Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

USE OF A COMBINED DISCRETE RATE AND POPULATION BALANCE SIMULATION FOR THE DESIGN AND OPTIMIZATION OF A HIGH SHEAR AGGLOMERATION PROCESS

David Krahl
David Milburn-Pyle

Kromite LLC
243 N Union Street, Suite 117
Lambertville, NJ 08530, USA

ABSTRACT

The application of a hybrid discrete rate and population balance simulation model to a high shear agglomeration process provides in-depth and dynamic insight into how a non-homogeneous product behaves across the entire process in addition to the individual process component level. By tracking individual product attributes throughout the process, the customer can size equipment, design control systems, and optimize process parameters, while reducing engineering development time and cost.

1 INTRODUCTION

The manufacture of granular or particulate products is difficult to understand and predict due to the complex nature of solid materials compared to that of liquids or gases. Particulate processes have risks of segregation, particle breakage, and complex flow behavior. The residence times of granules may vary as a function of their size. Additionally, agglomerated products may vary in composition based on particle size with larger granules having different levels of active ingredients than smaller granules. (Feeco 2019) All these issues make ensuring product quality and consistency difficult when developing new formulas or products. Process modeling can be used to understand the impacts of these risks, ensure product quality, and improve process operation.

Computational modeling is often used to predict process behavior and improve process operation. Many software suites have been developed to model industrial processes for a wide range of industries. However, predicting the output of a particulate process may require more sophisticated models than those of fluid processes. Population balance modeling has been used successfully in many industries to develop models of particulate processes including crystallization, granulations, and fluidized beds. These models typically examine granules and their properties as a function of particle size. Many models focus on the mechanisms of an individual process to predict the particle size distribution and material properties as a function of particle size. These models are valuable to understand and predict how material composition and yield will be affected by material properties and equipment settings.

When modeling a full particulate process however, it is sometimes useful to simplify fundamental mechanisms with assumptions and model the overall material flow to understand a complete process. In these cases the particle size distribution of the population balance method is still necessary to understand how particle size impacts material properties and recycle streams. Using the discrete rate method allows for the accurate modeling of residence times to produce a full description of the process over time.

In this work, a full manufacturing process of product created in a high shear agglomerator is modeled to understand the evolution of product attributes over time. The agglomeration mechanisms are not modeled specifically and most unit operations are treated as “black boxes” that perform transformations on particle sizes and compositions based on experimental data and assumptions. The material streams leaving each unit operation are tracked through the process between the individual pieces of equipment. By tracking the

978-1-7281-3283-9/19/$31.00 ©2019 IEEE

material streams of the process, a manufacturing plant can be described in enough accuracy to assist with the development of new formulas. With accurate modeling of residence times and material flow rates, the process can be evaluated for improvements to plant scheduling and operation.

The objective of this work is to produce a working model of the manufacturing facility with sufficient accuracy to simulate the production of a variety of products and process conditions. The completed model is used to evaluate the time required to transition between formulas and the impact of process variability on product attributes. Additionally, the model is used to improve process operations by testing various operator interventions in response to process upset conditions or demand scenarios.

A literature search yielded few previous papers on similar topics. This type of process and its control system have been studied (Mort et al. 2000) and discrete-rate simulation models have existed for two decades (Siprelle 1997), but attributes that can be assigned to the flow are relatively new (Imagine That 2019). While it is possible to add properties to flow without direct software support, it would be difficult to achieve in the context of the requirements of this model. Most population balance models of similar industrial processes focus on unit operations and not the whole process (Vesjolaja et al. 2018; Bárkányi et al. 2013), hence the limited references to related work.

Please note that this model was developed for, and with the direct support of, a global consumer goods manufacturer. Due to the confidential and competitive nature of the production process, no details of the customer or their representatives can be provided. Furthermore, any data included in this paper are strictly representative and for illustrative purposes only, and certain labels on screenshots have been redacted.

2 DESCRIPTION OF THE AGGLOMERATION PROCESS

Agglomeration is the process whereby fine particles of various materials are combined into larger pellets or granules. In its basic form, the process of high shear agglomeration starts with raw materials in a powder form, which are then combined with a binding agent (typically water) in one or more mixers such as pin or paddle mixers. The resultant granules are then dried in a rotary or fluid bed dryer, after which they are screened for size. The absence of a pelletizer in the agglomeration process means that granules leaving the mixer(s) will have a variance in size, mass, and composition. In the screening process, granules that are over-sized are sent to a mill to be broken down to either the correct size or smaller. Agglomerates that are under-sized - either coming from the mixer(s) or the mill - are recycled to the start of the process and reintroduced into the mixer(s) along with the raw materials. Correctly-sized or in specification granules are stored in silos for further processing and packaging. This process is shown below in Figure 1.

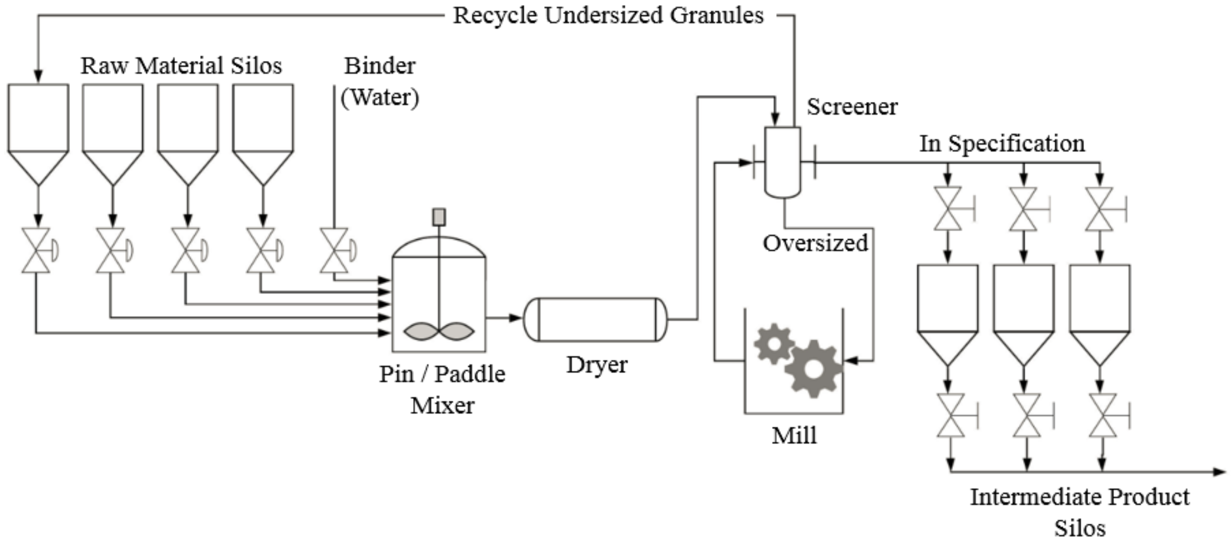

Figure 1: Agglomeration process schematic.

3 MODEL REQUIREMENTS

The client had a need to model an end-to-end agglomeration and batching process. Their specific requirements for the model were:

- Record the material composition (mass-based % of each raw material and binding agent) throughout the process in fixed time increments.
- Track product flow rates in all streams (raw materials, binding agent, mixer(s), dryer, screener, mill, recycle, batch silos) throughout the process at discrete time increments.
- Track material provenance, i.e. to be able to identify the granules from each mixer throughout the downstream process. Similarly, if over-sized granules were milled, the resultant granules needed to carry the attribute that they had been processed in the mill.
- Vary granule product composition (“recipe”) by changing the raw material flow rates into the mixer(s).
- Introduction of perturbation into the process, such as mixer performance (e.g. increased quantity of over- or under-sized granules than the equipment specification), changes in raw material quality (varying moisture content), and raw material input (variation due to flow control).

The client’s objectives for the model were:

- Verify initial process parameters and testing.
- Determine equipment sizing, especially recycle and intermediate-product silos.
- Design the process control logic, with focus on binding agent flow control, and intermediate and batch silo switching.
- Analysis of process fluctuations, and how the impact – defined as product “smearing” - is carried through the process given the product residence times within equipment or on intermediate conveyors. Specific focus was on process start-up, changes to raw material make-up, and equipment performance.

4 MODEL SOLUTIONS

The agglomeration process was modeled as a hybrid discrete rate/population balance simulation using ExtendSim v9.2. The choice of software / modeling tool was largely driven by the client, as ExtendSim was already part of their installed software suite. This software was well suited to the problem as it had native support for discrete-rate attributes that were used to track the constituents of the agglomerated material. In addition, the built-in database was useful for storing the material flow logs.

Discrete rate was chosen due to the bulk flow of material through the system and the need to process events at specific times such as product changeovers, equipment status changes, and control logic triggers (Krahl 2009). A continuous layer was included to simulate the continuous variability of material in the process.

In this case, attributes of flow and material composition were modeled by events occurring within the process – events such as a change in flowrate, or a change in a material % within the product composition. To achieve the levels of precision required within a reasonable simulation time, a discrete rate model was the best simulation technology for this application.

5 MODEL CONSTRUCTION

At first glance this appeared to be a straightforward system to model. However, the following challenges required a very detailed model to capture all of the aspects of the system behavior. In particular, detailed control logic was necessary to produce the required material quality.

5.1 Process variability

There are a number of sources of variability in the process. These pose challenges in producing a homogeneous product. Because this variability is continuous, it was necessary to sample random distributions at regular intervals. The following contribute to process variation:

- Moisture content: The moisture content of the raw material varies over time (e.g. bulk deliveries into the raw material silos), which impacts the flow rates of the raw materials and binder into the mixer(s).
- Raw material rate variability: The rate at which raw material arrives varies over time due to tolerances in conveyor-based flow control.
- Mixer performance varies producing different particle sizes: Due to wear, operating speed, and other factors, the mixer may produce different distributions of material size and composition over time.
- Composition of recycle stream: The makeup of the material leaving the mixer(s) varies by particle size. Larger granules will tend to have more of one ingredient than smaller granules. Because the small and large granules are unacceptable in the final product, they are recycled and either reduced in size (large granules) or re-agglomerated (small granules). This recycling with different component proportions affects the composition of the product as it moves through the system.

5.2 Mixer

The mixer(s) were each modeled as a series of 3 Merge blocks that combined the flow in proportions dynamically calculated by the simulated control system to account for the raw material "recipe" composition, the recycle of under-sized granules, and the binder flow (calculated off the moisture condition of the raw materials). Functionality was included to allow for a degradation of mixer performance (mass flow throughput), and also to vary the amount of recycled material (based on the recycle silo level).The mixer was naturally modeled.

A fundamental outcome of this type of agglomeration process is that the granules leaving the mixer(s) vary in size, composition, and mass, as can be seen in Figure 2. While this was not an issue with modeling the mixer(s), it has a significant impact on the downstream process.

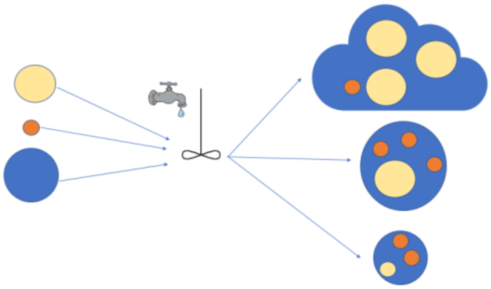

Figure 2: Illustration of agglomerate variability.

5.3 Dryer

As mentioned above, the nature of the high shear agglomeration process is that there is a variance in size, mass, and composition of the granules leaving the mixer(s). When introduced to a constant-airflow rotary or fluid bed dryer, these variations result in non-homogeneous material flow rates and some granules travel faster or slower than others through the dryer. This poses a challenge when modeling the residence time of material in the dryer and becomes significant when the process is in transition from one state to another.

To model this behavior, it was necessary to introduce a virtual screener between the mixer(s) and the dryer. While the physical screener is situated downstream of the dryer, in the simulation environment it was introduced at the exit from the mixer(s). Laboratory testing had determined a distribution of granule sizes exiting the mixer(s), which could be grouped into seven streams of increasing size, each with a different residence time in the dryer as shown in Figure 3. The dryer could then be modeled as having seven parallel streams, each with a Convey Flow block at operating at a different speed, representing the different speeds at which the larger or smaller granules were carried through the dryer. Each stream also simulated the drying process, as water was removed from the granules to a user set-point.

By modeling the behavior of each granule stream individually, it was possible to track the composition of the product leaving the dryer, and to identify the impact of process changes (product "smearing"). This was important in tracking the makeup of the final product, particularly during start-up or after a product change-over occurred.

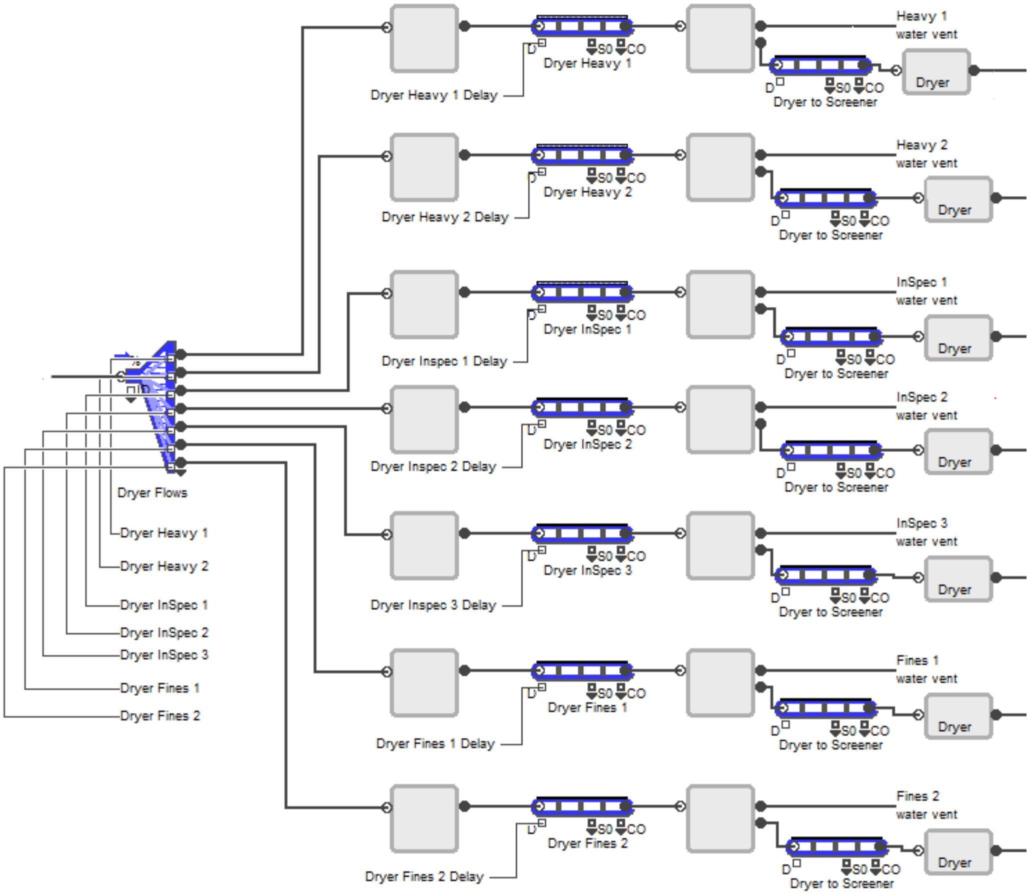

Figure 3: Dryer model.

5.4 Screener

The physical screener was modeled by mathematically recombining the flow of material based on the proportion of the various particle sizes as they exited the dryer and the mill. Of the seven particle sizes tracked, there were three different paths based on the particle diameter. Two particle sizes were classified as large, three as medium, and two as small. Only the medium-sized granules were useful for final production. The small granules are mixed back with the raw material and re-agglomerated and the large granules are processed through a mill to reduce their size.

5.5 Mill

The function of the mill is to reduce the size of over-sized granules to in-specification or to under-sized granules. The agglomerates leaving the mill return to the screener, where, along with product from the dryer, the correctly-sized granules pass to the intermediate product silos, and the under-sized granules are recycled. A key modeling criteria was that the over-sized granules are not broken down into smaller granules of the same composition as the original granule. The granules differ by size, mass, and composition, and hence will have different residence times within the mill. Like the dryer, the mill had to be modeled as having five individual product streams, with individual Convey Flow blocks. Note that the largest granules or "heavies" do not leave the mill. Holding tanks were added as a model validation feature. As the only size granules that were leaving the tank were "fines" and "inspecs", if these heavy tanks were not empty at the end of the simulation, there was a problem with the calculation. This can be seen in Figure 4 below.

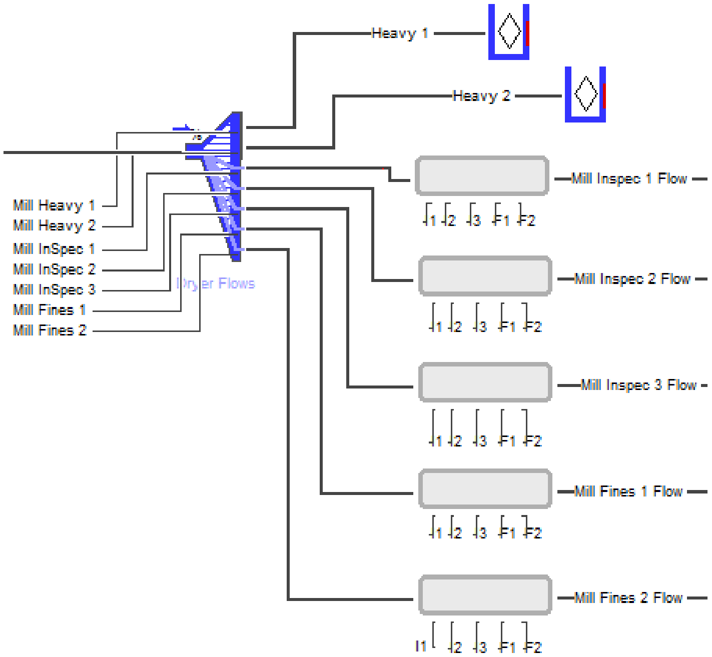

Figure 4: Mill model.

5.6 Control System

One of the model purposes was to assist in the design of the control system for the agglomeration process. There were four areas where the control system was addressed. These all contributed to the production of the proper quality of the final product.

5.6.1 Binder injection into the mixer(s)

Based on the moisture content of the raw materials and a moisture setpoint for agglomerate leaving the mixer(s), logic was included in the model to control the input flow rate of the binder into each mixer. This logic not only ensured the correct product mass flow downstream of the mixer(s), but also quantified the volume flow of binder required across varying process parameters. The logic was tested by adding a normal distribution of variability to the raw material moisture content and then monitoring the mixer exit conditions.

5.6.2 Moisture removal in the dryer

Just as the binder is added to the process in the mixer(s), water is removed in the dryer. Again, logic was added to determine the moisture to be removed in the dryer for each of the 7 product streams, such that the moisture levels matched a set point. This calculation provided insight to the volume of water to be removed – required for specifying the dryer size and performance – as well as setting the correct attribute values for the product downstream of the dryer.

5.6.3 Recycle silo level control

Depending on the performance of the mixer(s), the quantity of undersized granules being recycled will vary. If the distribution of granule sizes leaving the mixer(s) has more under- or oversized granules than design, the recycle flow rate will be greater than design, and this is managed through control of level in the recycle silo. Logic in the model controls the recycle silo level against a setpoint target level, by varying the amount of recycled granules that are reintroduced to the mixer(s).

5.6.4 Batch management

The final product from the production process was a blend of granules of two different material compositions (two "recipes"). This meant that the simulation needed the ability to switch between two preset raw material % inputs to the mixer(s). The corresponding downstream impact was the need for an intermediate product silo into which granules would be filled during the "recipe" change-over period. The model included time-based batching logic to account for changes in "recipe". The times for the logic were set by running the simulation and determining how long product had to be filled into the intermediate silo until the new product recipe had been carried through the entire process.

5.7 Attributes

The agglomeration process and the customer requirements resulted in a significant number of attributes to track. Assuming the process had three raw materials and a binding agent, the simulation has to then track the flow rate and % composition of the four materials. However, on exiting the mixer, the model diverges into seven streams of different size and composition. Furthermore, the model diverges again as over-size streams are milled to size. If a second mixer is allowed, the model is then tracking 28 combinations of the flow rate and four properties or 140 attributes and calculating the total streams for another 25 records. Table 1 shows the structure of the information stored per flow. These flows are tracked in multiple places in the model. Note that the values in this table are for illustrative purposes, the actual data is proprietary.

Table 1: Database representation of flow properties.

#	Granule Size	Mixer Number	Output from a Mill	Mass Flow Rate	Material 1 %	Material 2 %	Material 3 %	Water %
1	Total Stream	1	No	73.0	11.0	71.7	2.4	15.0
2	Size 1	1	No	2.3	5.2	77.7	2.1	15.0
3	Size 2	1	No	3.8	5.3	77.6	2.1	15.0
4	Size 3	1	No	25.1	10.8	72.1	2.1	15.0
5	Size 4	1	No	25.8	10.6	71.9	2.5	15.0
6	Size 5	1	No	5.0	10.3	72.3	2.4	15.0
7	Size 6	1	No	5.5	15.7	66.6	2.7	15.0

8	Size 7	1	No	5.5	15.7	66.6	2.7	15.0
9	Total Stream	2	No	73.0	11.0	71.7	2.4	15.0
10	Size 1	2	No	2.3	5.2	77.7	2.1	15.0
:	:	:	:	:	:	:	:	:
30	Size 5	2	Yes	1.3	10.3	72.3	2.4	15.0
31	Size 6	2	Yes	1.4	15.7	66.6	2.7	15.0
32	Size 7	2	Yes	1.4	15.7	66.6	2.7	15.0
33	Total of all Streams	-	-	183	11.0	71.7	2.4	15.0

A combination of database tables and flow attributes were used to track the proportion of each component as well as the previous path that the material had passed through. A total of 140 attribute values needed to be tracked for the material flow as it passed through the system. The flow attributes were used where there was a time delay for the material (conveyors) and the existing attribute structure could be utilized. In other places, a database table was used as it led to more efficient model performance.

A custom block was created to calculate the component proportions. This provided an easy-to-use interface to select the appropriate function and tables for a specific point in the model. Figure 5 shows the user interface that was developed to set, update, and append the properties of the flows. Database tables were used to store the proportions of the different materials. Once the block had been programmed, making the calculations at different points in the model required selecting database tables and the appropriate options for that point in the model. This reduced model building and verification time.

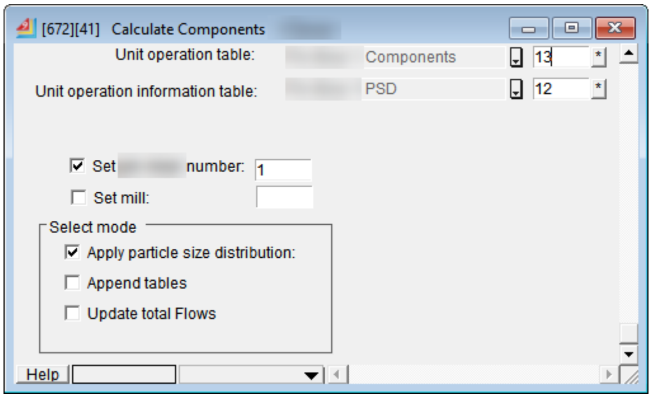
[672][41] Calculate Components [redacted]

Unit operation table: [redacted] Components 13

Unit operation information table: [redacted] PSD 12

☑ Set [redacted] number: 1

☐ Set mill:

Select mode

☑ Apply particle size distribution:

☐ Append tables

☐ Update total Flows

,

Figure 5: Block to calculate component properties.

A requirement was to log the status of the system at set intervals. This was achieved by adding equations, scheduled to calculate at every interval, that write the current flow status to a database table. A hierarchical block was created that encapsulated these equations and parameterized them such that only the particle size and mill status were entered to record the flows. This hierarchical block was placed at various locations in the model to record the flow status. Utilization of a single, parameterized hierarchical block in multiple locations improved the model construction time, reliability, and maintenance. Figure 6 shows the structure of the hierarchical block. Note the popup menus for selecting the granule size and mill status. In the final version of the model, five minutes was used as the calculation and recording interval.

This interval was deemed by the client to have the appropriate level of detail based on the performance of the actual process.

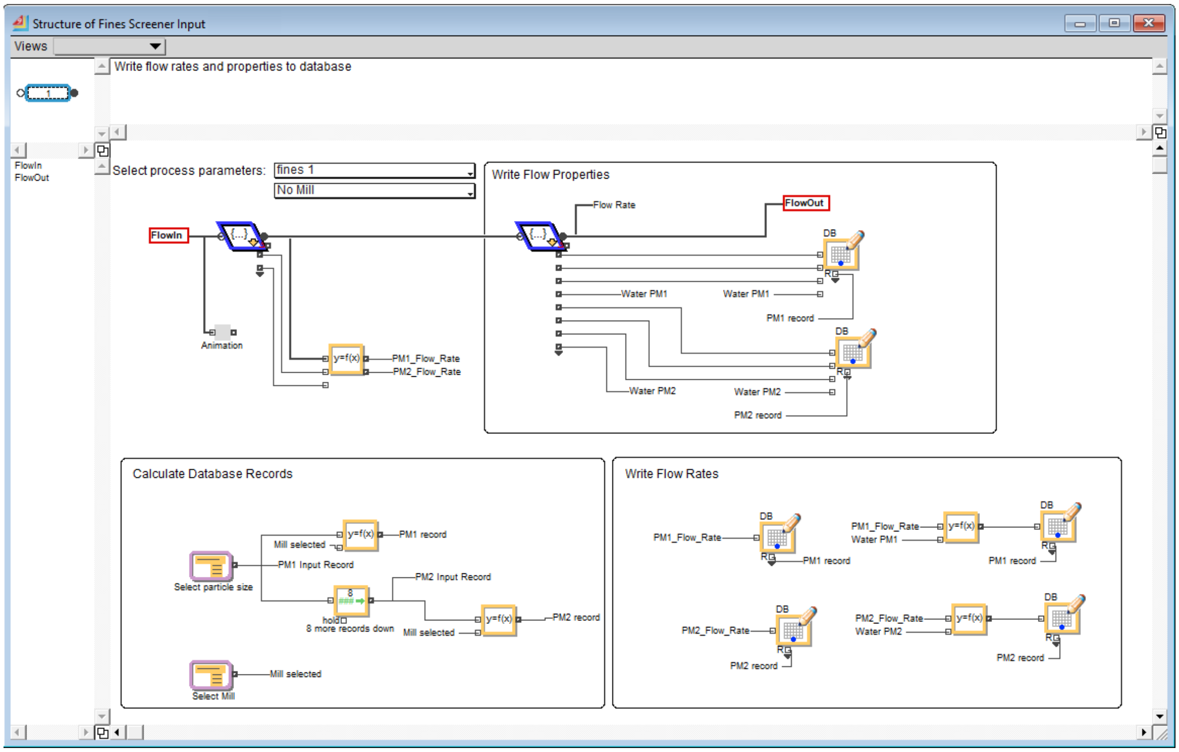

Figure 6: Hierarchical block writing material properties.

5.8 Model Accuracy and Performance

The client had previously developed a proof of concept model in Excel, and this, along with experimental data, was used to validate the ExtendSim results across the model given specific input and operating conditions. The ExtendSim model could then be run to analyze the process under more rigorous conditions, process upsets, and longer cycle times.

High sampling frequency and small changes in composition of material led to large number of layers of material in the tanks and Convey Flow blocks that needed to be tracked by the software. Performance was improved significantly by storing the component information in database tables and only using the built-in attributes when absolutely necessary. Additional performance improvements were gained by rounding the component proportions to fewer significant digits (5) so that fewer layers were necessary to track the material. The rounding level did not significantly affect the quality of the results.

When running the model for the “full” process – from raw material tanks to finished batched product, and with empty intermediate tanks and conveyors – the model takes ~15 minutes to simulate a 180 minute production cycle. However, when testing specific operating conditions and modes as specified in the customer’s design protocol, it was possible to reduce this time by only running the relevant processes within the model.

6 CONCLUSIONS AND RECOMMENDATIONS

The process model is capable of tracking all the material attributes throughout the process flow. Material streams from different mixers combine together and create the agglomerated product with the properties expected based on experimental work. The model is able to produce all necessary recipes with different raw material combinations. Variability in the process is investigated by examining the limiting steady state cases based on mixer yield and recipe range to determine which operational areas will lead to a process failure.

6.1 Dryer Residence Time

The effective modeling of the dryer residence time distribution is clearly demonstrated in Figure 7. The figure depicts the cumulative residence time of dryer based on the dryer inlet material feed. The smallest material exits the dryer first at 4.5 minutes. Each additional size bin then exits the dryer after the next specified interval of 3 minutes each. After 22.5 minutes all material is exiting the dryer and the total flow rate has achieved its final value. The series does not converge to 1 as a final value because the dryer removes moisture from the product. The difference between the final amount of material exiting the dryer versus the dryer input depends on the water added in the mixer recipe.

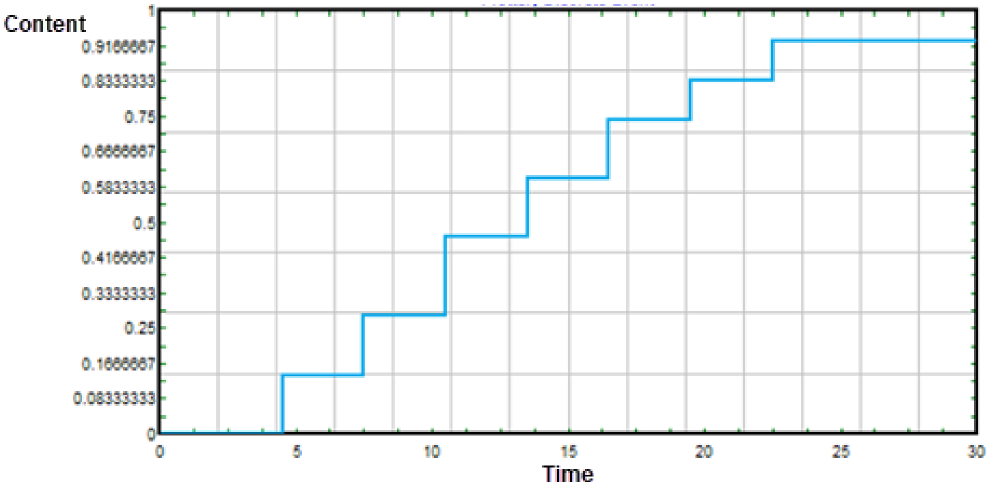

Figure 7: Dryer residence time cumulative distribution based on dryer material feed.

Validation of the dryer residence time is based off of experimental observations of different particle sizes. With a fixed dryer design, the residence time distribution cannot be changed without changing materials or particle size. The difference in particle size will impact recipe switches as materials from a set point change will smear in the dryer. Small particles from the new set point mix with large particles of the previous set point. This effect is innate to the process and must be accounted for in any plant operation scenario.

6.2 Recipe Transition

Transitioning from one mixer recipe to another will introduce a step change that must propagate through the process. The process will require time to converge the product stream when switching between recipes with different levels of active ingredients. Figure 8 describes the level of one active ingredient in product stream exiting the screener headed to material storage silos. The first recipe with high active ingredient at a level of 1 is run from time 0 to 300 minutes. At 300 minutes, the mixer is switched to a recipe with no active ingredient. The new recipe exits the screener several minutes later causing a rapid drop in the level of active. The active level then continues to decrease as the active material that was built up in the recycle loop is purged from the process.

The total transition time is the time between when the product no longer has enough material to be acceptable as the initial recipe and the time when the process has not purged out enough active in the recycle to be considered a zero active ingredient product. The transition time is dependent on the amount of recycle in the process, the magnitude of the set point change, and the particle size distribution of the process.

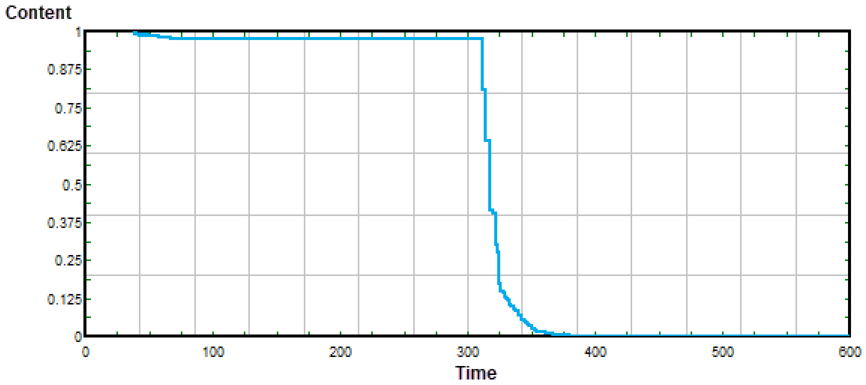

Figure 8: Active ingredient level change due to a recipe transition at 300 minutes.

Transition time is primarily impacted by the level of recycle in the fines tank. This is because the more recycle in the process, the more residual active ingredient that must be purged from the system. This effect is seen clearly in Figure 9. The amount of recycle is described in the minutes required to purge the process of all current recycle. The figures shows that the active percent in the product material stream decreases in several steps. The first step occurs just after 300 minutes. The amount of time spent at this active level is dependent on the recycle amount. During this time new recycle material is created with a lower amount of active ingredient which is fed into the process. The active level then decreases twice more before converging to a value close to zero.

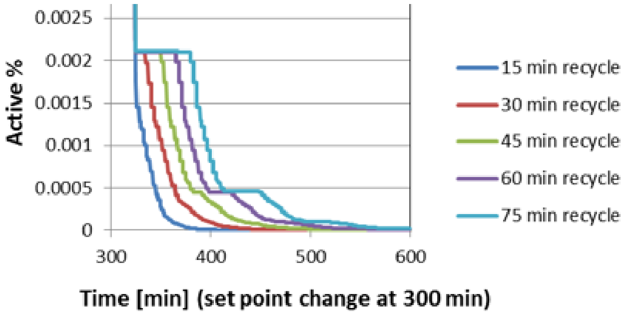

Figure 9: Difference in transition times based on the amount of recycle in the process.

The longer the transition time of the process, the more out of specification intermediate material is generated that must be reprocessed or discarded. Minimizing the transition time is important to reduce these operational losses. Tuning the recycle control loop on the recycle level is necessary to reduce the level in the silo as low as possible before a process transition occurs. Process transitions are required in order to provide a sufficient amount of product from two recipes to different packaging lines.

6.3 Final Comments and Recommendations

The model results allowed the client to understand their process in greater depth. The model effectively simulated process transitions and flow and compared well to experimental data. Although no individual unit operations were modeled mechanistically, using experimentally validated assumptions is sufficient to build a process model that can be evaluated for operational improvements. The client can continue to use the model to test new formulas and try different process control strategies to further minimize the transition time and waste product. All of this can be completed without spending resources on expensive experimental studies.

The population balance model approach was necessary in order to accurately model the process. The use of multiple particle size bins produced an accurate description of the dryer residence times. Additionally, the particle sizes define how the screener sorts streams to the mill, recycle, and finished product. Assigning material attributes to the process for each size bin enabled the tracking of active ingredients throughout the process. The level of active ingredient is instrumental in understanding the transition time of the process when alternating between recipes. The mill attribute also allows for a better description of the transition time by identifying particles in the product stream that are still high in active ingredient. Future work can investigate the case where mixers produce different recipes or have different yields.

In conclusion, this modeling approach provides an excellent understanding of a complete particulate process. The population balance approach was successfully implemented in the ExtendSim discrete rate software combining the power of a discrete event simulation with population balance modeling. This modeling approach is appropriate to model any full process that requires the tracking of many attributes moving through a process in time. The model is sufficiently robust to experiment with different operational modes of the manufacturing facility without conducting expensive trials. Although the fundamental mechanisms of individual unit operations are not modeled, an accurate description of the full process can be developed. This approach is recommended for any case in which process scheduling and optimization is more important that mechanistic understanding.

REFERENCES

FEECO. 2019. The Agglomeration Handbook. FEECO International, Inc. http://go.feeco.com/acton/media/12345/the-agglomeration-handbook, accessed April 8, 2019.

Imagine That. 2010. Evolution of ExtendSim. Imagine That, Inc. https://www.extendsim.com/products/advantages/history, accessed April 8, 2019.

Krahl, D. 2009. “ExtendSim Advanced Technology: Discrete-Rate Simulation”. In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johnson, A. Dunkin, and R. G. Ingalls, 333-338. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Mort, P. R, S. W. Capeci, and J. W. Holder. 2001. “Control of Agglomerate Attributes in a Continuous Binder-agglomeration Process”. *Powder Technology* 117(1–2):173-176.

Siprelle, A. J. and Phelps, R. A. 1997. “Simulation of Bulk Flow and High Speed Operations.” In *Proceedings of the 1997 Winter Simulation Conference*, edited by S. Andradottir, K. J. Healy, D. H. Withers, and B. L. Nelson, 706-710. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Vesjolaja, L., B. Glemmestad, and B. Lie. 2018. “Population Balance Modelling for Fertilizer Granulation Process.” In *Proceedings of The 59th Conference on Simulation and Modelling (SIMS 59)*, 95-102, Oslo Metropolitan University, Norway.

Bárkányi, A., S. Németh, and B. G. Lakatos. 2013. “Modelling and Simulation of a Batch Poly(Vinyl Chloride) Reactor”. *Chemical Engineering Transaction* 32:769-774.

AUTHOR BIOGRAPHIES

DAVID KRAHL is the CEO of Kromite and has over 30 years of experience in the simulation software industry. Specializing in discrete event simulation, he has worked with a wide range of industries including automotive, public sector, petro-chemical, defense, healthcare, manufacturing, and mining. Prior to joining Kromite, Dave was the Technology Evangelist at Imagine That, Inc. In this position, he promoted ExtendSim and simulation technology to a wide variety of industries. In past positions, Dave has performed development, consulting, technical sales, support, and training for a range of simulation software programs. Dave is one of the primary architects of ExtendSim’s libraries of model components. His particular interests include tools and techniques for troubleshooting simulation models, design of experiments, and the architecture of message-based discrete event simulation. Dave has a Master’s degree from Golden Gate University and a Bachelor’s degree from the Rochester Institute of Technology. He is also a Certified Analytics Professional (CAP). His email address is dkrahl@kromite.com. His website is http://www.kromite.com.

DAVID MILBURN-PYLE is a consultant at Kromite LLC. He holds an MBA (Finance and Management) from the Stern School of Business at New York University, and a BSc (Mechanical Engineering) from the University of Cape Town, South Africa. His focus is on process and supply chain modeling and optimization, data analytics, and decision support. His email address is dmilburn-pyle@kromite.com. His website is http://www.kromite.com.