Hierarchical Multiobjective Strategy for Particle-Size Distribution Control

Charles David Immanuel and Francis Joseph Doyle III
Dept. of Chemical Engineering, University of Delaware, Newark, DE 19716

A multiobjective optimization formulation is presented to track a desired particle-size distribution (PSD) trajectory in semi-batch emulsion polymerization. This is done in a hierarchical framework, controlling the individual processes of nucleation, growth, and coagulation that determine the PSD. An evolutionary technique, the Nondominated Sorting Genetic Algorithm, is employed as the optimization tool. Experimental implementation of these open-loop optimal recipes yield promising results, but simultaneously highlight the importance of feedback corrective action. Preliminary state/parameter estimation studies are performed, which show an ability to estimate the full distribution and related states from the available measurements (relative distribution, latex density, and feed rates), and also elucidate the requirements in the strategy to be employed for estimation. The potential for on-line feedback is examined using experimental data, in view of the irreversibilities and process limitations.

Introduction

The control of the full particle-size distribution (PSD) in emulsion polymerization is important due to the complex and multi-modal nature of the target distributions that are relevant for industrial applications. These include adhesives, paints and coatings, cosmetics, synthetic rubber, and so on, wherein the adhesion, rheological properties, drying characteristics, optical properties, and mechanical strength characteristics are critical. These properties are primarily determined by the PSD. The evolution of PSD is driven by three major phenomena: nucleation, growth, and coagulation. These phenomena strongly interact with each other, and also are influenced in complex ways by the available manipulated variables in the process, calling for advanced model-based control.

Although the modeling of the PSD in emulsion polymerization in a population balance framework has been pursued since the 1970s (Min and Ray, 1974), studies addressing the control of entire distributions are very recent. These coincide with advances in instrumentation that enable on-line and sufficiently rapid measurements of PSD, and also with advances in the computation techniques. Semino and Ray (1995b) demonstrated successful control of PSD in emulsion polymerization, based on their study on the controllability of distributions in different population balance systems (Semino and Ray, 1995a). Crowley et al. (2000) were the first to employ a detailed population balance model in the design of an optimal recipe for semi-batch emulsion polymerization of styrene, to attain a target PSD at the end of the batch. Flores-Cerrillo and MacGregor (2002) proposed a "mid-course" correction strategy to the open-loop optimal recipe to account for disturbances, wherein the recipe is re-computed at a single point in the middle of the batch. They used a PLS model, which reduces the re-computation to solving simple regression problems.

A related contribution is the batch-to-batch strategy proposed by Crowley et al. (2001) to address the computational feasibility and the process limitations which make in-batch feedback ineffective. They used a hybrid model for this purpose, with an information-rich and extrapolative fundamental model complemented with a PLS model to account for the coagulation events and for other model mismatch. They observed in simulations that convergence can be obtained within 3–4 batches in typical cases. Recently, studies in our group have addressed the issues of process sensitivities and the reachability of distributions and profiles (Wang et al., 2002; Immanuel and Doyle III, 2003a). The latter study presents experimental sensitivity results for vinyl acetate (VAc) and butyl acrylate (BuA) co-polymerization using nonionic surfac-

Correspondence concerning this article should be addressed to F. J. Doyle III at this current address: University of California, Chemical Engineering Dept., Santa Barbara, CA 93106.

tants. The findings suggest the use of the feedrates of surfactant and a monomer as the manipulated variables for the control of PSD, reveal the irreversibilities in the process, and suggest a possible hierarchical strategy in which the nucleation, growth, and coagulation events are regulated separately as part of the overall aim to control the full distribution.

In an earlier study (Immanuel and Doyle III (2002)), we extended the work of Crowley et al. (2000) to the semi-batch emulsion co-polymerization of vinyl acetate (VAc) and butyl acrylate (BuA). A notable distinction in this study is the use of a nontraditional, direct optimization technique—the Genetic Algorithm (GA)—as the optimization tool. This was motivated partly by the computational burden associated with the computation of gradients and Hessians. Although approximate techniques exist to circumvent these issues, the discontinuity in the process *vis-a-vis* the nucleation phenomenon and the nonconvexity observed in the underlying optimization problem by earlier studies motivate the application of a direct and global optimization technique.

In this article, a hierarchical strategy is presented for the optimization of a recipe to attain a target distribution in semi-batch emulsion polymerization. The system under investigation is vinyl acetate (VAc)–butyl acrylate (BuA) co-polymerization, with nonionic surfactants as emulsifiers. The difference between this study and the previous study (Immanuel and Doyle III, 2002) is that a novel hierarchical strategy is proposed here to overcome some of the limitations with direct control of PSD. Experimental implementation of these optimal recipes is presented. This is followed by an evaluation of the effectiveness of in-batch feedback to correct the discrepancies introduced by the unavoidable disturbances and uncertainties. This involves the evaluation of several state/parameter estimation strategies, and a determination of the requirements for successful state estimation. These strategies are employed in off-line feedback calculations, to investigate their effectiveness. This off-line evaluation effort is motivated by the delayed and sparse measurements, the irreversibilities in the process, and clear indications in other studies (Immanuel and Doyle III, 2003a; Wang et al., 2002) of possible cases wherein in-batch feedback would not be able to rectify a discrepancy in the distribution. The importance of state estimation in these feedback calculations is evinced.

Open-Loop Optimization Studies

In the semi-batch operation of emulsion polymerization reactors, the primary control problem is the regulation of PSD to a target at the end of the batch. One critical issue in this optimization is the choice of the objective function. The target distributions are multi-modal, usually with complex shapes of the modes. Crowley et al. (2000) examined the sensitivity to different objective functions, and settled on a multiobjective formulation, with one objective describing each mode, proposing a min-max optimization problem involving these two objectives. In the current work, bimodal targets with clearly separated modes are used, characterized by complex relations in terms of the number of particles in the two modes and the mass of the particles in the two modes. Such distributions are not captured adequately by any straightforward objective function. The choice of a proper objective function to capture all aspects of the distribution is a difficult task, and usually varies from target to target. Previous sensitivity studies (Immanuel and Doyle III, 2003a) suggest the benefit of regulating the nucleation, growth, and coagulation processes separately, and, hence, controlling the PSD in a hierarchical framework. An advantage in this framework is that it is applicable to all types of target distributions. Another advantage is that it facilitates the tracking of a target PSD trajectory along the course of the batch—say, to guide the distribution to the target along a path that avoids problematic regimes, or to monitor the evolution of the distribution and decide upon the appropriate feedback measure.

Figure 1 depicts the proposed hierarchical strategy. As observed in the previous study (Immanuel and Doyle III, 2002), coagulation is one of the most imporatant sources of irreversibility, and it is more difficult to manipulate coagulation to shape the distribution in emulsion polymerization. Thus, it is advisable to minimize coagulation to the extent possible. In following the trajectory of nucleation and growth, the idea of the control of instantaneous properties can be utilized (Kozub and MacGregor, 1992a; Urretabizkaia et al., 1994; Saldivar and Ray, 1997). Exploiting this idea, the trajectory of the nucleation rate can be controlled by regulating the profile of total particles, while the trajectory of the growth rate can be controlled by regulating the profile of solids content.

Thus, the problem of following a target PSD trajectory is transformed first into an equivalent problem of tracking the profiles of nucleation and growth rates, with coagulation minimized. This in turn is transformed into another equivalent problem of tracking the profiles of total particles and solids

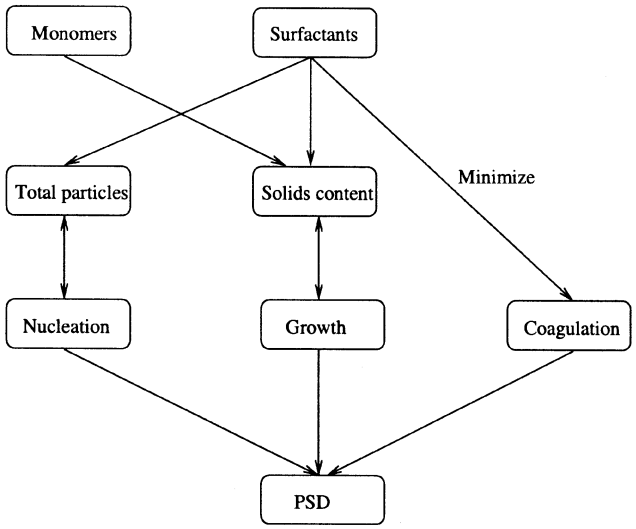

Figure 1. Proposed hierarchical strategy for the control of PSD, by regulating the nucleation, growth and coagulation events individually, in a multi-objective framework.

The target PSD trajectory is transformed into trajectories of the individual rates of the nucleation and growth events (with minimization of the coagulation events). Furthermore, the idea of the control of instantaneous properties is exploited to re-cast the trajectories of nucleation and growth rates as equivalent profiles of total particles and solids content.

content, with coagulation minimized. The transformed problem is a multiobjective optimization problem, in which the first objective (denoted θ_1) is the tracking error of the total particles, and the second objective (θ_2) is the tracking error of the solids content. These are defined as follows

$$\theta_1 = \sum_i w_i^1 \left(N_{p,i} - N_{p,i}^{\text{ref}} \right)^2 \tag{1}$$

$$\theta_2 = \sum_i w_i^2 \left(sc_i - sc_i^{\text{ref}} \right)^2 \tag{2}$$

where $N_{p,i}$ (sc_i) denote the value of total particles (solids content) at time instant i, and $N_{p,i}^{\text{ref}}$ and sc_i^{ref} are the corresponding target values at this time instant. w_i^1 and w_i^2 are appropriate weights. A third objective can also be considered for the end-point distribution error ($\theta_3(t_f)$), defined as follows

$$\theta_3(t_f) = \int_{r_{\text{nuc}}}^{r_{\max}} \left(W(r,t_f) - W_{\text{ref}}(r) \right)^2 dr \tag{3}$$

Here, $W(r,t_f)$ is the weight-averaged PSD at the end time t_f and $W_{\text{ref}}(r)$ is the target weight-averaged PSD, r_{nuc} and $r_{\max}$ being the minimum and maximum particle sizes in the population. θ_1 and θ_2 can be collapsed into a single weighted objective and the problem can be solved using the regular single-objective GA (Immanuel and Doyle III, 2002). However, these objectives conflict with each other, with improvement in one objective leading to a potential worsening of the other (due to the interaction between nucleation, growth, and coagulation (Immanuel and Doyle III, 2003a). Also, in order to alleviate the problems with the choice of a proper weight, a multiobjective formulation is employed.

Multiobjective optimization is very common in engineering applications. For example, in the model predictive control (MPC) algorithm, set point tracking error is minimized while simultaneously penalizing the manipulated inputs. Likewise, in optimizing the performance of chemical reactors, a typical pair of objectives is to maximize the conversion and minimize the byproducts. More generally, one is often faced with the tradeoff between an engineering objective vs. an economic objective (Luyben and Floudas, 1994). It is a common practice to combine these multiple objectives through suitable weights—by emphasizing one objective and relegating to a lower position the other. Although this is acceptable in most cases, there are a number of situations in which the objectives are in competition. This is particularly true in several polymerization applications, prompting the use of multiobjective strategies for these applications (Tsoukas et al. (1982); Choi and Butala (1991); Mitra et al. (1998); Bhaskar et al. (2000)). A common strategy used for multiobjective optimization is the ϵ-constraint method, wherein only one of the objective is chosen for the optimization, and all the other objectives are appended as inequality constraints in the optimization problem (Tsoukas et al., 1982; Choi and Butala, 1991; Luyben and Floudas, 1994). This problem is amenable to gradient-based optimization techniques. Thus, a series of optimization problems are solved to find a pareto-optimal set of solutions. Among these pareto solutions, improvement in one objective leads to a worsening of one or more of the other objectives, thereby rendering each of them equally acceptable or tradeoff solutions. A different technique that directly obtains the pareto curve, and has found wide applications, is the Nondominated Sorting Genetic Algorithm (NSGA) (Mitra et al., 1998; Bhaskar et al., 2000; Silva and Biscaia, Jr., 2001). This is based on an extension of the genetic algorithm to multiobjective problems.

NSGA is employed to solve the multiobjective problem in this study. The genetic algorithm is an example of a reverse engineered strategy: it is based on the natural evolution paradigm. The algorithm starts with a random set of solutions, each of which is evaluated with respect to the defined objective(s) and assigned fitness values. These solutions (members of a population) are reproduced into the next generation in proportion to their fitness values. The solutions (members) are allowed to evolve with respect to the desired characteristics by cross-over and mutation operations. The selection (reproduction), cross-over and mutation operations are repeated over several generations until they show no improvement in the objective function value. See Immanuel and Doyle III (2002) for more details on the specific algorithm employed in this work. As discussed previously, a multiobjective optimization problem produces solutions which are not necessarily optimal with respect to any of the objectives considered separately, but are the best tradeoff solutions. In multiobjective optimization, a solution is considered better than another only if it is better with respect to all the objectives. In this case, the better solution is assigned a superior rank and a higher fitness for selection. On the other hand, if one solution is better than another with respect to one objective, while the second is better with respect to a different objective, then the two solutions are considered nondominated or equivalent, and are called pareto solutions. Pareto solutions are assigned equal rank and fitness in the algorithm. Thus, if $\theta_{1,i} < \theta_{1,j}$ and $\theta_{2,i} < \theta_{2,j}$, where i and j represent two members of the population, then member i is superior to member j, and is assigned a better rank. If $\theta_{1,i} < \theta_{1,j}$, but $\theta_{2,i} > \theta_{2,j}$, then i and j are nondominated or noninferior or pareto solutions.

Figure 2 shows a block diagram of the NSGA algorithm employed in the study. As stated above, the algorithm examines the members of a population as they progress from generation to generation. Each generation is composed of N_{pop} members. The members of the first generation are obtained by a random sampling of the solution space. The members are characterized by the decision variables, which are the feedrates of the reagents at the various intervals of the batch (discretized as zero-order holds). The batch is divided into 11 intervals spanning 11 min each (determined by the frequency of the PSD measurements). These decision variables are represented as binary strings of length l_{chrome}, called the chromosomes. The chromosomal information is translated into actual feedrate values using the following relationship

$$\frac{u_i - u_{\min}}{u_{\max} - u_{\min}} = \frac{\sum_{k=1}^{l_{\text{chrome}}} u_{bk} 2^{k-1}}{\sum_{k=1}^{l_{\text{chrome}}} 2^{k-1}} \tag{4}$$

where u_i is the reagent feedrate in interval i, $u_{\min}$ and $u_{\max}$ account for the range of values for the input, and u_{bk} is the

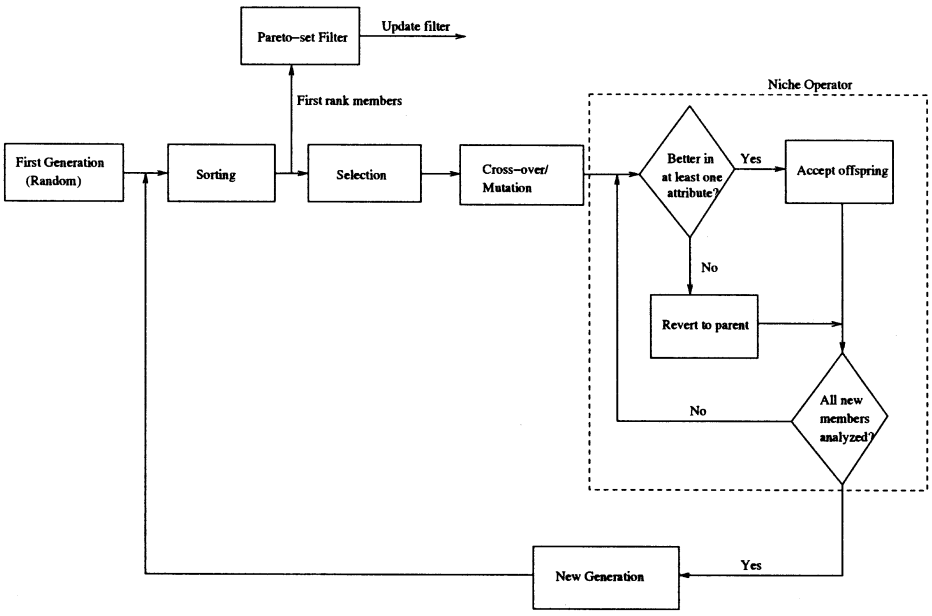

Figure 2. NSGA algorithm.

The algorithm works with a population of candidate solutions, and enables an iterative improvement in the solutions by combining different solutions (cross-over) and altering some of them (mutation).

binary value of the string at position k. The random members of the first generation are evaluated by solving the underlying first-principles model, and determining the values of the defined objective functions. The model comprises the population balance representing the PSD (Eq. 5), material balances for the various reagents and species, and partitioning calculations for surfactants and monomers

$$\frac{\partial}{\partial t}F(r,t)+\frac{\partial}{\partial r}\left(F(r,t)\frac{dr}{dt}\right)=\Re_{\text{nuc}}\delta(r-r_{\text{nuc}})+\Re_{\text{coag}} \quad (5)$$

In the above equation, $F(r,t)$ is the particle density function defined such that $F(r,t)dr$ is the moles of particles of the size between r and $r+dr$. $\Re_{\text{nuc}}$ represents the rate of nucleation and $\delta(r-r_{\text{nuc}})$ is the dirac delta function, which is unity at $r=r_{\text{nuc}}$ and zero elsewhere, assuming that the nucleation occurs at a fixed minimum size r_{nuc}. $\Re_{\text{coag}}$ incorporates terms accounting for the loss and gain of particles of size between r and $r+dr$ due to coagulation. The partial derivative with respect to r in the lefthand side Eq. 5 represents the growth event, where dr/dt is the size-dependent growth rate.

Upon determining the objective function values of each member, the members are sorted into pareto sets using the nondomination concept described above, and are assigned ranks (the box labeled "sorting" in the block diagram—Figure 2). All members having the same rank are assigned equal fitness, defined as follows (Silva and Biscaia, Jr., 2001)

$$fit_i=\frac{N_r-k+1}{SS} \quad (6)$$

where

$$SS=\frac{1}{N_{\text{pop}}}\sum_{k=1}^{N_r}(N_r-k+1)P_{sk} \quad (7)$$

In the above equations, N_{pop} is the size of the population, N_r is the total number of ranks in the generation (or the highest rank), P_{sk} is the size of rank k (the number of members assigned that rank), and fit_i is the fitness of any member i which is assigned a rank k.

The members are reproduced into the next generation in proportion to their fitness values, emulating the concept of the survival of the fittest ("Selection" box in the block diagram). The "pareto-set filter" block is introduced, and all members assigned the first rank are stored in this filter. In the new generation, some of the members are chosen at random and exposed to cross-over and/or mutation operations, subject to certain probabilities p_{cross} and p_{mut}, respectively. In effecting a cross-over, a certain location of the binary string is determined at random, and each of the binary strings of the two members are swapped after this location. In effecting the mutation operation, one of the chromosomes characterizing the member (representing the feedrate of a particular reagent in a particular interval) is chosen at random, and the mutation location of the string is also determined at random. The binary value that represents the mutation location is flipped from 1 to 0 or vice versa. The members of the new generation are evaluated. If a new member is better than its parent in at least one objective, the new member is accepted into the current generation. If not, the parent member is sent into the current generation. This concept is called the Niche operator, introduced in the NSGA technique (for example, Silva and Biscaia, Jr., 2001) to avoid a genetic drift. The members of the current generation are sorted and ranked. The members assigned rank 1 in the current generation are appended to those already present in the pareto-set filter. The members in the filter themselves are ranked, and only the first ranked members are retained. The current genera-

tion is again subjected to selection, and the entire operation is repeated for a certain number of generations. At the end of any generation, the pareto-set filter has the current pareto-optimal solutions. These solutions form a classical pareto tradeoff curve.

Two-objective formulation — coagulation-free model

In this section, a two-objective optimization problem is considered (based on the objectives pertaining to the profiles of total particles θ_1 and solids content θ_2). A detailed population balance model is utilized, which assumes that the coagulation events are negligible (Immanuel et al., 2002). Orthogonal collocation on finite elements is used to discretize the population balance equation (a hyperbolic partial differential equation) (Eq. 5). Forty finite elements of varying width are used to span the domain of particle sizes, with 5 collocation points per element. This results in a system comprising 161 differential algebraic equations (DAEs) (see Immanuel et al. (2002) for details on the discretization). These DAEs are combined with other DAEs representing material balances and partitioning equations, forming a complete system of 330 DAEs. This complete system is solved using the FORTRAN-based DAE-solver DDASSL. The decision variables for the optimizer are the feedrates of surfactant solution and VAc monomer. These are discretized as zero-order holds (11 intervals spanning 11 min each), as mentioned previously. No surfactant or VAc feed is allowed towards the end of the batch, and the target profile of solids content reflects the necessity of diminishing growth rates towards the end of the batch, to prevent the distribution from evolving beyond the end-point. Specifically, the feedrate of the surfactant solution (22.7% by weight aqueous solution) in the first five intervals and the VAc monomer in the first seven intervals are the decision variables. The feed rates of these reagents in the other intervals are set equal to zero. The NSGA parameters employed are $N_{pop} = 25$, $l_{chrome} = 24$, $p_{cross} = 0.7$, and $p_{mut} = 0.001$. Figure 3 shows the sorting of some of the members corresponding to the first generation. The solution #1 is better than the rest with respect to both the objectives, and is assigned rank 1. The solutions marked #2 are paretos (nondominated) to each other. Hence, these are both assigned rank 2, and so on. Solution #1 is stored in the pareto-set filter after the first generation. The operations of selection, cross-over and mutation, and the updating of the pareto-set filter are repeated generation after generation. The inset in Figure 3 shows the pareto-set filter at the end of five generations. It has four solutions, each of which have bettered solution #1 from the first generation. Table 1 lists the values of the target variables corresponding to these four pareto-optimal solutions. It can be seen that pareto #1 is the *best* with respect to the tracking of the profile of total particles (θ_1), while pareto #4 is the *best* with respect to the tracking of the profile of the solids content (θ_2). This table also provides the value of $\theta_3(t_f)$ (end-point PSD tracking error) corresponding to each of these pareto solutions.

Table 1. Objective Function Values Corresponding to the Four Pareto-Optimal Solutions

Pareto #	θ_1	θ_2	$\theta_3\ (t_f)$
1	1.3428	0.0325	5.63×10^{-4}
2	1.3852	0.0311	6.97×10^{-4}
3	2.1088	0.0296	1.5×10^{-3}
4	3.4663	0.0085	1.45×10^{-4}

One can see from this table, and also from Figure 4, that pareto #4 gives the best match to the end-point PSD. The decrease in the value of θ_2 from pareto #1 to pareto #3 is more than countered by the increase in the value of θ_1, resulting in the degradation of the match in the end-point PSD $\theta_3(t_f)$. However, the drop in θ_2 between paretos #3 and #4 results in a substantially better match of the end-point PSD in spite of the increase in θ_1. The nonuniform trend in the value of θ_3 suggests regions in the pareto-curve where a different objective is dominant, indicative of a multiobjective problem. Figure 5 shows a comparison of the feed profiles of

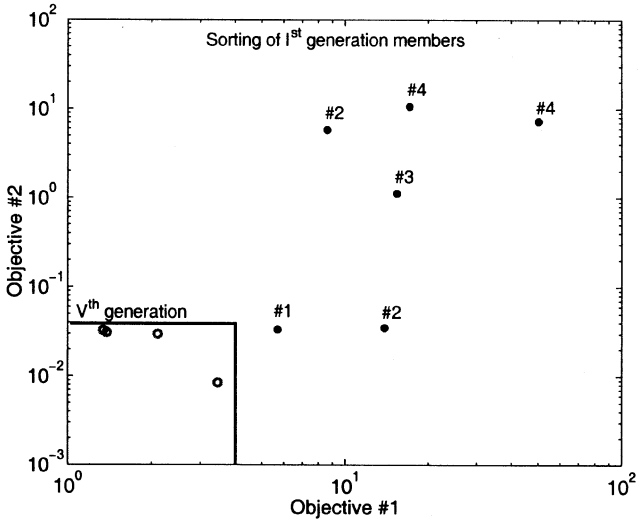

Figure 3. Evolution of the pareto-optimal solutions from generation I to generation V.

The inset shows the pareto-set filter corresponding to the fifth generation.

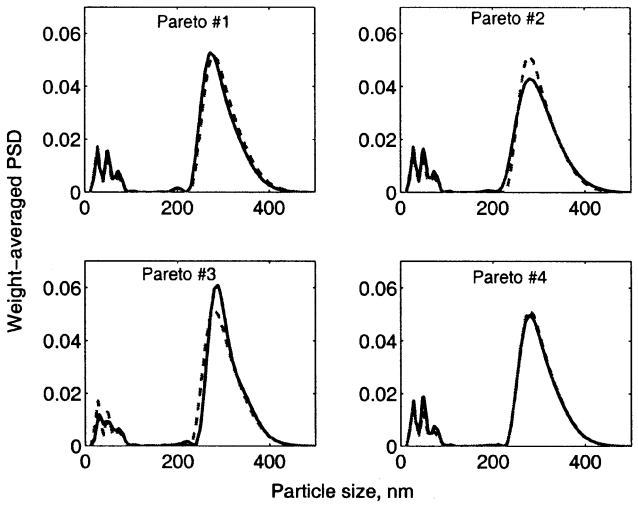

Figure 4. Performance of the pareto-optimal solutions obtained at the end of five generations, with respect to the end-point PSD target.

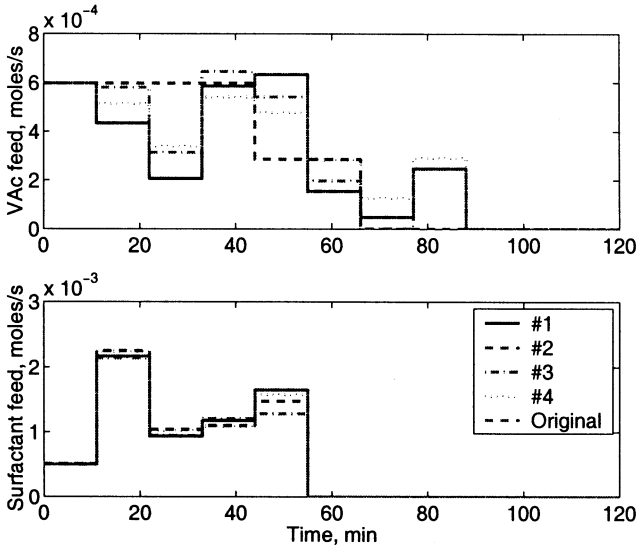

Figure 5. Comparison of the optimal feed profiles characterizing the four pareto solutions, with the target (original) recipe.

the surfactant solution and VAc monomer corresponding to the four nondominated solutions with the actual feed profiles that were originally used to generate the target distribution.

Among the four pareto-optimal solutions obtained from the NSGA, one can choose the solution closest to the utopia point for implementation in the process. The utopia point is the ideal point with respect to the objectives, and in this case corresponds to $\theta_1 = 1.3428 = {}^{\min}_{i}\theta_{1,i}$ and $\theta_2 = 0.0085 = i^{\min}\theta_{2,i}$, where i refers to the members in the pareto-set filter (see Table 1). Alternatively, one can utilize other process information to discriminate among the pareto solutions. For instance, one can use the value of $\theta_3(t_f)$ corresponding to these solutions to choose the solution to be implemented. In light of the values of $\theta_3(t_f)$ for the four solutions in Table 1, the user would likely choose pareto #4 for implementation.

The pareto solution #4 was implemented in an experimental reactor system. Figure 6 shows the evolution of the PSD over the course of the batch corresponding to this recipe. Figure 7 shows a comparison of the estimates of the profiles of total particles and solids content from the experimental

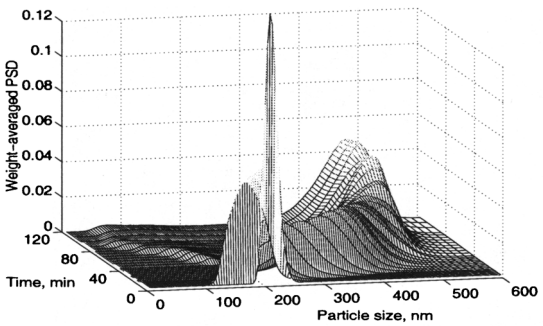

Figure 6. Evolution of the (weight-averaged) PSD in the experimental implementation of pareto #4.

Input profiles are those depicted in Figure 5.

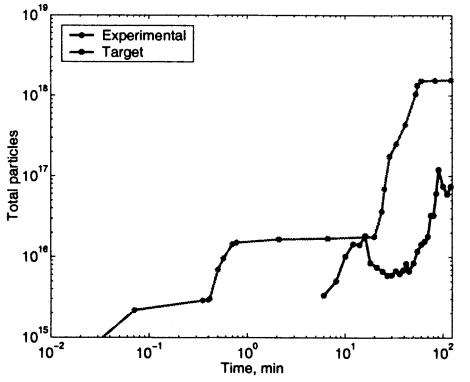

(a) comparison of the profile of total particles

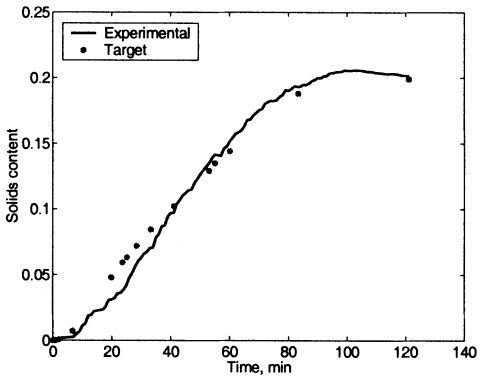

(b) comparison of the profile of solids content

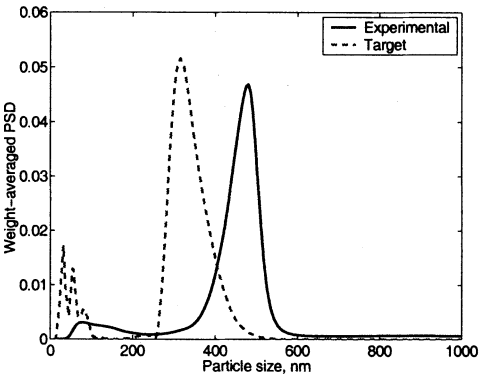

(c) comparison of the end-point PSD

Figure 7. Comparison of the experimental results observed on implementing pareto #4 (from Table 1—Two-objective formulation) with the target.

data, with the target profiles. It also depicts the end-point PSD obtained experimentally and the target end-point PSD. The experimental data for the total particles (Figure 7a) shows that the primary nucleation event is quite delayed (a longer induction period), and also shows the presence of a coagulation event, that result in a lower number of particles at early times. This, in turn, results in larger values of the average number of radicals/particle and monomer concentrations inside the particles, resulting in larger growth rates. Hence, the

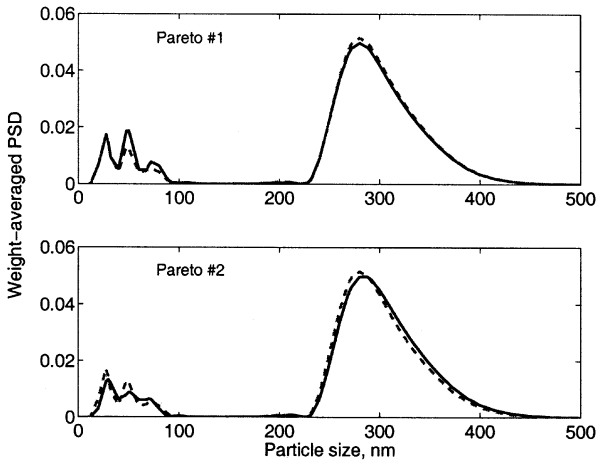

Figure 8. Performance of the pareto-optimal solutions obtained at the end of five generations in the three-objective problem, with respect to the end-point PSD target.

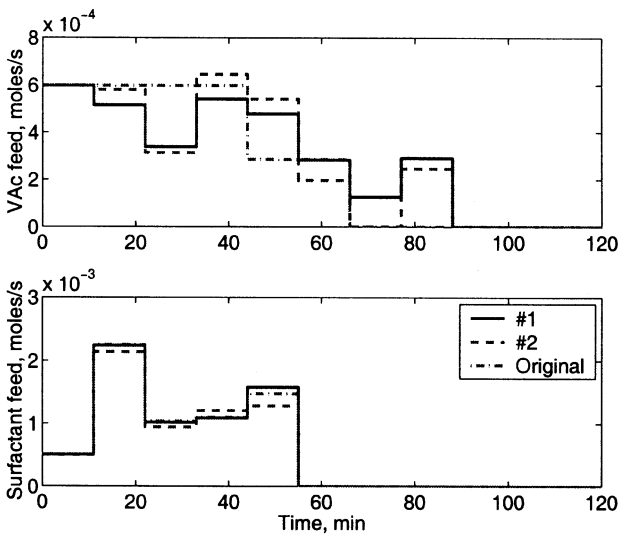

Figure 9. Comparison of the optimal feed profiles characterizing the two pareto solutions with the target recipe—three-objective formulation.

difference in the size of the larger mode, relative to the target (Figure 7c).

Three-objective formulation — coagulation-free model

In the previous results, the profiles of nucleation and growth rates that characterize the desired PSD trajectory were re-cast as equivalent profiles of total particles and solids content (exploiting the idea of the control of instantaneous properties). Viewing this approach from a different perspective, total particles and solids content are related to the moments of the distribution. Consequently, tracking the profiles of total particles and solids content might not guarantee the attainment of a desired multimodal distribution with complex shapes. In light of this, the third objective on the end-point weight-averaged PSD ($\theta_3(t_f)$) is explicitly included into the formulation, and the complete three-objective formulation implicit in Figure 1 is solved. The nonuniformity in the value of $\theta_3(t_f)$ in the pareto solutions (Table 1) with respect to θ_1 and θ_2 also supports this formulation. Figures 8 and 9 show the results pertaining to the three-objective optimization problem. In this case, the algorithm generates two nondominated solutions (after five generations), one of which is the same as that obtained in the earlier two-objective formulation (pareto #4 in Table 1). Table 2 lists the objective function values corresponding to these two solutions. For implementation purpose, one could select the solution which gives the maximum surface coverage to the particles with surfactants—based on the hypothesis that larger surface coverage leads to lower coagulation events. Alternatively, the user could employ other process information (not built into the mathematical framework) in deciding the which solution to implement.

Convergence tests

It is well known that, in most cases, a Genetic Algorithm usually produces a solution in the vicinity of the global optimum. This solution can be improved using a local optimization technique. Alternatively, one has to employ a very large number of generations to obtain the optimum. In the case of a multiobjective optimization, the solution(s) obtained (either a pareto-optimal set of nondominated solutions or a unique solution) are not necessarily optimal in a formal sense, but are the *best* tradeoff solutions. The actual optimal point with respect to the different objectives (considered separately) is called a *utopia* point, to reflect its ideality. This is the implicit tradeoff in optimizing multiple objectives separately. However, in the current case, *the solutions obtained are very close to the targets*, leaving little room for further refinement. Also, previous studies with GA, aimed at optimizing a single target, show that the solutions converge (to within an accuracy dictated by the choice of objective function) in about *three* generations (of a total of *eight* generations) (Immanuel and Doyle III, 2002). Thus, in this NSGA application, the search was restricted to 5 generations. However, two measures were undertaken to confirm that the solutions had converged to the "optimal" ones. The first was to run the algorithm for a larger number of generations, specifically 10. Another was to parameterize the input space, and perform a systematic search (contrasted with the facilitated random search that GA signifies)—by solving a combinatorial problem. In these cases, different solutions were obtained, but these were paretos to the ones found already. Table 3 lists all the solutions from the different strategies. In summary, five generations were found to be adequate to develop the optimal pareto curve.

Three-objective formulation — coagulation-inclusive model

In the previous studies, a coagulation-free model was utilized in performing the optimization. Coagulation events were

Table 2. Objective Function Values Corresponding to the Two Pareto-Optimal Solutions—Three-Objective Formulation

Pareto #	θ_1	θ_2	θ_3
1	3.4663	0.0085	1.45×10^{-4}
2	2.3652	0.0279	2.94×10^{-4}

Table 3. Objective Function Values Corresponding to the Pareto-Optimal Solutions Obtained from Different Strategies

Pareto #	θ_1	θ_2	$\theta_3(t_f)$	Technique
1a	3.4663	0.0085	1.45×10^{-4}	NSGA—
2a	2.3652	0.0279	2.94×10^{-4}	5 Generations
1b	2.0651	0.02758	4.28×10^{-4}	
2b	2.05693	0.02738	5.48×10^{-4}	NSGA—
3b	2.0096	0.02789	3.37×10^{-4}	10 Generations
4b	1.11708	0.03884	3.11×10^{-3}	
1c	9.9468	0.0238	1.22×10^{-4}	Systematic
2c	2.3881	0.0140	4.73×10^{-4}	search

minimized by setting constraints on the surface coverage of the particles with surfactants. Such an approach was required because of the computational burden associated with the complete coagulation-inclusive model. This strategy was revised due to the excessive coagulation events in the experimental implementation of the optimal recipe, both in the current study (see Figure 7a), and in the earlier study based on single objective optimization (Immanuel and Doyle III, 2002). A new solution technique has been developed (Immanuel and Doyle III, 2003b), which enables very efficient computation. Thus, the coagulation-inclusive model presented in Immanuel et al. (2003a) is combined with the complete three-objective strategy. A population size $N_{pop} = 25$ is adopted, and the algorithm is executed for 24 generations. The algorithm produces two pareto-optimal solutions (shown in Table 4), which track the target profiles of total particles and solids content closely, and also attain the target end-point PSD. One of these solutions is shown in Figure 10. The input profiles that are determined by the algorithm match the target input profiles rather closely. Experimental implementation of this recipe is shown in Figure 11. Clearly, there is no evidence of a strong coagulation event in the experimental distribution (Figure 11c). However, there is a discrepancy between the target distribution and the experimental distribution. This is attributed to model and parametric uncertainties, and input disturbances.

Feedback Analysis

The process of evolution of PSD in emulsion polymerization is quite complex, comprising the phenomena of nucleation, growth and coagulation. The mechanisms that characterize each of these sub-processes in turn are complex and not completely well characterized. As newer processes are developed to address environmental impacts or to produce better products, they add to the complications and to the unknown mechanisms in the model. For example, the use of ionic surfactants results in undesired properties in the final products relative to their water-resistance characteristics. Thus, the ionic surfactants are being replaced with nonionic ones. However, the latter result in considerable alterations in the nucleation and coagulation phenomena (and indirectly on the growth phenomenon) due to their tendency to partition into the dispersed phases. Since these aspects and mechanisms are still subjects of active research, the model developed at this point has associated uncertainties (Immanuel et al., 2002, 2003a). There are also a number of parameters in the model that are uncertain.

Although the experimental implementation of the open-loop optimized recipe produced encouraging results, there is

Table 4. Objective Function Values Corresponding to the Two-Pareto-Optimal Solutions—Three-Objective Formulation, Using a Coagulation-Inclusive Model

Pareto #	θ_1	θ_2	θ_3
1	7.4774	0.0057	1.122×10^{-6}
2	7.5131	0.0056	1.069×10^{-6}

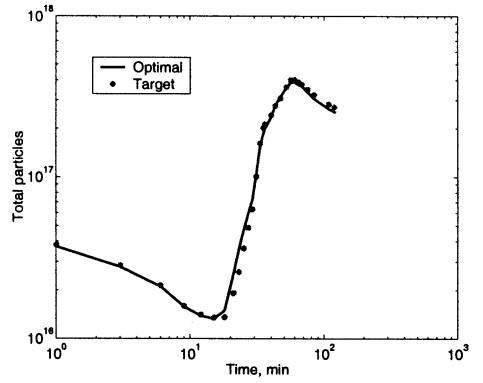

(a) profile of total particles

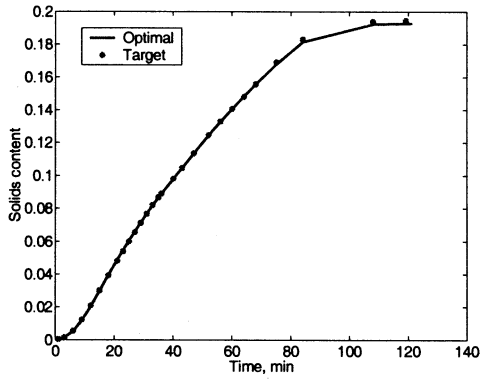

(b) profile of solids content

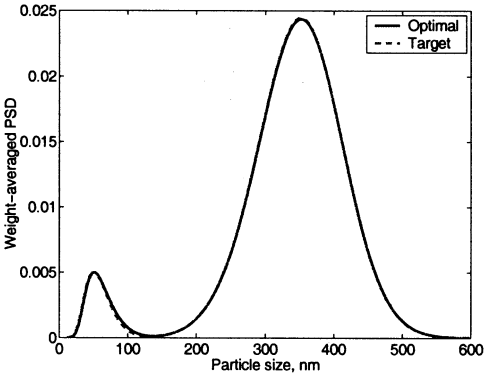

(c) end-point PSD

Figure 10. Comparison of NSGA-generated "optimal" solution with the target, utilizing a coagulation-inclusive model to perform the optimization.

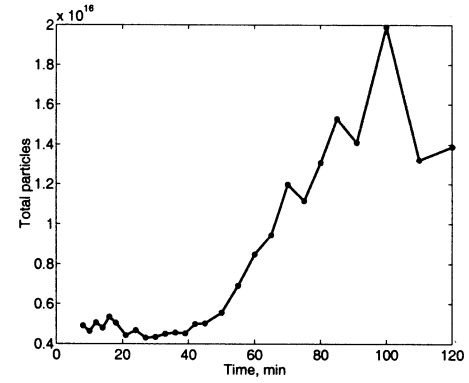

(a) profile of total particles

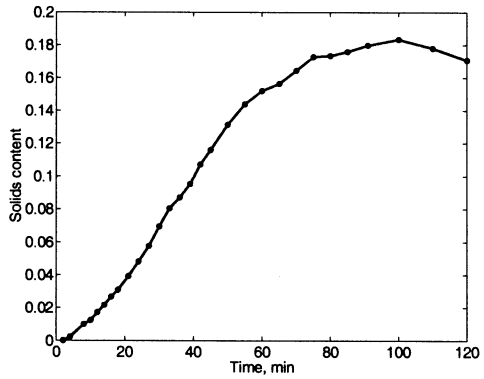

(b) profile of solids content

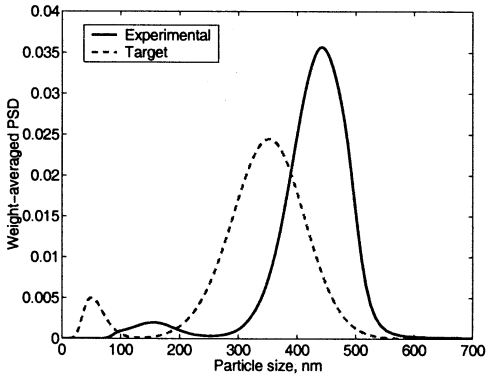

(c) end-point PSD

Figure 11. Experimental implementation of the open-loop optimized recipe corresponding to the solution shown in Figure 10, and comparison of the obtained end-point PSD with the target.

a clear need for improvement through feedback. However, there are also indications that in-batch feedback may not be effective to account for disturbances in semi-batch emulsion polymerization (Immanuel and Doyle III, 2003a; Wang et al., 2002). For example, in the experimental implementation presented in Figure 7, the primary nucleation event is delayed considerably, and also there is a coagulation event. The resultant decrease in the number of particles caused the primary mode to grow to a much larger size than desired. This is due to the interactive effect of the nucleation phenomenon (decrease in the number of particles) on growth. To correct this error (in a relative sense—and, hence, preserve the shape of the distributions), one might need an earlier secondary nucleation event (of a reduced magnitude) and also a drop in the growth rate. These highlight an opportunity for improvement with feedback, subject to measurement limitations. Feedback information on the total particles (or the free surfactant concentration relative to the critical micelle concentration) could enable rectification of these errors. However, it depends on whether these states are observable without the PSD measurements, as the very first PSD measurement is not available until at least 12 min into the batch. In general, the PSD measurements are both sparse and delayed, which could potentially make them of a reduced utility to correct the rapid and irreversible processes. Also, there are cases in which a strong coagulation event could occur, leaving an irreparable impact on the distribution (see Immanuel and Doyle III, 2002) for an example in which coagulation causes a substantial alteration on the distribution). This is mainly a concern due to the randomness associated with shear-induced coagulation.

The previously described attributes challenge the application of control for this process. Generally, there are two feedback corrective measures that have been proposed—in-batch feedback (Flores-Cerrillo and MacGregor, 2002), and batch-to-batch feedback (Crowley et al., 2001). The in-batch feedback approach presented by Flores-Cerrillo and MacGregor (2002) is based on a single point re-computation. This requires a determination of the most effective time to perform the re-optimization. As described in the case of Figure 7 (reduced nucleation at early times), the cascaded effects of the errors motivate early corrective action. On the other hand, corrective action taken based on too few measurements may lead to an erroneous action influenced by faulty measurements. Hence, it is of interest to identify the most opportune time to perform the first corrective action. Also, a single recomputation may not be effective, and one might need a full nonlinear model predictive control (MPC) formulation in a receding horizon framework to correct the errors. There are several efficient solution techniques that have been developed (Mahoney and Ramkrishna, 2002), which can render a full nonlinear MPC feasible. However, it is important to determine whether physical limitations can be overcome with on-line feedback subject to the process constraints. If these strategies prove ineffective, it is better to employ the batch-to-batch feedback strategy presented by Crowley et al. (2001), with the possibility of an improvement in subsequent batches.

The present study evaluates these tradeoffs with a two-fold objective. One is the evaluation of different state estimation strategies to determine the configuration that leads to a reasonable convergence of the estimates in the available time. The second objective is to employ these estimation techniques to address the important question of the effectiveness of feedback. This study is carried out off-line using actual experimental data.

Biased state estimation strategy

As mentioned earlier, successful feedback measures rely on the availability of successful and robust estimation techniques

—which take the measurements from the process and reconcile them with the model to estimate the actual process states. The on-line measurements available to monitor the process of the evolution of PSD in the system considered here are (Immanuel et al., 2003b):

(1) The relative PSD measurements, available typically 11–12 min from a Capillary Hydro-Dynamic Fractionator (CHDF).

(2) Latex density measurements from a densitometer.

(3) Feedrate measurements (from flowmeters and load cells).

In typical MPC applications using linear models on continuous processes, a classical approach to correcting for disturbances and model mismatch is to bias the outputs against the measurements, and hold the bias constant for the rest of the prediction horizon (Garcia and Morari, 1982; Garcia and Morshedi, 1986). A similar idea can be used in this semi-batch process in a receding horizon framework.

This approach is an extension of the strategy employed in Immanuel et al. (2003b). In that case, estimation was accomplished using sensor measurements combined with process calculations performed using simple steady-state models. It was seen that latex density and feedrate values could be used to infer the conversion of the monomers and the solids content. Combining this information with the relative PSD measurements from the CHDF, the absolute PSD can be obtained, and all the particle-related states including the total particles can be calculated. A major implication of this study is the ability to observe all important process states. Dynamic state estimation is possible if the above estimation technique is combined with a dynamic model. The dynamic model presented in Eq. 5 has the following generic structure

$$\begin{aligned} \dot{x} &= f(x, p, y, z, u) \\ y &= g(x, p) \\ z &= h(y) \end{aligned} \tag{8}$$

The above model is obtained by a model reduction strategy applied to a finite-dimensional approximation of the infinite-dimensional partial differential equation model. Here, the vectors y and z represent the inertial manifold or the pseudo-steady-state variables. A sequential solution technique is employed. The vector x represents the states of the system, p represents the parameters, and u represents the inputs. The model reduction itself is obtained from process considerations and the underlying mechanisms. The vector z comprises the PSD at discrete points and other outputs such as the total particles and solids content. The vector y comprises the faster modes such as the average number of radicals/particle corresponding to the discrete points of the PSD, and the partitioning equations. They also comprise the individual rates of nucleation, growth, and coagulation. The states x comprise the slower modes, mainly corresponding to the material balance equations. Thus, this model essentially has a similar hierarchical strategy to the one proposed in this article for open-loop recipe optimization (Figure 1). It produces solutions that closely match the rigorous solutions obtained from solving a finite-dimensional system resulting from orthogonal collocation on finite elements-based discretization of the population balance equation. A significant advantage is that it results in very efficient computation. Additional details on the solution technique can be found in Immanuel and Doyle III (2003b). At each discrete time instant k when measurements (including the PSD) are available, the process calculations are performed (Immanuel et al., 2003b), and the variables in z are forced to the calculated values $z_{\text{calc},k}$ (akin to an output bias estimate (Garcia and Morari, 1982; Garcia and Morshedi, 1986). The solution is then continued in a receding horizon implementation. Thus, the dynamic model takes the following generic form

$$\begin{aligned} \dot{x} &= f(x, p, y, z, u) \\ y &= g(x, p) \\ z &= h(y) \quad \text{or} \quad z = z_{\text{calc},k} \end{aligned} \tag{9}$$

Figure 12 compares the biased estimation results with the open-loop data. As seen in Figure 12b, at early times, the biasing results in an even larger error in the nucleation rate and the number of particles compared to the open-loop simulations. The very first data point corresponding to 6 min shows 3.8×10^{14} particles—not seen in the plot in Figure 12b—that is considerably lower than the predicted value. Biasing against this value results in raising the predicted value to $\sim10^{17}$ at this time compared to $\sim1.5\times10^{16}$ in the open-loop case.

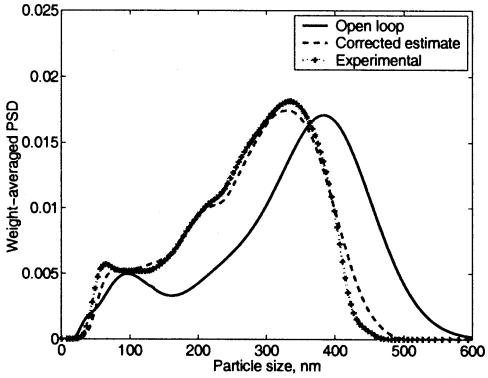

(a) end-point distribution

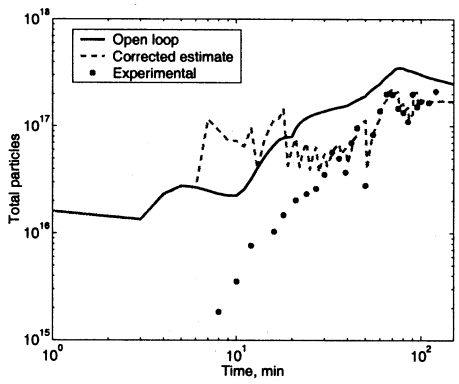

(b) profile of total particles

Figure 12. Biased estimation strategy without parameter updating, showing the convergence of the results beyond 30 min.

However, after about 30 min, the predicted profile closely follows the observed data (and the estimated end-point PSD and the experimental PSD match closely—Figure 12a).

Re-Optimization of Inputs Based on Biased Estimation. The experimental results from the open-loop optimized recipe using the coagulation-inclusive model (Figure 11) are utilized for this study. Note that more frequent PSD data–than the typical 11 min—are utilized at the early times, based on off-line sampling. The very first successful measurement in this experiment corresponds to the sample taken at 8 min, which would be available at approximately 19–20 min (due to delay in measurements). Earlier measurements are not possible owing to the low particle concentrations (below 0.5%) that cannot be detected by the instrument (Immanuel et al., 2003b). The model prediction is biased against the measured/calculated value at 8 min ($z = z_{\text{calc}}$) and the new end-point PSD, and the new trajectories of total particles and solids content are estimated. Based on this estimate, the receding horizon problem is re-solved to find the inputs that would drive the solution back to the target trajectory. The inputs from the third interval forward (22 min) are re-optimized. In general, for batch processes, one needs to delay the re-optimization until a sufficient number of measurements are available, to obtain a good estimation of the process disturbances and the future course of the controlled variables. This aspect is particularly critical in complex processes such as the one under study here, which is faced with irreversible effects and very intricate and often ill-categorized mechanisms. We first demonstrate the problem associated with performing the re-optimization with too few data points, and, hence, with erroneous estimates of process trends. Figure 13 shows a biased estimate of the end-point distribution based on the 8-min sample, which shows much larger size particles in the larger (primary) mode relative to the target, and a much larger ratio of the number of particles in the smaller mode to that in the larger mode, relative to the target. This is mainly due to an estimate of a reduced nucleation rate at early times, which causes larger growth rates at these early times. It also causes an earlier and larger secondary nucleation event.

In the face of the error in the nucleation rate that has occurred, the end-point distribution targeting suggests that one should not follow the original trajectory of total particles and solids content, which represent a particular nucleation and growth rate trajectory. Instead, these trajectories need to be re-computed. For example, the decrease in the particles nucleated in the primary mode would necessitate a decrease in the particles nucleated in the subsequent nucleation event to preserve the relative magnitude of the distributions. In addition, it might necessitate an earlier secondary nucleation event, for instance, to preserve the distance between the modes. Thus, the simplest and most straightforward method to perform on-line optimization is to solve a single objective optimization problem, based only on the end-point distribution (on a weight-averaged basis—as defined in $\theta_3(t_f)$). In this case, there is no explicit target for the profiles of total particles and solids content. The single-objective algorithm presented in the previous study (Immanuel and Doyle III, 2002) is utilized for this purpose. Since the first PSD measurement is available at about 19 min, the surfactant feed in intervals 3–5 min, and the feed of both the monomers (VAc & BuA) in intervals 3–7 min are re-computed. Note that it is

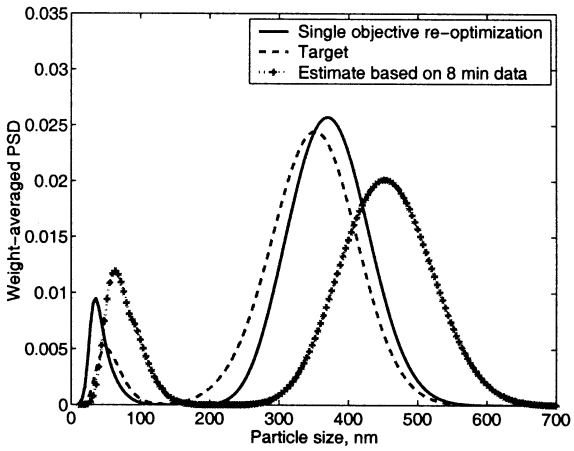

Figure 13. End-point PSD obtained with re-optimized inputs, employing a single-objective re-optimization, using a single PSD measurement based on a sample taken at 8 min into the batch.

essential to include the feed rate of BuA as a decision variable. Figure 14 presents the re-optimized input profiles, and Figure 13 depicts the corresponding end-point PSD. It is seen that the size of the particles are reduced considerably, and the larger mode is brought closer to the corresponding mode in the target. Also, the secondary nucleation event is reduced, and the secondary mode is brought closer to that in the target. However, the two modes are farther apart from each other in the solution than in the target. This implies that, to correct for the reduced primary nucleation event, the secondary nucleation event must have started before 15 min, or the growth rate must be reduced at very early times (recall that the first PSD data are not available until 19 min). The re-optimized solution has a significantly reduced solids content ($\sim$10%) at the end of the batch. Further receding horizon optimization as the batch proceeds (a full nonlinear MPC) could potentially improve upon these results, as more data become available and more information about the experiment can be inferred.

In the above cases, the simple biased-estimation strategy (analogous to a constant disturbance) was used. This predicts a secondary nucleation event that starts by approximately 15 min, and also is stronger than in the nominal case. However, a closer book at Figure 11a reveals that the secondary nucleation event does not happen until about 40 min in the actual case. This clearly highlights a *shortcoming of basing the feedback calculation on too few measurements*. This also suggests the ineffectiveness of the biased estimate in most cases, particularly without a parameter update. To elaborate upon the latter aspect first, the model predicts a larger primary nucleation event compared to the actual experimental occurrence. This is probably due to the errors in the surfactant-partitioning parameters and the critical micelle concentration (cmc) value (discounting the effect of inhibitors such as residual dissolved oxygen in the initial mixture—which might occur in spite of the purging of the reaction mixture with nitrogen gas). For example, either the cmc value used in the model is lower than the actual value, or the partitioning parameters

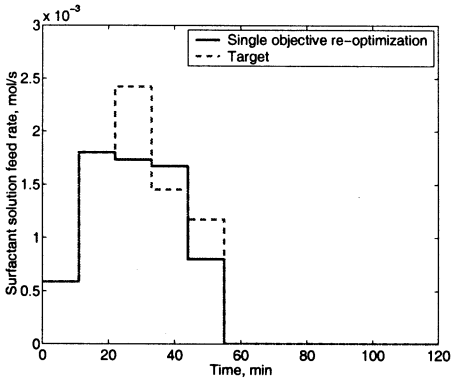

(a) re-optimized profile of surfactant feed rate

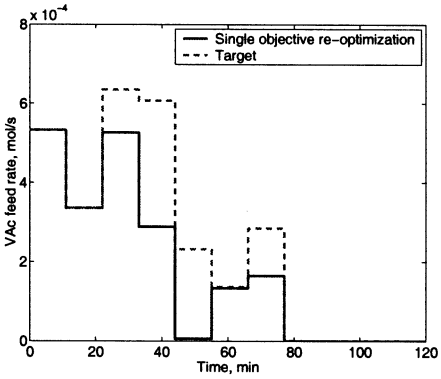

(b) re-optimized profile of vinyl acetate feed rate

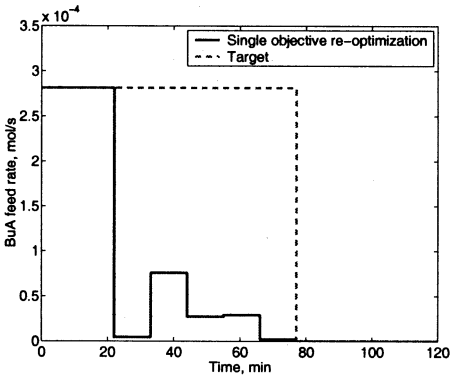

(c) re-optimized profile of butyl acrylate feed rate

Figure 14. Re-optimized inputs corresponding to the results in Figure 13.

are such that they predict a larger free surfactant concentration in the aqueous phase. One or both of these uncertainties result in predicting a larger concentration of the micelles, and a larger micellar nucleation rate. Now, when the distribution is updated at 8 min based on the measured value, the total particles and the total surface area are reduced (to the values calculated from the measurements). This results in an even larger free surfactant concentration than in the nominal case, and, hence, the prediction of an earlier and larger secondary nucleation event. This clearly demonstrates the need for a parameter update in addition to output biasing. There are situations in which a pure biasing of the outputs alone (without parameter updating) is effective in driving the simulations to the experimental results (Figure 12 presented above); however, one needs to employ enough data points. Since the biasing method is not consistent in all cases, a more formal estimation strategy is considered next.

Luenberger observer

There are several studies reported in the literature addressing the issue of state estimation for semi-batch polymerization processes including emulsion polymerization, using lumped parameter models. These include the works of Dimitratos et al. (1989, 1991), Kozub and MacGregor (1992b), Liotta et al. (1997), and Astorga et al. (2002). These studies used an extended Kalman filter with parameter augmentation of the states. Kozub and MacGregor (1992b) suggested the need to update the initial condition in addition to the parameters (reiterative extended Kalman filter), along the lines of the proposition in Jazwinsky (1970). They suggest a full nonlinear optimization strategy utilizing all the data points (available up to the current time), to determine the initial condition. In contrast to the rich literature dealing with lumped parameter models, there are no state estimation studies reported using models dealing with the full PSD in emulsion polymerization.

In the present work, for the purposes of observer design, the model is augmented with the actual outputs (measurements) $w = q(z)$, where w represents the relative (weight-averaged) PSD discretized at different points in the size domain, and lumped variables such as total particles and solids content. The auxiliary variables z (250×1) account for the absolute PSD at more refined intervals than the outputs w (25×1), based on the finite-element discretization chosen (Immanuel and Doyle III, 2003b). Thus, the vector z gives the actual absolute distribution, while the outputs w represent the actual measurements. Three of the states x (M_1 and M_2—the two monomers, and R_w—the aqueous phase initiator-derived radicals) are chosen for update in the estimator framework, and the rest are allowed to evolve in open loop. The disturbances and noise are not incorporated at present. To render the estimator robust in all cases, a combined state-parameter estimator is configured. A discrete (static) parameter update is performed for the cmc, Γ_∞ (one of the surfactant-partitioning parameter), and k_{p11} (propagation rate constant for polymer of type 1—with a VAc end group, with monomer 1—VAc). The structure of the continuous-discrete observer is given by (Astorga et al., 2002)

$$\begin{aligned}
\dot{\hat{x}} &= f(\hat{x}, \hat{p}, \hat{y}, \hat{z}, u) + L_i(w_k - q(\hat{z})) \\
\hat{y} &= g(\hat{x}, \hat{p}) \\
\hat{z} &= h(\hat{y}) \\
\dot{\hat{p}} &= 0 \\
\hat{p}^-_{k+1} &= \hat{p}^+_k \\
\hat{p}^+_{k+1} &= \hat{p}^-_{k+1} + L^p(w_{k+1} - q(\hat{z}))
\end{aligned} \tag{10}$$

(The character $\hat{}$ denotes estimates of the corresponding variable.) Here, L_i (over a time interval i) is a Luenberger gain, designed such that the eigenvalues of $(A - L_i CD)$ lie in the LHP. Here, A is the Jacobian of $f(x, p, y, z, u)$ at $\tilde{x}$, C is the Jacobian of $h(g(x, p))$ at $\tilde{x}$, and D is the Jacobian of $q(z)$ at $z = h(g(\tilde{x}, p))$, with $\tilde{x}$ being a representative point in the time interval i. w_k is the measured output value at time instant k (held constant until another measurement is available). The states are corrected using combinations of the error in the weight-averaged PSD at different discrete points of the distribution. Accordingly, the three non-zero rows of the matrices L_i (corresponding to the states M_1, M_2 and R_w) have non-zero elements at different points to catch the error in the PSD at different discrete points. Although the tuning of L_i is done using the Luenberger concept, the estimator has a Kalman-like structure. The gain matrices L_i are calculated at the start in this off-line study, such that the convergence criterion is satisfied along the entire batch, by dividing the entire batch into intervals spanning 10 min each. L^p is the fixed gain used to update the parameters. The superscript $^-$ denotes the value of the parameter before updating —an infinitesimal time instant before the measurement, and the superscript $^+$ denotes the value after updating using the current measurement. Essentially, the parameters are updated at each time instant a measurement is available, and are held constant until the next measurement. The error in the lumped outputs—total particles, and solids content (which represent moments of the full distribution—are used to update the parameters, that is, only those corresponding rows of w are utilized.

Figure 15 illustrates the estimated profiles of total particles and solids content, and the end-point PSD, with the experimental observations. The estimation scheme captures the delay in the secondary nucleation event and also reduces the number of particles nucleated during the secondary nucleation event (compared to the open-loop case). The parameter updating also corrects the overall growth rate, and pushes the particles closer to the experimental end-point PSD, by correcting k_{p11}. However, this results in an increased error in the solids content. The discrepancy seen in the end-point PSD is mainly due to the irreversible effects during the first sample time in the batch. This suggests the need for not only a parameter update as one progresses with the batch (receding horizon), but *a need to update the initial conditions* (and simulation of the entire batch after each data is available). Since the model simulates a larger number of particles at early times, there is a reduced growth rate which accounts for the discrepancy in the size of the larger mode. Although the size can be corrected by adjusting the parameters, the spread of the distribution is dictated by the early nucleation rate and cannot be corrected in a receding horizon formulation (there is no input to manipulate the competitive growth—see Immanuel and Doyle III (2003a)). Also, the relative magnitude of the two modes cannot be corrected as the number of particles nucleated in the primary mode cannot be taken away in a receding horizon formulation.

Figure 16 shows the improvement in the estimates as more samples are utilized. As one would expect, these results indicate a better convergence as more data are utilized. The estimates utilizing the samples up to 55 min converge completely with the one corresponding to using all the samples (and no better estimates are available beyond this time point). Even at an earlier time, say 21 or 33 min, the estimation is quite close to convergence. These earlier estimates capture the delaying of the nucleation, and also partially capture the reduced secondary nucleation event. Although it may not be advisable to base the re-computation of the inputs on the sample at 8 min alone, or even on samples up to 12 min, one can choose a time between 21–33 min to perform the *first* re-computation. This could potentially be early enough to

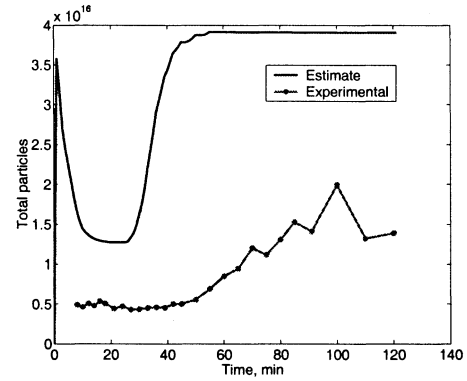

(a) profile of total particles

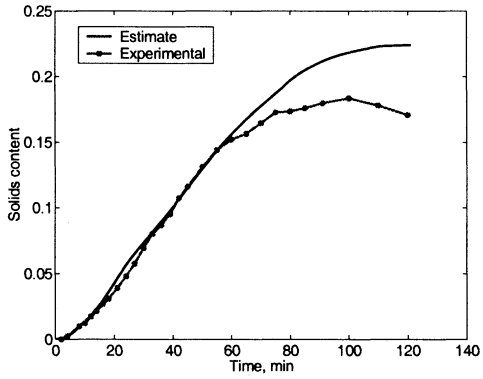

(b) profile of solids content

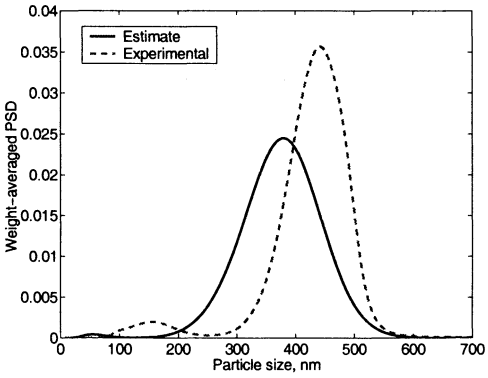

(c) end-point distribution

Figure 15. Comparison of the estimates based on a Luenberger estimator with parameter updating, with the experimental data.

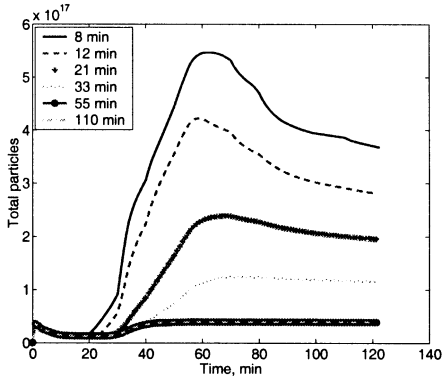

(a) profile of total particles

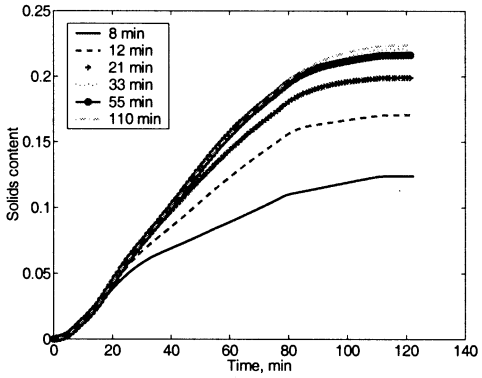

(b) profile of solids content

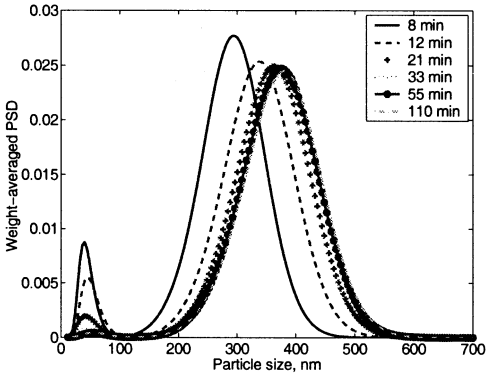

(c) end-point distribution

Figure 16. Improvement and convergence of the estimates based on Luenberger estimator, with availability of additional data.

prevent any irreversible effects on the distribution while simultaneously being a reliable estimate.

Re-optimization Based on Luenberger Observer. Figures 17 and 18 show the results corresponding to a re-computation of the inputs using the data sampled up to 33 min (available at approximately 44 min). This is based on a single objective re-optimization, namely, minimization of the deviation in the end-point weight-averaged PSD ($\theta_3(t_f)$). The surfactant feed in intervals 5–7, and the feed of both monomers in intervals

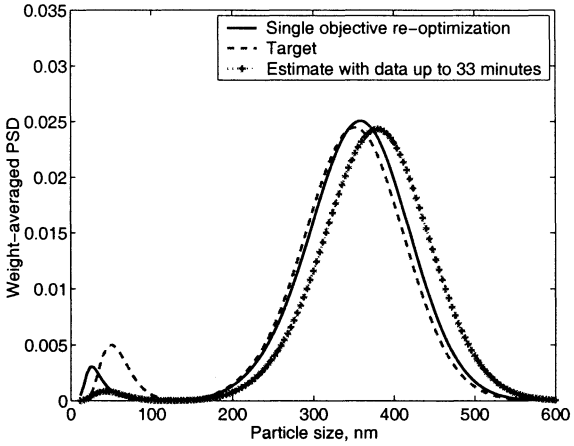

Figure 17. Single-objective re-optimization of the inputs, based on a Luenberger estimator with parameter updating, utilizing data up to 33 min into the batch.

5–8 are re-computed. The inputs in this case drive the solution closer to the target, relative to the estimate (Figure 17). The error in the smaller mode is partly due to limitations with the objective function choice that was discussed earlier —with no single objective function having the capability to describe all aspects of the distribution. In this case, the surfactant feed in the fifth interval hits the upper constraint, and that in the sixth interval is close to the upper constraint, to cause a larger nucleation event (Figure 18a)—in contrast to the case based on a single measurement (Figure 14a), wherein the surfactant feedrates were reduced in all the intervals. The monomer feeds in intervals 5–7 are reduced significantly to affect the growth rates and the particle size (Figures 18b and 18c). After the sixth interval, nucleation is reduced, and the optimizer attempts to raise the growth rate to accommodate the smaller mode within the target.

Observability issues based on faster measurements

The potential benefit of feedback control that was described previously is predicated on the observability of key states at early batch times. An evaluation of that observability is performed as follows. The parameters (alone) are updated based on the solids content measurements until the time when the first PSD measurement is available (L_i is a null matrix in Eq. 11). Thereafter, the batch is run in open loop.

Figure 19 shows the estimated profiles of total particles and solids content, and the estimated end-point PSD, when solids content samples up to various cut-off times are utilized to update the parameters. The estimates corresponding to the previous state-parameter estimator (utilizing all PSD measurements) are also shown in these figures (labeled as "Original estimate"). It is seen that using all solids content measurements through 21 min (which is approximately the time when the first PSD measurement is available) gives very close estimates to the ones obtained from the earlier estimation strategy. The estimates are improved when the early parameter updating (based on the solids content measurements) is combined with the later state/parameter correction in the

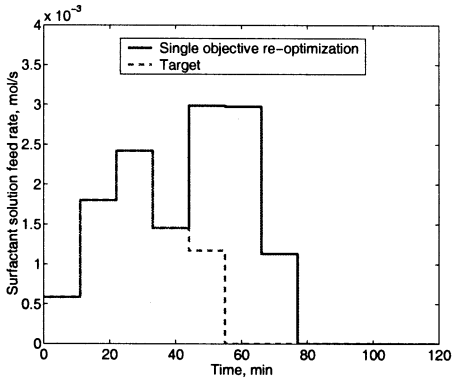

(a) profile of surfactant feed rate

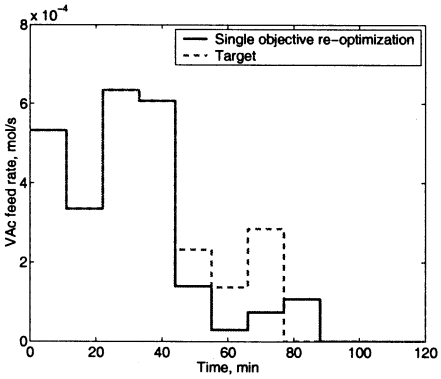

(b) profile of vinyl acetate feed rate

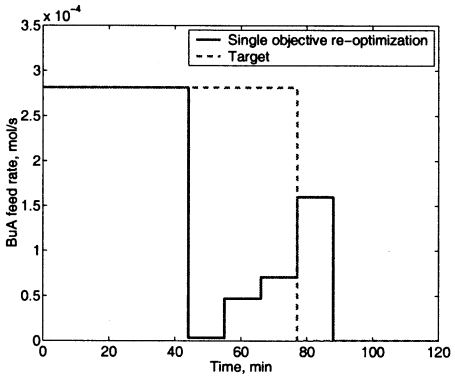

(c) profile of butyl acrylate feed rate

Figure 18. Inputs corresponding to the results presented in Figure 17, re-optimized from 44 minutes into the batch.

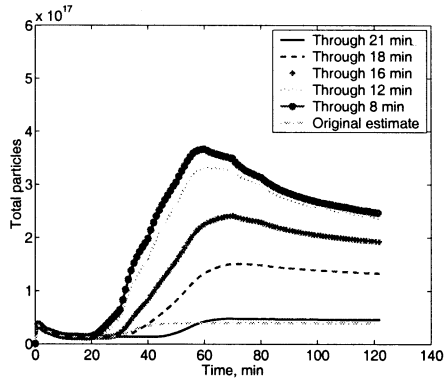

(a) profile of total particles

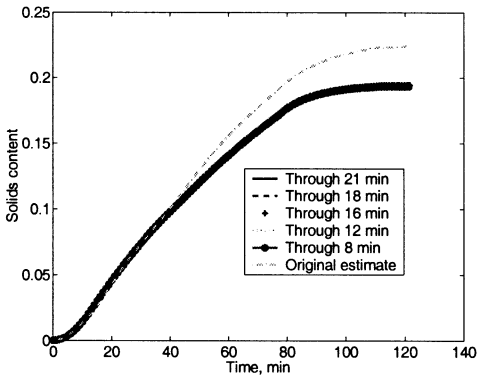

(b) profile of solids content

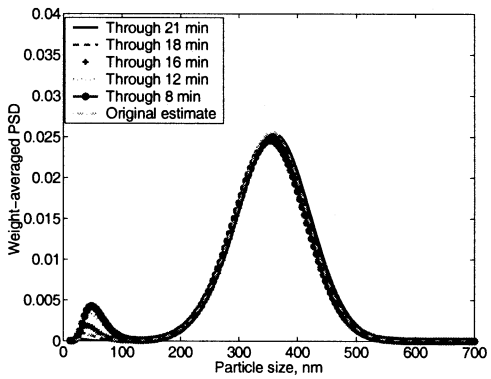

(c) end-point distribution

Figure 19. Estimates based on discrete parameter updating using the early measurements of solids content, compared with the estimates based on the Luenberger continuous-discrete state/parameter estimator.

Parameter updating is stopped after the time shown in the legend of the various cases.

continuous-discrete formulation that was presented in the previous subsection. These results show improvements particularly in the profile of solids content. However, the major impact is the ability to observe the process trends at very early times of the batch. These results indicate the potential to extract information on the total particles and the nucleation rates even at times before the availability of PSD data. This could substantially improve feedback control with the early prevention of irreversible effects.

Summary

This article details results on the control of PSD in semibatch emulsion polymerization. The primary contribution of this article is the proposal of a novel strategy for recipe design that attains a target PSD at the end of a batch. The

proposed method is hierarchical in which the desired PSD is produced by manipulating the individual rates of nucleation, growth, and coagulation over the course of the batch. The ideas of the control of instantaneous properties and partial control are invoked in implementing the technique. The problem of the attainment of a desired PSD is decomposed into one of tracking profiles of nucleation and growth rates, with coagulation minimized. This in turn is re-cast as an equivalent problem of the tracking profiles of total particles and solids content, with coagulation minimized. While the tracking of the two profiles constitute two specific objectives, a third objective on the full distribution itself is also added to the formulation. The Nondominated Sorting Genetic Algorithm (NSGA) is used as the optimization technique to solve the resultant multiobjective problem, and to produce the pareto-optimal solutions. The new strategy is effective in identifying solutions that match the target very closely, compared to the previous single-objective formulation which was faced with the inability of defining a suitable objective function in terms of the full distribution. Experimental implementation produces results that are close to the target. The explicit incorporation of the coagulation events in the model avoids any conspicuous coagulation events in the results. The results, however, show differences attributed to the inevitable implementational and model uncertainties, thereby motivating the use of feedback control.

An analysis of several state estimation and feedback strategies is presented. The underlying on-line state estimation is very complex, consisting of multirate characteristics and delayed measurements. In the current study, the estimation was performed off-line. It was found that a combined state/parameter estimation strategy is essential. The estimation was performed using a Luenberger observer with a continuous state update and a discrete parameter update. Furthermore, it was observed that a pure receding horizon estimation would not be very effective, and an update of the initial conditions also is essential as each data point becomes available (at least for the early portion of the batch). (However, there are also cases in which a simple biased estimator, even without a parameter update, is effective in converging to the experimental data.) It is also important to account for noise and disturbances. The last aspect was not pursued in the current study. Another observation that has a profound implication for feedback control is the ability to influence the states on the basis of the solids content values (from the more frequent densitometer and load cell measurements). The early solids content measurements can be utilized to get estimates of total particles and free surfactant concentration values. This would enable an earlier feedback correction. The general aim of feedback control of PSD (as addressed in this article) was to regulate the timing, duration, and magnitude of the secondary nucleation event and the overall growth rates to thereby match the target PSD on a relative basis. These early estimates will substantially aid in regulating both the secondary nucleation event and growth. However, they can also potentially aid in regulating the primary nucleation event (and, hence, attain the target PSD on an absolute basis). However, this is only true for negative errors in the number of particles. Also, the shape of the primary mode may not be correctable, as nucleation and growth are inter-related (and not sequential) processes.

The state estimation studies show the need to delay the re-computation until approximately 4–6 data points are available. The re-optimized results are encouraging in their ability to drive the solutions closer to the target. The implementation of a full nonlinear MPC beyond this point would be a desirable strategy. At a minimum, re-computation at several instances would be needed (as opposed to a single re-computation). A statistical fault detection tool (Flores-Cerrillo and MacGregor, 2002) might be a possibility to compensate for measurement noise. A batch-to-batch re-computation (Clarke-Pringle and MacGregor, 1998; Crowley et al., 2001; Lee et al., 2002) could complement in-batch feedback for the control of PSD in semi-batch emulsion polymerization. An off-line parameter/state estimator similar to the one considered in this study can be employed, utilizing all data points available from previous batches. These parameter estimates can be utilized to re-optimize the inputs for the subsequent batches, using the straightforward methods presented here.

Acknowledgments

The authors express their thanks to Dr. Cajetan Cordeiro of Air Products and Chemicals, Inc., Allentown, PA, Dr. Yang Wang of the University of California at Santa Barbara, and Dr. Babatunde A. Ogunnaike of DuPont Co. for their invaluable comments on this work. The authors also acknowledge funding from the University of Delaware Competitive Fellowship, the Office of Naval Research, and the Process Control and Monitoring Consortium at the University of Delaware.

Literature Cited

Astorga, C. M., N. Othman, S. Othman, H. Hammouri, and T. F. McKenna, "Nonlinear Continuous-Discrete Observers: Application to Emulsion Polymerization Reactors," *Cont. Eng. Prac.*, **10**, 3 (2002).

Bhaskar, V., S. K. Gupta, and A. K. Ray, "Multiobjective Optimization of an Industrial Wiped-film PET Reactor," *AIChE J.*, **46**(5), 1046 (2000).

Choi, K. Y., and D. N. Butala, "An Experimental Study of Multiobjective Dynamic Optimization of a Semibatch Copolymerization Process," *Polym. Eng. Sci.*, **31**, 353 (1991).

Clarke-Pringle, T., and J. K. MacGregor, "Optimization of Molecular-Weight Distribution Using Batch-to-Batch Adjustments," *Ind. Eng. Chem. Res.*, **37**, 3660 (1998).

Crowley, T. J., C. A. Harrison, and F. J. Doyle III, "Batch-to-Batch Optimization of PSD in Emulsion Polymerization using a Hybrid Model," *Proc. American Control Conf.*, Arlington, VA, Omnipress, Madison, WI, 981 (2001).

Crowley, T. J., E. S. Meadows, E. Kostoulas, and F. J. Doyle III, "Control of Particle Size Distribution Described by a Population Balance Model of Semibatch Emulsion Polymerization," *J. Proc. Cont.*, **10**, 419 (2000).

Dimitratos, J., C. Georgakis, M. S. El-Aasser, and A. Klein, "Dynamic Modeling and State Estimation for an Emulsion Copolymerization Reactor," *Computers Chem. Eng.*, **13**(1/2), 21 (1989).

Dimitratos, J., C. Georgakis, M. S. El-Aasser, and A. Klein, "An Experimental Study of Adaptive Kalman Filtering in Emulsion Copolymerization," *Chem. Eng. Sci.*, **46**(12), 3203 (1991).

Flores-Cerrillo, J., and J. F. MacGregor, "Control of Particle Size Distribution in Emulsion Semi-batch Polymerization using Mid-course Correction Policies," *Ind. Eng. Chem. Res.*, **41**, 1805 (2002).

Garcia, C. E., and A. M. Morshedi, "Quadratic Programming Solution of Dynamic Matrix Control," *Chem. Eng. Commun.*, **46**, 73 (1986).

Garcia, C. E., and M. Morari, "Internal Model Control 1. A Unifying Review and Some New Results," *IEC Process Des. and Dev.*, **21**, 308 (1982).

Immanuel, C. D., and F. J. Doyle III, "A Sensitivity Approach to

Reachability Analysis for Particle Size Distribution in Semi-batch Emulsion Polymerization," *Ind. Eng. Chem. Res.*, in press (2003a).
Immanuel, C. D., and F. J. Doyle III, "Computationally-Efficient Solution of Population Balance Models Incorporating Nucleation, Growth and Coagulation," *Chem. Eng. Sci.*, in press (2003b).
Immanuel, C. D., and F. J. Doyle III, "Open-loop Control of Particle Size Distribution in Semi-batch Emulsion Co-Polymerization using a Genetic Algorithm," *Chem. Eng. Sci.*, **57**(20), 4415 (2002).
Immanuel, C. D., C. F. Cordeiro, S. S. Sundaram, and F. J. Doyle III, "Population Balance PSD Model for Emulsion Polymerization with Steric Stabilizers," *AIChE J.*, (2003a).
Immanuel, C. D., C. F. Cordeiro, S. S. Sundaram, E. S. Meadows, T. J. Crowley, and F. J. Doyle III, "Modeling of Particle Size Distribution in Emulsion Co-Polymerization: Comparison with Experimental Data and Parametric Sensitivity Studies," *Comp. Chem. Eng.*, **26**(7-8), 1133 (2002).
Immanuel, C. D., T. J. Crowley, E. S. Meadows, C. F. Cordeiro and F. J. Doyle III, "Evolution of Multimodal Particle Size Distribution in Vinyl Acetate/Butyl Acrylate Emulsion Copolymerization," *J. Polym. Sci. Polym. Chem.*, **41**(14), 2232 (2003b).
Jazwinsky, A. H., *Stochastic Processes and Filtering Theory*, Academic Press, New York (1970).
Kozub, D. J., and J. F. MacGregor, "Feedback Control of Polymer Quality in Semibatch Copolmerization Reactors," *Chem. Eng. Sci.*, **47**, 929 (1992a).
Kozub, D. J., and J. F. MacGregor, "State Estimation for Semi-Batch Polymerization Reactors," *Chem. Eng. Sci.*, **47**, 1047 (1992b).
Lee, K., J. H. Lee, D. R. Yang, and A. W. Mahoney, "Integrated Run-to-Run and on-line Model-Based Control of Particle Size Distribution in Semi-Batch Precipitation Reactor," *Comp. Chem. Eng.*, **26**(7-8), 1117 (2002).
Liotta, V., C. Georgakis, and M. S. El-Aasser, "Real-time Estimation and Control of Particle Size in Semi-Batch Emulsion Polymerization," *Proc. American Control Conf.*, Albuquerque, NM, Omnipress, Madison, WI, 1172 (1997).
Luyben, M. L., and C. A. Floudas, "Analyzing the Interaction of Design and Control—1. A Multiobjective Framework and Application to Binary Distillation Synthesis," *Comp. Chem. Eng.*, **18**(10), 933 (1994).
Mahoney, A. W., and D. Ramkrishna, "Efficient Solution of Population Balance Equations with Discontinuities by Finite Elements," *Chem. Eng. Sci.*, **57**, 1107 (2002).
Min, K. W., and W. H. Ray, "On the Mathematical Modelling of Emulsion Polymerization Reactors," *J. M. S.-Rev., Macromol. Chem.*, **C11**(2), 177 (1974).
Mitra, K., K. Deb, and S. K. Gupta, "Multiobjective Dynamic Optimization of an Industrial Nylon 6 Semibatch Reactor Using Genetic Algorithm," *J. Applied Polym. Sci.*, **69**, 69 (1998).
Saldivar, E., and W. H. Ray, "Control of Semicontinuous Emulsion Copolymerization Reactors," *AIChE J.*, **43**(8), 2021 (1997).
Semino, D., and W. H. Ray, "Control of Systems Described by Population Balance Equations—I. Controllability Analysis," *Chem. Eng. Sci.*, **50**(11), 1805 (1995a).
Semino, D., and W. H. Ray, "Control of Systems Described by Population Balance Equations—II. Emulsion Polymerization with Constrained Control Action," *Chem. Eng. Sci.*, **50**(11), 1825 (1995b).
Silva, C. M., and E. C. Biscaia Jr., "Genetic Algorithm Development for Multiobjective Optimization of Polymerization Reactors," *Second Pan American Workshop on Process Systems Engineering*, Guaruja, Brazil (2001).
Tsoukas, A., M. Tirrell, and G. Stephanopoulos, "Multiobjective Dynamic Optimization of Semibatch Copolymerization Reactors," *Chem. Eng. Sci.*, **37**(12), 1785 (1982).
Urretabizkaia, A., G. Arzamendi, M. J. Unzue, and J. M. Asua, "High Solids Content Emulsion Teropolymerization of Vinyl Acetate, Methyl Methacrylate and Butyl Methacrylate. II. Open Loop Composition Control," *J. Polym. Sci.: Part A: Polym. Chem.*, **32**, 1779 (1994).
Wang, Y., C. D. Immanuel, S. Sundaram, C. F. Cordeiro and F. J. Doyle III, "Reachability Studies for Particle Size Distribution in Emulsion Polymerization," AIChE Meeting, Indianapolis (2002).

Manuscript received Oct. 2, 2002, and revision received Mar. 4, 2003.