Choosing an Operating Policy for Seeded Batch Crystallization

Jeffrey D. Ward, Duncan A. Mellichamp, and Michael F. Doherty
Dept. of Chemical Engineering, University of California, Santa Barbara, CA 93106

DOI 10.1002/aic.10808

Determination of the optimal temperature (or supersaturation) trajectory for a seeded batch crystallizer is a well studied problem. The majority of researchers find that the optimal trajectory has a greater supersaturation at the end of the batch. A significant minority of researchers, however, using very similar models and different but equally justifiable objective functions, find that the optimal trajectory has the greatest supersaturation at the beginning *of the batch. Still other researchers have suggested that crystallization processes should be operated with* constant *supersaturation. These apparent conflicts are resolved into different equivalence classes of objectives and corresponding operating policies. It is shown that the equivalence classes arise as a result of an inherent tradeoff for seeded batch crystallizers involving secondary nucleation: Early growth operating policies tend to minimize the* number *of nuclei formed because supersaturation is greatest when the mass of crystals is least. However, those nuclei which are formed early have an opportunity to grow to a larger size and compete with the growth of the seed crystals. By contrast, late growth operating policies tend to produce more nuclei because supersaturation is greatest when the mass of crystals is greatest. However, nuclei formed late in the batch do not have a chance to grow as large, and the total mass of nucleus-grown crystals is small. Which of these outcomes is preferred depends on the particular application.* © 2006 American Institute of Chemical Engineers *AIChE J,* 52: 2046–2054, 2006
Keywords: crystallization, process control, process synthesis

Introduction

Crystallization is a very common unit operation in the chemical and pharmaceutical industries for separation and purification, and for the production of solid phase products. The driving force behind crystallization is the supersaturation of the solute in solution. Crystallization is often conducted in batches rather than continuously because this allows for greater flexibility and shorter process development times. A critical question in the design and operation of batch crystallization processes is how the supersaturation should vary with time during the batch.

Relative supersaturation (also referred to simply as "supersaturation" in this article) is defined as

$$S = (C - C_{\mathrm{sat}})/C_{\mathrm{sat}} \tag{1}$$

where C is the actual concentration of solute in the solution, and C_{sat} is the saturation concentration. In a cooling crystallizer, supersaturation is achieved by cooling the solution over time and C_{sat} is a function only of temperature for a particular solute-solvent pair. Determination of the optimal temperature (or supersaturation) trajectory for a seeded batch crystallizer is one of the most extensively studied path optimization problems in chemical engineering (after batch reaction and batch distillation). Beginning with the pioneering work of Jones,[1] more than a dozen researchers have used a variety of techniques (including sequential quadratic programming, optimal control

---

Correspondence concerning this article should be addressed to M. F. Doherty at mfd@engineering.ucsb.edu.

theory, and stochastic methods) to determine the optimal trajectory. A summary of the literature is presented in the section on choosing an objective function.

The majority of researchers have found that the optimal temperature trajectory is one which leads to an increase in supersaturation near the end of the batch (here termed "late growth"), and this has been suggested as a generally valid heuristic.[2] However, a significant minority of researchers have found a result which in some sense is almost exactly opposite: that the optimal policy has an increased supersaturation at the beginning of the batch (here termed "early growth"). Still other researchers have indicated that, at least to a first approximation, the optimal way to operate a seeded batch crystallizer is to maintain constant supersaturation.[3] In spite of the large number of publications in this field, these discrepancies apparently remain unresolved.

The format of this article is as follows: First the development of a mathematical model for seeded batch crystallization is reviewed. Then, the subtleties associated with the selection of an objective function are discussed, and objectives used by previous authors are summarized. A theory is presented to account for discrepancies between the results found by different researchers based on differences between objective functions used by them. Finally, the theory is tested by optimizing a case study process using six different objective functions. It is shown that the theory presented in this article correctly predicts the qualitative shape of the optimal trajectory in every case.

Formulation of the General Optimization Problem via the Method of Moments

This section summarizes the development of a mathematical model which has been explored by many authors in order to determine the optimal temperature (or supersaturation) trajectory for a seeded batch crystallizer. Similar developments have been presented elsewhere, for example, Randolph and Larson[4] and Chung et al.[5]

Population balance equations

A general statement of the population balance for a batch crystallizer with one characteristic growth dimension is

$$\frac{\partial f}{\partial t} + \frac{\partial (fG)}{\partial L} = B \tag{2}$$

where f is the population density function for the entire crystallizer, G is the linear growth rate of the crystals (m/s), and B is the rate of nucleation (#/kg s).

As written, Eq. 2 is difficult to solve even numerically. There are a number of different ways that the population balance can be reformulated into a series of equations more amenable to solution. An excellent review is given by Ramkrishna.[6] In this work, the method of moments is employed, not because it is the most accurate, robust or versatile, but merely because it is the method which has been by far the most popular with researchers historically. One of the primary goals of this work is to interpret the results of previous researchers.

Equation 2 can be converted to a series of ordinary differential equations by the well-known method of moments,[7] where the ith moment of the crystal-size distribution is defined as

$$\mu_i = \int_0^\infty L^i f(L)\,dL \tag{3}$$

where μ_i has units of m^i/kg.

If the nuclei are assumed to have essentially zero size when formed, the result is

$$\frac{d\mu_{T_0}}{dt} = B \tag{4}$$

$$\frac{d\mu_{T_i}}{dt} = i\mu_{T_{i-1}}G \tag{5}$$

where μ_{T_i} is the ith moment of the total crystal population (including crystals grown from seeds, as well as crystals grown from nuclei which form during the course of the batch). Some authors consider objectives which involve only properties of the seed-grown crystals or only properties of the nucleus-grown crystals. Equations 4 and 5 can be written exclusively for the seed-grown crystals

$$\frac{d\mu_{S_0}}{dt} = 0 \tag{6}$$

$$\frac{d\mu_{S_i}}{dt} = i\mu_{S_{i-1}}G \tag{7}$$

Finally, the moments of the nucleus grown crystals at any point in time can be determined by subtracting the moments of the seed-grown crystals from the corresponding moments of the total crystal population

$$\mu_N = \mu_T - \mu_S \tag{8}$$

Kinetics of birth and growth

For size-independent growth, growth is typically assumed to be proportional to the supersaturation raised to some power

$$G = k_g S^g \tag{9}$$

where k_g and g are empirically determined parameters.

It is well known that for practical industrial crystallization systems, the supersaturation must be kept low enough that primary nucleation is effectively suppressed, so that new particles are formed almost exclusively by secondary nucleation. The most often-used expression for secondary nucleation is

$$B = k_b \mu_{T_3} S^b \tag{10}$$

where k_b and b are empirically determined parameters. The rate of formation of new particles depends on both the supersaturation and mass (volume) of solid crystals in the crystallizer.

There are several ways of manipulating the supersaturation during the course of a batch, including decreasing the temperature, evaporating solvent, and adding antisolvent. The analysis presented in this article is applicable regardless of the method used to manipulate the supersaturation. However, the most common method in both industry and the literature is by cooling. Therefore, results will be illustrated for this case.

The (relative) supersaturation was defined in Eq. 1. The model is complete when the concentration of solute (C) is determined by material balance

$$\frac{dC}{dt} = -3\rho_c G \mu_{T_2} \tag{11}$$

Boundary conditions and objective functions

The initial condition on the moments of the crystal size distribution is given by the properties of the seeds. In order for the problem to be well posed and industrially relevant, it must be possible to specify a production rate constraint for the process. For a seeded batch crystallizer, this corresponds to specifying a time for the batch and a final value of the total mass of crystals produced. These conditions are expressed by

$$t = t_f \quad \mu_{T_3} = \mu_{T_{3,f}} \tag{12}$$

Finally, it is necessary to choose an objective function J. Although a small number of researchers have considered path objectives, in this work attention is restricted to objective functions which depend only on the properties of the crystal size distribution at the end of the batch:

$$J = J(\mu_{T,f}, \mu_{S,f}, \mu_{N,f}) \tag{13}$$

Thus, the optimization problem is

$$\min_{T(t)} J \tag{14}$$

subject to the mathematical model given by Eqs. 4–12, and possibly additional constraints which may enter into the formulation depending on the particular system under consideration. The remainder of this article explores how the choice of the objective function J affects the optimal temperature (and supersaturation) trajectory.

Early Growth and Late Growth

The exact supersaturation (or temperature) trajectory which solves Eq. 14 will depend on the values of the parameters which are unique to each particular system, as well as the objective function. However, it is argued in this work that many of the supersaturation trajectories found by different authors can be broadly classified into one of two groups: "early growth" and "late growth." Early growth trajectories are characterized by the fact that supersaturation (and, therefore, growth rate) is greatest at the beginning of the batch, and is reduced toward the end of the batch. Late growth trajectories are characterized by the feature that supersaturation is increasing during the batch and is greatest at the end of the batch. For a programmed cooling crystallizer, the slope of the temperature trajectory is greatest in an absolute sense at the beginning of the batch for an early growth policy and greatest at the end of the batch for a late growth policy.

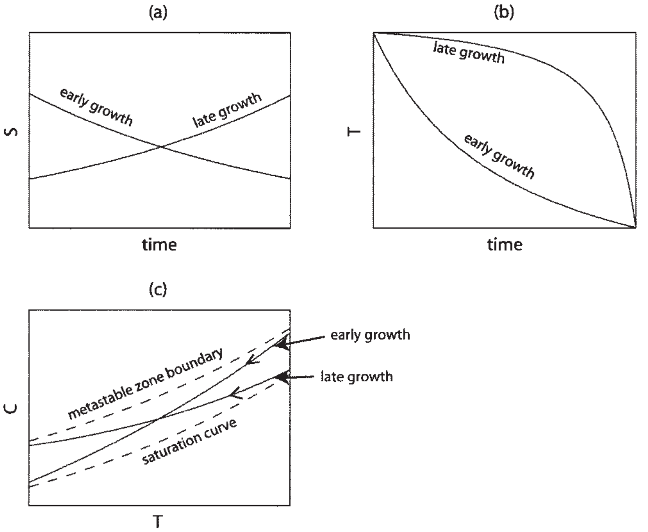

Figure 1. Sketches of various trajectories for early and late growth operating policies. (a) Supersaturation vs. time, (b) temperature vs. time, and (c) concentration vs. temperature.

Figure 1 shows sketches of three different representations of trajectories for both early growth and late growth operating policies. In Figure 1a, supersaturation is shown decreasing with time for an early growth policy, but increasing with time for a late growth policy. In Figure 1b, the temperature decreases most rapidly at the beginning of the batch for an early growth policy, but decreases most rapidly at the end of the batch for a late growth policy. In Figure 1c, for an early growth policy the solute concentration is closest to the metastable limit at high-temperatures (early in the batch), whereas for a late growth policy it is closest to the metastable limit at low-temperatures (late in the batch). Note that the diagrams in Figure 1 are exaggerated to emphasize the difference between policies and do not correspond to any particular physical system.

Between the extremes of early growth and late growth is a policy of constant growth, in which the supersaturation is constant with time. Constant growth trajectories are not shown in Figure 1 for clarity, but this would correspond to a horizontal line in Figure 1a. Note that it would not necessarily correspond to a straight line (linear cooling) in Figure 1b.

So called "natural cooling," in which the batch is merely allowed to cool from an elevated temperature to room-temperature by free convection according to Newton's law of cooling, is an extreme example of an early growth operating policy. It results in a very rapid drop in the solution temperature early in the batch. This may cause the concentration to exceed the metastable envelope and produce a shower of nuclei by primary nucleation. Such an extreme operating policy is likely to produce an undesirable result. What are here termed early growth and late growth operating policies include only policies where supersaturation is controlled to be within the metastable zone at all times so that nucleation is by secondary mechanisms only.

Choosing an Objective Function

There is a great deal of subtlety associated with selecting an objective function for seeded batch crystallizers. Part of the problem, as several researchers have noted, is that it may be difficult to correlate the cost of downstream processing to the crystal-size distribution. As a result, it has not generally been practical to use process economics to formulate an objective. Consequently, researchers have been forced to identify a surrogate objective function based on some property of the crystal-size distribution. (Recent advances suggest that it may be possible to develop and use an economic relationship in the near future. Seymour[8] presents two case studies in which processes with continuous crystallization are designed and costed, including operating costs. Wibowo et al.[2] summarize correlations which are useful for the analysis of unit operations that are typically found downstream from a crystallizer.)

The following is a list of four common objectives for the crystal size distribution at the end of the batch, each of which seems reasonable:

1. minimize the "amount" of nucleus-grown crystals,
2. maximize the "average size" of the total crystals,
3. minimize the "variation in size" of the total crystals,
4. grow the seed crystals as "large" as possible.

Two additional considerations must be kept in mind when translating these word objectives into mathematical objectives. The first is that it is important to decide whether the objective should involve only the seed-grown crystals, only the nucleus-grown crystals, or both. The second is that there is ambiguity associated with terms such as those marked in the list of objectives with quotation marks. The ambiguity arises because it is not clear which moments of the crystal-size distribution should be used in the mathematical objective. For example, in objective 1, minimizing the "amount" of nucleus grown crystals could mean minimizing the *number* of nucleus-grown crystals (μ_{N_0}) or the total volume (or mass) of nucleus grown crystals (μ_{N_3}). Although less common, minimizing μ_{N_1}, μ_{N_2} or μ_{N_4} are also valid mathematical representations of objective 1. Likewise, the term "average size" in objective 2 is ambiguous. The quantity μ_{T_1}/μ_{T_0} is the number average size, while μ_{T_4}/μ_{T_3} is the weight average size. Again by analogy to objective 1, the quantities μ_{T_2}/μ_{T_1} and μ_{T_3}/μ_{T_2} are less common but equally valid measures of the "average size" of the total crystals. The variability in the crystal size distribution (objective 3) can be defined several ways: $\sqrt{\mu_{T_0}\mu_{T_2}/\mu_{T_1}^2 - 1}$ or $\sqrt{\mu_{T_5}\mu_{T_3}/\mu_{T_4}^2 - 1}$ depending on whether one is concerned with the variability in the low moments or the high moments. In summary, terms such as "amount," "average size" and "variation in size" can be defined with equal validity using either low moments of the crystal-size distribution or high moments. It is shown in subsequent sections that the choice of whether the objective should involve low moments of the crystal size distribution or high moments has a much greater impact on the optimal supersaturation trajectory than does the choice among objectives 1–4.

For reference, Table 1 summarizes results from twelve authors spanning more than 30 years starting with the seminal work of Jones[1] and extending to the present day. The criteria for inclusion in Table 1 are that the authors must have used the method of moments to model the population balance equations and considered secondary nucleation (that is, their model must have been similar to the one developed in this article and in particular they must have used Eq. 10 as the expression for secondary nucleation); and they must have attempted to determine an optimal supersaturation (or temperature) trajectory based on some property of the moments at the end of the batch. Although a reasonable effort was made to identify and include every study which meets these criteria, some may have been left out inadvertently. Nevertheless, Table 1 represents a broad survey of the literature in this field to date. Although it is not shown in the table, it is worth noting that these authors have used a variety of numerical methods to tackle the optimization problem, including optimal control, sequential quadratic programming (with the continuous variable discretized using a linear spline) and stochastic methods, such as simulated annealing.

For each work, Table 1 shows the exact mathematical objective function, which word objective(s) 1–4 the authors were attempting to represent with their mathematical objective, whether they were concerned with low or high moments of the crystal-size distribution, and whether their resulting temperature (or supersaturation) trajectory could be classified as either late growth or early growth. There is some inherent subjectivity associated with classifying trajectories as late growth or early growth. The classification was based either on statements by the authors that supersaturation was greatest at the beginning or end of the batch, or by examination of the supersaturation (or temperature) trajectories published in the work. If reasonable doubt remained as to the appropriate classification of a particular result, it was marked "other." Some trajectories in this category exhibited nearly constant supersaturation; some showed an increase in supersaturation at both the beginning and at the end of the batch; some had a shape different from any of these descriptions.

Looking at Table 1, several features stand out. First, many different mathematical objective functions have been used by different authors over the years. Second, of the 14 different trajectories which could be classified as either early or late growth, the majority (nine) were late growth, but a significant minority (five) were early growth. No correlation between the word objective used in a study and the optimal operating policy that was found is discernable: examples can be found where each of the four word objectives results in a late growth policy, and there are examples of word objectives 2 and 3 resulting in early growth policies. By contrast, there is a strong correlation between the optimal operating policy and whether the authors were concerned with the high moments of the crystal size distribution or the low moments. In 10 out of 12 possible cases, authors who were concerned with low moments of the crystal-size distribution found early growth operating policies to be optimal while authors who considered high moments of the crystal-size distribution found late growth operating policies to be optimal. The apparent exceptions to this pattern are the paper by Jones[1] and the article by Rawlings et al.[11] However, unlike the other studies both of these works used objective functions which considered *only* the seed-grown crystals. The importance of this distinction will be discussed later. Allowing for this, the results given in Table 1 can be summarized as follows:

Table 1. Summary of the Literature on the Optimal Control of Batch Crystallizers

Authors	Mathematical Objective (minimize)	Word Objective	Moments Considered	Operating Policy	Notes
Jones[1]	$-\mu_{S_1}$	4	Low	Late growth	
Chang and Epstein[9]	$-\frac{\mu_{T_1}}{\mu_{T_0}}$	2	Low	Early growth	
Chang and Epstein[9]	$-\mu_{T_3}$	—	High	Early growth	a
Chang and Epstein[9]	$\frac{\mu_{T_2}}{\mu_{T_0}} - \left(\frac{\mu_{T_1}}{\mu_{T_0}}\right)^2 - \alpha \frac{\mu_{T_1}}{\mu_{T_0}}$	2, 3	Low	Early growth	
Eaton and Rawlings[10]	$\frac{\mu_{N_3}}{\mu_{S_3}}$	1, 3	High	Late growth	
Rawlings et al.[11]	$-\mu_{S_1}$	4	Low	Late growth	b
Miller and Rawlings[13]	$\frac{\mu_{N_3}}{\mu_{S_3}}$	1, 3	High	Late growth	
Chung et al.[5]	$-\frac{\mu_{T_4}}{\mu_{T_3}}$	2	High	Late growth	
Chung et al.[5]	$\sqrt{\frac{\mu_{T_2}\mu_{T_0}}{\mu_{T_1}^2} - 1}$	3	Low	Early growth	
Chung et al.[5]	$\frac{\mu_{N_3}}{\mu_{S_3}}$	1, 4	High	Late growth	
Ge et al.[14]	$-0.00081 \frac{\mu_{T_4}}{\mu_{T_3}} + 0.025 \frac{\mu_{T_3}}{\mu_{N_3}} + 0.10 \sqrt{\frac{\mu_{T_2}\mu_{T_0}}{\mu_{T_1}^2} - 1}$	1, 2, 3, 4	Both	Late growth	
Feng and Berglund[15]	$-\frac{\mu_{T_4}}{\mu_{T_3}}$	2	High	Late growth	
Ma and Braatz[17]	$\frac{\mu_{T_1}}{\mu_{T_0}}$	2	Low	Early growth	c
Zhang and Rohani[18]	$\frac{\mu_{T_4}}{\mu_{T_3}} + 0.0005 \sqrt{\frac{\mu_{T_5}\mu_{T_3}}{\mu_{T_4}^2} - 1}$	2, 3	High	Other	
Rawlings et al.[16]	μ_{N_3}	1	High	Late growth	
Choong and Smith[19]	$-\frac{\mu_{T_4}}{\mu_{T_3}}$	2	High	Other	
Choong and Smith[19]	$\sqrt{\frac{\mu_{T_5}\mu_{T_3}}{\mu_{T_4}^2} - 1}$	3	High	Other	

[a] This objective is not meaningful in the context of this work, because it is assumed here that the third moment of the total crystal size distribution is constrained to be constant.

[b] Miller and Rawlings do not identify the exact mathematical objective function in their paper; the reported objective function is based on a personal communication from Rawlings.[12] As mentioned by Rawlings in the personal communication and discussed earlier, when the objective concerns only the seed crystals, the exact nature of the objective function is not important, because the PSD of the seed crystals depends only on the total amount of crystal growth which occurs during the batch.

[c] Ma and Braatz consider a system of crystals with two characteristic lengths. Thus their actual objective function is $-\mu_{01}/\mu_{00}$.

- Authors who considered *only* word objective 4 (maximize terminal seed-grown crystal size) always find late growth policies to be optimal.
- Authors who considered word objectives 1–3 and were concerned with the higher moments of the crystal size distribution also found late growth policies to be optimal.
- Authors who considered word objectives 1–3 and were concerned with the lower moments of the crystal-size distribution found early growth policies to be optimal.

The remainder of this article will develop specific insight as to why these results were obtained.

Insight from Suboptimal Supersaturation Trajectories

In order to investigate the influence that the choice of objective function has on the optimal supersaturation trajectory, the problem was "reverse engineered" by subjecting the crystallizer model to a series of predetermined supersaturation trajectories. Ideally, the expression for the supersaturation trajectory can be parameterized by a variable which when adjusted over an appropriate range causes the supersaturation trajectory to transition smoothly from early growth to constant growth to late growth. A function which meets these criteria and is employed in this work is

$$S(t) = \alpha e^{\beta t/t_f} \tag{15}$$

where α and β are adjustable parameters. If $\beta < 0$ then the supersaturation trajectory is one of early growth; if $\beta = 0$ then the supersaturation trajectory is one of constant growth; if $\beta > 0$ then the supersaturation trajectory is one of late growth. For any given value of β, the parameter α is adjusted so that the

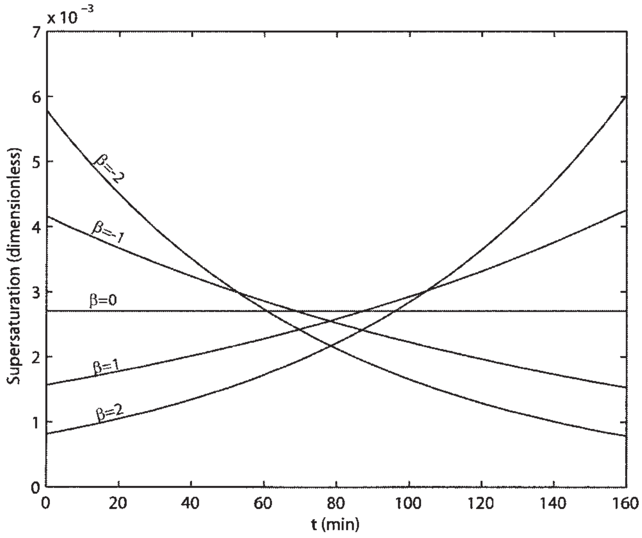

Figure 2. Supersaturation trajectory parameterized by β.

total amount of growth is sufficient to meet the production rate requirement. Figure 2 illustrates the supersaturation trajectory parameterized by five different values of β.

Figure 3 shows the *final* value of the zeroth through fourth moments of the nucleated, seed-grown and combined (seed-grown + nucleus-grown) crystals versus the parameter β for the case study system of potassium nitrate in water. Further details about this case study are presented in the section which gives optimal temperature trajectories for various objectives. It is important to realize that this plot *does not* show the time evolution of the crystal-size distribution—all values of the moments are at the end of the batch. Values are plotted vs. the parameter β of Eq. 15. Negative values of β correspond to an early growth strategy, while positive values of β correspond to a late growth strategy. The discrepancy between the findings of different authors in terms of the optimal supersaturation trajectory can be explained entirely in terms of this figure. However, it is important to make a few observations first.

First, both the zeroth moment of the seed crystals and the third moment of the combined crystals are constant. The zeroth moment of the seed-grown crystals is constant because by definition no additional seed crystals are formed during the batch, therefore, the number of seed-grown crystals at the end of the batch is the same as the number of seeds at the start and does not depend on the supersaturation trajectory. The third moment of the combined crystals is constant because this is the constraint imposed on the system. For each value of β in Eq. 15, the parameter α was adjusted so that the total volume of crystals produced by the end of the batch would be the same.

Second, the lower moments of the crystal-size distribution are influenced much more by the choice of supersaturation trajectory than are the higher moments, and the zeroth moment is influenced by far the most. For comparison, the "variability" of each moment is defined as the maximum value of that moment minus the minimum value divided by the maximum value expressed on a percent basis

$$\%\ \text{variability} = \left(\max_{-2\le\beta\le2} \mu_{T_i} - \min_{-2\le\beta\le2} \mu_{T_i} \right) \Big/ \left(\max_{-2\le\beta\le2} \mu_{T_i} \right) \times 100\% \quad (16)$$

Table 2 shows the variability of the first five moments of the total (combined) crystal-size distribution. Note that the third moment is not listed because the variability in the third moment is zero because this value is fixed. One consequence of the high variability of the zeroth moment is that an objective that combines more than one moment, such as $-\mu_{T_1}/\mu_{T_0}$ or $\mu_{T_2}\mu_{T_0}/\mu_{T_1}^2$ will tend to be dominated by the zeroth moment.

It is now possible to explain the disparate results of Table 1.

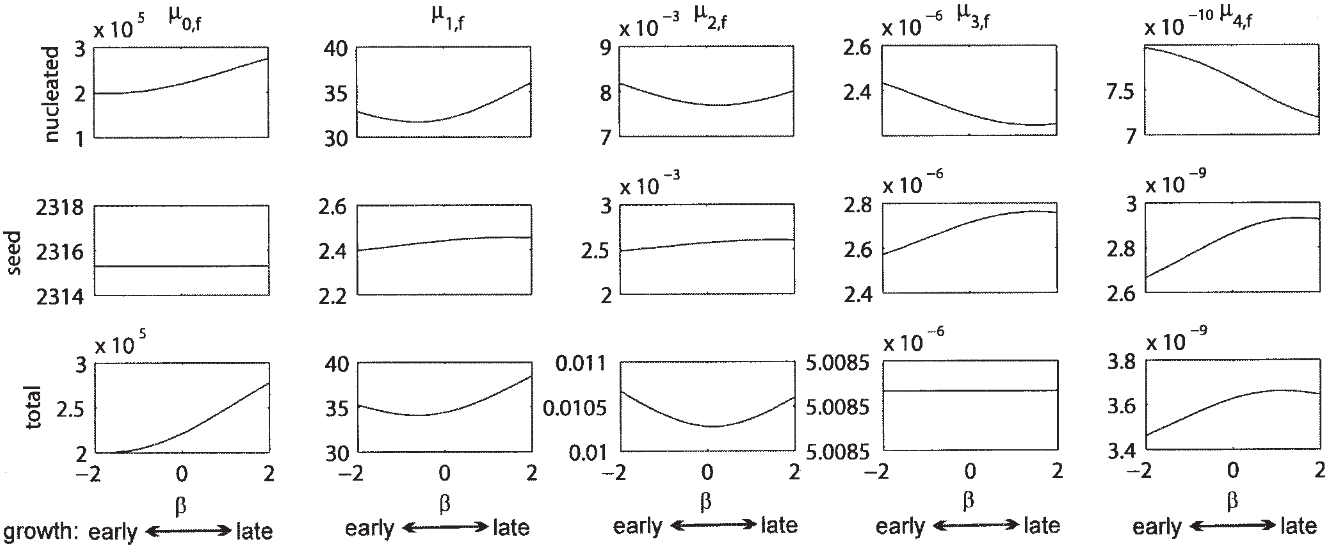

Figure 3. Final values of the first four moments of the crystal-size distribution of nucleated crystals, seed-grown crystals, and the sum of these as a function of the parameter β.

Negative values of β correspond to an early growth policy, while positive values of β correspond to late growth policies. Note that these plots do *not* show the time evolution of the moments of the crystal-size distribution. Values of the moments are reported only at the end of the batch.

Table 2. % Variability in Moments 0–4 (as Defined in Eq. 16)

Moment	% Variability
μ_{T_0}	28
μ_{T_1}	11
μ_{T_2}	3.7
μ_{T_4}	5.4

It is important to bear in mind the critical tradeoff that one encounters when selecting an operating policy for seeded batch crystallizers with secondary nucleation:

Early growth operating policies minimize the *number* of nuclei formed, because supersaturation is highest when the mass of crystals is least, namely at the beginning of the batch. However, those nuclei that *are* formed have a chance to grow to a larger size because they are formed early in the batch. By contrast, late growth operating policies produce a greater number of nuclei because the supersaturation is high at a time when the mass of crystals is great (at the end of the batch). However, nuclei that are formed late in the batch do not have a chance to grow to as large a size as those formed early in the batch.

The identification of this tradeoff and its consequences are the most important results of this work. Evidence for the validity of this way of thinking about the design of a crystallizer operating policy is shown in Figure 3.

Consider the first row of figures, which shows how the final value of the moments of the nucleus-grown crystals change as a function of the parameter β. Plausible objectives for the crystal-size distribution include minimizing any of the moments 0–3. Minimizing the zeroth moment corresponds to minimizing the number of nuclei formed, and minimizing moments 1–3 corresponds to minimizing the "total size" of the nucleus-grown crystals in some sense, either minimizing the total length, area or volume of nucleus-grown crystals. It is clear from Figure 3, however, that the optimal operating policy changes substantially based on which moment of the nucleus-grown crystals is considered. If one is concerned with the zeroth moment, then an early growth policy is best. For the third moment a late growth policy is best. In between these extremes, if the first or second moments are of concern, then an operating policy closer to constant growth is best.

Minimizing the "total size" of the nucleus-grown crystals is only one plausible objective. Another plausible objective is *maximizing* the *average* size of the total quantity of crystals. This choice corresponds to maximizing a ratio of consecutive moments such as μ_{T_1}/μ_{T_0} (number average size) or μ_{T_4}/μ_{T_3} (weight mean size).

Earlier it was stated that objectives which involve only properties of the seed-grown crystals always lead to late growth operating policies. This result can also be understood in terms of Figure 3, and the critical tradeoff identified in this section. With the set of assumptions employed in this work, the evolution of the crystal-size distribution of the seed-grown crystals is deterministic in the sense that it depends only on the total amount of growth which has occurred, that is $\int G(t)dt$. One can imagine that the nucleus-grown crystals form and grow completely independently from the seed-grown crystals. However, the nucleus-grown crystals do interact with the seed-grown crystals in one important way: through the production rate constraint. If the total mass of nucleus-grown crystals is allowed to become large, this process will "steal" mass away from the seed-grown crystals: at the end of the batch a large fraction of the total mass of crystals will comprise the nucleus-grown crystals, and, therefore, a smaller amount will comprise the seed-grown crystals. However, the evolution of the seed-grown crystals is deterministic. As a consequence, *any* objective involving only the seed-grown crystals, such as maximizing the amount or average size of the seed grown crystals, based on *any* moment of the seed-grown crystal-size distribution, is mathematically equivalent to minimizing the total volume (mass) of the nucleus-grown crystals. And it has already been shown that such an objective always leads to a late growth operating policy.

As an alternative way of understanding the different effects caused by early growth and late growth operating policies, consider Figure 4 which shows the time evolution of the zeroth and third moments of the total crystal-size distribution for an early growth ($\beta = -2$) and a late growth ($\beta = 2$) operating policy. For the zeroth moment, the early growth operating policy produces more crystals early in the batch, but by the end of the batch the late growth operating policy produced more new crystals because the supersaturation is greatest at the end of the batch when the total mass of solid crystals is greatest. For the third moment, at the end of the batch both operating policies have produced the same volume of crystals because of the production rate constraint. The early growth operating policy generates more volume at the beginning of the batch, and the late growth operating policy produces more volume at the end of the batch.

In summary there are two types of objectives that lead to late growth operating policies:

- Objectives involving the higher moments of either the nucleus-grown crystals or the total quantity of crystals.
- Objectives involving *any* moments of the seed-grown crystals.

By contrast, only one type of objective leads to early growth operating policies:

- Objectives involving the lower moments of either the nucleus-grown crystals or the total quantity of crystals.

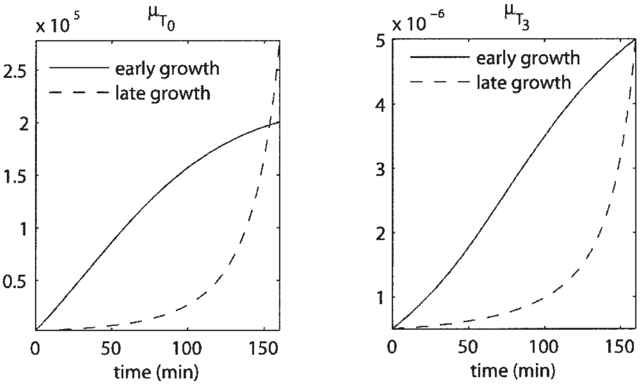

Figure 4. Time evolution of the zeroth and third moments of the total crystal size distribution with early- and late growth operating policies.

Table 3. Objective Functions Considered for Optimization

Objective Function (Minimize)	Prediction
$-\mu_{T_4}/\mu_{T_3}$	Late growth
μ_{T_0}	Early growth
$\sqrt{\mu_{T_2}\mu_{T_0}/\mu_{T_1}^2 - 1}$	Early growth
$-\mu_{S_1}/\mu_{S_0}$	Late growth
$-\mu_{T_1}/\mu_{T_0}$	Early growth
$-\mu_{N_3}/\mu_{S_3}$	Late growth

Table 5. Moments of the Crystal Size Distribution for the Seeds

Moment	Value	Units
0	2.31×10^3	#/kg
1	1.39	m/kg
2	8.34×10^{-4}	m^2/kg
3	5.01×10^{-7}	m^3/kg
4	3.01×10^{-10}	m^4/kg

Optimal Temperature Trajectories for Various Objectives

In order to provide additional support for the arguments presented in this article, a case study process was optimized using a variety of objective functions listed in Table 3. Table 3 also shows the prediction, based on the results developed in the article, of whether the objective should lead to a late growth or early growth operating policy. Numerical values (listed in Table 4) are taken from a case study published by Chung et al.[5] on the crystallization of potassium nitrate from water. Table 5 gives the moments of the crystal-size distribution for the seeds. In order to reduce the problem from an infinite-dimensional path optimization problem to a finite dimensional problem, the temperature trajectory was fit to a linear spline with nine data points (following Chung et al.). Chung et al. then solved the optimization problem by sequential quadratic programming. However, it was found in this study that for at least some of the objectives, the sequential quadratic programming solver in MATLAB (fmincon) converged to local minima from some starting conditions. For some objectives, the trajectories obtained from different starting points were similar, but for other objectives they were markedly different and some gave substantially inferior values of the objective function.

In order to rule out the possibility of converging to local minima, the optimization problem was solved by gridding the parameter space using a grid size of 0.5°C and evaluating the objective function at every feasible grid point. This method is in the spirit of the class of global optimization techniques known as interval analysis.[20,21] It is well known that the disadvantage of interval analysis is that it may be computationally expensive in some circumstances. However, there are two features of this problem which make interval methods an attractive alternative. First, the dimensionality of the problem is not prohibitively high. Second, it was desired in this work to investigate six different objective functions for the same set of constraints. Integrating the differential equations which constitute the crystallizer model was (by far) the most time-consuming part of the evaluation of the objective function. Using an interval method, it was possible to evaluate all six objective functions at each grid point after having integrated the differential equations only once at that grid point. This feature resulted in significant computational savings. For the stated grid size, there were 23,940 feasible trajectories and it took approximately eight hours on a workstation with a 2.2 GHz processor to evaluate the objective functions for every feasible trajectory.

Table 4. Values of Parameters in Case Study

Variable	Name	Value	Units
ρ_c	Density of crystal	2.11×10^3	kg/m^3
T_{max}	Maximum temperature	32	°C
T_{min}	Minimum temperature	22	°C
R_{max}	Maximum rate of T change	0	°C/min
R_{min}	Minimum rate of T change	−0.1	°C/min
k_b	Nucleation parameter	4.6401×10^{11}	$\#/m^3$ s
b	Nucleation parameter	1.78	Dimensionless
k_g	Growth parameter	1.1612×10^{-4}	m/s
g	Growth parameter	1.32	Dimensionless

The saturation concentration C_{sat} is given by:

$$C_{sat} = 0.1286 + 0.00588T + 0.0001721T^2$$

where C_{sat} has units of kg KNO_3/kg H_2O and T is in °C.

The results of the optimization are shown in Figure 5. Where a late growth operating policy is predicted to be optimal, the computed optimal temperature trajectory decreases as rapidly as possible at the end of the batch, whereas when an early

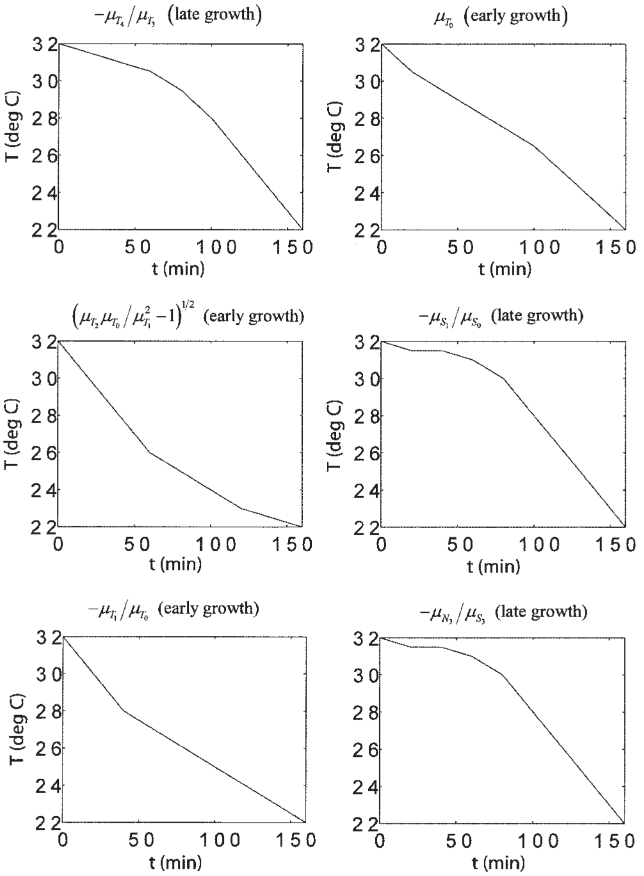

Figure 5. Optimal temperature trajectories for the potassium nitrate case study corresponding to six different objective functions.

growth policy is predicted to be best the computed temperature trajectory drops more rapidly at the beginning of the batch.

Conclusions

In concluding this article, it is worthwhile to put these results into context and to discuss how they might be used by the process engineer. Previous authors have suggested that as an operating heuristic the supersaturation should be kept constant during a batch run. They have demonstrated that this policy results in a substantially improved crystal-size distribution compared to natural cooling. Constant relative supersaturation is a very reasonable operating policy—it is expected to provide a vast improvement in CSD compared with an uncontrolled operating policy, such as natural cooling. A constant supersaturation operating policy may be "good enough" in many practical cases. However, the results presented in this paper allow the process engineer to make an additional refinement to the constant supersaturation operating policy and achieve a further improvement in CSD without much additional effort. With our approach there is no need to develop a sophisticated mathematical model of the process, to regress process data so as to identify parameters in the model, and to optimize the model using some sophisticated numerical method. All of this work has been done "in advance" in order to provide the intuition in this article.

An engineer could make use of these results in the following manner: In the very early stages of the process design, the engineer may discover by trial and error that a constant relative supersaturation of 10^{-2} is much too high and immediately produces a shower of nuclei, while a constant relative supersaturation of 10^{-4} is much too low and causes no perceptible crystal growth after many hours, but that a constant relative supersaturation of 10^{-3} allows for reasonably steady crystal growth with a tolerable rate of formation of nuclei. At that point, using the results presented in this article, the engineer can make an immediate improvement to the operating policy without much additional effort. If the concern is the number of small particles, on the next batch the engineer could try an early growth operating policy, in which the supersaturation is decreased linearly from (say) 1.1×10^{-3} to 0.9×10^{-3}. Conversely, if the concern is about minimizing the total volume of the nucleus-grown crystals, or maximizing the total volume or average size of the seed-grown crystals, then a late growth operating policy might be applied, gradually increasing the relative supersaturation from 0.9×10^{-3} to 1.1×10^{-3}. In this way, much of the residual benefit of further optimizing the relative supersaturation trajectory (beyond a constant relative supersaturation trajectory) can be realized with only a small fraction of the effort required to conduct a rigorous optimization.

It is also important to note the limitations of this method. All of the work presented in this article is dependent on the validity of the crystallizer model presented earlier, and in particular on the expression for secondary nucleation given in Eq. 10. If some aspect of this model is invalid, and in particular if the expression for the rate of nucleation is invalid, the results presented in this paper may not hold. For example, if nucleation were to occur through a primary rather than secondary mechanism, or if it were to depend on some moment of the total crystal-size distribution other than the third, or if it did not depend on the supersaturation, then the results presented in this paper would have to be modified accordingly. The crystallizer model developed earlier is the one most commonly used by researchers, and one contribution of this work is that it is able to explain and place into context a great deal of previous work by other researchers. The reader is merely cautioned that the results presented here (and elsewhere) are no better than the assumptions they are based on.

Acknowledgment

This work was partially supported by the National Science Foundation (Grant No. CTS-0411630).

Literature Cited

1. Jones AG. Optimal operation of a batch cooling crystallizer. *Chem Eng Sci.* 1974;29:1075–1087.
2. Wibowo C, Chang WC, Ng KM. Design of integrated crystallization systems. *AIChE J.* 2001;47:2474–2492.
3. Fujiwara M, Nagy ZK, Chew JW, Braatz RD. First-principles and direct design approaches for the control of pharmaceutical crystallization. *J Process Contr.* 2005;15:493–504.
4. Randolph AD, Larson MA. *Theory of Particulate Processes.* Academic Press: San Diego; 1988.
5. Chung SH, Ma DL, Braatz RD. Optimal seeding in batch crystallization. *Can J Chem Eng.* 1999;77:590–596.
6. Ramkrishna D. *Population Balances.* Academic Press: San Diego; 2000.
7. Hulburt HM, Katz S. Some problems in particle technology—a statistical mechanical formulation. *Chem Eng Sci.* 1964;19:555–574.
8. Seymour CB. *Conceptual design of dedicated batch versus continuous processing in multireaction step processes.* University of Massachusetts at Amherst; 1995. Ph.D. Dissertation.
9. Chang CT, Epstein MAF. Identification of Batch Crystallization Control Strategies Using Characteristic Curves. In: *Nucleation, Growth and Impurity Effects in Crystallization Process Engineering;* Epstein MAF, ed. AIChE: New York; 1982.
10. Eaton JW, Rawlings JB. Feedback control of chemical processes using on-line optimization techniques. *Comput Chem Eng.* 1990;14:469–479.
11. Rawlings JB, Miller SM, Witkowski WR. Model identification and control of solution crystallization processes: A Review. *Ind Eng Chem Res.* 1993;32:1275–1296.
12. Rawlings JB. Personal communication; 2005.
13. Miller SM, Rawlings JB. Model identification and control strategies for batch cooling crystallizers. *AIChE J.* 1994;40:1312–1327.
14. Ge M, Wang QC, Chiu MS, Lee TH, Hang CC, Teo KH. An effective technique for batch process optimization with application to crystallization. *Trans IChemE.* 2000;78:99–106.
15. Feng LL, Berglund KA. ATR-FTIR for determining optimal cooling curves for batch crystallization of succinic acid. *Cryst Growth Des.* 2002;2:449–452.
16. Rawlings JB, Sink CW, Miller SM. Control of Crystallization Processes. In: *Handbook of Industrial Crystallization.* 2nd ed. Myerson AS, ed. Butterworth Heinemann: Boston; 2002.
17. Ma DL, Braatz RD. Robust identification and control of batch processes. *Comput Chem Eng.* 2003;27:1175–1184.
18. Zhang GP, Rohani S. On-line optimal control of a seeded batch cooling crystallizer. *Chem Eng Sci.* 2003;58:1887–1896.
19. Choong KL, Smith R. Optimization of batch cooling crystallization. *Chem Eng Sci.* 2004;59:313–327.
20. Lin Y, Stadtherr MA. Advances in interval methods for deterministic global optimization in chemical engineering. *J Global Optim.* 2004;29:281–296.
21. Hanson E, Walster GW. *Global Optimization Using Interval Analysis.* 2nd ed. New York: Marcel Dekker; 2004.

Manuscript received July 19, 2005, and revision received Jan. 24, 2006.