

A DISCRETE-EVENT SIMULATION TO EXPLORE DISAGGREGATION OF BIOTECHNOLOGY RESEARCH AND DEVELOPMENT WORKFLOWS

Susan S.M. Hanson
Noah Mecikalski
Alex Tobias

Jack Morris
Neal Wagner
Rebecca S. Widrick

The MITRE Corporation
7515 Colshire Drive
McLean, VA 22102, USA

The MITRE Corporation
202 Burlington Road
Bedford, MA 01730, USA

Damon Bayer

The MITRE Corporation
2280 Historic Decatur Road
San Diego, CA 92106, USA

ABSTRACT

Research and development (R&D) of biotechnology products is an iterative process typically characterized by a monolithic workflow in which a single organization takes a project from start to finish through many complex operations. This paper presents a discrete-event simulation methodology to explore an alternative disaggregated workflow in which R&D is managed by a single organization but individual operations are distributed among multiple organizations. This methodology is applied to a protein engineering R&D process to compare the monolithic and disaggregated workflows over a range of conditions and scenarios. Based upon a set of assumed parameters, results identify conditions favorable to either workflow and provide a first indication that the industry's trend towards disaggregation may lead to improvements in development timelines. The methodology also provides a foundation for decision support tools that enable decision-makers to manage biotechnology R&D projects.

1 INTRODUCTION

Within the field of biotechnology (biotech), protein engineering aims to improve the properties of valuable proteins such as enzymes and antibodies for industrial, medical, agricultural, or other applications. This is effected by changing the amino acid sequences of existing proteins and measuring the effects of these mutations on end properties important for the application (Engqvist and Rabe 2019). Because there is effectively an infinite number of possible mutations that could be made, and because it remains challenging to predict the precise effect of mutations on the parameters of interest, these R&D efforts tend to follow iterative cycles where improved mutants are discovered and then further built upon in subsequent cycles, until the desired performance is achieved or the project runs out of time or resources (Liao et al. 2022).

This iterative process, paired with human cognitive biases and external pressures for positive outcomes, can commonly result in R&D costs exceeding manager predictions (Lovallo and Kahneman 2003). Additionally, biotech R&D is typically monolithic, characterized by vertical integration in which a single organization develops a product from concept to final production. This workflow requires that an organization be proficient at most if not all operations to provide a likelihood of success. The monolithic workflow

is not resilient to contingencies, does not encourage specialization, and does not promote the open and fair competition necessary to foster technological advances in the field. Looking to the semiconductor industry and other industries that have disaggregated their vertical processes and developed into new horizontal markets (Camelot Management Consultants 2011; Regan and Heenan 2010; Walsh 2006), can we expect that the same transition will occur in biotech?

There are several lines of evidence which indicate that biotech R&D is trending towards disaggregation. In recent years, the biotech industry has experienced a number of pressures similar to those experienced by other industries (e.g., the chemical, pharmaceutical, automotive, semiconductor, and aerospace industries) in advance of their disaggregation including the reduction in capital investment (Hodgson 2023), pressures to reduce organizational headcounts and costs (Sadovi 2022), the emergence of Contract Research Organizations and Contract Manufacturing Organizations (Bleys et al. 2022), and the establishment of cloud laboratories like Strateos and Emerald Cloud Lab (Armer et al. 2023).

This paper proposes a discrete-event simulation (DES) methodology to address this question. We consider a monolithic workflow, as is currently standard practice, as well as an alternative "disaggregated" workflow in which a single organization conceptualizes and manages product development, but where labor and resources are distributed among several organizations. This workflow disaggregation is hypothesized to improve the speed, cost, resilience, fungibility, and interoperability of biotech R&D. Parameter estimates and overall validation of the analysis are provided by biotech subject matter experts (SMEs).

Biotech product R&D workflows differ from other development processes because they are nonlinear, characterized by repetitive cycles in which operations that yield suboptimal results are expected and necessary to learn and adjust procedures. This process, commonly referred as "Design, Build, Test, and Learn" (Liao et al. 2022), requires managers to re-estimate project timelines and costs multiple times over the course of a single product development R&D program. The combination of complex process flow and uncertainty makes the DES modeling approach a more appropriate fit for this problem when compared to other agent-based or system dynamics modeling approaches (Borshchev and Filippov 2004; Babulak and Wang 2008). Additionally, the nature of iterative improvement in the problem makes DES more suitable than a simpler Markovian alternative (Standfield et al. 2014).

We propose that this DES framework could also be the foundation for tools that aid managers from the onset to capture and manage the uncertainties of the biotech product R&D environment. Such tools would provide more reliable forecasting of product development timelines and costs, recommend services whose operational characteristics align with the consumer's objectives, and provide insights into portfolio management such as which projects to continue and which to abandon to cut losses.

1.1 Related Work

Process-oriented modeling techniques such as DES have been used in various areas of systems biology and healthcare. A DES model is used to estimate DNA-protein binding times in (Ghosh et al. 2007). Multiple studies employed DES models to capture patient flow times in health clinics (Jun et al. 1999). In (Temple and Fone 2002), a DES model is used to compare methods for identifying patients with high risk of heart disease. Similar to the focus of this paper, a DES approach was used to capture disaggregated production workflows in the semiconductor industry (Godding et al. 2007). Although process modeling has become a standard tool in healthcare (Günel and Pidd 2010) and industrial manufacturing (Babulak and Wang 2008), we encountered no studies that apply process modeling to end-to-end biotech product R&D processes. This paper aims to address this gap by using DES to capture the Design, Build, Test, Learn process that characterizes the development of biotech products and to explore the potential benefits and trade-offs of monolithic and disaggregated workflows.

2 MODELING METHODOLOGY

Leveraging the DES framework, we developed a software application, hereafter referred to as the Biotech R&D model or “BiRD model,” designed to explore various process workflows for R&D of biotech products. BiRD provides a quantitative capability that supports assessment, design optimization, and what-if analysis. Specifically, the model uses DES to compare various configurations of monolithic and disaggregated workflows with the aim of identifying workflows and conditions that support more rapid R&D project timelines and lower project costs, as well as increased resilience to disruptions. The model is implemented using the Python-based DES framework SimPy (Lünsdorf et al. 2023).

Workflows are specified using a directed graph whose states capture distinct product R&D operations that are executed sequentially, with stochastic “loopback” probabilities to return to the same or earlier operations, until the final desired product is manufactured. The workflow graph’s edges represent transitions between these operations. An example workflow is shown in Figure 2. Table 1 gives brief descriptions of the model inputs associated with each state and edge in the graph. Each input can also be dimensioned by an attempt number, allowing distributions to vary over time and capture the impact of multiple attempts at completing the operation (i.e., improvements by learning from experience). We note that the times, measured in calendar days, are inclusive of weekends, holidays and other expected downtime.

Table 1: Workflow graph inputs.

Input	Entity	Description
Dwell Time	State	Distribution of execution time (days) for workflow operation
Cost	State	Distribution of cost paid for a successful workflow operation (USD)
Transition Probabilities	Edge	Probability adjacency matrix for transitions from one state to another
Handoff Times	Edge	Distribution of time (days) to transition from current state to next state

To simulate the workflow of a single product, BiRD follows the control flow shown in Figure 1. A simulated product begins at an initial state, waits a random dwell time drawn from a user-specified distribution for the state, randomly determines its next state – including whether it progresses or loops back to a previous state based on a set of user-specified transition probabilities, waits a random handoff time drawn from a user-specified distribution associated with transitions from the current state to the next state, and repeats these steps until the product reaches a terminal state. After conducting many iterations to achieve product completion or failure, the relevant metrics are saved and displayed in a user interface.

For this initial analysis, product workflows do not include queue times to begin an operation, so the results for any workflow represent the best case. Not shown in Figure 1 are costs paid for each state’s operation. The model’s preliminary representation of cost is not the focus of this paper, but SME guidance suggests that development costs will be highly proportional to development times because variable costs - such as labor and supplies costs - predominate in biotech R&D.

2.1 A Biotech Product R&D Process Workflow: Protein Engineering

An example biotech R&D workflow is a protein engineering process, which is defined by the set of nodes and edges depicted in Figure 2. Nodes, commonly referred to as states or operations throughout this paper, are represented by the labeled boxes. Edges, represented by the arrows, indicate the flow from one state to another. Blue arrows indicate the forward flow of the process while red arrows indicate backwards flow.

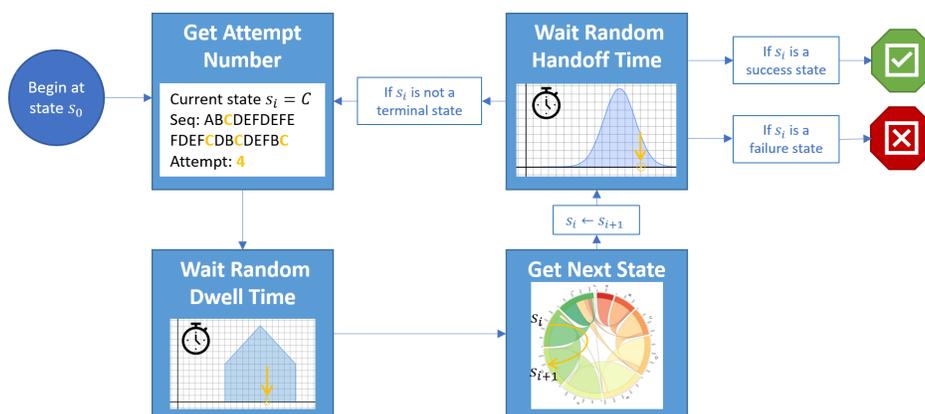


Figure 1: BiRD model control flow.

When a state does not successfully complete, it may repeat the same state or revert to a previous state in the process (self-directed edges are not shown in the figure for clarity). States A and H in the figure are the starting and goal states, respectively. Each state has transition probabilities associated with each edge leaving the state and a dwell time distribution indicating the duration of a state’s operation. Edges may be associated with handoff times which specify the transition time from one state to the next. By SME recommendation, handoff time on backward edges is assumed to be negligible for the purposes of these experiments. Each state may also include a probability of entering a failure state where the process is terminated before reaching the success state. However, these failure states are not used in our experiments. Based on SME guidance, dwell and handoff time probabilities are represented by a modified triangular distribution where that the minimum and maximum values are provided as input and occur with some non-zero probability and where a moderate preference is given to the average.

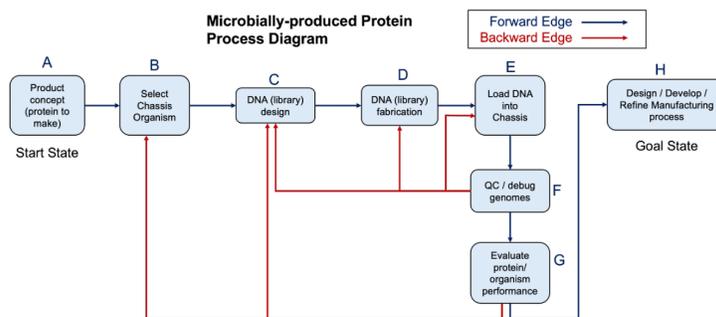


Figure 2: Process for engineering a microbially-produced protein.

2.2 Model Verification and Validation

The BiRD model captures the interaction of complex and uncertain dynamics present in a real biotech product R&D process. The goal of verification is to ensure that the model implementation produces behavior that matches that of the intended model. To achieve this, sensitivity analysis was performed to quantify 1) how many simulation iterations must be executed to achieve convergence of outputs, 2) how small and large variations to input settings affect the magnitude and direction of change in outputs, and 3) how varying input settings across wide ranges and in different combinations affect the ordinality of outputs. All verification experiments are conducted using the protein engineering process described in Section 2.1 and specific scenarios as described in Section 3. For model validation, we relied on SME assessments that

the minimum, maximum, median and average outcomes for total development time and total number of operations were within reasonable ranges based on the input provided.

2.2.1 Output Convergence

Identifying the number of iterations to achieve a stable outcome is critical to the experiments in Section 3. Figure 3 gives the mean and standard deviations of the convergence results for three representative scenarios. The intention is not to compare the performance of these workflows, but rather to determine the number of iterations necessary to reach output convergence for representative scenarios. For all scenarios, including other product types not presented in this paper, greater than 1000 iterations is sufficient to see a convergent result for both the mean and standard deviation as given by the relatively flat line we see after 1000 iterations. It is important to note that not all probabilistic models converge; some diverge and some oscillate between two or more convergent states (Dekking et al. 2005). Run times on a personal laptop (Dell 11th Gen Intel(R) Core (TM) with 16.0 GB of RAM) are roughly 0.65 seconds for 100 iterations, 6.7 seconds for 1000 iterations, and 29.94 seconds for 5000 iterations.

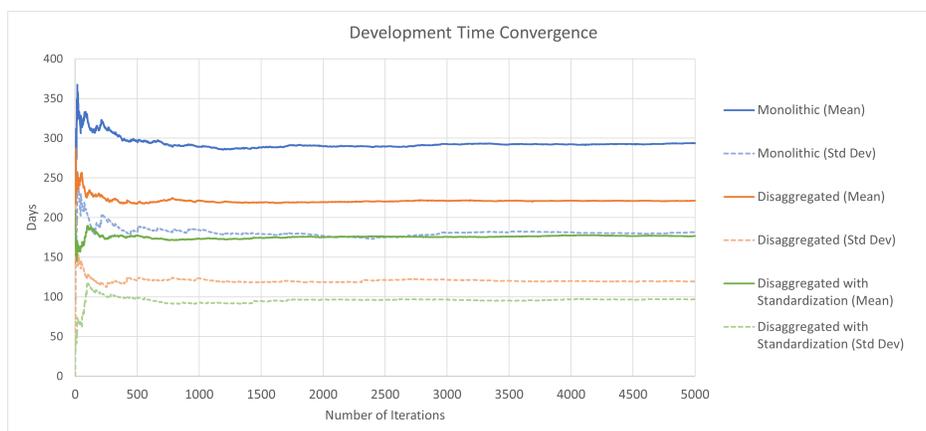


Figure 3: Convergence test results.

2.2.2 Parameter Verification

For the remaining verification experiments, we determined the following desired model behaviors: 1) relatively small changes to any single input (e.g., forward state transition probability, dwell time) resulted in proportionally small changes to appropriate outputs in the appropriate direction, 2) relatively large changes to any single input resulted in proportionally large changes to appropriate outputs in the appropriate direction, and 3) both small and large non-offsetting changes to combinations of inputs produced the expected ordinality of outputs. Offsetting changes, referring to two or more changes whose dynamics pull an output in opposite directions such that the effect on the output is not trivially obvious, are not explored. Outputs checked are total development time and total number of operations completed.

3 SIMULATION EXPERIMENTATION

Two experiments are conducted to evaluate the expected total development time of proposed biotech product workflows. These experiments use the workflow for engineering a microbially-produced protein product shown in Figure 2. For Experiment 1, we compare three scenarios for developing a protein product as posited by biotech SMEs, which include 1) a monolithic workflow, 2) a disaggregated workflow, and 3) a disaggregated workflow with additional specification standardization to improve the transfer of information. For Experiment 2, we conduct an experimental design for determining thresholds that need

to be met or exceeded by a disaggregated scenario to improve upon the total development time of the traditional monolithic workflow. For the purposes of these experiments, an improved workflow is one that has a shorter development time than the traditional monolithic workflow.

3.1 Experiment 1: Protein Engineering Workflow Analysis

Experiment 1 is intended to show the expected differences between three SME-provided scenarios, exploring the effects of disaggregation and disaggregation with standardization. The primary scenario differences include: expected 1-2 day handoff times between operations for the monolithic workflow, compared to 2-3 day handoff times for the disaggregated workflows due to longer exchange times between different organizations. Disaggregated scenario dwell times are about 23% faster on average (ranges between 0%-67% faster) and are less variable than monolithic due to added efficiencies gained through leveraging specialists. Disaggregated scenario forward transition probabilities are lower during earlier attempts but improve at a faster rate (due to more room for improvement) and eventually approach (but do not exceed) monolithic transition probabilities. On average, the success rates for the forward transition probabilities in the disaggregated scenario are about 1.5% lower than the monolithic scenario (ranges between 0%-7% lower). The standardization process for the disaggregated workflow is expected to decrease the likelihood of having to redo an operation and increase the chance of success. Thus, the disaggregated with standardization workflow is identical to the disaggregated workflow with the exception of a 10% decrease in self-loopback probabilities, resulting in success rates for forward transitions that are, on average, 28% and 26% better than disaggregated and monolithic, respectively.

Each workflow is simulated 5000 times to obtain a distribution of development time. The results can be seen in Figure 4, and summary statistics can be seen in Table 2 where **M** is Monolithic, **D** is Disaggregated, and **S** is Disaggregated with Standardization. On average, the total development time of the monolithic workflow is 294 days. The disaggregated scenario has a mean total development time of 221 days, an improvement of 73 days. The disaggregated with standardization case results in a mean total development time of 177 days, a further 44 days saved. Based on respective 25% and 40% improvements in mean development time of the posited disaggregated and disaggregated with standardization workflows, this presents the first indication that disaggregation could offer an effective alternative to the industry's standard practice. Of additional note is the effect of disaggregation and standardization on reducing the standard deviation and shrinking the heavy tails of the probability distribution function.

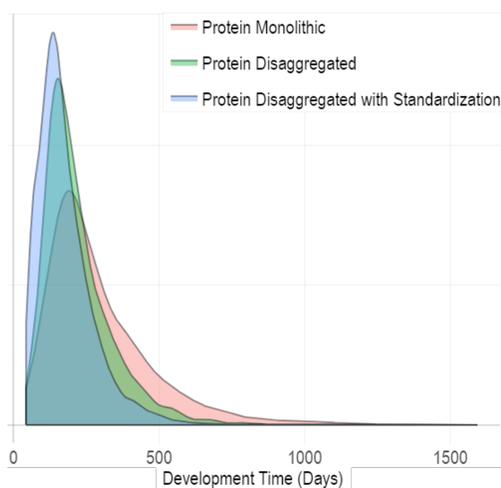


Figure 4: Development time probability density.

Table 2: Summary statistics.

	M	D	S
Minimum	46	47	43
5 th Perc.	97	87	63
20 th Perc.	160	133	104
Median	242	190	154
Mean	294	221	177
80 th Perc.	409	299	239
95 th Perc.	647	454	363
Maximum	1593	1314	912
Std. Dev.	181	119	97

3.2 Experiment 2: Identifying Conditions that Support Disaggregated and Monolithic Workflows

Experiment 2 explores a range of system conditions to identify combinations of conditions where an alternative workflow (e.g., disaggregated) would, and would not, be an improvement over the traditional monolithic workflow. We follow the widely used Design of Experiments (DoE) methodology to systematically formulate and conduct a full factorial experimental design (Montgomery 2013).

The experimental design uses the protein monolithic workflow as the baseline, which will serve as the midpoint of the design. The DES model inputs (DoE factors) that are varied for the experiments can be seen in Table 3. Given the large number of input variables, varying each independently would result in an untenably large number of experiments. Thus, variables are unilaterally varied as a percentage of baseline data and the number of iterations is reduced from 5000 to 2000. The data in Section 2.2.1 indicates that 2000 iterations should be sufficient for a convergence of outputs.

Table 3: Full factorial design of experiments factors and levels.

Success Rates (% of Monolithic Data)	Dwell Times (% of Monolithic Data)	Handoff Times (Days)
80%	120%	2-3
85%	115%	1-2
90%	110%	
95%	105%	
100%	100%	
105%	95%	
110%	90%	
115%	85%	
120%	80%	

(Worse/Better than monolithic workflow)
 (Roughly aligns with data for the disaggregated workflows)

Dwell times are multiplied by the scaling factor denoted by the experiment. For example, if the level is 120%, then the dwell time parameters (minimum and maximum) are multiplied by 1.2. Success probabilities are scaled with respect to their distance to 100% success. For example, they will be 20% closer to one, rather than grow by 20%, thus avoiding issues where the probability of success is greater than 100%. The remaining non-successful probability is distributed among all possible backflows, proportional to their original probabilities. This process is repeated across all states and attempts. Excursion 1 uses the 1-2-day handoff to determine the sensitivity of the baseline to changes in the dwell times and probabilities, while Excursion 2 uses the 2-3-day handoff to explore the threshold improvements needed to make a disaggregated scenario more attractive than the monolithic.

The measured DoE response variable is the change in average total development time when compared to the baseline result of 294 days. The full results of both excursions can be seen in Figure 5. A positive value (red) indicates the combination of settings results in a longer total development time. A negative value (blue) indicates a shorter total development time. The x-axis indicates the success rate and the y-axis indicates the dwell time as levels applied to the baseline data. The black lines show simple thresholds where development times show improvement over the baseline result.

In Figure 5, the left chart shows the results for Excursion 1 where the midpoint, indicated by the blue star, has a difference of zero to show the baseline scenario. We see that it is much easier to increase the development time than decrease it, which makes sense intuitively. We also see that the threshold line and

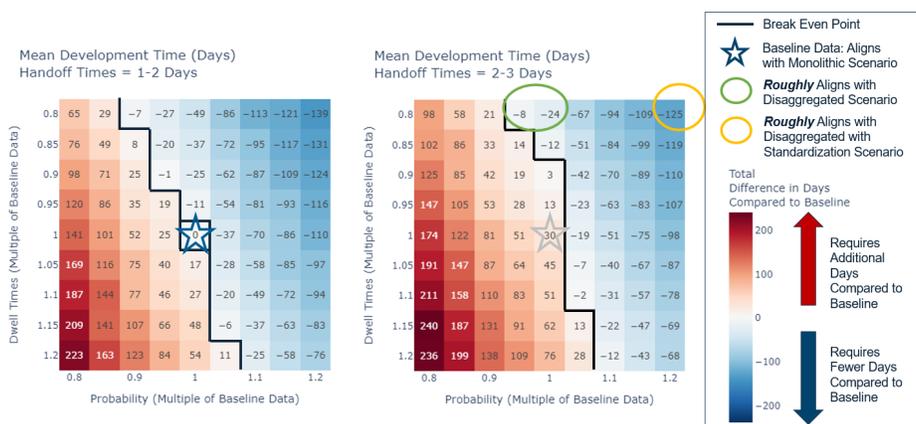


Figure 5: Design of experiments results.

absolute changes are not symmetric, indicating that offsetting changes have non-trivial interactions that can be further explored by operations research techniques.

The right chart contains the results for Excursion 2 which has an extra day of handoff time per handoff (due to shipping, coordination or some other delay between operations). The gray star highlights this simple change results in a development time increase of 30 days on average. We use the chart to identify how much improvement via combinations of reduced dwell times and/or increased success rates of moving forward is required to compensate for the added delays. For example, if dwell times remained the same, the break even point requires achieving a 0% to 5% improvement in all success rates from the baseline (this can be seen as we move to the right of the gray star). Conversely, if success rates remained the same, the break even point requires achieving a 10% to 15% decrease in all dwell times (this can be seen as we move up from the gray star). In addition to finding break-even points, the resulting DoE data helps highlight the trade-offs between improving combinations of dwell time and transition probabilities which in turn can inform where efforts should be focused to improve existing and alternative workflows. The DoE indicates that disaggregation could be beneficial if the providers are more skilled in their particular field and can provide similar or better processing times. However, those savings could be neutralized if managing the requirements and expectations of transferring R&D between providers adds additional delays.

The green and yellow circles in Figure 5 respectively indicate the approximate placement of the SME-posed disaggregated and disaggregated with standardization scenarios from Section 3.1. However, note that these are rough estimates due to rigidity of DoE levels as compared to the finer grained SME inputs. The individual states in the protein process have different magnitudes of dwell time and transition probabilities, and the SME posited improvements were not uniform. For this reason we can only roughly estimate where in the parameter space the disaggregated workflows would fall.

4 RESULTS AND LIMITATIONS

In Section 3.1, we find that disaggregated workflows – with and without standardization – engineer a protein product on average 25-40% faster than the monolithic workflow. Similar results are seen for other biotech product types not presented in this paper. This presents a first indication that disaggregation could offer an alternative to the industry’s standard practice.

However, it is not evident that this improved performance will hold as model assumptions are relaxed. First, the input data assumes that disaggregation and standardization offer an improvement to dwell times and success rates, but the magnitude of this improvement directly relates to the size and diversity of participating service providers, factors which are highly uncertain. Second, the results found in this paper represent best-case outcomes. In particular, the model assumes that the provider of a given state’s service immediately begins work on a product once the handoff is completed when, in reality, the provider

likely balances many projects and may only begin work on a product as its production capacity allows. Thirdly, as shown in 3.2 the accuracy in the baseline success probabilities can significantly affect the expected performance. All of these assumption relaxations have the potential to hinder the performance of disaggregated workflows. When considering the sample protein engineering workflow analyzed herein, this paper indicates that disaggregation decreases timelines by approximately 73-117 days. Therefore, when improving model fidelity, the disaggregated workflow's performance may degrade by up to 73-117 days before it no longer affords a meaningful alternative to the traditional monolithic workflow. Evaluating whether this margin is exceeded will require the development of additional modeling capabilities.

Moreover, the derived utility between disaggregated and monolithic workflows is multifaceted. Our paper is single-focused on the duration of product development, however consumers also care about cost and quality, and their balance of priorities will vary. Though the current model contains a preliminary form of cost, developing these new provider-focused modeling capabilities presents an excellent opportunity to integrate a much more comprehensive form of cost modeling that will help mediate consumer-provider interactions and demonstrate other impacts of disaggregation.

Lastly, the BiRD model covers a single product in a wider ecosystem and does not measure the broader benefit to new entrants who are unable to execute a traditional workflow. This is another potential avenue of model enhancements.

5 SUMMARY

The proposed BiRD model provides a method to evaluate a complex workflow system with multiple paths and outcomes. Its application to the biotech domain represents a first-of-its-kind DES analysis of end-to-end biotech product R&D workflows and their disaggregation. The model demonstrates a M&S capability that could provide multiple benefits to biotech development practitioners. In a field where predicting project timelines and success rates often goes awry, the model could help add systematic analysis to project planning and management, underscoring that even modest improvements in the time required or success probability of an iteratively revisited state can have a significant impact on overall project timelines.

The reported experiments evaluate a small set of scenarios over a broad range of system conditions. The example process examined here, with its assumed parameters, has a best case potential to shave months from overall project timelines if its operations are disaggregated. Further, the disaggregated scenario assumptions could degrade considerably before it would be outperformed by the existing monolithic practice. When exploring a broad range of system conditions, experiments indicate that dynamic interactions between handoff times, dwell times, and operation success rates do hold significant sway over system outcomes. Notably, through a DoE of model inputs, this paper finds that there are areas of opportunity for a disaggregated workflow. Modest time improvements to biotech operations, paired with standardization improvements to compensate for handoff delays, provides an opportunity space that should be further explored.

BiRD exhibits key strengths in its flexibility and generality, its understandability, and its computational tractability. Its underlying DES framework is well suited to enable additional refinements and expansions to better inform the development and utilization of a networked biotechnology ecosystem.

6 FUTURE WORK

This paper is intended to be a first attempt at simulating a disaggregated, networked biotechnology ecosystem. However, the current early-stage design of the BiRD model is limited to perfect provider availability, a single product workflow, and is primarily focused on development time as an objective. Because of these, currently simulated products never experience bottleneck delays or choose from selections of providers – both of which are core mechanics of a disaggregated biotech product ecosystem. To address these gaps, the next iteration of the model would 1) use a multi-objective approach of both development time, product quality, and cost, 2) simulate multiple different products at different points in their workflow and 3) facilitate interactions between ecosystem participants through scarcity and diversity.

Further analysis would consider the impacts of new entrants on a biotech ecosystem and dive deeper into identifying critical workflow assumptions and their impact on outcomes. Pairing this expansion with a revisit to Experiment 1 and broadening the scope to treat each state independently, we would create a richer dataset to explore and allow us to better identify critical parameters in the workflows. This could also enable a project manager to assess the value of outsourcing a particular state by estimating expected changes to dwell times, success rates and handoff delays.

Further development as a decision aid would explore the effects on product workflows when managers are exposed to more information about providers and the biotech ecosystem as a whole.

6.1 Example Case: Supply Chain Disruptions

To illustrate the value of additional development, we consider the impact of supply chain disruptions with tests of an updated version of the model with preliminary forms of provider capacity constraints. With parameters consistent, we simulated the same protein engineering workflow in a networked biotechnology ecosystem while constraining the total service capacity for each state to 15, 10, and 5 jobs at a time. We also considered a scenario in which all states had 15-job capacity with the exception one bottleneck state with a reduced 5-job capacity. Products were randomly initiated according to an exponential arrival rate with a mean of five days (i.e., on average, one new protein product is introduced every 5 days), and competed for services over a duration of several years.

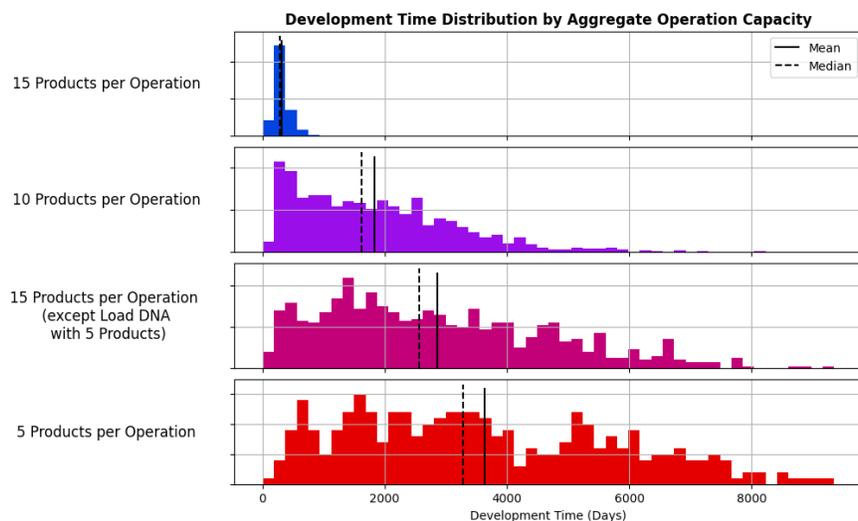


Figure 6: Impact of capacity constraints on development time.

As shown in Figure 6, a production capacity of 15 products per state operation resembles results of the former model in which no product encounters wait times. However, it is clear that decreasing the production capacity quickly escalates development time well beyond what is considered viable. Moreover, it appears that the presence of even a single bottleneck state (e.g., "Load DNA") can significantly disrupt the whole supply chain, nearly as much as limiting all states by the same degree.

Though these efforts are preliminary, they demonstrate the potential for additional model developments to enable the exploration of biotech product R&D workflow disaggregation in terms of its core purposes. The experiments in this paper indicate the capability of disaggregated workflows to produce sample biotech products at a quickened pace. However only by considering production within an interacting network can we fully illuminate qualities that arise through system interactions such as resilience, competition, innovation, and democratization.

6.2 The Model as a Project Management Tool

Across industries, projects frequently exceed their budgets, miss their deadlines, and fail to deliver. Combinations of uncertainty, manager cognitive biases, and external pressures for positive outcomes, lead to project managers rarely anticipating the true cost and duration of projects (Lovallo and Kahneman 2003). Work in biotech often requires an iterative “Design-Build-Test-Learn” approach in which managers cyclically learn from experience and adjust procedures until success can be achieved (Liao et al. 2022). With this uncertainty and a lack of standardization and reproducibility, biotech project R&D process management remains a challenge (El Karoui et al. 2019).

Future versions of the model can provide more systematic analysis to project planning by enabling forecasts of project development timelines and budgets. In its current form, the proposed model captures a full distribution of outcomes and the impacts of failures and delays, underscoring that even marginal improvements in cycle times and success rates can cause large improvements to overall project timelines.

Moreover, with the additional development detailed above, the model could serve as the backbone of a fully-fledged project management tool. In addition to timeline forecasting, the model could provide data-driven service provider recommendations to consumers, scheduling and production-maximizing assistance for providers, evaluations of proposed workflows and services, and portfolio management for all participants in a distributed biotechnology ecosystem.

ACKNOWLEDGMENTS

This work was funded by MITRE’s IR&D program; we thank MITRE Fellow Dr. Marc Salit and Ms. Jenny McDaniel as additional experts for providing their inputs and commentary on this work. Approved for Public Release (Case 23-1440). Distribution unlimited. ©The MITRE Corporation. All rights reserved.

REFERENCES

- Armer, C., F. Letronne, and E. DeBenedictis. 2023. “Support Academic Access to Automated Cloud Labs to Improve Reproducibility”. *PLoS Biology* 21(1):e3001919.
- Babulak, E., and M. Wang. 2008. “Discrete Event Simulation: State of the Art”. *International Journal of Online and Biomedical Engineering* 4(2):60–63.
- Bleys, J., E. Fleming, H. Mirman, and L. The. 2022. “CROs and Biotech Companies: Fine-tuning the Partnership”. Technical report, McKinsey & Company. <https://www.mckinsey.com/industries/life-sciences/our-insights/cros-and-biotech-companies-fine-tuning-the-partnership>, accessed 12th April 2023.
- Borshchev, A., and A. Filippov. 2004. “From System Dynamics and Discrete Event to Practical Agent Based Modeling: Reasons, Techniques, Tools”. In *Proceedings of the 22nd International Conference of the System Dynamics Society*. July 25th-29th, Oxford, England.
- Camelot Management Consultants 2011. *Pharma’s Future with a Disaggregated Supply Chain*. <https://www.chemanager-online.com/en/news/pharmas-future-disaggregated-supply-chain>, accessed 12th April 2023.
- Dekking, F. M., C. Kraaikamp, H. Lopuhaä, and L. Meester. 2005. *A Modern Introduction to Probability and Statistics: Understanding Why and How*. London: Springer.
- El Karoui, M., M. Hoyos-Flight, and L. Fletcher. 2019. “Future Trends in Synthetic Biology—A Report”. *Frontiers in Bioengineering and Biotechnology* 7(175).
- Engqvist, M. K. M., and K. S. Rabe. 2019. “Applications of Protein Engineering and Directed Evolution in Plant Research”. *Plant Physiology* 179:907–917.
- Ghosh, P., S. Ghosh, K. Basu, and S. Das. 2007. “Modeling Protein-DNA Binding Time in Stochastic Discrete Event Simulation of Biological Processes”. In *2007 IEEE Symposium on Computational Intelligence and Bioinformatics and Computational Biology*, 439–446. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Godding, G., H. Sarjoughian, and K. Kempf. 2007. “Application of Combined Discrete-event Simulation and Optimization Models in Semiconductor Enterprise Manufacturing Systems”. In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1729–1736. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Günal, M. M., and M. Pidd. 2010. “Discrete Event Simulation for Performance Modelling in Health Care: a Review of the Literature”. *Journal of Simulation* 4(1):42–51.
- Hodgson, J. 2023. “2022 - Toughing Out the Trough”. *Nature Biotechnology* 41(2):170—173.

- Jun, J. B., S. H. Jacobson, and J. R. Swisher. 1999. "Application of Discrete-event Simulation in Health Care Clinics: A Survey". *Journal of the Operational Research Society* 50:109–123.
- Liao, X., H. Ma, and Y. J. Tang. 2022. "Artificial intelligence: a Solution to Involution of Design-Build-Test-Learn Cycle.". *Current Opinion in Biotechnology* 75:102712.
- Lovullo, D., and D. Kahneman. 2003. "Delusions of Success". *Harvard Business Review* 81(7):56–63.
- Lünsdorf, O., S. Scherfke, and P. Grayson. 2023. *SimPy: Discrete Event Simulation for Python*. Revision 2973dbe7 ed. Read the Docs. <https://simpy.readthedocs.io/en/latest/>, accessed 12th April 2023.
- Montgomery, D. C. 2013. *Design and Analysis of Experiments*. Hoboken, New Jersey: Wiley.
- Regan, M. C., and P. T. Heenan. 2010. "Supply Chains and Porous Boundaries: The Disaggregation of Legal Services". *Georgetown Law Faculty Publications and Other Works* 78(11-04):2142–2159.
- Sadovi, M. W. 2022. "Layoffs Part of Top-down, Bottom-up Cost-cutting Strategy, Surface CFO Says". <https://www.cfodive.com/news/layoffs-part-top-down-bottom-up-cost-cutting-strategy-surface-cfo-says/636627/>, accessed 12th April 2023.
- Standfield, L., T. Comans, and P. Scuffham. 2014. "Markov Modeling and Discrete Event Simulation in Health Care: A Systematic Comparison". *International Journal of Technology Assessment in Health Care* 30(2):165–172.
- Temple, J. M. F., and D. Fone. 2002. "Implementing National Guidance on Prevention of Coronary Heart Disease: Clinical Governance and Computer Simulation Modelling". *British Journal of Clinical Governance* 7.1:27–33.
- Walsh, K. 2006. "Reverse Disaggregation - How Silicon IP Will Change the Semiconductor Supply Chain". <https://www.design-reuse.com/articles/13320/reverse-disaggregation-how-silicon-ip-will-change-the-semiconductor-supply-chain.htm>, accessed 12th April 2023.

AUTHOR BIOGRAPHIES

SUSAN S.M. HANSON is an Operations Research Analyst at The MITRE Corporation (MITRE). She holds a M.E. and B.S. in Applied Operations Research (OR) from Cornell University and has worked on numerous OR projects across the Army, Air Force, Joint Staff, DHS, and FAA, and has co-authored a Macroeconomic Model of Federal and Commercial Spectrum Usage. Her email address is shanson@mitre.org.

NOAH MECIKALSKI is an Operations Research Analyst at MITRE. He holds a B.S. in Industrial and Systems Engineering from Auburn University. He has applied OR techniques to areas such as critical infrastructure analysis, aviation scheduling, and failure rate modeling. His email address is nmecikalski@mitre.org.

ALEX TOBIAS is a Synthetic Biology Principal Researcher at MITRE. He holds PhD and B.Eng. degrees in Chemical Engineering from the California Institute of Technology and McGill University, respectively. He has worked as a Chemical and Biological Research Engineer in industry and for the US Army. His email address is atobias@mitre.org.

JACK MORRIS is an Operations Research Analyst at MITRE. He holds a B.S. in Applied Mathematics from William & Mary and completed an S.M. in MIT's Technology & Policy Program in 2022. While at MITRE, he has applied OR techniques within a range of areas including energy system resilience, transportation and logistics, and synthetic biology. His email address is morrisj@mitre.org.

NEAL WAGNER is an Operations Research Analyst at MITRE. He has published ~ 50 peer-reviewed papers in areas of modeling and simulation, optimization, and time series prediction, holds a patent on mathematical modeling for cybersecurity, and has worked on numerous research projects for multiple government agencies. He has also worked in industry and as a university faculty. He holds PhD, MS, and BA degrees from University of North Carolina. His email address is nwagner@mitre.org.

REBECCA S. WIDRICK is an Operations Research Analyst at MITRE and US Air Force Veteran. She holds a B.S in Mathematics from Rensselaer Polytechnic Institute and an M.S. in OR from the Air Force Institute of Technology. She specializes in stochastic modeling and applied statistics. In her ten years as an OR Analyst, she has worked across several industries including the Department of Defense, healthcare, and transportation. Her email address is rwidrick@mitre.org.

DAMON BAYER was a summer Operations Research Analyst at MITRE. He is pursuing a PhD in Statistics at University of California, Irvine specializing in Bayesian methods, and holds a M.S. and B.S in Math and Statistics from South Dakota University. His email address is damon@damonbayer.com.