overall organization in terms of experts' time. To avoid the manual, costly, and time-consuming iteration cycles, the SMDD organization seeks to exploit rigorous, mathematical optimization based approaches to this portfolio-wide decision-making problem. The ultimate goal

[1] https://www.lilly.com

that the SMDD organization wants to achieve is the development and adoption of a *decision support system* that shall enable each person in all teams to make good decisions regarding each product development process, and very importantly, to support a collaborative planning culture.

In this paper, we present the scope of portfolio-wide modeling and optimization at the SMDD organization within Eli Lilly, the creation of capabilities we developed to address this scope, and its application to benefit Eli Lilly's portfolio of small and medium molecules. The overall goal is to be in position to quantitatively negotiate the costs, delays or other weighted objectives that a set of decisions bring to the whole portfolio, and to create a computational infrastructure that largely automates the process of planning R&D activities and allocating resources to this purpose. The decision support system that we developed is based on a mixed-integer linear programming (MILP) approach, mapping the complex scheduling problem Lilly faces to a rich-featured, multi-mode resource-constrained project scheduling problem (MRCPSP).

The MRCPSP is an archetypal optimization problem with broad application in planning and process-level enterprise decision problems. Several extensions of the archetypal problem have been developed to accommodate practical applications (Hartmann and Briskorn 2010), including sales force deployment, natural resource development projects, R&D and production scheduling (Varma et al. 2007). Motivated by specific industrial settings, Voß and Witt (2007) study the production planning problem at a steel manufacturer by incorporating setup states into the MRCPSP model, while Tiwari et al. (2009) consider the customer training projects in the telecommunication industry by extending the MRCPSP to account for service quality. Bartels and Zimmermann (2009) study a project scheduling problem motivated in the automotive industry, proposing an MRCPSP with minimum and maximum time lags as well as renewable and cumulative resources. There also exist applications of MRCPSP models in cargo transportation (Lorenzoni et al. 2006, Hill et al. 2019) and healthcare facility operations (Riise et al. 2016, Poppenborg and Knust 2016), among many other settings.

The MRCPSP was first introduced in the context of optimizing pharmaceutical research projects by Kolisch and Meyer (2006), where the authors proposed two heuristics for solving the problem. Other works later implemented MILP models based on the MRCPSP within a heuristic simulation-optimization framework for addressing R&D pipeline management (Varma et al. 2008, Zapata et al. 2008). While some studies viewed the single product development as a resource constrained problem (Jain and Grossmann 1999), other studies on portfolio-wide optimization used stochastic programming methods (Colvin and Maravelias 2008, 2009, 2011). For example, Colvin and Maravelias

(2011) addressed the similarities shared between R&D planning and the single-mode RCPSP, while they used a multi-stage stochastic programming framework to account for the success or not of various clinical trials. For our purposes, we adapt the well-known discrete-time model for MRCPSP by Talbot (1982) to our context, extending the literature model in multiple ways to accommodate applicable rich features.

In the remainder of the paper, we define the problem at hand in more detail, discuss our model, and provide computational results to elucidate the tractability of the computational framework we have implemented. Furthermore, we illustrate how an operational decision-making tool of this sort can be used in a more tactical fashion to provide insights about organizational limitations and potential improvements, as well as the organization's amenability to cope with major disruptive events.

Problem Definition

When a business decision is made to pursue the next phase of development of a particular experimental drug (also referred to as *molecule*), and up to the point where a sufficient quantity of the drug product is delivered to the clinical trials allowing the latter to commence, researchers at Lilly need to pursue a sequence of R&D activities. These include, among others, activities to develop the Active Pharmaceutical Ingredient (API), activities to determine its solid form, activities to develop the drug product formulation, studies to determine the Relative Bio-Availability (RBA), the manufacturing campaign to produce the product, product stability studies, as well as to prepare and file paperwork with the U.S. FDA and other international regulatory authorities. In doing so, a number of resources available to Lilly R&D are put to action, including personnel of given specializations, such as process chemists, chemical engineers, drug product formulators and analytical chemists, as well as leading laboratory equipment and space, and available budgets, among others. All these resources need to be appropriately synchronized, in order to ensure that all molecules within the portfolio are developed on time and within specific business-driven deadlines.

Even though time-to-market is the most critical factor in the pharmaceutical industry (Shah 2004), since first-to-market players can have an advantage in the market share (Cha and Yu 2014), the company still needs to ensure consistent quality standards, which is fundamental to a successful launch. The goal of the Chemistry, Manufacturing, and Controls (CMC) development is to develop commercial processes to produce quality product in a safe, economical and environmentally friendly

manner within desired timelines.

For each drug product, among the massive volume of activities that are involved in the CMC development, there is a fair amount of information that needs to be transferred from one to another to reduce the risk that is inherent to the development process. Due to the need for such information flow, there exist complicated interdependencies between activities that need to be accounted for in any feasible development plan. For example, the API stability studies must wait for enough quantity to be manufactured during the API campaign before they can commence, while the drug product stability studies should have an overlap with the drug product manufacturing campaign to calibrate the drug product design. Moreover, important to capture information flows exist also between CMC and other R&D activities. In fact, after drug discovery, all stages of the drug development life cycle involve CMC. For example, the preclinical phase requires analytical activities to validate the product, the clinical stage needs additional characterization of the drug product, while after the clinical trials conclude, the process needs to be appropriately scaled-up to ensure the same specifications are met under commercial manufacturing. In fact, for those candidate molecules that survive multiple phases of clinical trials, certain development activities get executed several times with increasing and/or more detailed scope.

Model Development

Lilly aims to optimize its work plans and resource allocations, given its activity inter-dependencies, resource constraints, and commercial deadlines. We model the problem of Lilly's portfolio-wide scheduling and resource allocation as a mixed-integer linear programming formulation that is based on a multi-mode resource-constrained project scheduling problem. The MRCPSP considers a *project* with a number of non-preemptive *activities* to be scheduled, each of which comes with a set of execution *modes*. The latter constitute different, non-dominated combinations of duration and resources consumption via which the activity can be carried out. The goal of the problem is to determine a feasible activity schedule and associated execution modes for all activities along a predefined, discrete-time horizon, which can often be an especially challenging task due to the tightly limited *resources* and the existence of *precedence* relations between activities.

We remark that, in order to capture various realistic constraints associated with the application setting of focus, our optimization model deviates from archetypal MRCPSP models that exist in the literature, extending them in many ways that are described later in the paper. In summary, we

highlight that our model generalizes the concept of activity precedence, allowing for various types of partial activity timing overlaps, considers the presence of resources of mixed type (i.e., resources with both intensive and extensive consumption limits), as well as supports various business motivated objectives. More specifically, the model can be configured either towards the minimization of the net present total cost of development, towards the minimization of tardiness against the imposed (soft) deadlines, or towards some suitable combination of these two objectives. Furthermore, unlike the archetypal MRCPSP definition that focuses exclusively on the minimization of project makespan, our model can be configured to instead minimize the makespan of a subset of the project, which is especially useful when there arises a business priority to accelerate the development of a specific molecule in the portfolio. Next, we describe each component of the optimization model and the associated data collection process.

Project, Activities, Resources

We consider the portfolio of molecules under development to be the *project* and all required R&D activities for each such molecule to be the set of *activities* to be scheduled. A planning horizon for the development of the portfolio is specified and is suitably discretized in monthly periods. Personnel of various different specialties are regarded as *renewable resources* (in terms of full-time equivalent units, FTE). The API and drug product manufacturing facility's total capacity is also considered as a renewable resource, allowing us to limit the number of manufacturing activities (a.k.a., "campaigns") that can be ongoing at any point in time. Unlike renewable resources, which are limited on a per-period basis, *non-renewable resources* are defined to have a limited cumulative availability across the entire project. In our setting, we do not define any strictly non-renewable resources, but we do consider the available budget for carrying out the activities associated with developing the portfolio (a.k.a., "direct spend") as a *doubly-constrained resource*. The latter is to be interpreted as a resource for which both per-period as well as cumulative limitations apply, which enables us to model applicable budgetary and accounting restrictions. Cumulative limitations on the direct spend resource consumption can be defined for both the complete horizon or appropriately selected sub-horizons. We use the concept of sub-horizons, for example, to accommodate cash flow availabilities that are applicable on a calendar year basis.

Activity Timing

We have defined the activity network depicted in Figure 1 to be pertinent to the development of a novel pharmaceutical product. In this network, each activity is represented by a rectangle, and it is color-coded according to a domain-specific classification: red for API development; blue for drug product development; purple for solid state determination; green for enhanced API activities; gray for RBA studies; and brown for regulatory activities. This classification broadly covers the major aspects of the R&D process in this context. We highlight that not all depicted activities are relevant to be carried out by all molecules, and that the diagram as depicted represents merely a superstructure of possible activities for each molecule. To that end, the diamonds represent high-level decisions that scientific leadership at Lilly has already made (before the planning task addressed in this paper is performed) so as to specify the development path of each molecule in the portfolio, as appropriate. We also remark that a particular molecule may have to perform multiple passes through this network, as it undergoes separate phases of its multi-year development process.[2]

The arcs in Figure 1 represent the existence of some type of *precedence relationship* between the linked activities. Our concept of such a relationship is generic, and it can be viewed as a relative activity timing restriction with maximum and minimum time lags, allowing for two activities to also partially overlap (see Figure 2). Beyond the standard precedence relationships that imply a hard zero lag timing restriction, where the successor activity may begin as long as the predecessor activity has concluded, we identify three additional types of partial overlap restrictions, namely the *finish-start*, *finish-finish*, and *start-start* timing restrictions with minimum and maximum lag (Demeulemeester and Herroelen 1997). The finish-start relationship type requires some prescribed amount of lag between the conclusion of the predecessor activity and the beginning of the successor activity. Similarly, the finish-finish (start-start) precedence relationship requires lag between the conclusions (beginnings) of two adjacent activities. These overlaps are imposed to account for necessary technical information sharing between the executors of different activities in the context of the scientific work that is carried out. Finally, in order to link different development stages for each molecule, arcs across activities from two separate instantiations of the depicted network are also defined, designating also their applicable precedence type.[3] Consequently, the "start" and "end"

[2] In general, each pass of a given molecule will follow a different path and/or the modes available for each activity box in the diagram will be different to reflect the molecule's stage of development (e.g., pre-clinical, first-in-human, first registered dose, first submission).

[3] Usually, these cross development phase precedences are of finish-start type, reflecting the need for time to assess if the molecule should proceed to the next phase (e.g., via the performance of a clinical trial), or reflecting material flows across various campaigns (e.g., one API campaign feeding multiple drug product campaigns).

nodes do not reflect the initiation or conclusion of development of a given molecule, rather represent the signal that a molecule should enter its next phase of development and that the molecule has completed the current phase, respectively.

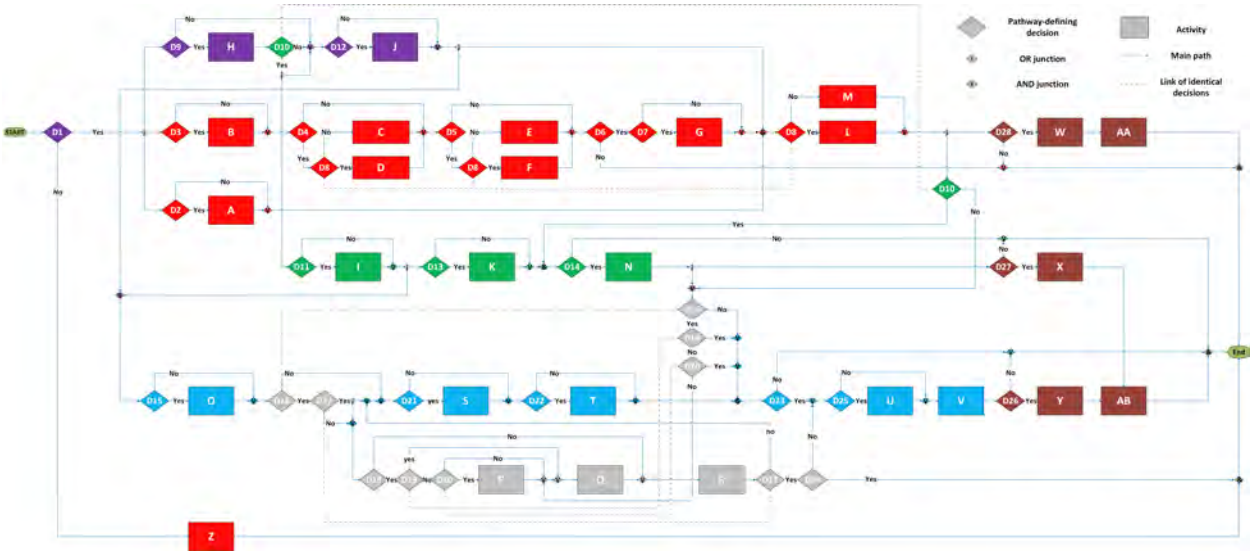

Figure 1: Activity network superstructure for pharmaceutical R&D (adapted from Viswanath et al. (2020)).

(a) Standard (zero lag)

(b) Finish-start with min and max lag

(c) Finish-finish with min and max lag

(d) Start-start with min and max lag

Figure 2: Four types of timing restrictions.

Lilly provides all information related to the activities to be scheduled in the portfolio as part of the model input. Spreadsheets store the specific types of precedence relationships between activities. Whenever a technological requirement for information sharing exists between two activities, a lag time interval is provided in the corresponding cell. Example of such entries are shown in Tables 1, where the activity names in the header column and header row correspond to the predecessors and successors in the generalized precedence relations, respectively. An "∞" entry in the time interval indicates the absence of an upper bound restriction.[4] For clarity, let us focus on the left-most spreadsheet. It imposes that activity B needs to start immediately after activity A finishes, while

[4]From this perspective, a standard precedence relationship can be viewed as simply a finish-start type relationship with minimum lag time equal to zero and maximum lag time equal to ∞.

activity C can only start after one month has passed since the conclusion of activity B. Similarly, the data in the other spreadsheets implies that activity C needs to have finished three to nine months before activity D finishes, while activity E may only start after six months have passed since activity B did. Finally, we remark that these spreadsheets are considered to be master documents and used as default input for most molecules under development, while molecule-specific deviations from these entries are possible, as applicable.

Act	A	B	C	D	E	...
A		[0,0]				
B			[1,∞]			
C						
D						
E						
...						

(a) Finish-start precedence

Act	A	B	C	D	E	...
A						
B						
C				[3,9]		
D						
E						
...						

(b) Finish-finish precedence

Act	A	B	C	D	E	...
A						
B					[6,∞]	
C						
D						
E						
...						

(c) Start-start precedence

Table 1: Example spreadsheets with generalized precedence data.

Activity Modes

For each activity in the network, we define a set of *modes* by convoluting the various applicable "scientific modes" with their allowable range of resources consumption trade-off. The scientific modes that are available to a specific molecule usually depend on its phase of development. As the drug development process moves forward, the development activities tend to require more comprehensive work to avoid outcomes of failure at later stages; as a consequence, the scientific modes are generally more expensive for later stages. Given an applicable scientific mode, there usually exist duration/resource and resource/resource trade-offs. For example, a particular activity could be performed by three scientists working on it for three months, or alternatively, by one scientist working for nine months. In such a case, we will split the scientific mode into multiple "execution modes," which are then referenced in our optimization model.[5]

Lilly provides a master spreadsheet that contains raw data of scientific modes for each activity in the network. These data include the development phase each mode is related to, the total direct spend and total FTE-months required, as well as an FTE "envelope," where the minimum and maximum duration and minimum and maximum FTE fractions for each FTE type are included. Then, at a preprocessing step, we will discretize these duration and FTE consumption intervals and generate a set of non-dominated execution modes that respect the total FTE-months requirement, while rounding up their duration to the nearest integer value. We allow at most four execution

[5]Execution modes are equivalent to the concept of "mode" in standard MRCPSP literature.

modes for each originally provided scientific mode.

Deadlines and Time Windows

Typically, the model is configured to include the development of certain molecules of interest until market launch, by inserting the typical sequence of activities, given impending clinical readouts and probabilities of clinical trial success. As an example, CMC readiness for rapid acceleration of assets following a positive clinical signal can be facilitated by building "planning to launch" activity networks, which include development (information generation) and material supply activities that are contingent on the positive clinical signal.

In this context, deadlines can be imposed as input data to the model for a number of critical activities that the business has set completion targets for (e.g., product launch activities). These deadlines could be imposed as hard constraints that could never be violated, or as constraints relaxed with slack variables to be penalized in the objective (soft deadlines). In the latter case, Lilly defines the penalties (weighting factors) that apply. Optionally, both hard and soft deadlines may also be associated with a per-period gain to be accrued, if the activity concluded in advance of its stated deadline. These gains are meant to offset penalties for violating other soft deadlines. In general, we do not restrict the defined deadlines to apply against the completion time of an activity; we rather generalize this concept and, when applicable, impose a deadline for partial completion of an activity. For example, stability studies for API and drug products are expected to start immediately after material production and release, while the regulatory expectation of demonstrating twelve months of API and drug product stability prior to a new drug application submission adds a back-end constraint to the activity, given submission deadlines. Moreover, predefined time windows for activities may also be added to the model by fixing binary variables associated with activity start times. This feature is useful, for example, to address cases where development cannot start until budgets are issued.

Horizon, Sub-horizons, and Rolling Horizon

We formulate the mathematical model based on a uniform discrete-time representation in which a ten-year planning *horizon* is divided into one month intervals. This representation forces the beginning and ending of an activity to be associated with calendar months, which from an operational perspective, constitutes an acceptable approximation for the Lilly organization. With the ten-year horizon as the superset, we also define (possibly overlapping) *sub-horizons* along which the doubly-

constrained resource consumptions are limited. Often, the sub-horizons coincide with calendar or fiscal years, as budgets for molecule assets in the portfolio are appropriated accordingly.

At this stage, uncertainty in input data (e.g., clinical trial results, budget availability, resource reduction due to employee attrition, management guidance for rapid acceleration/deceleration of certain molecules) is mitigated by implementing a *rolling horizon* approach, where the deterministic problem is solved iteratively, advancing the optimization horizon forward once a month. Therefore, the input data collected by Lilly is updated monthly according to the portfolio progress, and an updated optimal schedule is determined and broadcast to the organization, while the activities scheduled for the near-term are committed upon. In order to accommodate this rescheduling approach, an *activity status* is designated for each activity. The activity status field admits the values *completed*, *committed*, or *to be scheduled*. The latter are effectively the activities that are subject to optimization. In contrast, a firm decision about mode and exact timing of the committed activities has already been made, and thus those are no longer amenable to rescheduling. Note that the set of committed activities includes also the ones that are presently ongoing. From an implementation point of view, such activities can easily be handled in the optimization model by fixing binary variables. Finally, the completed status is reserved for activities that have already occurred in the past and are now concluded. Although they do not need to be explicitly referenced in the scheduling model, these activities might still be necessary to account for due to their generalized precedence relationships with activities to be scheduled. Evidently, as the rolling horizon advances each month, there is a gradual shifting of activities from a "to be scheduled" to "committed" and finally to "completed" status.

Mixed-Integer Linear Programming Model

This portfolio-wide scheduling and resource allocation problem for R&D activities is formulated as a mixed integer linear model (see Appendix), which represents an extended multi-mode resource-constrained project scheduling problem. The primary decision variables correspond to the start time and execution mode of each activity, while the constraints enforce generalized precedence relations, resource limitations and deadline fulfillment. Given Lilly's application setting, we propose multiple potential objectives that align with various SMDD business goals, such as to minimize the net present cost of portfolio development,[6] minimize total deadline tardiness, accelerate the scheduling of some critical activity (e.g., accelerate the launch of a priority molecule), or maximize resource

[6]This cost might include direct spend as well as indirect charge related to resource utilization.

utilization, to name but a few. Stakeholders can also be presented with *optimal* plans that negotiate between these different objectives (e.g., *Pareto frontier* solutions).

Implementation and Test Results

In order to test the performance of our MILP-based decision support system, we instantiated the model using realistic inputs from part of the current SMDD portfolio, and we conducted a number of motivated case studies. The model was implemented and solved using the C++ advanced programming interface of the IBM ILOG CPLEX Optimizer 12.9, while all computational experiments were conducted on four threads of an Intel Xeon CPU E5-2689v4 @3.10GHz processor with 64GB of RAM. All optimizations were conducted down to zero optimality gaps.

Baseline Case Study

Based on the portfolio-wide data, we first generate our baseline case study. SMDD requires the schedule to be optimized across a ten-year planning horizon, which is discretized in monthly periods. Developing a portfolio of 17 different molecule products can be mapped to a total of 135 activities, including 32 campaigns. On average, each activity is associated with 6.69 execution modes. Furthermore, in order to account for the capacity of the manufacturing facilities to which the campaigns are outsourced, we limit the number of concurrent campaign activities to not exceed 4.[7] We also consider 4 distinct types of personnel (FTE: full-time equivalents) as renewable resources, while we consider the direct spend across product classes (sub-portfolios) as resources of mixed type with both per-period as well as horizon-wide availability limits. Notably, these availability limits need not remain constant throughout the horizon, a feature that we use to model the real life practice of budgets being gradually released over time. More specifically, three distinct budgets are taken into account, each one limiting the spend against the development of 5 or 6 molecule products. The first two budgets become available in two parts, once at the beginning of the horizon and once after the one- or two-year mark. The last budget is allocated at the beginning, but comes with the additional restriction of having to be consumed within the five-year mark (see blue lines in Figure 5).

The primary optimization objective chosen for the baseline case study is to minimize the total deadline tardiness, which considers deadline-specific, per-month penalties (weights). Per-month

[7] This constraint can be easily implemented as a renewable resource of which each campaign activity consumes one unit.

gains (negative weights) can also be considered for the earliness of certain activities. This objective is hereafter referred to as "total delays." Among the 135 activities, 7 hard deadlines and 54 soft deadlines are imposed. Minimizing the net present total direct spend was also considered as a secondary objective, in order to break possible ties. This was generically implemented by resolving the optimization model, constraining the primary objective function to be equal to its optimal value.[8] In defining the secondary objective, an appropriate per-period (i.e., monthly) interest rate was taken into account. The total net present spent is hereafter referred to as "total cost," and it is reported in arbitrary monetary units (m.u.).

The optimal Gantt chart resulting from solving the baseline case is presented in Figure 3, where we use different colors to represent each molecule product. The blue and red thick bars represent the soft and hard deadlines, respectively, as applicable for the activities in the corresponding rows. The optimal total delays for the baseline test instance is 18 months, while all activities in the portfolio can be completed within 70 months. We also present in Figure 4 the personnel utilization profiles for the four types of FTE considered. We observe that, with the exception of type 1, the availability of personnel is restricting, as there exist certain periods in which the latter are utilized up to their availability levels. Finally, we present in Figure 5 the cumulative and per period consumption of sub-portfolio budgets, revealing that all three budgets are limiting as a non-renewable resource.

[8] In this approach, the second optimization run benefited by utilizing the first optimal solution as an initial incumbent.

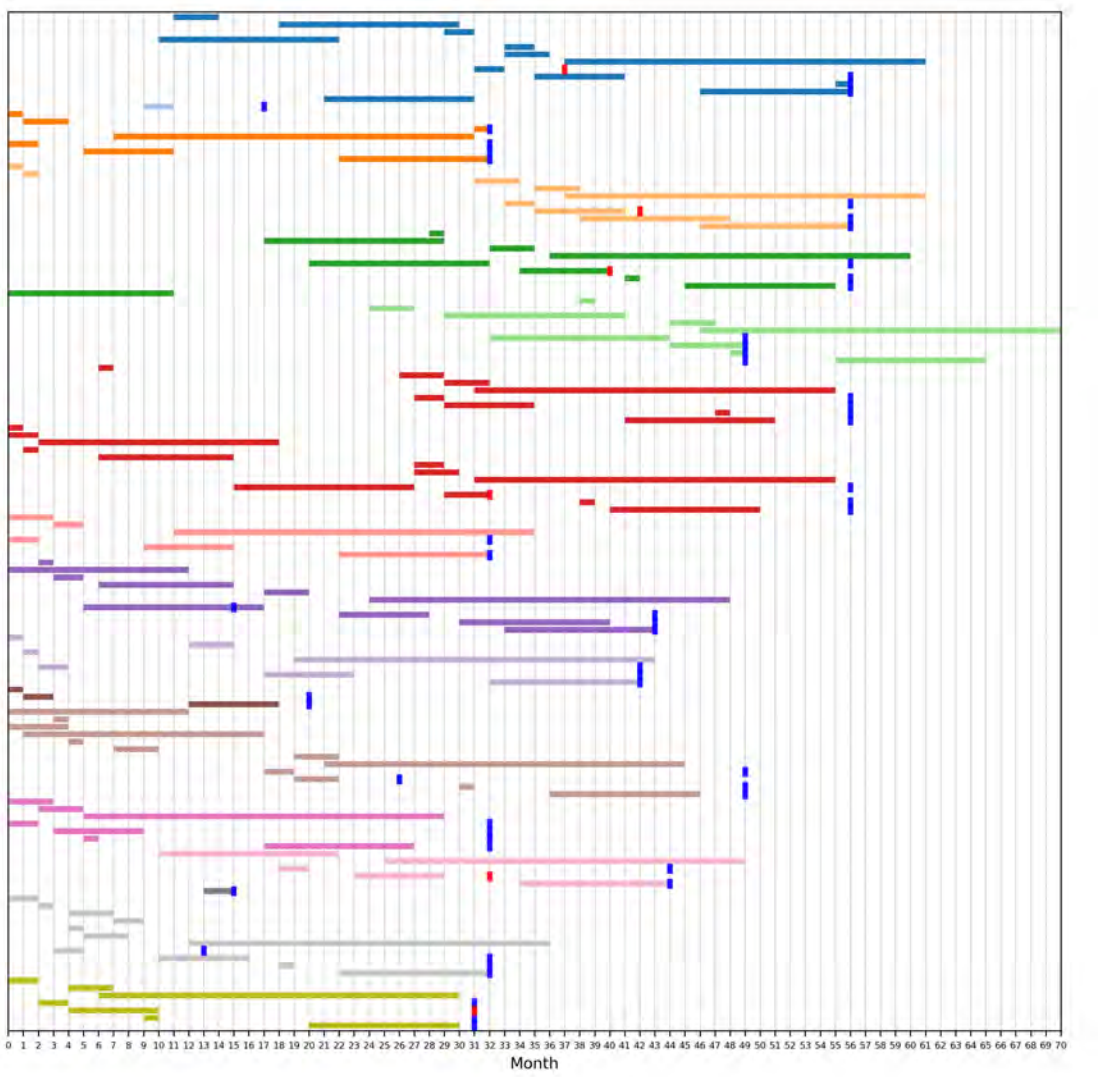

Figure 3: The optimal Gantt chart for the baseline case study.

In order to provide a reference for the computational effort, in Table 2 we report some details on the input size, the size of the resulting MILP formulation (i.e., number of non-fixed variables[9] and number of constraints), as well as the two optimal objective values and the CPU time it took the solver to close the gap. Factoring in the frequency at which Lilly is expected to be updating its input to this planning problem, the CPU times required are deemed to be relatively short, demonstrating that the MILP approach can be a viable means for operational decision-making in this setting.

[9]In general, a number of binary variables can be fixed to zero after application of the well-known *Critical Path Method* (Kelley 1961), which we applied before optimizing our models.

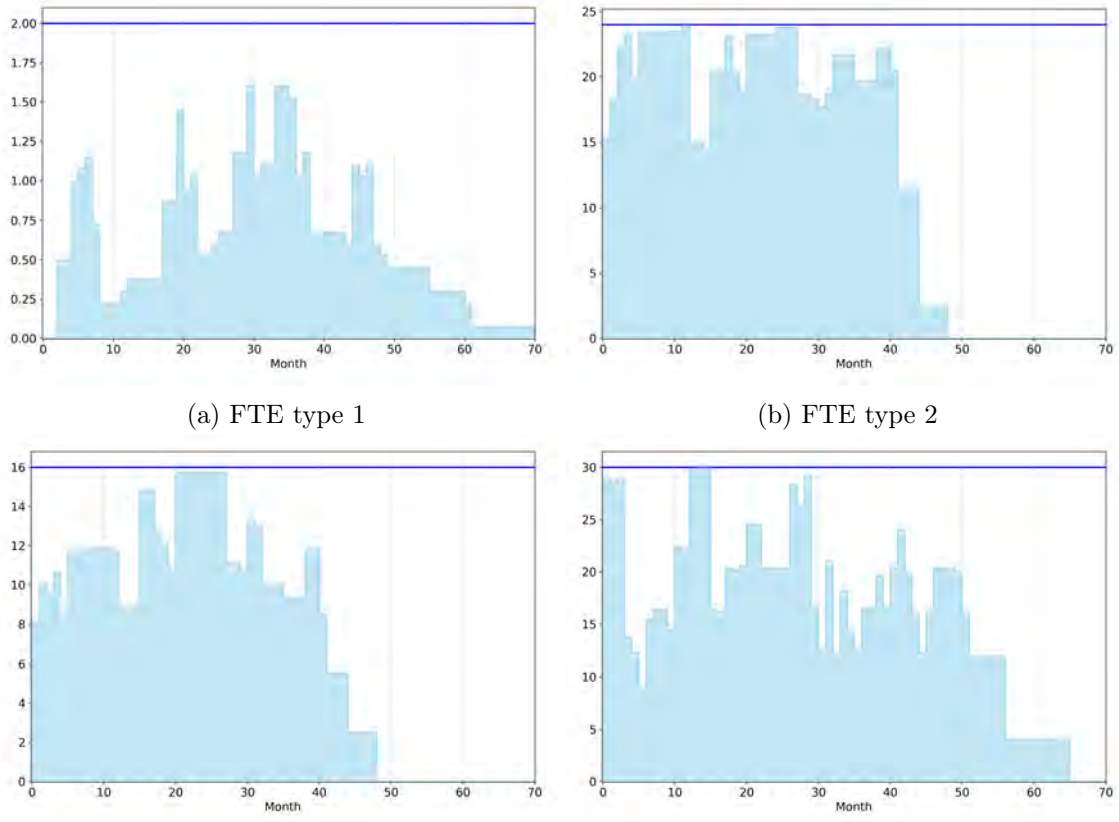

(a) FTE type 1

(b) FTE type 2

(c) FTE type 3

(d) FTE type 4

Figure 4: FTE utilization profiles (in normalized FTE units) at the optimal solution.

Table 2: Computational results for the baseline instance.

# of Activities	# of Modes per Activity	# of Periods	# of Bin. Variables	# of Cont. Variables	# of Constraints	Total Delays (months)	CPU Time (s)	Total Cost (m.u.)	CPU Time (s)
135	6.69	120	79,857	115	1,491	18	202.6	30.480	201.2

The Pareto Frontier

In order to better highlight the effect that the commitment to minimize total delays has on the total cost, we used the ε-constrained method (Mavrotas 2009) to perform bi-objective optimization of our baseline case and produce Pareto optimal solutions that negotiate these two objectives. The Pareto frontier is thus presented in Figure 6. As expected, by allowing an increase in overall tardiness, the net present spend required decreases. However, it is noticeable from the "step-like" shape of this Pareto frontier that multiple-month delay increases are sometimes necessary to drive down the total cost further. In particular, increasing the total delays beyond 24 months does not

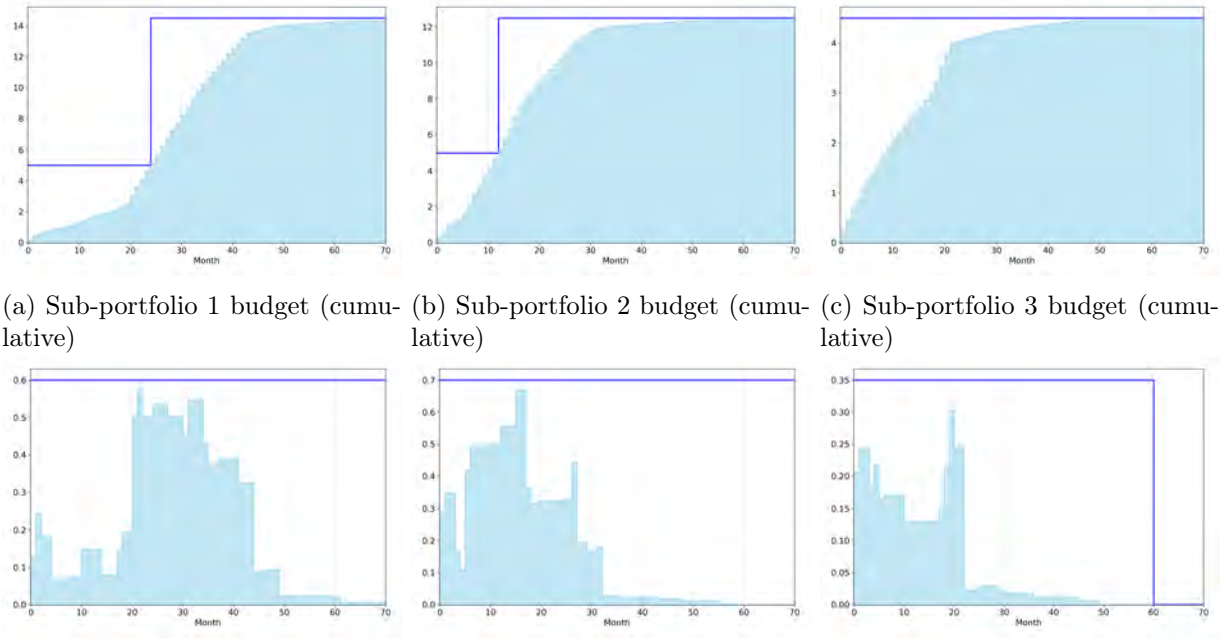

(a) Sub-portfolio 1 budget (cumulative)

(b) Sub-portfolio 2 budget (cumulative)

(c) Sub-portfolio 3 budget (cumulative)

(d) Sub-portfolio 1 budget (per period)

(e) Sub-portfolio 2 budget (per period)

(f) Sub-portfolio 3 budget (per period)

Figure 5: Cumulative and per period budget consumptions (in monetary units) at the optimal solution.

translate to any significant cost savings.

Impact of Concurrent Campaigns

Here, we seek to gain some insights about the impact that the limited infrastructure to perform campaign activities has on the optimal schedule. To this end, we focus on the constraint that limits the number of concurrent campaigns, and by varying the latter and re-optimizing, we establish the sensitivity of the minimum net present spend on this quantity. This analysis provides valuable information for SMDD management to determine whether additional investments to expand the campaign capacity are warranted.

Figure 7 presents the results of this sensitivity analysis. We observe that, by allowing 5 concurrent campaigns (i.e., one more than in the baseline case), the optimal total cost decreases by approximately 3%.[10] However, any additional increase in our capacity to conduct campaign activities has marginal to no additional benefit. Consequently, this suggests that investing in additional capacity to run up to 5 (instead of only up to 4) concurrent campaigns should be further explored to debottleneck overall portfolio development.

[10]Note that, for this analysis, the total allowable delays were fixed to their optimum value, namely 18 months.

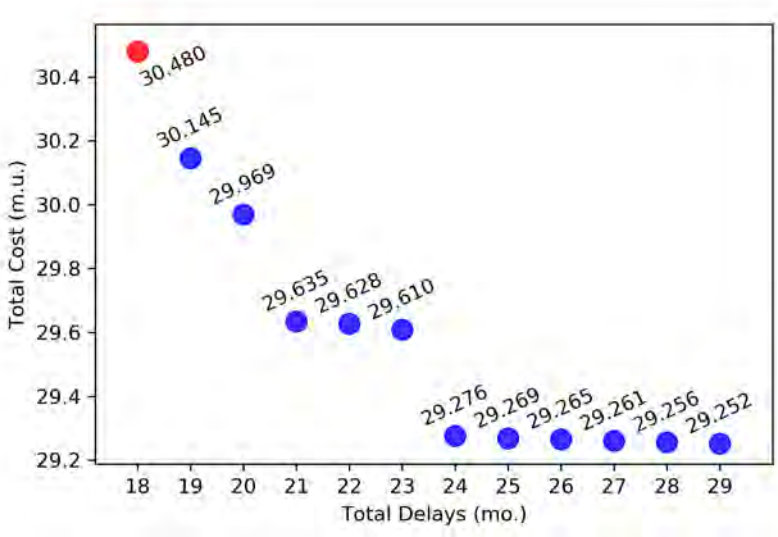

Figure 6: The Pareto frontier for the primary and secondary objectives (the original baseline solution is marked in red).

Benefit From Increasing Available Resources

Recognizing that even a one-month acceleration of launching a product to the market may potentially translate to significant value for the company, SMDD is also interested at identifying the impact that additional available FTE or budget allocations could have on further minimizing the total delays. It is worthy to note that deadlines are generally imposed against critical portfolio development activities, and hence, reducing delays of completing the latter directly translates to accelerating the time to launch final products. Recall that, by solving the baseline instance, we have determined that the minimal total delays is 18 months when resource availabilities are at their current levels. Below we present the results of the sensitivity analysis.

Figure 8a corresponds to the case of increasing the amount of available FTE. In each run, the postulated increase is implemented at time period 12 onwards (i.e., one full year after the current time), reflecting the fact that some lead time is necessary for new personnel hiring to take effect. We remark that we allow for the possibility of all four FTE types to increase, as long as the total increase does not exceed the amount reported in the horizontal axis.[11] The plot reveals a step-like behavior, alluding to the fact that typically a single-unit FTE increase is not sufficient to debottleneck the organization enough for a critical activity to be pulled earlier in the horizon. Evidently, such a

[11]This was implemented simply by adding (i) slack variables to indicate the required FTE availability increase for each type, and (ii) a constraint to limit the sum of those slack variables.

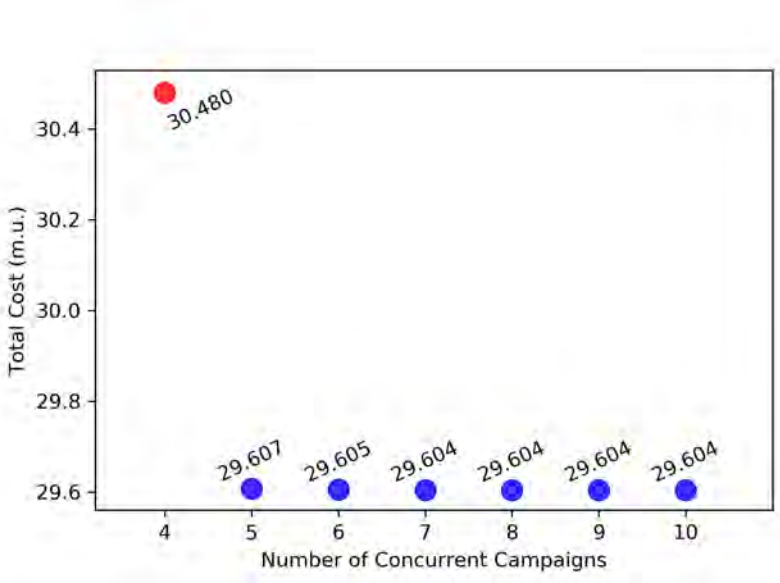

Figure 7: Optimal net present spend as a function of the number of allowable concurrent campaigns (the original baseline solution is marked in red).

quantitative analysis can assist Lilly at identifying bottlenecks in the current portfolio's resourcing plans and make more informed decisions regarding human resource planning, including identifying specific FTE types that would be more helpful to be added in the present effort.

In Figure 8b, we present optimal solutions for the case where budgets can be increased. Here, we allow for the possibility of all three budget groups to be increased, as long as the total increase does not exceed the amount reported, while the per period availabilities are increased by 10% of the required increase for the corresponding budget group, as determined by the optimizer. Similarly to the case of FTE increase, budget increases are only factored in as available at time period 12 onwards, since budget increase requests usually require some lead time until they are approved. We observe that, in the range between 0.3 and 0.9 m.u. of total budget increase, a rapid decrease of total delays down to 15 months can be achieved. However, no further reductions appear to be possible by increasing allocated budgets. In addition, no noticeable deviation is observed in terms of the secondary objective (total cost) across all the solutions depicted in both plots, which can be attributed to little shifting of modes as resource availabilities are perturbed.

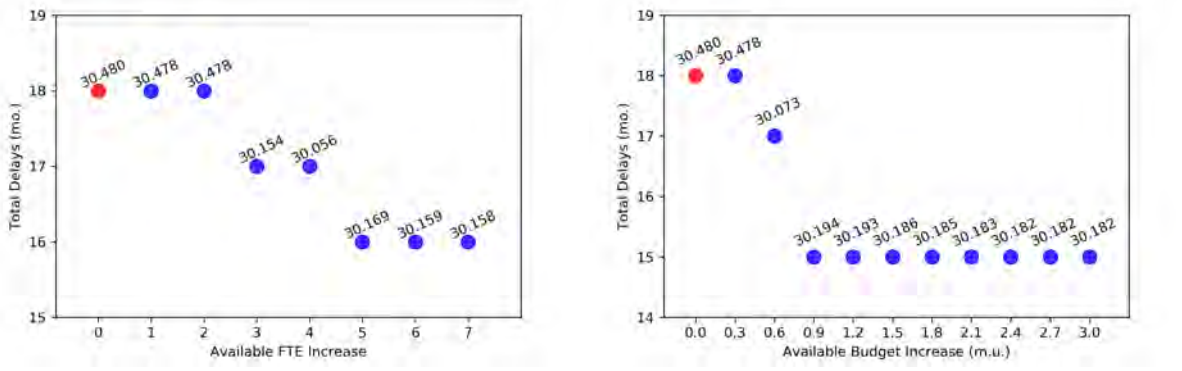

(a) Additional FTE

(b) Additional budget

Figure 8: Minimal total delays as a function of resource availability increases (the original baseline solution is marked in red).

Sudden Portfolio Size Increase

In practice at a pharmaceutical R&D organization, there often arise opportunities for new products to enter the portfolio. These could be new products, developed organically within the company, or products in later stages of development that were onboarded from other companies as part of intellectual property acquisition (a.k.a., *drop-in assets*). In the following, we test how our mathematical model performs in light of such significant changes in input data.

More specifically, we study three augmented instances featuring an expanded portfolio compared to that considered in the baseline case study. The instances are generated by starting with the original portfolio and then duplicating the molecule products whose development is subject to each of the three budget groups (duplication of only one group in each new instance). All applicable precedence relationships and activity data are kept the same. The only exception pertains to the deadlines of activities associated with drop-in assets, which are pushed back by 18 months compared to the original data, in order to reflect a more realistic expectation for the development timeline of these newly introduced products. While the FTE resources will now be shared and rebalanced across the augmented portfolio, it is assumed that the drop-in assets will not compete with existing products for budget resources, rather be subject to their own budget (a mere duplicate of the original applicable budget data). Moreover, we fix the development plan to the baseline case solution for the first 3 months in the horizon, avoiding in this manner substantial changes to near-term FTE allocations.

Table 3 reports the relevant computational results, following the same format as Table 2. It can

be observed that the time to solve the model under the total delays minimization objective increases for all three drop-in scenarios, and so does the time to solve under the cost minimization objective in two out of three runs. Undoubtedly, as the number of to-be-scheduled activities increases and the resources become more limiting for the portfolio, the time to solve such optimization problems to optimality tends to increase. However, for the input sizes we considered in this work, the computational times remain reasonable enough to support decision-making at the planning level using the proposed MILP based approach.

Table 3: Computational results for the drop-in assets study.

	# of Activities	# of Modes per Activity	# of Periods	# of Bin. Variables	# of Cont. Variables	# of Constraints	Total Delays (months)	CPU Time (s)	Total Cost (m.u.)	CPU Time (s)
Baseline	135	6.69	120	79,857	115	1,491	18	202.6	30.480	201.2
Drop-in 1	154	5.99	120	77,591	140	1,599	58	2762	44.059	931.4
Drop-in 2	154	6.20	120	82,737	139	1,595	37	1,227	46.255	1,360
Drop-in 3	151	6.11	120	79,115	139	1,590	23	2,421	36.121	180.3

Integration with Real-life Decision-making

To best utilize this new decision-making and analytical capability presented in this work, Lilly developed a digital tool and also implemented a new work process to utilize it. The tool can be configured to tackle several business-driven scenarios, demonstrating the capability of the mathematical optimization based method to support decisions regarding portfolio development.

Decision-making Tool

One of the major challenges to full digitization of the decision-making process is managing the data. Information retrieval from source systems, such as PDF report documents or lab notebooks, is manual, expensive, error-prone and slow. Lilly's SMDD organization is evaluating the development and adoption of digital technologies to promote *FAIR* (Find, Access, Inter-operable, Reuse) principles that will enable the company to better integrate product data, ranging from starting materials to packaged drug products, and to digitally link the decisions to commission work so as to generate data with the risk of not meeting product requirements in the absence of data. Arguably, this is a challenge that requires multiple linked layers of semantic ontologies.

Ontologies are an explicit formal specification of the terms in the domain and the relations

among them, which can store and retrieve data that share many structural similarities (Gruber 1993). Not only do ontologies bring great efficiency in the tool development itself, they also enable the R&D department to have consistent, reliable, and well linked data, which further enhances the linkages among all stakeholders in different functions and maximizes collaboration, both internally and externally. We refer to a previous publication for details regarding the ontological framework developed to support the utilization of our tool (Viswanath et al. 2020).

The data integration process at SMDD greatly facilitated the development and implementation of the decision-support tool. Based on an Excel spreadsheet template that keeps every related R&D activity for each asset in the current portfolio, a framework aware of all technological requirements is applied to automatically retrieve the applicable activity network for each specific asset from the superstructure of Figure 1. Similarly, the bottom-up timing and resourcing requests as well as the top-down budgetary requirements and deliverable dates for each R&D activity are obtained from other ontology layers.

The user interface (see relevant screenshot in Figure 9) admits user-friendly, drag-and-drop style data entry from a number of Excel spreadsheets. Automated processes are applied to roll the horizon forward each month, and although most inputs remain relatively stable from month to month, every piece of data can also be manually updated on an as-needed basis. A notable case would be the path-defining decisions for the activity networks of drop-in assets.

The data is then passed along to instantiate the optimization model via a compact JSON-formatted file that is also automatically generated. Once the model is solved, the solution is appended to the JSON file, which would then store both the input and output of the optimization run. The essential information for the optimized schedule could then be extracted to other file formats (e.g., Excel spreadsheet reports) that are more familiar to the target users of this tool. In addition, the target users are given a quick graphical view of the Gantt chart of the optimal schedule as well as the FTE requirements for each type of personnel and the budget consumptions for the different groups of activities (see Figures 10 and 11).

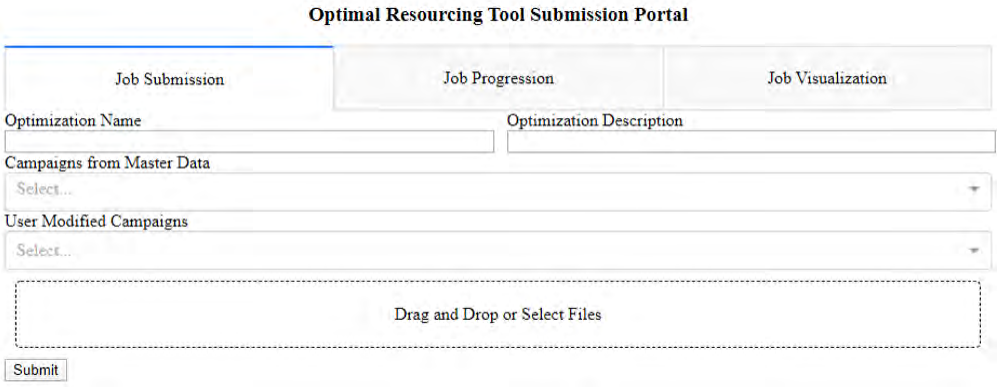

Figure 9: Example screenshot of the user interface.

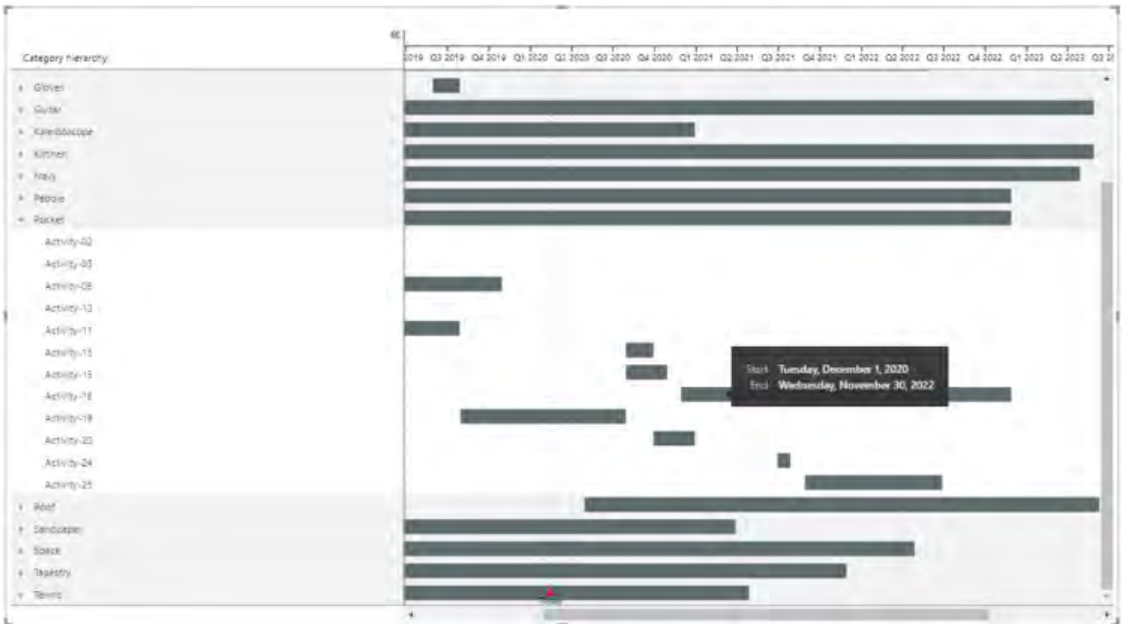

Figure 10: Example Gantt chart visualization.

Figure 11: Example resource consumption profiles visualization.

Scenario-driven Analyses

As demonstrated also in the previous section, one may define appropriate optimization workflows to elicit answers to a series of intriguing "what-if" questions, which are regularly of interest to ask, and which hence have been sufficiently automated within the tool. These include:

- **Asset acquisition or discontinuation.** Portfolio management decision problems are complicated by both the opportunistic entry and the failure-driven exit of candidate products over time, and this feature is especially prominent in the pharmaceutical industry. Given that portfolio product selection is determined by stakeholders at the strategic level, our mathematical model, which targets operational planning-level decisions, treats the portfolio as given input. However, the acquisition or discontinuation of assets arises frequently, and whenever it happens, current scheduling plans are subject to adjustment according to these strategic decisions. *What if we are tasked with the development of one additional late-phase oncology drug?*

- **Human resource planning.** Sudden changes in personnel availability can have an impact on the organization's ability to address its portfolio development mandate. To that end, management is generally interested at understanding how such changes affect the optimal scheduling and resourcing plan. This could be relevant in a scenario for planned recruitment,

where new hires will come on board and be ready to contribute to the current portfolio. Another case of personnel availability change occurs when a major project suddenly appears and causes sudden shifts in priorities, refocusing individual employees to the new project. Such a scenario leads to a reduction in the overall FTE availability for developing the regular portfolio that is under automated scheduling with our tool. *What if the three analytical chemists retiring next year are not adequately replaced?*

- **Manufacturing capacity disruption.** The organization maintains certain outsourcing facilities to perform the API and drug product campaigns, and it makes sure that scheduled activities will not overload these facilities in time. However, the capacity planning for those facilities, which is outside the authority of the SMDD organization, is often subject to variations. In order to understand the system's ability to cope with changes in outsourcing facilities, we are regularly interested at considering the scenario where the manufacturing site is subject to unexpected short-term shutdowns (e.g., due to bad weather). *What if no new campaigns can start over the next six months?*

- **Access to new technology.** Researchers at Lilly are constantly developing more convenient and cost-effective technologies and workflows. These innovations can reduce the resources required to carry out some R&D activities. By quantitatively analyzing various relevant scenarios, we can gain insights into the tactical decisions concerning the adoption and the timing of implementation of these new capabilities. *What if we upgraded our analytical lab with this new piece of equipment?*

Conclusions

The Small Molecule Design and Development (SMDD) organization of Eli Lilly and Company is constantly tasked with pursuing the development of a large pharmaceutical product portfolio. SMDD desired to implement a systematic approach to utilize its human resources and allocated budgets in an optimal manner that negotiates the complex and unintuitive tradeoffs between various system constraints and the sequence of activities that have to be performed. To this end, we developed an MILP optimization framework to address their needs, as well as developed an associated decision-support tool to enable application of this framework for activity scheduling and resource planning in an operational setting. This tool can also be employed at the tactical level to

quantify the impact of postulated structural changes in the system and to understand the ability of the organization to cope with possible disruptions.

The utilization of the decision-making tool we developed as part of this industry-academia collaboration can greatly enhance SMDD productivity. The numerical case studies demonstrated the application of the tool on realistic portfolio problems as well as how it can be utilized to assess the impact of various unexpected scenarios, including sudden portfolio expansion, resource availability changes and manufacturing bottlenecks. Importantly, the tool facilitates transparent collaborations between individual decision units within SMDD, improving the overall quality of the master portfolio plan. SMDD has embraced the outcomes of the computational case studies reported in this paper and has now started using the optimization framework in a pilot capacity that will allow us to collect additional feedback towards further improving the tool.

Acknowledgments

The authors would like to acknowledge multiple stakeholders within the Small Molecule Design and Development organization from Eli Lilly and Company for their continuous support and input into this effort.

Appendix. MILP Formulation

Notation

Sets

$\mathcal{A}$	Set of activities to be scheduled
$\mathcal{A}^{ds} \subseteq \mathcal{A}$	Set of activities with soft deadlines
$\mathcal{A}^{dh} \subseteq \mathcal{A}$	Set of activities with hard deadlines
$\mathcal{E} \subseteq \mathcal{A} \times \mathcal{A}$	Set of activity pairs with some form of precedence relationship
	(Note: $\mathcal{E} \equiv \mathcal{E}^g \uplus \mathcal{E}^s \uplus \mathcal{E}^f$)
$\mathcal{E}^g \subseteq \mathcal{E}$	Set of activity pairs with restricted delay between them
	(includes pairs governed by standard precedence, i.e., $[0, \infty)$ delay)
$\mathcal{E}^s \subseteq \mathcal{E}$	Set of activity pairs with restricted start time differences
$\mathcal{E}^f \subseteq \mathcal{E}$	Set of activity pairs with restricted finish time differences
$\mathcal{M}_i$	Set of available modes for activity $i \in \mathcal{A}$
$\mathcal{R}^\rho$	Set of renewable (i.e., with intensive availability) resources
$\mathcal{R}^\delta$	Set of doubly-constrained (i.e., with both intensive and extensive availabilities) resources
$\mathcal{N}_r^{sub}$	Set of sub-horizons for limiting extensive availability of resource $r \in \mathcal{R}^\delta$
$\mathcal{T}$	Set of time periods
$\mathcal{T}_{rk}^{sub} \subseteq \mathcal{T}$	Set of time periods of sub-horizon k for resource $r \in \mathcal{R}^\delta$

Indices

i, j	$1 \ldots \lvert\mathcal{A}\rvert$
m	$1 \ldots \lvert\mathcal{M}_i\rvert$
r	$1 \ldots \lvert\mathcal{R}^{\rho/\delta}\rvert$
k	$1 \ldots \lvert\mathcal{N}_r^{sub}\rvert$
t	$1 \ldots \lvert\mathcal{T}\rvert$

Parameters

$\left[g_{ij}^L, g_{ij}^U\right]$	Interval of time periods after activity i finishes in which activity j must start (for all $(i,j) \in \mathcal{E}^g$)
$\left[s_{ij}^L, s_{ij}^U\right]$	Interval of time periods after activity i starts in which activity j must start (for all $(i,j) \in \mathcal{E}^s$)
$\left[f_{ij}^L, f_{ij}^U\right]$	Interval of time periods after activity i finishes in which activity j must finish (for all $(i,j) \in \mathcal{E}^f$)
b_{imr}^{ρ}	Consumption of renewable resource r by activity i at mode m per time period
b_{imr}^{δ}	Total consumption of doubly-constrained resource r by activity i at mode m
B_{rt}^{ρ}	Availability of renewable resource r at time period t
$B_{rt}^{\delta\rho}$	Availability of doubly-constrained resource r at time period t
$B_{rk}^{\delta\nu}$	Total availability of doubly-constrained resource r during sub-horizon k
p_{im}	Duration of activity i in mode m
π_i^+	Gain for earliness of activity $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$
π_i^-	Penalty for lateness of activity $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$
δ_i	Due period (deadline) of activity $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$
n_{im}^d	Number of periods of activity $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$ in mode $m \in \mathcal{M}_i$ that must have been completed by the due period δ_i (for activities that must have been fully completed, use the value of p_{im})
α	Interest rate per time period
H	Horizon ($= \lvert\mathcal{T}\rvert$)
$r^* \in \mathcal{R}^{\delta}$	Doubly-constrained resource whose total consumption needs to be minimized (objective 2 only)

Binary variables

x_{imt}	1, if activity i starts in mode m at time period t; 0, otherwise
y_i^s	1, if activity i finishes before its due period δ_i; 0, otherwise (for all $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$)

Continuous variables

$\tilde{s}_i^+$	Earliness of activity $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$
$\tilde{s}_i^-$	Lateness of activity $i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds}$

Equations

Objectives (primary and secondary)

$$\min_{x_{imt},\tilde{s}_i^+,\tilde{s}_i^-,y_i^s} \quad \sum_{i\in\mathcal{A}^{ds}} \left(\pi_i^- \tilde{s}_i^- - \pi_i^+ \tilde{s}_i^+\right) + \sum_{i\in\mathcal{A}^{dh}} \left(-\pi_i^+ \tilde{s}_i^+\right) \tag{1}$$

$$\min_{x_{imt},\tilde{s}_i^+,\tilde{s}_i^-,y_i^s} \quad \sum_{i\in\mathcal{A}} \sum_{m\in M_i} \sum_{t\in\mathcal{T}} \left\{ \sum_{t'=t}^{t+\lceil p_{im}\rceil-1} \frac{b^{\delta}_{imr^*}/\lceil p_{im}\rceil}{(1+\alpha)^{t'}} \right\} x_{imt} \tag{2}$$

Sequencing Constraints

$$\text{s.t.} \sum_{m\in M_i} \sum_{t\in\mathcal{T}} x_{imt} = 1 \qquad \forall\, i\in\mathcal{A} \tag{3}$$

$$\sum_{m\in\mathcal{M}_j} \sum_{t\in\mathcal{T}} t\; x_{jmt} \geq \sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} \left(t + \lceil p_{im} + g_{ij}^L\rceil\right) x_{imt} \qquad \forall\, (i,j)\in\mathcal{E}^g \tag{4}$$

$$\sum_{m\in\mathcal{M}_j} \sum_{t\in\mathcal{T}} t\; x_{jmt} \leq \sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} \left(t + \lfloor p_{im} + g_{ij}^U\rfloor\right) x_{imt} \qquad \forall\, (i,j)\in\mathcal{E}^g \tag{5}$$

$$\sum_{m\in\mathcal{M}_j} \sum_{t\in\mathcal{T}} t\; x_{jmt} \geq \sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} \left(t + \lceil \min\left\{s_{ij}^L, p_{im}\right\}\rceil\right) x_{imt} \qquad \forall\, (i,j)\in\mathcal{E}^s \tag{6}$$

$$\sum_{m\in\mathcal{M}_j} \sum_{t\in\mathcal{T}} t\; x_{jmt} \leq \sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} \left(t + \lfloor s_{ij}^U\rfloor\right) x_{imt} \qquad \forall\, (i,j)\in\mathcal{E}^s \tag{7}$$

$$\sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} (t + \lfloor p_{im}\rfloor)\; x_{imt} \geq \sum_{m\in\mathcal{M}_j} \sum_{t\in\mathcal{T}} \left(t + \lceil p_{jm} - f_{ij}^U\rceil\right) x_{jmt} \qquad \forall\, (i,j)\in\mathcal{E}^f \tag{8}$$

$$\sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} (t + \lceil p_{im}\rceil) x_{imt} \leq \sum_{m\in\mathcal{M}_j} \sum_{t\in\mathcal{T}} \left(t + \lfloor p_{jm} - \min\left\{f_{ij}^L, p_{jm}\right\}\rfloor\right) x_{jmt} \quad \forall\, (i,j)\in\mathcal{E}^f \tag{9}$$

Resource Consumption Constraints

$$\sum_{i\in\mathcal{A}} \sum_{m\in\mathcal{M}_i} b^{\rho}_{imr} \sum_{\substack{t'\in\mathcal{T}\cap \\ \{t-\lceil p_{im}\rceil+1,\ldots,t\}}} x_{imt'} \leq B^{\rho}_{rt} \qquad \forall\, t\in\mathcal{T},\ \forall\, r\in\mathcal{R}^{\rho} \tag{10}$$

$$\sum_{i\in\mathcal{A}} \sum_{m\in\mathcal{M}_i} \frac{b^{\delta}_{imr}}{\lceil p_{im}\rceil} \sum_{\substack{t'\in\mathcal{T}\cap \\ \{t-\lceil p_{im}\rceil+1,\ldots,t\}}} x_{imt'} \leq B^{\delta\rho}_{rt} \qquad \forall\, t\in\mathcal{T},\ \forall\, r\in\mathcal{R}^{\delta} \tag{11}$$

$$\sum_{i\in\mathcal{A}} \sum_{m\in\mathcal{M}_i} \sum_{t\in\mathcal{T}} \left\{ \sum_{\substack{t'\in\mathcal{T}^{sub}_{rk}\cap \\ \{t,\ldots,t+\lceil p_{im}\rceil-1\}}} \frac{b^{\delta}_{imr}}{\lceil p_{im}\rceil} \right\} x_{imt} \leq B^{\delta\nu}_{rk} \qquad \forall\, k\in\mathcal{N}^{sub}_r,\ \forall\, r\in\mathcal{R}^{\delta} \tag{12}$$

Activity Deadline Constraints

$$\sum_{m \in M_i} \sum_{t \in \mathcal{T}} \left(t + n^d_{im}\right) x_{imt} + \tilde{s}^+_i - \tilde{s}^-_i = \delta_i \qquad \forall i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds} \tag{13}$$

$$0 \leq \tilde{s}^+_i \leq H y^s_i \qquad \forall\, i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds} \tag{14}$$

$$0 \leq \tilde{s}^-_i \leq H(1 - y^s_i) \qquad \forall\, i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds} \tag{15}$$

$$y^s_i = 1 \qquad \forall\, i \in \mathcal{A}^{dh} \tag{16}$$

Variable Integrality

$$y^s_i \in \{0, 1\} \qquad \forall\, i \in \mathcal{A}^{dh} \cup \mathcal{A}^{ds} \tag{17}$$

$$x_{imt} \in \{0, 1\} \qquad \forall\, t \in \mathcal{T},\ \forall\, m \in \mathcal{M}_i,\ \forall\, i \in \mathcal{A} \tag{18}$$

In the above model, the primary objective function (1) minimizes total penalties for missing soft deadlines, while also accounting for credits due to meeting (soft or hard) deadlines with slack. The secondary objective function (2) minimizes the net present cost associated with the total consumption of a doubly-constrained resource $r^* \in \mathcal{R}^\delta$. In our context, r^* is chosen to be the total direct spend.

Constraints (3) ensure that every activity is executed, while constraints (4)–(9) enforce the applicable (complete, partial, or otherwise restricted) precedence relations among activities. Constraints (10) limit the consumption of renewable resources within their availabilities in each period. Similarly, constraints (11) limit the consumption of doubly-constrained resources within their own maximum availability per period. Constraints (12) enforce that the total consumptions of doubly-constrained resources do not exceed their respective total availabilities in each applicable sub-horizon. Constraints (13)–(15) help us encode the earliness and lateness of activities that have a deadline, whereas constraints (16) enforce that hard deadlines cannot admit any lateness. Finally, constraints (17) and (18) declare the binary nature of the corresponding model variables.

References

Bartels JH, Zimmermann J (2009) Scheduling tests in automotive r&d projects. *European Journal of Operational Research* 193(3):805–819.

Cha M, Yu F (2014) Pharma's first-to-market advantage. *Retreived from https://www. mckinsey. com/industries/pharmaceuticals-and-medical-products/ourinsights/pharmas-first-to-market-advantage* .

Colvin M, Maravelias CT (2008) A stochastic programming approach for clinical trial planning in new drug development. *Computers & Chemical Engineering* 32(11):2626–2642.

Colvin M, Maravelias CT (2009) Scheduling of testing tasks and resource planning in new product development using stochastic programming. *Computers & Chemical Engineering* 33(5):964–976.

Colvin M, Maravelias CT (2011) R&d pipeline management: Task interdependencies and risk management. *European Journal of Operational Research* 215(3):616–628.

Demeulemeester EL, Herroelen WS (1997) A branch-and-bound procedure for the generalized resource-constrained project scheduling problem. *Operations Research* 45(2):201–212.

Gruber TR (1993) A translation approach to portable ontology specifications. *Knowledge acquisition* 5(2):199–220.

Hartmann S, Briskorn D (2010) A survey of variants and extensions of the resource-constrained project scheduling problem. *European Journal of Operational Research* 207(1):1–14.

Hill A, Lalla-Ruiz E, Voß S, Goycoolea M (2019) A multi-mode resource-constrained project scheduling reformulation for the waterway ship scheduling problem. *Journal of scheduling* 22(2):173–182.

Jain V, Grossmann IE (1999) Resource-constrained scheduling of tests in new product development. *Industrial & Engineering Chemistry Research* 38(8):3013–3026.

Kelley JE (1961) Critical-path planning and scheduling: Mathematical basis. *Operations Research* 9(3):296–320.

Kolisch R, Meyer K (2006) Selection and scheduling of pharmaceutical research projects. *Perspectives in modern project scheduling*, 321–344 (Springer).

Lorenzoni LL, Ahonen H, de Alvarenga AG (2006) A multi-mode resource-constrained scheduling problem in the context of port operations. *Computers & Industrial Engineering* 50(1-2):55–65.

Mavrotas G (2009) Effective implementation of the ε-constraint method in multi-objective mathematical programming problems. *Applied mathematics and computation* 213(2):455–465.

Poppenborg J, Knust S (2016) Modeling and optimizing the evacuation of hospitals based on the mrcpsp with resource transfers. *EURO Journal on Computational Optimization* 4(3-4):349–380.

Riise A, Mannino C, Burke EK (2016) Modelling and solving generalised operational surgery scheduling problems. *Computers & Operations Research* 66:1–11.

Shah N (2004) Pharmaceutical supply chains: key issues and strategies for optimisation. *Computers & chemical engineering* 28(6-7):929–941.

Talbot FB (1982) Resource-constrained project scheduling with time-resource tradeoffs: The nonpreemptive case. *Management Science* 28(10):1197–1210.

Tiwari V, Patterson JH, Mabert VA (2009) Scheduling projects with heterogeneous resources to meet time and quality objectives. *European Journal of Operational Research* 193(3):780–790.

Varma VA, Pekny JF, Blau GE, Reklaitis GV (2008) A framework for addressing stochastic and combinatorial aspects of scheduling and resource allocation in pharmaceutical r&d pipelines. *Computers & Chemical Engineering* 32(4-5):1000–1015.

Varma VA, Reklaitis GV, Blau GE, Pekny JF (2007) Enterprise-wide modeling & optimization—an overview of emerging research challenges and opportunities. *Computers & Chemical Engineering* 31(5-6):692–711.

Viswanath S, Guntz S, Dieringer J, Vaidyaraman S, Wang H, Gounaris C (2020) An ontology to describe small molecule pharmaceutical product development and methodology for optimal activity scheduling. *Under Review* .

Voß S, Witt A (2007) Hybrid flow shop scheduling as a multi-mode multi-project scheduling problem with batching requirements: A real-world application. *International journal of production economics* 105(2):445–458.

Zapata JC, Varma VA, Reklaitis GV (2008) Impact of tactical and operational policies in the selection of a new product portfolio. *Computers & Chemical Engineering* 32(1-2):307–319.