Annual Reviews in Control 35 (2011) 99–110

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

An industrial and academic perspective on plantwide control

James J. Downs [a], Sigurd Skogestad [b,*]

[a] Eastman Chemical Company, Kingsport, TN 37662, USA
[b] Dept. of Chemical Engineering, Norwegian University of Science and Technology, N-7491 Trondheim, Norway

ARTICLE INFO

Article history:
Received 13 January 2010
Accepted 7 January 2011

Keywords:
Process control
Control structure design
Plantwide control
Inventory control
Throughput

ABSTRACT

The purpose of this paper is to stress the importance of selecting the right plantwide control structure and the need for a formalized approach that can encompass the many issues that arise in plantwide control design. Since the concept of process control design based on a holistic view of the process came about, the variety of procedures and approaches to the design problem has illustrated the difficulty of a unified approach. Using examples, the need and advantages of using a systematic approach based on considering the plant economics are highlighted. The examples deal with disturbance rejection, throughput maximization and economic optimization of plants consisting of parallel units.

1. Introduction

Industry uses a variety of approaches to accomplish plantwide control design. The range of tools used spans from engineering judgment to the applications of complex model based algorithms. Over the last 40 years the field of research in this area has attacked this design problem on various levels. Larsson and Skogestad (2000) provide a good review of the various approaches. Design heuristics based on experience, design rules based on case studies, algorithms for objective function minimization, etc. have all contributed to the improvement of how designs can be accomplished (e.g., Chen & McAvoy, 2003; Douglas, 1988; Downs, 1992; Konda, Rangaiah, & Krishnaswamy, 2005; Kookos & Perkins, 2002; Luyben, Tyreus, & Luyben, 1998; Narraway & Perkins, 1993; Skogestad, 2004; Vasbinder & Hoo 2003; Ward, Mellichamp, & Doherty, 2006; Zheng, Mahajanam, & Douglas, 1999). However, the complex nature of the problem and the various depths to which it needs to be solved have resulted in a design procedure that is difficult to piece together from the various approaches that have been put forth. This is not a new issue and almost 20 years ago the "Tennessee Eastman challenge problem" (Downs & Vogel, 1993) was put forward so that various approaches could be tested against each other. Nevertheless, today in industry, much of the research in this area has still not gained sufficient acceptance to have a profound influence. The purpose of this paper is to stress the need for a formalized yet simple approach that can encompass the many levels that arise in plantwide control design.

* Corresponding author.
E-mail addresses: jjdowns@eastman.com (J.J. Downs), skoge@ntnu.no (S. Skogestad).

1367-5788/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.arcontrol.2011.03.006

In Section 2, industrial aspects of plantwide control design are discussed and two examples illustrate how industry has approached the plantwide control problem. Although the final control structure may be good, the approach and criteria to arrive at it may vary. This points to a need for a more formal procedure which is presented in Section 3. In Section 4, the inclusion of plantwide economic variables is presented and illustrated in Section 5. The paper concludes that the formal approach presented is a step in the direction of helping to organize the design procedure for plantwide control. This paper also illustrates the application of the formal procedure to more complex examples that illustrate that plantwide design involves many issues and one-criterion approaches may not be sufficient.

2. Status in industry

The traditional approach for Eastman for designing process control strategies for chemical plants has been to set production rates using the process feed rates and then to design automatic control systems around each unit operation sequentially through the process. For processes with significant in-process inventory, including buffer tanks, and limited recycle, this approach can be used successfully. However, as processes become more complex and at the same time have less in-process inventory, the design of a plant-wide control strategy becomes a more important part of the overall process control design problem. The interrelation of the plant-wide control strategy with the process chemistry and economics requires both control theory and also process knowledge. It has become apparent that the design of plant-wide control strategies involves not only the development and application of process control theory but also, in a more fundamental sense, the development of a methodology one uses to approach the plant-wide control problem.

While we usually think about material balance and energy balance equations applying to a unit operation, they also apply to whole processes and to entire chemical complexes. The time it takes to accumulate and deplete inventories may be longer for large processes or chemical complexes, but the laws of accumulation and depletion of material hold nonetheless. Whereas for a steady-state process, we assume the rate of accumulation of each component to be zero, the fact that the control system must ensure that to be the case is often overlooked. The manipulation of flows, utilities, and the readjustment of process operating conditions to maintain a balance of material and energy entering and leaving an entire process are some of the overriding priorities for the control system (Buckley, 1964). The material balance must be maintained not only from an overall viewpoint but also for each component in the system.

While traditional control theory can be used to approach the control problem as, "Given a process described by a model of the form…", the plant-wide control problem requires much more in the development of the problem statement itself. It is not intuitively obvious at the outset what the underlying control problems are – much less how they should be solved. As researchers have begun to explore the plant-wide control area, the application of methods and techniques as applied to case studies (e.g., Luyben et al., 1998) has elucidated issues that are difficult to quantify and are in need of further discussion and research.

Despite the ever-increasing incentive, segregation of the process design and control tasks is still common. Two contributing factors to this segregation are: (1) the difficulty of changing from the historical approach of fixing the process design before the control engineer becomes involved, and (2) the difference in the thought pattern of design and control engineers. In addition it can be costly and time consuming to address controllability and operability in a rigorous way at the design stage. The common notion is that process economics are solely determined by the steady-state process design. While the nominal steady-state design point is very important, it loses its distinction if one is unable to maintain plant operation at the desired operating point. Design decisions are often based on a nominal steady-state analysis with little consideration of disturbances, changes in active constraints, controllability, process and product variability, or plant-wide control issues. The basic thought pattern in the design stage usually follows the form, "Given these conditions, create a design to perform this function" (design question), as opposed to, "Given this design, how well will it perform its intended function?" (rating question). As existing plants are pushed to produce greater throughputs, an additional question becomes important, "Given this plant, how can I maximize profit?" (optimization question). In summary, we need a plantwide control system that implements in practice the operation envisioned at the steady state design stage. There is a link between process design and control here, in that the process design team should, in addition to the plant design and its nominal operating data, provide optimal operating data also for expected future changes, including the expected location of the bottleneck when the throughput increases.

Current industrial practice is usually focused on unit operation control. This viewpoint emanates from the overriding issue of reliable operation. These unit control strategies are simple and understandable by operators and engineers alike and lead to operations that when "sick" can usually be healed without the capabilities of experts. This approach has worked reasonably well for many years. Furthermore, the high costs of building new facilities have led to more retrofits and plants producing products that they were not designed to produce. As plants are campaigned to produce a wider variety of product specifications, control strategies that are simple and perhaps applicable to many different operating points are strongly desirable.

This current design practice is being challenged as process economics drive toward fewer new designs and more operation of existing facilities in new ways. Techniques for plant wide process control design are needed (1) that result in processes that are operated in near optimal fashion while not employing complex control technology, for example, real-time optimization (RTO), and (2) that do not require the care and feeding of control experts. Several approaches that address the attainment of optimal operation of plants while not requiring implementation of complex, perhaps difficult to understand control systems, have emerged. Two of these, self optimizing control design (Skogestad, 2000) and operational strategies based on process chemistry (Ward, Mellichamp, & Doherty, 2004; Ward et al., 2006) have found particular appeal at Eastman.

The importance of being able to discriminate how process variables need to behave to achieve optimal operation is fundamental when designing plantwide strategies. Often the underlying unit operation strategies can be kept simple and usually single-input single-output (SISO) while the overall plant wide strategy is optimizing plant operation in a more natural fashion. This approach has a wide appeal when plant reliability and control system understandability are required. Each of these approaches builds into the control system a natural "self-optimizing" that is part of a normal operation. Contrasted with the centralized approach (e.g., RTO) of using models to determine an optimum and then driving a process to that optimum point, "self-optimizing" strategies designed in from the bottom, provide simplicity, robustness and reliability.

From the start-up the primary objective for a new plant is to achieve nameplate (nominal) capacity in a reliable and predicable way. Often the need for optimization of plant operations comes after the facility has been operational for a few years. By this time top-down optimization strategies can be implemented, provided the plant has a good regulatory control system. If the optimization strategy is counterintuitive, then operator understanding can suffer. Most control engineers have experienced the difficulty of keeping in service control strategies that, while driving the process to the correct economic conditions, does so in an unusual or difficult to understand fashion.

The importance of having plantwide control strategies that are optimizing in a natural, fundamental way can have long term effects. Operator training and understanding during the early years of plant operation sets thought patterns for years to come. When the need for plant optimization arises, the basic building blocks of how the control system automatically drives plant operation are in place. The process optimizer at this time may only have to make small adjustments to a process that is close to optimum already. The trick, of course, is that these strategies must be basically simple and for the most part SISO. Our experience from Eastman is that for plants where "self-optimizing" regulatory control strategies have been built in from the beginning, we have been successful with process optimization projects that have been undertaken. On the other hand, for older processes which have control strategies not designed with optimization in mind, we may struggle for years working to gain operator acceptance to a new strategy. Even the simple idea of setting process throughput at a place other than the process feed can become a difficult endeavor.

Example 1 (*Changing the production throughput manipulator TPM for an esterification plant*). Eastman operates many processes that have produced chemicals for over 50 years. Esterification chemistry is well known and has been a workhorse for the company. Traditionally, units were designed with the process throughput set at the feed to the process. Control systems consisted of pneumatic single input/single output controllers that were difficult to change and had a long operating history. As production rates increased over the years due to demand growth and incremental process improvement, the original plantwide strategy would become limiting. A flowsheet of a typical esterification process, consisting of a reactor with a distillation column, followed by an extraction column and a distillation column is shown in Fig. 1. The first distillation column separates water, ester, and alcohol from the unreacted acid, The extraction column washes the unreacted alcohol from the ester product. The final distillation column removes the remaining unreacted alcohol, which is recycled to the extraction column. The flowsheet shows the original plant had with the throughput manipulator (TPM) located at the reactor feed.

In the late '70's and early '80s' Eastman benefited from implementing a change in the TPM location on numerous plants. Early adoption of this significant change was difficult because of (1) an ingrained mindset toward needing process feeds constant, (2) operator understanding of an "inventory-to-feed" strategy, and (3) the difficulty of reversing the control decision using pneumatic hardware. Today at Eastman, the notion of setting the TPM at a location other than the process feeds is common and is driven by variability propagation and ease of operation requirements. The benefits of choosing the best location for the TPM have also become realized in our capital design process.

For the esterification process the first change was to move the TPM from the process feed rate to the distillate flow rate leaving the first distillation column as shown in Fig. 2. This strategy worked well for many years because many of the disturbances entering the reactor were directed away from the more sensitive extraction/distillation separation portion of the process.

Later, as the extraction step became the process bottleneck, it became evident that its behavior as a function of organic feed rate was very nonlinear. This nonlinearity stemmed from the fact that increasing organic feed rate resulted in an increasing composition of the alcohol taken from the extractor to the final distillation column. The increase in distillate rate (recycle stream R in Fig. 2) needed to remove the alcohol from the final product would aggravate the situation by increasing the feed rate to the extractor. The point at which the process would enter this "windup" varied with the amount of unreacted alcohol reaching this part of the process. This windup in the recycle loop is similar to Luyben's "snowball effect" (Luyben, 1994), but the cause in our case is a limitation in mass transfer rate whereas in Luyben's case it is a limitation in reaction rate. For this process, the windup condition usually took

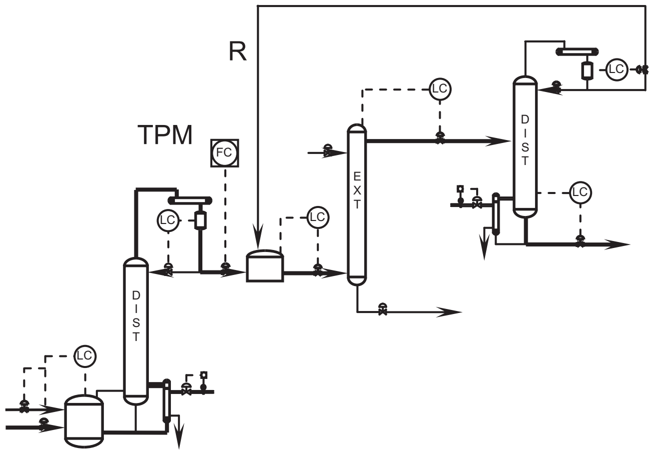

Fig. 2. Esterification process. Inventory control strategy with column 1 distillate rate used as the TPM.

12–24 h to get fully engaged. This made it difficult for operators to confidently set the production rate. In addition, what may be a maximum and stable rate today might result in the windup condition tomorrow. The outcome of this uncertainly resulted in operations setting a lower than optimum production rate to guarantee process stability.

A further improvement in locating the TPM occurred when it was relocated to be the feed to the extraction system (Fig. 3). Obviously, this eliminated variability from propagating to the extractor, but more importantly, it resulted in a self regulating system that avoids the windup should the operator set the TPM too high. In particular, if the TPM is set too high and excess alcohol leaks to the final distillation system, then less feed is drawn from the front end of the process, and the extractor, while not at the optimum feed rate, does remain stable. This situation is quite recoverable by operators who note that production rates have fallen, and realize that they have set the extractor feed rate too high. We found that the operators were capable of optimizing the operation once fear of setting the extractor feed too high was removed.

The strategy in Fig. 2, and in particular the original strategy in Fig. 1, were unforgiving in that once the throughput was set too high, it resulted in flooding in the extractor and distillation columns and several hours of lost production. With the final control strategy in Fig. 3, the ability to experiment with the throughput without the penalty of passing this "point of no return", gave

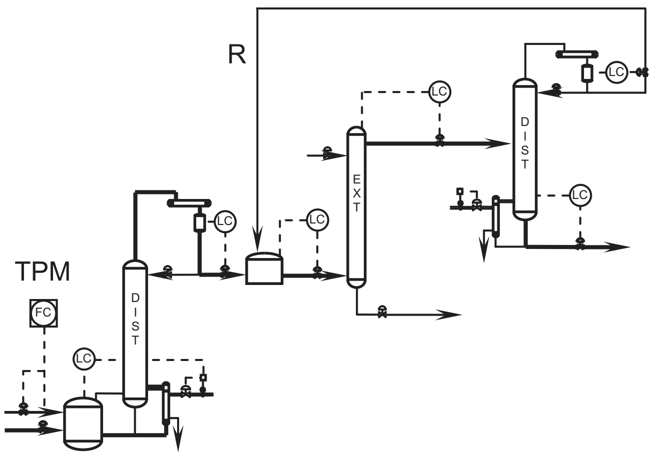

Fig. 1. Esterification process. Original inventory control strategy with column 1 feed rate used as the TPM.

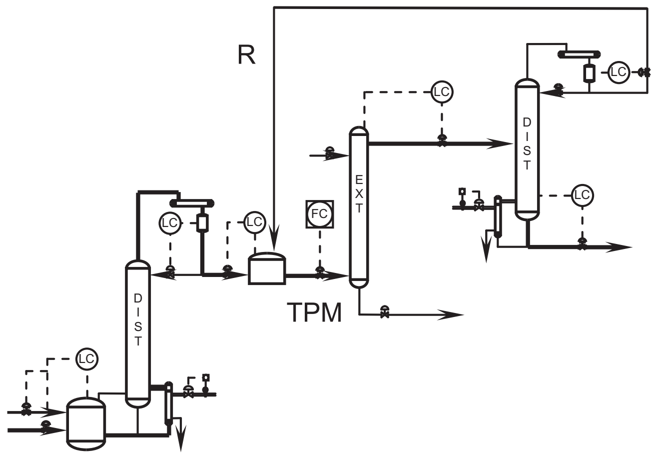

Fig. 3. Esterification process. Final inventory control strategy with extractor feed rate used as the TPM.

operators confidence in the control system to recover if they ended up pushing rates too high, and made them operate the process closer to its maximum throughput.

The basis for relocating the TPM in Figs. 2 and 3, was mainly process insight as described above. However, it agrees with more general recommendations. First, note that the capacity of the extraction column is the bottleneck of the process, so locating the TPM here agrees with the rule of Skogestad (2004) of moving the TPM to the bottleneck unit. Second, note that the TPM is located inside the recycle system, so it agrees with Luyben's rule (1994), 1997, 1998) of fixing a flow inside the recycle system. One justification for Luyben's rule is that setting a flow in the recycle loop avoids the slow dynamics of the recycle system (e.g., Morud & Skogestad, 1996) that otherwise makes it difficult to have tight control, especially using manual control.

Example 2 (*Control strategy for a liquid–liquid extraction process*). During the control design phase, the engineer may chose from a variety of criteria (or rules) and the criterion chosen is usually based upon engineering judgement. The importance of the criterion choice is often not appreciated. The objective of this example is to illustrate the design criterion of propagating disturbances to insensitive location, which has been successfully used in many applications in Eastman. The resulting control strategy can then be compared with those obtained using a more methodical approach.

Consider the extraction process in Fig. 4 where acid is transferred from the water/acid feed (F) to the extract (E) by use of a solvent (S). The remaining water is the raffinate byproduct (R). Apart from maintaining stable operation, the main control objective is to have small variations in the extract product composition.

The total inventory is self-regulated by overflow of extract, but the interface level (component inventory) does not self regulate. How should this inventory be controlled? Two alternatives are shown in Fig. 5.

Strategy I. Let aqueous feed F control interface level (with constant outflow R).
Strategy II. Let aqueous outflow R control interface level (with constant feed F).

Both of these structures have been used for extraction control in various services including the example given here. Obviously, both structures work and give the same result if everything is constant (no disturbances). How do the two strategies differ when there are disturbances? To understand the difference we ask the question: "Where does the disturbance go"?

Let x denote the acid fraction, and consider variation (disturbance) in the acid feed fraction x_F by ±1% (30 ± 1%). We would like that the resulting variation in the extract product to be reduced by at least a factor 3, that is, the variation in x_E should be less than

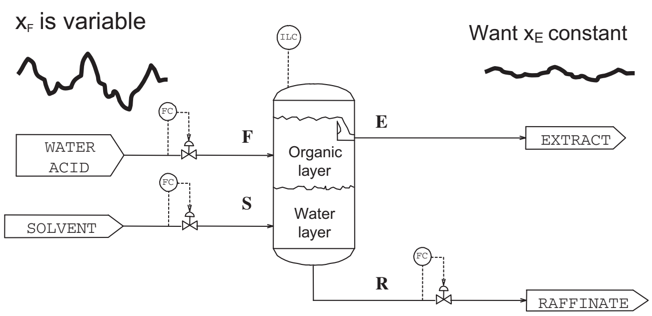

Fig. 4. Liquid–liquid extraction process without control of interface level.

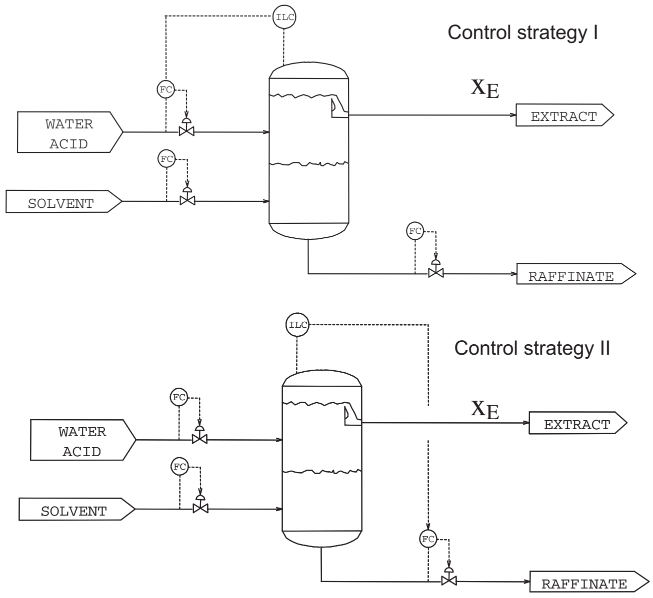

Fig. 5. Alternative control strategies for liquid–liquid extraction.

±0.33%. For strategy I, the resulting variation in the acid composition of the extract product (x_E) is ±0.856% (21.4 ± 0.856%) and for strategy II it is reduced to ±0.506% (21.4 ± 0.506%), which is still significantly higher that the desired variation. For details see the mass balances in Table 1. The mass balances were derived by assuming that we for the extract phase E have acid/water = 3 (equilibrium), and that the raffinate R is pure water (see Table 2).

In summary, for strategy I with R constant, the variation in x_F results in variations in the feed flow, F, and in x_E (with gain 0.856), while for strategy II with F constant, the variation in x_F results in variations of the outlet flow, R, and to a lesser extent in x_E (with gain 0.506). Strategy II is the preferred strategy of the two since the main objective is to have small variations in x_E. This example illustrates how different inventory strategies may result in process variability being transferred to different parts of the process.

The criterion of propagating disturbances to insensitive locations gives good insight and has given good designs for many cases. However, for more complex problems and for less experienced engineers a more systematic approach is needed. We will discuss this below where we introduce control strategies III, IV and V.

3. A plantwide control design procedure

This section is largely based on the plantwide procedure of the second coauthor (Skogestad, 2004). No matter what approach we use, the following decisions need to be made when designing a plantwide control strategy:

Decision 1. What to control? Selection of controlled variables (CVs) to achieve

a. Good steady-state performance (economics), and
b. "Stable" operation with little dynamic drift (including selecting CVs related to inventories).

Decision 2. Where to set the production rate? Placement of throughput manipulator (TPM).

Decision 3. How to control the inventories? How to pair the loops? That is, selection of a control configuration that interconnects CVs and MVs.

Table 1
Mass balances (mol/s) for extraction process: Base case (x_F = 30%), and with disturbance (x_F = 31%) for control strategies I–IV.

	Feed, *F*				*S*	Extract, *E*				Raffinate, *R*			
	Base	I	II	III, IV	All	Base	I	II	III, IV	Base	I	II	III, IV
Water	70	70.568	69	66.774	0	10	10.568	10.33	10	60	60	58.67	56.77
Acid	30 (30%)	31.705 (**31%**)	31 (**31%**)	30 (**31%**)	0	30 (21.4%)	31.705 (22.3%)	31 (21.9%)	30 (21.4%)	0	0	0	0
Solvent	0	0	0	0	100	100	100	100	100	0	0	0	0
Total	100	102.273	**100**	96.77	100	140	142.273	141.33	140	60	**60**	58.67	56.77

Table 2
General mass balances for extraction process.

	Acid/water feed (*F*)	Solvent feed (*S*)	Extractproduct (*E*)*	Raffinate byproduct (*R*)
Water	F_W	0	$F_A/3$	$F_W - F_A/3$
Acid	F_A	0	F_A	0
Solvent	0	F_S	F_S	0
Total	$F = F_W + F_A$	$S = F_S$	$E = (4/3)F_A + F_S$	$R = F_W - F_A/3$

* Acid/water = 3 in extract phase.

Often in industrial practice all issues are considered simultaneously without making formal decisions that answer the above three questions. For the extraction process in Example 2 the need for good extract composition raised the question of how best to control the aqueous inventory. This naturally leads one to consider the same issues on a broader, plantwide scale. To be effective, a more systematic procedure is helpful.

The plantwide control structure design procedure of Skogestad (2004) consists of the following seven steps:

I. Top-down part

Step 1. Define operational objectives (economics) and constraints
Step 2. Identify degrees of freedom (MVs) and optimize operation for important disturbances (offline analysis)
Step 3. Select primary (economic) controlled variables (**Decision1a**)
Step 4. Select location of throughput manipulator. (**Decision 2**)

II. Bottom-up part

Step 5. Structure of regulatory control layer (including inventory control)
  a Select secondary ("stabilizing") CVs (**Decision 1b**)
  b Select "pairings" between CVs and MVs (**Decision 3**)
Step 6. Structure of supervisory control layer (decentralized, MPCs?) (Related to **Decisions 1a and 3**)
Step 7. Structure of (and need for) optimization layer (RTO) (Related to **Decision 1a**)

The top-down part (steps 1–4) is mainly concerned with economics and steady-state considerations are often sufficient. Dynamic considerations are required for steps 4–6.

Steps 1 and 2 involve *analysis* of the optimal operation of the plant, and should form the basis for the actual *decisions* in steps 3–7. A detailed analysis in steps 1 and 2 requires that one has available a steady-state model and that one performs optimizations using for the given plant design ("rating mode") for various disturbances. As mentioned in Section 2, where we discussed the link between process design and process control, this should ideally be done at the design stage where the steady-state model is usually available. However, this is often not done in industrial practice, because the model used for design may not be suitable or available, the working relationship between the design and control functions may be weak, or there may not be time to perform this analysis.

Nevertheless, one should at least perform a simplified engineering version of steps 1, 2, and 3 where one thinks through the economics of the present and future operation with aim of using process insight to propose which variables to control, keep constant, from a steady-state economic point of view. In particular, a good engineer can often easily identify the "active constraints" that the control system should maintain. That is, where should one optimally stay at maximum or minimum values of flow, temperature, pressure, composition, etc?

Simplified step 1–3. Identify degrees of freedom and main disturbances. Based on process insight, select variables to keep constant at steady-state in order to achieve close-to-optimal economic operation (in spite of disturbances). (**Decision 1a**)

There have been many applications of the above design procedure, e.g., see Araujo, Govatsmark, and Skogestad (2007), but most of them on academic problems. There exist several other procedures for plantwide control (e.g., Luyben et al., 1998), but they focus mainly on the bottom-up part, and in particular on step 5. However, making good decisions in step 5 can be difficult without having first gone through the top-down plantwide economic analysis in steps 1–4.

Step 4 (location of TPM) was addressed in Example 1, and this issue is further discussed in the recent PhD thesis by Aske (2009); see also Aske, Strand, and Skogestad (2008).

The focus of the rest of this paper is on step 3 (economic CVs). In this respect it is important to notice that the best control structure may vary, and, depending on market conditions, there are two main modes of operation:

Mode I. Maximize efficiency (for a given throughput).

With a given throughput (production rate), the value of the products is usually known, and provided there are degrees of freedom left after satisfying the constraints (specifications), the economic objective is to minimize the use of utilities, maximize raw material yield, and to minimize waste treatment costs. These and other issues that increase specific production costs are the same as maximizing the efficiency. As discussed in Section 2 on the industrial status, the control system for a new plant is usually set up to handle this mode of operation well.

Mode II. Maximize throughput (with production rate as a degree of freedom).

When market demand is good and product prices are high, the profit is maximized by running the plant at maximum throughput. In fact, the first thing that the operation people usually focus on after startup of a new plant is to increase throughput because the opportunities for extra profit in mode II are usually much larger than in mode I. In spite of this there is usually no effort during the design phase to design a control system that can operate at maximum throughput. Operation at maximum throughput usually corresponds to using all degrees of freedom to satisfy active constraints. There will be a bottleneck somewhere in the plant against which operation at maximum throughput will run. Trying to increase the throughput will result in infeasible operation in the bottleneck unit. The maximum flow through the bottleneck unit is then an active constraint, and operation in mode II should be focused on keeping this flow at its maximum (Aske, 2009). This results in the rule of locating the TPM at the bottleneck unit (Skogestad, 2004).

As already discussed, the location of the TPM for the esterification plant in Example 1 agrees with this rule,

We now want to apply the systematic plantwide procedure to the extraction example, where we previously used the design criterion that disturbances should be propagated to insensitive locations.

Example 3 (*Application of the systematic design procedure to liquid–liquid extraction process in Example 2*). The process is very simple and there are no degrees of freedom left for steady-state economic optimization once the specifications are satisfied. Thus, for the systematic procedure we use the simplified approach for selection of controlled variables.

Simplified step 1–3. Identify degrees of freedom (MVs) and main disturbances and based on process insight, select primary controlled variables(**Decision 1a**).

The extract product flow (E) is on overflow, so there are 3 MVs that can be used for control; the two feed flows (F and S) and the raffinate R. However, at steady state there are only 2 degrees of freedom because the interface level, which has no steady-state effect, needs to be controlled. Further, the throughput is assumed to be given (mode I), which consumes another degree of freedom. We are then left with only 1 steady-state degree of freedom, and thus need to decide on 1 "economic" CV. From process insight it is important to maintain a constant product composition (x_E) so we decide that this should be controlled. There are then no degrees of freedom left for economic optimization.

Decision 1a*:* The acid product composition x_E should be kept constant. The "economic" CV is therefore $CV_1 = x_E$.

Step 4. Select location of throughput manipulator (TPM) (**Decision 2**).

The TPM location is important both for economic reasons in terms of minimizing the backoff (Aske et al., 2008) and because it dictates the structure (pairing) of the inventory control system (Buckley, 1964; Price & Georgakis, 1993). The throughput is typically located at the main feed, but could generally be anywhere in the process. Since the two proposed control strategies both have a given (constant) solvent feed flow, this implies that the solvent feed S is the throughput manipulator (**Decision 2**). Indeed, this is the optimal location for this process because the regeneration of solvent is the capacity bottleneck.

Step 5. Structure of regulatory control layer (including inventory control)

Decision 1b: The total inventory is self-regulated by overflow, but also the interface level between the two liquid phases must be controlled. Thus, CV_2 = interface level.

We must next decide *how* to control the interface level, that is, decide on the pairing. With solvent feed rate S as the TPM, we have left two candidate MVs: Feed F and outflow R. The main issue for regulatory (level) control is usually dynamics, and from this point of view there does not seem to be any significant difference between the two choices. Another issue for regulatory control is to avoid saturation of the MV, and this tells us that we should prefer the largest flow, which is the feed F. However, one should also think ahead to step 6, which is the structure of the supervisory layer. Here, the concern is to control acid composition in the extract ($CV_1 = x_E$) which depends directly on the feed F, but only indirectly, through the effect of other controllers, on the outflow R. Thus, from a pairing point of view, we would like to "save" F for the supervisory layer.

Decision 3. Use R to control the interface layer ($MV_2 = R$), which is the same as control strategy II in Fig. 5.

Step 6. Structure of supervisory control layer

Decision 3, continued. The remaining $MV_1 = F$ is used to control acid composition in the extract ($CV_1 = x_E$). The final control structure is shown as strategy III in Fig. 6. It is the same as strategy II but with an additional composition loop to keep the extract composition x_E constant.

Strategy III will meet the control objectives, but it assumes that the extract product composition x_E can be measured ($CV_1 = x_E$), which may not be possible in practice.

Self-optimizing strategies. Are there any ways of keeping x_E approximately constant, without having to measure it? One option is to estimate x_E using a model and available measurements ("soft sensor"), but this is a bit complicated. Another option is to find something else to "control" (keep constant), which indirectly leads to small variations in x_E. In fact, this is what we attempted in the previously proposed strategies where we selected.

Strategy I: Keep $CV_1 = R$ constant (and use F to control the interface level)

Strategy II: Keep $CV_1 = F$ constant (and use R to control the interface level)

However, both of these strategies gave too large variations in the product composition x_E; we found $\Delta x_E/\Delta x_F = 0.856$ for strategy I and $\Delta x_E/\Delta x_F = 0.506$ for strategy II, which are larger than the specification $\Delta x_E/\Delta x_F < 0.33$.

Can we do better? Is it possible to find a simple strategy (maybe a combination of strategies I and II) that gives $\Delta x_E/\Delta x_F = 0$? Yes, it is!

Strategy IV: Keep the flow difference

$$CV_1 = F - R$$

at a constant value (while at the same time adjusting F and R to control the interface level). One possible implementation is shown in Fig. 7.

Why does strategy IV give constant composition x_E? Controlling the interface level (which indirectly depends on the feed composition x_F) closes the material balance at steady-state. From the total steady-state material balance we have $E = S + (F - R)$ so by keeping $F - R$ constant, we have that the flow E remains constant (because the throughput S is constant) and assuming equilibrium it follows

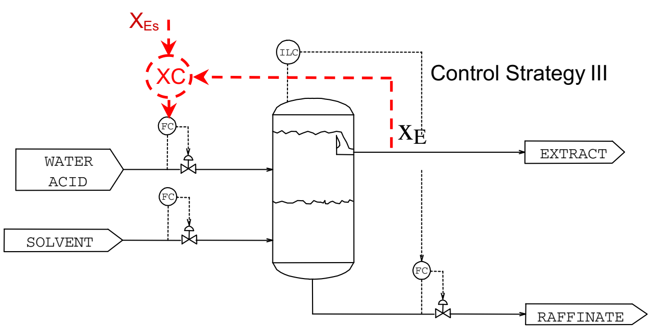

Fig. 6. Control strategy for extraction process resulting from systematic procedure, assuming on-line composition measurement.

that the composition of E must remain constant (again assuming S constant).

Strategy IV is a special case of a "self-optimizing" measurement combination, as discussed below. In fact, since we have $n_d = 1$ disturbance (x_F) and $n_u = 1$ steady-state degrees of freedom, we have from the nullspace theorem (Alstad & Skogestad, 2007) that self-optimizing control can be obtained by controlling a combination of $n_d + n_u = 2$ independent measurements. The flows (MVs) R and F are here candidate "measurements", so a possible controlled variable is $CV_1 = h_1F + h_2R$. In general the values for h_1 and h_2 can be found in a systematic manner using the *nullspace theorem*, and as we will shown later this yields $h_1 = 1$ and $h_2 = -1$ (strategy IV). For this particular example we arrive at the same result using process insight, but in general this may not be so easy and the systematic approach is preferable.

Im strategy IV, the setpoint $(F - R)_s$ may be adjusted based on (infrequent) measurements of the product composition x_E. In particular, we will need to increase $(F - R)_s$ if the throughput increases. However, by process insight (or looking at the model equations) we have that all flows should be scaled by the throughput to keep x_E constant, and this leads the improved strategy V.

> *Strategy V:* Measure the solvent federate S (measured disturbance) and keep the variable
>
> $$CV_1 = (F - R)/S$$

at a constant value, see Fig. 8. To correct for measurement errors, drift and other disturbances, the setpoint for $(F - R)/S$ may be adjusted based on (infrequent) measurements of the product composition x_E. Note that control strategy V is "self-optimizing" (giving constant product composition x_E) with respect to both disturbances in feed composition (x_F) and solvent federate (S), whereas control strategy IV only handles disturbances in feed composition (x_F). In the next section, we will show to find self-optimizing strategies in a systematic manner.

4. Selection of economic (primary) CVs

In the above example, we found that the relative flow difference $CV_1 = (F - R)/S$ is a good primary CV. How do we find good primary (economic) CVs in a systematic manner (step 3)?

We make the standard assumption here that a steady-state analysis is sufficient for studying the economics. The question is: How can we turn optimization into a setpoint problem? What should we control (CV) to get close-to-optimal operation (with minimum cost J)? The obvious CV is the gradient of the cost function, $J_u = dJ/du$, which should be zero at the optimum point, independent of disturbances. However, before we look at this and other ideas, let us look in a bit more detail in steps 1–3 in the proposed procedure for selecting CVs.

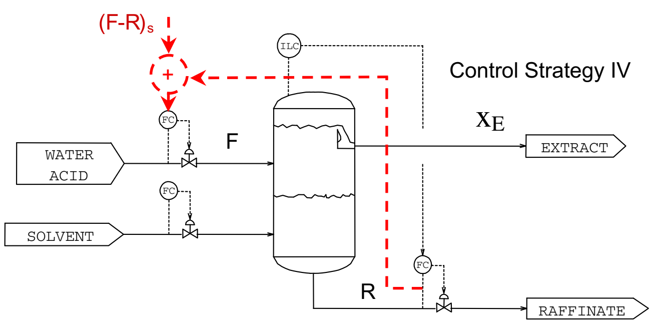

Fig. 7. "Self-optimizing" control strategy for extraction process.

> *Step 1.* Define operational objectives (cost J) and constraints

In many cases a simple economic cost is used:

$$\text{Profit} = -J$$
$$= \text{value products} - \text{cost feeds} - \text{cost utilities (energy)}$$

Other operational issues, such as safety and environmental impact are usually formulated as constraints. For cases with good market conditions we often have a constrained optimum and the cost function can be simplified to $J = -\text{TP}$ (mode II, maximum throughput).

Other cost functions are also possible. For example, consider the extraction process. Here, the optimum is to keep a constant product composition x_E, but we cannot control it directly because there is no online measurement. We therefore want to control something else that gives *indirect control* of the primary output (Hori, Skogestad, & Alstad, 2005). The cost function is then $J = (x_E - x_{ES})^2$.

> *Step 2.* Identify degrees of freedom and optimize operation for various disturbances.

One approach for optimization is to use a steady-state flowsheet simulator, if available, and optimize the operation with respect to the degrees of freedom for various disturbances. However, this may be very time consuming and in many cases, simpler models and approaches may be used. Typical "disturbances" include feed composition, feed rate, reaction rate constants, surroundings, values of constraints and prices.

> *Step 3.* Select primary (economic) controlled variables

The issue here is to select the primary (economic) controlled variables (CVs). That is, for what should we use the (steady-state) degrees of freedom? What should we control?

1. *Control active constraints*. The active constraints come out of the analysis in step 3 or may in some cases be identified based on physical insight. The active constraints should be selected as CVs because the optimum is not "flat" with respect to these variables. Thus, there is usually a significant economic penalty if we "back off" from the active constraints, *so tight control of the active constraints is usually desired*. Poor control makes it necessary to "back off" from the active constraint and this gives a loss. If a constrained optimization method is used for the optimization, then we can find the magnitude of the loss from the Lagrange multiplier λ, and we have Loss = $\lambda * \text{backoff}$.

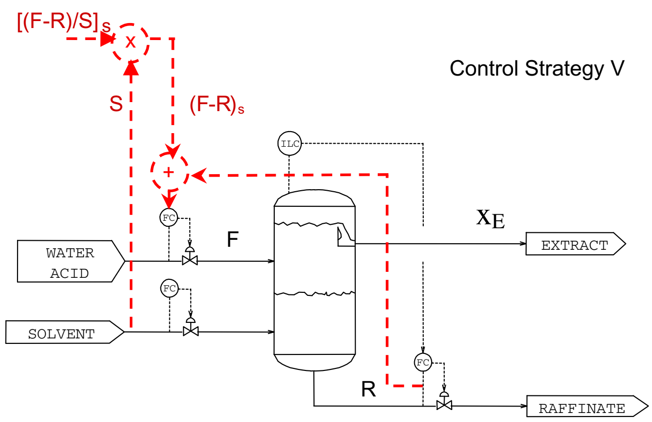

Fig. 8. Improved "self-optimizing" control strategy for extraction process.

Specifically, in mode II the feed rate should be adjusted to keep the bottleneck unit operating at its active constraints. Any back-off from the active constraints will reduce the flow through the bottleneck unit and give a loss in feed flow (production) which can never be recovered.

1. *Identify "self-optimizing" variables related to the (possibly) remaining unconstrained degrees of freedom*. These are (seemingly) "magic" variables which when held constant result in close-to-optimal operation (with a small loss), in spite of the presence of disturbances. The term "magic" is used to signify that the choice may have a significant effect on the economics (loss), and that it is not generally obvious what a good choice is. A good self-optimizing variable should give a "flat" optimum, which means that tight control of these variables is usually not required (as opposed to the active constraints). Note that the different self-optimizing variables must be found for each region of active constraints.

There are two main possibilities for selecting self-optimizing CVs:

1. Select single measurements as CVs (however, it is difficult to find single measurements in a systematic manner, so one must often use the "brute force" approach).
2. Use measurements combinations as CVs (here, methods exist to find optimal combinations; see below).

To identify good candidate measurements for a controlled variable, c, we may use the following four requirements (Skogestad, 2000):

Requirement 1. Its optimal value is insensitive to disturbances (so that the optimal variation Δc_{opt} is small).
Requirement 2. It is easy to measure and control accurately (so that the implementation error n is small).
Requirement 3. Its value is sensitive to changes in the manipulated variable, u; that is, the gain, G, from u to c is large (so that even a large error in controlled variable, c, results in only a small error in u. Equivalently, the optimum should be 'flat' with respect to the variable, c.
Requirement 4. For cases with two or more controlled variables, the selected variables should not be closely correlated.

All four requirements should be satisfied. For example, for a marathon runner, the heart rate may be a good "self-optimizing" controlled variable (to keep at constant setpoint). Let us check this against the four requirements. The optimal heart rate is weakly dependent on the disturbances (requirement 1) and the heart rate is easy to measure (requirement 2). The heart rate is relatively sensitive to changes in power input (requirement 3). Requirement 4 does not apply since this is a problem with only one unconstrained input (the power).

In addition to the above requirements, some systematic approaches to evaluate and find good "self-optimizing" CVs (especially associated with the unconstrained degrees of freedom) are:

Approach 1 (*Brute force*). Conceptually, the simplest approach for finding candidate CVs is the "brute force" approach where one considers the economic loss imposed by keeping a candidate set of CVs constant when disturbances occur (rather than re-optimizing their values),

$$\text{Loss} = J(\text{CV} = \text{constant}, d) - J_{\text{opt}}(d)$$

The term "brute force" is used because one must do a separate evaluation of each candidate set of CVs. The "brute force" approach is the most general and exact method, but also the most time consuming method because there are essentially an infinite number of possible CVs (at least if measurement combinations are included) that can be suggested, and for each of them we need to do computations to find the cost for each disturbance.

The "brute force" approach was essentially what we initially tried with strategies I and II for the extraction process, where we evaluated the change in product composition ($\Delta x_E/\Delta x_F$) resulting from a disturbance in feed composition.

Approach 2 (*Use analytic expressions or insight about the optimum*). This is not a general approach, but it may be very effective for cases where it works. One useful method is to start from the fact that at the optimum the gradient of the cost J with respect to the degrees of freedom should be zero:

$$\text{At optimum : Gradient} = J_u = dJ/du = 0$$

These are also known as the necessary condition of the optimum (NCO) (Srinivasan, Bonvin, & Visser, 2003). It seems obvious that the gradient CV = J_u is the "ideal" self-optimizing variable (Halvorsen & Skogestad, 1997; Halvorsen, Skogestad, Morud, & Alstad, 2003), However, it may be difficult to obtain the expression for J_u or it may depend on non-measured variables.

Approach 3 (*Optimal measurement combinations and nullspace method*). For measurement combinations, the nullspace theorem (Alstad & Skogestad, 2007) can be used to find optimal measurement combinations in a simple manner. This is an improvement of approach 2 for the linear case, where effectively unmeasured variables are eliminated from the expression for J_u. To use the nullspace theorem, we first need to indentify the measured variables that we want to combine and these are collected in the vector y. The number of y's (n_y) must be equal to (or greater than) the number of independent variable at steady-state (n_u) plus the number of disturbances (n_d) that we want to be optimal with respect to, that is, $n_y \geqslant n_u + n_d$. Next, we need to obtain the optimal sensitivity matrix

$$\mathbf{F} = dy^{\text{opt}}/dd$$

Each column in $\mathbf{F}$ expresses the optimal change in the y's when the independent variable (u) is adjusted so that the system remains optimal with respect to the disturbance d. If we have a model of the process, then it is in principle straightforward to obtain $\mathbf{F}$ numerically. The nullspace theorem now says that the optimal controlled variable is given by the measurement combination

$$c = \mathbf{H}y$$

where the matrix $\mathbf{H}$ satisfies

$$\mathbf{HF} = 0.$$

Thus the optimal combination $\mathbf{H}$ is in the left nullspace of $\mathbf{F}$, which is the reason for calling it the nullspace method. Any number of disturbances (d) may in theory be included, but we since we need $n_y \geqslant n_u + n_d$, we must then include more measurements in y. This requirement is not needed if we use the more general "exact local method" (Alstad, Skogestad, & Hori, 2009; Halvorsen et al., 2003) which can also include the effect of measurement noise, which is neglected in the simpler nullspace method. An analytical expression for $\mathbf{H}$ for the "exact local method" is given by Alstad et al. (2009).

Example 4 (*Nullspace method for extraction process; continuation of Examples 2 and 3*). Previously, in Example 3, we found from process insight that $c = (F - R)/S$ is a "self-optimizing" variable with respect to disturbances in feed composition ($d_1 = x_F$) and solvent feedrate ($d_2 = S$); see control strategy V. We now want to derive this using the systematic nullspace method. For this particular example, the nullspace method may seem more complicated, but in general it is not.

For the extraction process, there is one independent variable at steady-state ($n_u = 1$), so we want to find a single CV, $c = \mathbf{H}y$. We consider two disturbances ($n_d = 2$) so from the condition $n_y \geqslant n_u + n_d = 1 + 2 = 3$, we need to combine at least three measurements. We choose to use three flows (which are all measured),

$$y = \begin{pmatrix} F \\ R \\ S \end{pmatrix}$$

The controlled variable combination we are looking for is then

$$c = \mathbf{H}y = h_1 y_1 + h_2 y_2 + h_3 y_3 = h_1 \Delta F + h_2 \Delta R + h_3 \Delta S$$

where we have introduced the symbol Δ to show clearly that the nullspace method gives linear measurement combinations in terms of deviation variables. For the extraction example, optimality is defined as keeping the product composition (x_E) constant. Using the model, the sensitivity matrix $\mathbf{F}$ becomes

$$\mathbf{F} = \begin{pmatrix} (dF/dx_F)_{x_E} & (dF/dS)_{x_E} \\ (dR/dx_F)_{x_E} & (dR/dS)_{x_E} \\ (dS/dx_F)_{x_E} & (dS/dS)_{x_E} \end{pmatrix} = \begin{pmatrix} k & F^*/S^* \\ k & R^*/S^* \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} k & 1 \\ k & 0.6 \\ 0 & 1 \end{pmatrix}$$

The first column in $\mathbf{F}$ says that the feed flow F and the raffinate flow R have the same optimal sensitivity (k) with respect to the disturbance in x_E, whereas the sensitivity in the solvent feed S is zero because it is a disturbance. It is only the ratio between the column elements in $\mathbf{F}$ that matters, so the numerical value of the constant k is not important. The second column in $\mathbf{F}$ expresses that for disturbances in S the product composition is kept constant if F/S and R/S are constant. The numerical values follow since at the nominal point (denoted *), we have $F^* = 100$, $S^* = 100$ and $R^* = 60$. From the nullspace method, the optimal combination matrix $\mathbf{H}$ satisfies $\mathbf{HF} = 0$, and we find

$$\mathbf{HF} = (h_1 h_2 h_3) \begin{pmatrix} k & F^*/S^* \\ k & R^*/S^* \\ 0 & 1 \end{pmatrix} . = \Big(\underbrace{h_1 k + h_2 k}_{=0} \quad \underbrace{h_1 F^*/S^* + h_2 R^*/S^* + h_3}_{=0} \Big) = 0$$

We may always select one element in $\mathbf{H}$ freely, so setting $h_1 = 1$ we obtain from the first equation

$$h_2 = -h_1 = -1$$

and from the second equation

$$h_3 = -h1F^*/S^* - h_2 R^*/S^* = -(F^* - R^*)/S^* = -0.4.$$

Thus, the optimal measurement combination is

$$c = \mathbf{H}y = h_1 \Delta F + h_2 \Delta R + h_3 \Delta S = \Delta F - \Delta R - \frac{F^* - R^*}{S^*} \Delta S$$
$$= \Delta F - \Delta R - 0.4 \Delta S$$

As expected, the nullspace method yields a linearized version of the truly optimal nonlinear variable combination, which we found from physical arguments is $c = (F - R)/S$.

Approach 4 (*Maximum gain rule*). The maximum gain rule (Halvorsen et al., 2003) says that one should control "sensitive" variables with a large scaled gain $|G|$/span(CV). This captures two main concerns:

1. *The optimal value of the CV should be approximately constant* (independent of disturbances), that is, span(CV) should be small.
2. *The CV should be sensitive to changes in the unconstrained degrees of freedom* (to ensure a flat optimum), that is the gain $G = \Delta\text{CV}/\Delta\text{MV}$ should be large.

The maximum gain rule can be derived from the exact local method (with noise) by making some assumptions, which are generally satisfied except for some ill-conditioned cases (Halvorsen et al., 2003). An important advantage of the maximum gain rule is the insight that one should control variables that are sensitive with respect to the inputs.

5. Optimal operation of parallel units

Let us consider an important problem, often encountered in industrial practice. During the life of production of a product, a company often expands capacity as demand grows. Early plant design may involve process designs based on incomplete data as time to market drives commercialization timelines. Once operation begins, improvements and changes are made in operating conditions, equipment designs and process topology. When capacity expansion takes places the new capacity may come simply by adding equipment to the existing process or by construction of a parallel plant. The new plant is seldom run in a "stand alone" fashion, but instead may share some unit operations with the existing facility. As expansion continues, the complexity of the topology among the plants can lead to plantwide control problems.

In its simplest form, consider a number of plants operating in parallel, each of differing ages, and each with its own efficiency and yield relationships that are dependent on throughput. How should we optimally load each plant to achieve a target production while minimizing the total production cost J? We can derive useful result from the necessary optimality condition $J_u = 0$. We derive that, provided the total production rate is given, it is optimal to load the units such that we have *equal marginal costs in all units* (which corresponds to $J_u = 0$). This is a well-known fact but a simple proof is nevertheless enlightening.

Proof. To derive this result, consider *n independent parallel units* with a given total load (e.g., given total feed). Let the total cost be $J = \Sigma J_i$ and let the total feed (or some other limited load for the units) be fixed, $F = \Sigma F_i$. The necessary conditions of optimality is that $J_u = \delta J/\delta u = 0$ where u in this case is the vector of feed rates F_i. Since the total feed is fixed, there are $n - 1$ independent degrees of freedom F_i, and we assume these are the F_i's for $n - 1$ first units (and for unit n we have $F_n = F - \sum_{i=1}^{n-1} F_i$) The units are assumed to be independent which means that the cost in unit i, J_i, depends only on the flow into unit i, F_i. However, note that when we make a change in F_i, we also need to change F_n, and we have $dF_n = -dF_i$. The optimality condition $\delta J/\delta F_i = 0$ for variable F_i then becomes

$$\delta J/\delta F_i = \delta(J_1 + J_2 + \cdots J_i + \cdots J_n)/\delta F_i = \delta(J_i + J_n)/\delta F_i$$
$$= \delta J_i/\delta F_i - \delta J_n/\delta F_n = 0$$

or $\delta J_i/\delta F_i = \delta J_n/\delta F_n$. Since this must hold for all i units, we have proved that one should operate such that the *marginal cost* $\delta J_i/\delta F_i$ is the same in same units. □

There has been many applications of this criterion. Urbanczyk and Wattenberger (1994) applied it to the maximization of oil production of wells that produce both oil and gas, but where the total gas handling capacity is fixed (limited). In their application J_i is the oil production and F_i is the gas production in well i, and the idea is to operate the wells such that $\delta J_i/\delta F_i$ is the same for all wells; that is, by increasing the gas production by a given amount δF_i one gets the same benefit in terms of extra oil production δJ_i in all wells.

Good self-optimizing variables are then the difference in marginal cost between the units (which should be zero). Below we discuss two industrial applications from Eastman.

Example 5 (*Operation of parallel refining systems*). Eastman received an energy efficiency award for 2002 from the American Chemistry Council for its application of advanced control to optimally load three parallel refining systems. Each system consists of four distillation columns used to refine crude reactor product. The application uses process data to establish operating costs for processing material from crude reactor effluent to saleable product. *The allocation of load to each system is determined by adjusting the feed rates to each train to achieve equal marginal refining costs*.

Example 6 (*Syngas production in parallel furnaces*). For many years Eastman produced synthesis gas by reacting methane and steam in reforming furnaces. The process consisted of 15 furnaces operated in parallel, see Fig. 9. The effluent gas from the furnaces was combined as feed to three carbon dioxide removal systems. The product syngas from the three carbon dioxide removal systems was combined to form a single product gas used in downstream chemical production. The 15 reforming furnaces, constructed over the span of three decades, each had different energy efficiency characteristics as well as different yield performance as technology advanced. In addition, the three carbon dioxide removal systems were of varying efficiency and performance. Newer systems were better instrumented, had valves that performed better, and had on-line analytical measurements. At any time, there were one to three furnaces down for routine maintenance.

The optimum operation of the plantwide system to coordinate pressure and production among the interconnected gas flow network was a significant challenge. The simple objective of matching production of syngas with consumption often ended up varying the production rate on the newest furnace because it could most gracefully handle the needed changes. From an optimization point of view this approach usually resulted in the most efficient units not being operated at their maximum rates.

Often with units in parallel, a reasonable "self-optimizing" strategy is to operate with the same outlet conditions (temperatures or compositions) of all parallel units. At least this avoids mixing losses, and indeed this would have been a good strategy if the reactors were identical. However, in general, and certainly for this example it is more economical to operate each furnace differently based on its particular efficiency and yield profile and then ensure that the combination stream met the total stream specifications. In particular, the newer more efficient furnaces were able to produce a much purer product for the same cost as the older units producing a much less pure product. The purity of the product from each furnace was a relatively weak function of feed rate. The final layer of complexity arises from the efficiency of the carbon dioxide removal system. Each system was connected to a designated set of furnaces so that it was beneficial to operate furnaces linked to the better performing carbon dioxide removal system.

The optimization layer to coordinate the total process production and the allocation of that production to various parallel units was complicated by the presence of crossover lines. These lines added operational flexibility but created an ever increasing complexity of the optimization problem. Local MPC controllers for furnace operation and supervisory control for the carbon dioxide removal systems allowed for near optimal operation at the local level illustrated in Fig. 10. Overall optimization was approached by production loading strategies and coordination using a supervisory MPC controller. As solutions to this problem were developed, it became clear that technology to guide us on the appropriate degree of decentralization was sparse. Developing a centralized system with all the CV's and MV's in the same MPC was problematic due to the routine on-line/off-line operation of the furnaces. Being able to gracefully add and remove systems from the overall control system was critical to success. In addition, measurement reliability often resulted in some furnaces being operated in "local" mode; i.e., not connected to the centralized MPC. The eventual control system needed to be developed and commissioned in reasonable time, needed to be implemented on available hardware, needed to be understood by plant operating staff, and had to be maintainable as process improvements were made. This led to a decentralized strategy choice as shown in Fig. 11.

These examples illustrate the complex nature of an industrial plantwide control problem. The use of a formalized procedure can make known improved strategies that may go undetected

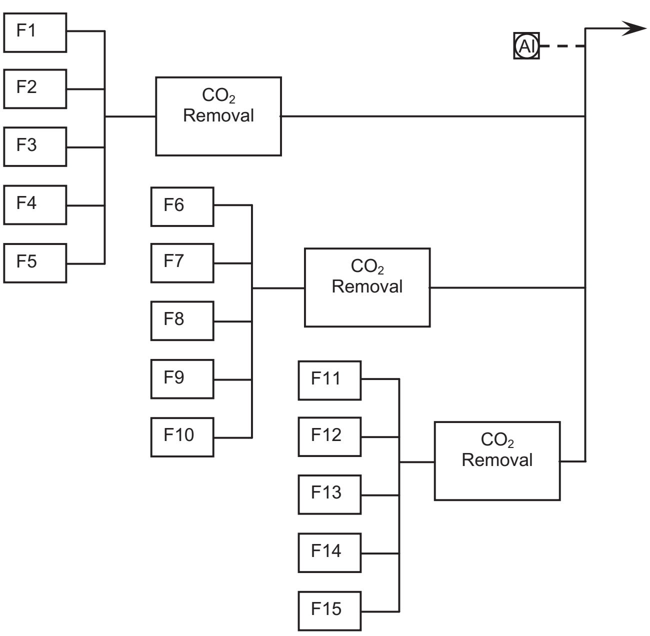

Fig. 9. Syngas process with 15 furnaces and three CO_2 removal systems.

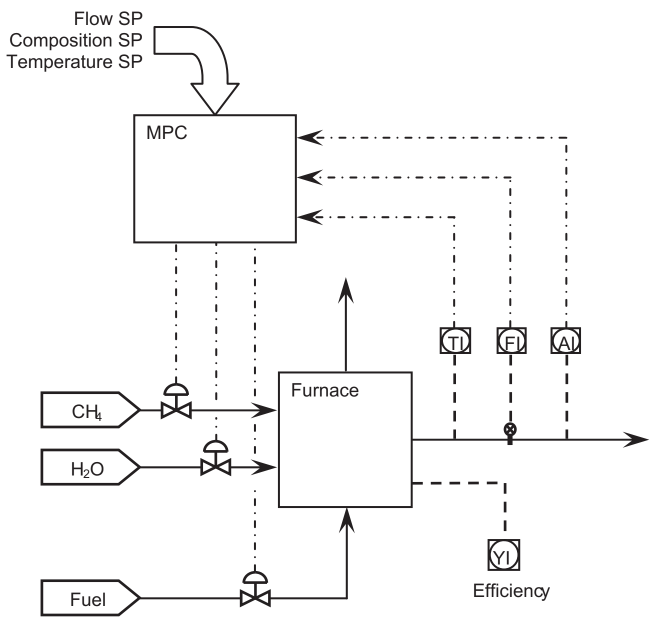

Fig. 10. Individual syngas furnace control with three manipulated variables and three controlled variables.

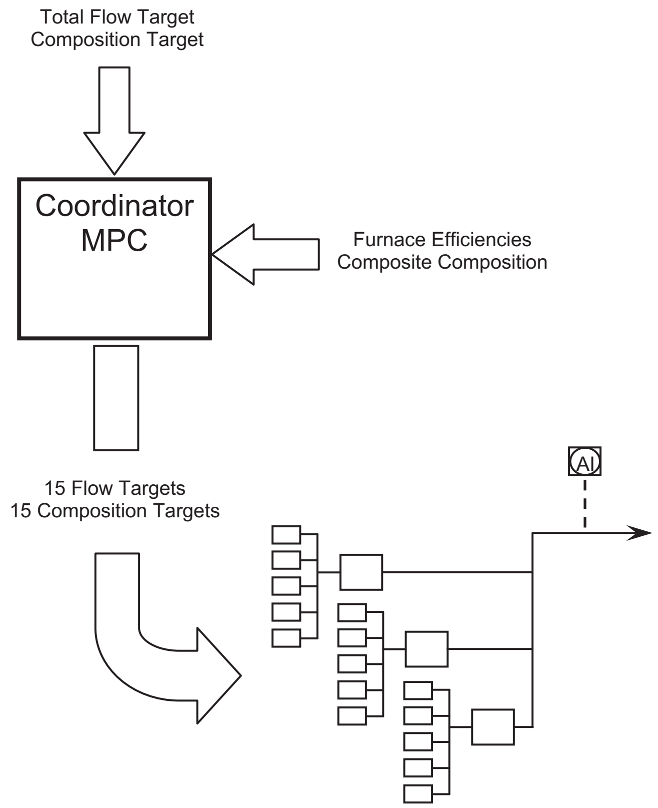

Fig. 11. Coordination MPC supervising 15 local furnace MPC controllers.

when using only one or two design criteria for guidance. The ability to weave in practical issues that complicate implementation is paramount. Using a formalized procedure can help unscramble the vast array of decisions that can overwhelm designers and cause them to continue reliance upon a unit operation focus.

6. Discussion

The design of plantwide control strategies can be seen from two viewpoints. These are (1) the design of control strategies for the regulation, including that of plant material and energy inventories, and (2) the design of control strategies for process economic optimization. The design of inventory control strategies determines the manipulated variables that remain for process economic optimization. Concepts are needed that guide the design of the inventory control strategy such that the design of the process economic optimization strategy is made easy. It is clear that if a good job is done during the design of the inventory control strategy, such as setting the TPM near the bottleneck, then the remaining process economic optimization strategy design is made easier. On the other hand, if the inventory control strategy results in key optimization variables being far away from available manipulated variables, then strategies for optimizing process economics will be difficult if not impossible to implement.

The emphasis upon placement of the TPM for a process has long been recognized as a key decision in the resulting inventory control strategy. It is becoming more evident that this decision also determines the difficulty of the remaining process economic optimization strategy design. Techniques to determine self-optimizing control variables can be effectively and easily employed if the variables available to optimize the process have good dynamic linkage with their manipulated variable counterparts.

The examples also illustrate that the application of optimization from a top down viewpoint may guide one to select manipulated variables that should remain free for economic optimization while other should be used for inventory control. The formalization of a procedure to organize the design of these two phases of control includes the concepts of: (1) TPM location within the process, (2) control of unit operation process variables against their local constraints, and (3) the development of measurement combinations whose control implies nearness to the economic optimum.

The application of plantwide control design procedures for new plants is certainly an obvious direction of growth. However, the redesign of plantwide control structures for existing plants has been shown to be very beneficial. The known locations of process bottlenecks, known market conditions and product demands, and the operating nuances of a running process all make the plantwide design procedure more understandable and manageable. Using a procedure to determine alternate control structures can lead to new ideas for control that may have been missed for existing processes. As noted, the migration from tried and true control, but inferior, control strategies to new and unfamiliar strategies can be difficult.

7. Conclusion

Since the concept of process control design based on a holistic view of the process came about, the variety of procedures and approaches to the design problem have illustrated the difficulty of a "one size fits all approach". The examples presented illustrate the application of a few industrial design approaches. A more formal design procedure is presented and it was shown on industrial examples to yield in a systematic way the same control strategies that were obtained with other more ad-hoc criteria. The importance of addressing process economics in the control design procedure is discussed and the industrial need to run plant at their maximum feed rate (mode II) is emphasized. The use of a plantwide design procedure that incorporates and organizes the variety of concerns and technical issues in this important area is demonstrated.

References

Alstad, V., & Skogestad, S. (2007). Null space method for selecting optimal measurement combinations as controlled variables. *Industrial and Engineering Chemistry Research, 46*(3), 846–853.

Alstad, V., Skogestad, S., & Hori, E. S. (2009). Optimal measurement combinations as controlled variables. *Journal of Process Control, 19*, 138–148.

Araujo, A. C. B., Govatsmark, M., & Skogestad, S. (2007). Application of plantwide control to the HDA process. I Steady-state and self-optimizing control. *Control Engineering Practice, 15*, 1222–1237.

Aske, E. M. B. (2009). Design of plantwide control systems with focus on maximizing throughput. PhD Thesis, Norwegian University of Science and Technology (NTNU).

Aske, E. M. B., Strand, S., & Skogestad, S. (2008). Coordinator MPC for maximizing plant throughput. *Computers and Chemical Engineering, 32*, 195–204.

Buckley, P. S. (1964). *Techniques of process control*. New York: Wiley (Chapter 13).

Chen, R., & McAvoy, T. J. (2003). Plantwide control system design: Methodology and application to a vinyl acetate process. *Industrial and Engineering Chemistry Research, 42*(20), 4753–4771.

Douglas, J. M. (1988). *Conceptual design of chemical processes*. New York: McGraw-Hill.

Downs, J. J. (1992). Distillation control in a plantwide control environment. In W. L. Luyben (Ed.), *Practical distillation control* (pp. 413–439). Van Nostrand Reinhold.

Downs, J. J., & Vogel, E. F. (1993). A plant-wide industrial process control problem. *Computers and Chemical Engineering, 17*, 245–255.

Halvorsen, I. J., & Skogestad, S. (1997). Optimizing control of Petlyuk distillation: Understanding the steady-state behavior. *Computers and Chemical Engineering, 21*(Suppl.), S249–S254.

Halvorsen, I. J., Skogestad, S., Morud, J. C., & Alstad, V. (2003). Optimal selection of controlled variables. *Industrial and Engineering Chemistry Research, 42*(14), 3273–3284.

Hori, E. S., Skogestad, S., & Alstad, V. (2005). Perfect steady-state indirect control. *Industrial and Engineering Chemistry Research, 44*(4), 863–867.

Konda, N. V. S. N. M., Rangaiah, G. P., & Krishnaswamy, P. R. (2005). Plantwide control of industrial processes: An integrated framework of simulation and heuristics. *Industrial and Engineering Chemistry Research, 44*(22), 8300–8313.

Kookos, I. K., & Perkins, J. D. (2002). An algorithmic method for the selection of multivariable process control structures. *Journal of Process Control, 12*, 85–99.
Larsson, T., & Skogestad, S. (2000). Plantwide control: A review and a new design procedure. *Modeling, Identification and Control, 21*, 209–240.
Luyben, W. L. (1994). Snowball effects in reactor separator processes with recycle. *Industrial and Engineering Chemistry Research, 33*, 299–305 (Res. 34(11), 3885–3898).
Luyben, W. L., Tyreus, B. D., & Luyben, M. L. (1997). Plantwide control design procedure. *AIChE Journal, 43*(12), 3161–3174.
Luyben, W. L., Tyreus, B. D., & Luyben, M. L. (1998). *Plantwide process control*. New York: McGraw-Hill.
Morud, J., & Skogestad, S. (1996). Dynamic behaviour of integrated plants. *Journal of Process Control, 6*, 145–146.
Narraway, L. T., & Perkins, J. D. (1993). Selection of control structure based on economics. *Computers and Chemical Engineering, 18*, S511–S515.
Price, R. M., & Georgakis, C. (1993). Plantwide regulatory control design procedure using a tiered framework. *Industrial and Engineering Chemistry Research, 32*, 2693–2705.
Skogestad, S. (2000). Plantwide control: The search for the self-optimizing control structure. *Journal of Process Control, 10*, 487–507.
Skogestad, S. (2004). Control structure design for complete chemical plants. *Computers and Chemical Engineering, 28*(1–2), 219–234.
Srinivasan, B., Bonvin, D., & Visser, E. (2003). Dynamic optimization of batch processes – III. Role of measurements in handling uncertainty. *Computers and Chemical Engineering, 27*(1), 27–44.
Urbanczyk, C. H., & Wattenberger, R. A. (1994). Optimization of well rates under gas coning conditions. *SPE Advanced Technology Series, 2*(2), 61–68.
Vasbinder, E. M., & Hoo, K. A. (2003). Decision-based approach to plantwide control structure synthesis. *Industrial and Engineering Chemistry Research, 42*, 4586–4598.
Ward, J. D., Mellichamp, D. A., & Doherty, M. F. (2004). Importance of process chemistry in selecting the operating policy for plants with recycle. *Industrial and Engineering Chemistry Research, 43*, 3957.
Ward, J. D., Mellichamp, D. A., & Doherty, M. F. (2006). Insight from economically optimal steady-state operating policies for dynamic plantwide control. *Industrial and Engineering Chemistry Research, 45*, 1343.
Zheng, A., Mahajanam, R. V., & Douglas, J. M. (1999). Hierarchical procedure for plantwide control system synthesis. *AIChE Journal, 45*(6), 1255–1265.

Dr. **James J. Downs** (PE) is an Engineering Fellow and manager of the Advanced Controls Technology group at Eastman Chemical Company. He has 33 years of experience in the design, startup, and support of industrial processes. His work has included the design of advanced control systems for reaction systems, distillation processes, polymer processes, extraction and other separations processes, gas handling systems and compressors, and other unit operations. He developed regulatory and overall plant-wide operation strategies for numerous capital expansion projects combining plant design and control system design techniques. He has implemented and applied to chemical process operations infinite horizon MPC, on-line Kalman filters, and on-line process optimization technology. His current research interests include plant-wide control strategy design, plant-wide process optimization, and the process design/process control interface. He was recognized by the AIChE for his contribution to the profession by receiving the CAST Computing Practice Award for 1996. His wife Patricia and he have been married 34 years and have two sons, ages 24 and 27. He is an instrument flight instructor and his hobbies include flying, traveling, and biking.

Sigurd Skogestad received his Ph.D. degree from the California Institute of Technology, Pasadena, USA in 1987 and has since then been a full professor at Norwegian University of Science and Technology (NTNU), Trondheim, Norway. He was Head of the Department of Chemical Engineering from 1999 to 2009. He is the principal author, together with Prof. Ian Postlethwaite, of the book "Multivariable feedback control" published by Wiley in 1996 (first edition) and 2005 (second edition). The book "Chemical and Engineering Process Engineering" was published by CRC Press in 2009. His research interests include the use of feedback as a tool to make the system well-behaved (including self-optimizing control), limitations on performance in linear systems, control structure design and plantwide control, interactions between process design and control, and distillation column design, control and dynamics. His other main interests are mountain skiing (cross country), orienteering (running around with a map) and grouse hunting. His wife Anne-Lise and he have been married 36 years and have four children.