An Industrial and Academic Perspective on Plantwide Control

James J. Downs*. Sigurd Skogestad**

**Eastman Chemical Company, Kingsport, TN 37662, USA*
(e-mail: jjdowns@ eastman.com).
*** Dept. of Chemical Engineering, Norwegian University of Science and Technology, N-7491 Trondheim, Norway*
(email: skoge@ntnu.no)

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 have illustrated the difficulty of a unified approach. Using examples, a formal design approach is presented to help put in context the need and advantages of using such an approach. The examples deal with disturbance rejection, throughput maximization and economic optimization of plants consisting of parallel units.

Keywords: process control, control structure design, plantwide control, inventory control, throughput

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., Downs, (1992), Narraway and Perkins (1993), Luyben (1998), Zheng, Mahajanam and Douglas (1999), Kookos and Perkins (2002), Chen and McAvoy (2003), Vasbinder and Ho (2003), Skogestad (2004), Konda et al. (2005), Ward et al. (2006)). 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 and Vogel, 1993) was put forward so that various approached could be tested against each other. Nevertheless, today in industry, much of the research in this area has still not gained traction to have the profound influence possible. 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.

In section 2 industrial aspects of plantwide control design are discussed and two examples illustrate how industry may approach the plantwide control problem using a single criterion for design guidance. This points to the 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 plantwide design involves many issues and one-criteria approaches may not be sufficient.

2. STATUS IN INDUSTRY

The traditional approach for designing process control strategies for chemical plants has been to set production rates by setting process feed rates and then to design automatic control systems around each unit operation sequentially through the process. For processes having significant in-process inventory and not too much in the way of recycles, 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 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 is one 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 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 design point. Design decisions are often based on steady-state analysis without consideration of 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).

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, can result in more reliable operation.

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 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 et al., 2004, 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 SISO while the overall plant wide strategy is optimizing plant operation in a more natural fashion. This approach has 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 normal operation. Contrasted with the centralized approach of using models to determine an optimum and then driving a process to that optimum point, optimization designed in from the bottom up provides the important robustness and reliability component.

From start-up the primary objective for a new plant is to achieve nameplate capacity in a reliable and predicable way. Often times 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. We can all attest to the uphill battle to achieve routine usage of a control system 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 is that for plants where “self-optimizing” regulatory control strategies have been build 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 endeavour.

Example 2.1 - 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. Units that were built 50 years ago were typically 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. The original plant had the standard scheme with the throughput manipulator (TPM) located at the feed as illustrated in Figure 2-1.

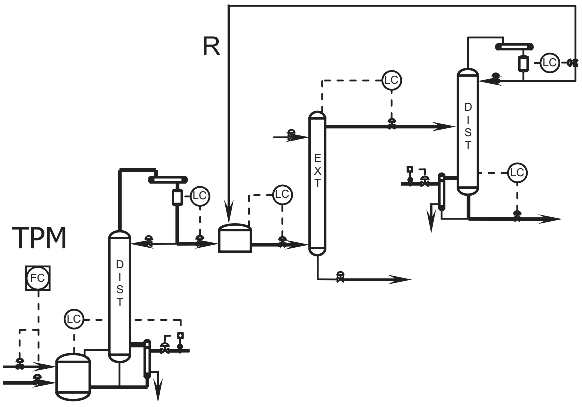

Figure 2-1. Esterification process. Inventory control strategy with column 1 feed rate used as the TPM.

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 Figure 2-2. This strategy worked well for many years because many of the disturbances entering the reactor were directed away from the more sensitive separation portion of the process. The extraction step of the process was intended to wash unreacted alcohol from the ester product. As the extraction step became the process bottleneck, it became evident that its behaviour as a function of organic feed rate was very nonlinear. This nonlinearity stemmed from the fact that increasing organic

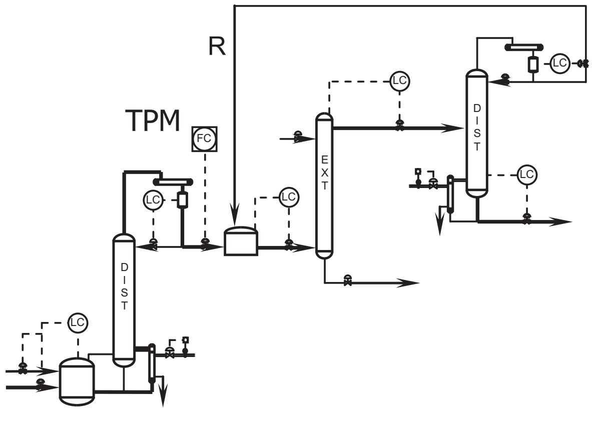

Figure 2-2. Esterification process. Inventory control strategy with column 1 distillate rate used as the TPM.

feed rate resulted in an increasing composition of the alcohol taken from the extractor to the final distillation column. The increase in distillate rate needed to remove the alcohol from the final product would aggravate the situation by increasing the feed rate to the extractor (stream "R" in Figure 2-2). 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 12-24 hours 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 (Figure 2-3). 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, take note of the system response to the extra distillate flow recycled to the extractor. Namely, it results in less flow being 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 principle that proved most useful is the idea that the optimum did not lie against or close to a process cliff. The original strategy was very unforgiving once the process was pushed too far. Extractor flooding, loss of liquid/liquid immiscibility, and flooding of the final distillation column meant several hours of lost production. The ability to experiment with the process without the penalty of passing this “point of no return”, gave operators confidence in the control system to recover if they ended up pushing rates too high.

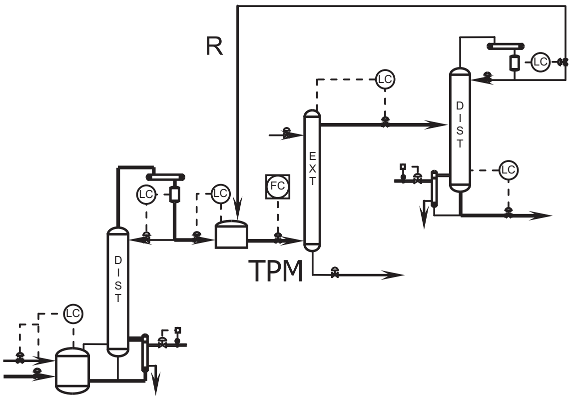

Figure 2-3. Esterification process. Inventory control strategy with extractor feed rate used as the TPM.

Example 2.2 - Control strategy for a liquid-liquid extraction process: During the control design phase one may chose from a variety of criteria to drive the control strategy design 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 how the choice of a design criterion that aims to *propagate disturbances to insensitive locations* results in a particular design. The resulting control strategy can then be compared with those obtained using a more methodical approach.

Consider the extraction process in Figure 2-4 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 product (R). 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 Figure 2-5.

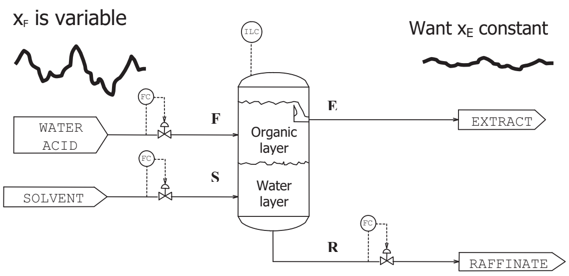

Figure 2-4. Liquid-Liquid Extraction Process

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%). 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 ±0.506% (21.4 ±0.506%). For details see the mass balances in Table 1. Thus, strategy II is the preferred strategy of the two if the objective is to have small variations in extract composition, x_E.

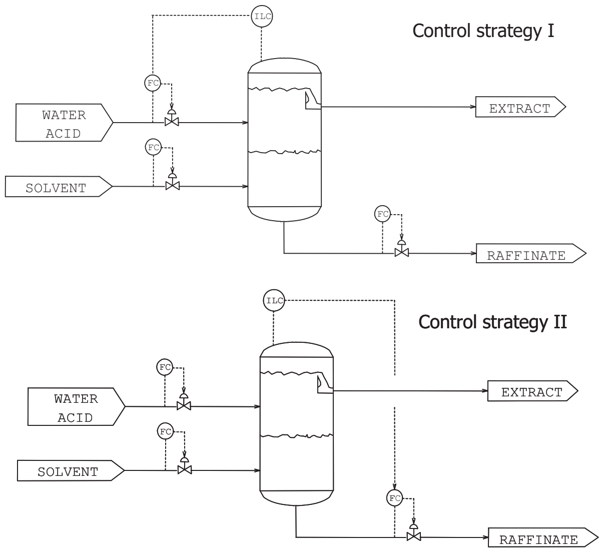

Figure 2-5. Alternative Control Strategies for Liquid-Liquid Extraction

In summary, for strategy I 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 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 if the main objective is to have small variations in x_E. This example suggests that different inventory strategies may result in process variability being transferred to portions of the process that are insensitive to variation or portions in which variability is harmful. The idea of propagating disturbances to insensitive locations gives good insight and can result in good designs. However, for more complex problems and for less experienced engineers a more systematic approach is needed.

	Feed, F				S	Extract, E				Raffinate, R			
Case	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	0	140	142.273	141.33	140	60	**60**	58.67	56.77

Table 1. Mass balances for extraction process: Base case ($x_F = 30\%$) and with disturbance ($x_F = 31\%$) for control strategies I, II, III, and IV. (Assumption: Equilibrium relationship Acid/Water = 3 in extract, E.)

3. A PLANTWIDE CONTROL DESIGN PROCEDURE

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.

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 2.2 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 to 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 to 7. A detailed analysis in steps 1 and 2 requires that one has available a steady-state model and that one performs optimizations using the model for various disturbances. This is often not done in industrial practice. 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 et al. (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 to 4.

Step 4 (location of TPM) was addressed in Example 2.1, and this issue is further discussed in the recent PhD thesis by Aske (2009); see also Aske et al. (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. Changes in production rate are considered a disturbance.

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 capacity 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).

The esterification plant in Example 2-1 is a case of operating in mode II with the extraction section being the bottleneck.

Example 3-1 - Application of the design procedure to Example 2-2: The design criterion for Example 2-2 was that disturbances should be propagated to insensitive locations. At this point we want apply the more systematic plantwide procedure. The process is very simple, so we use the simplified approach for selection of controlled variables (there are no degrees of freedom left for economic optimization once the specifications are satisfied).

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 location of the TPM influences the structure (pairing) of the inventory control system in Step 5. The throughput is often located at the main feed, but could generally be anywhere in the process. Since the two proposed control strategies both have a constant solvent feed flow, we assume here that the solvent feed S is the throughput manipulator ***(Decision 2).***

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. With solvent feed rate S as the TPM, we have left two candidate MVs: Feed F and outflow R. The main issue for regulatory 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 ($CV_1 = x_E$) which depends directly on the feed F but only indirectly on the outflow R. Thus, we would like to "save" F for the supervisory layer.

Decision 3. Use R to control the interface layer ($MV_2 = R$). This gives inventory control in the direction of flow, which is normal with the throughput set at the feed.

Step 6. Structure of supervisory control layer

Decision 3, continued. The remaining $MV_1 = F$ is used to control acid composition ($CV_1 = x_E$). The final control structure is shown as strategy III in Figure 3-1.

Note that we assumed that the product composition x_E can be measured ($CV_1 = x_E$), but this may not be possible in practice. We then need to find something else to "control" (keep constant). This is what we indirectly did 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 give undesired 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. It is possible to add a supervisor layer, where one adjusts R (strategy I) or F (strategy II) such that x_E is kept constant. This modification to strategy II is shown as *Strategy III* in Figure 3-1.

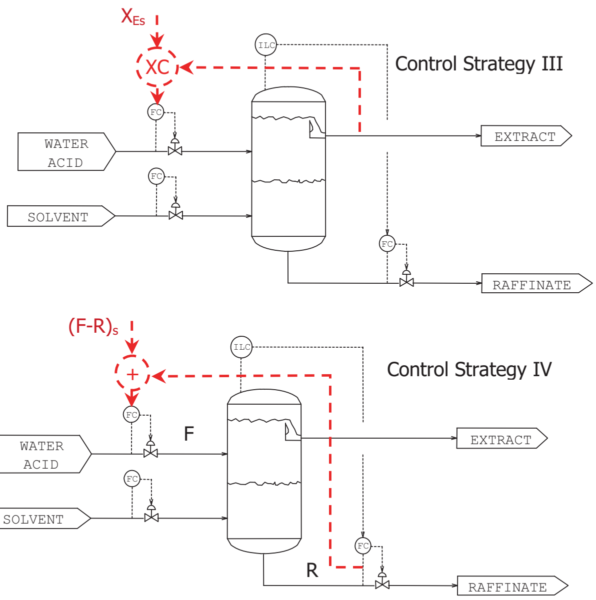

Figure 3-1. Self-Optimizing Strategies for Extraction Process

However, assume there is no online measurement of the extract composition x_E. One option would then be to estimate x_E using a model and available measurements ("soft sensor"), but this is a bit complicated. 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 $CV_1 = F - R$ constant (while at the same time adjusting F and R to control interface level). One possible implementation is shown in Figure 3-1.

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 material balance we have $E = S + (F\text{-}R)$ 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 that the composition of E must remain constant (again assuming S constant). If the throughput S varies (disturbance) then all flows should be scaled by S to keep x_E constant, so by process insight we derive that an "improved strategy IV" is to keep the variable $CV_1 = (F\text{-}R)/S$ 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 and 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 $CV_1 = h_1 F + h_2 R$, where in general the optimal h_1 and h_2 can be found from the nullspace theorem. In this example, we found by process insight that the optimal choice is $h_1 = 1$ and $h_2 = -1$ (strategy IV).

4. SELECTION OF ECONOMIC (PRIMARY) CVs

In the above example, we found that the flow difference F-R is a good primary CV. How do we select 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? The issue is to find some "magic" variable, c, to keep constant. The obvious "magic" variable 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 idea, let us look in a bit more detail in Steps 1 to 3 in the proposed procedure for selecting economic CVs.

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

In many cases a simple economic cost is used:

Profit = - J = value products – cost feeds – cost utilities (energy)

Other operational issues, such as safety and environmental impact are usually formulated as constraints. For cases with good marked conditions we often have a constrained optimum and the cost function can be simplified to J = - 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 this is not possible, even at steady-state, because there is no online measurement. We therefore want to control something else that gives *indirect control* of the primary output (Hori et al., 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 is to use a steady-state flowsheet simulator, if available, to optimize operation (with respect to the degrees of freedom) for various disturbances. 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 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.*

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.

2. Identify "self-optimizing" variables related to the (possibly) remaining unconstrained degrees of freedom. These are "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).

To identify good candidates 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),

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

The term "brute force" is used is 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:

At optimum: Gradient = $J_u = dJ/du = 0$

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

Approach 3 - Exact local method and optimal measurement combinations. The details are found in Halvorsen et al (2003), Alstad and Skogestad (2007) and Alstad et al. (2009). For the case single measurements as CVs, this is a "local" version of the brute force approach. However, the evaluation is much more efficient. In addition, the "nullspace method" can be used to find truly optimal measurement combinations, as was done in strategy IV for the extraction process.

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 = ΔCV / ΔMV should be large.

The maximum gain rule can be derived from the exact local method by making some not too serious assumptions. An important advantage of the maximum gain rule is the insight that it gives.

5. OPTIMAL OPERATION OF PARALLEL UNITS

Let us return to an important problem, often encountered in industrial practice. During the life of production of a product, a company often times 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, improved operating conditions, equipment designs, and process topology emerge. 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 plant wide 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 costs? 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 $J_u = 0$).

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

$$\begin{aligned}\delta J/\delta F_i &= \delta(J_1 + J_2 + \dots J_i + \dots J_n)/\delta F_i = \delta(J_i + J_n)/\delta F_i \\ &= \delta J_i/\delta F_i - \delta J_n/\delta F_n = 0\end{aligned}$$

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.* **End proof.**

Urbanczyk and Wattenbarger (1994) applied this criterion 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 of this idea.

Example 5-1 – Operation of parallel refining systems: Eastman received an industry award 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. Based on operating costs, process operation limits, and utility availability, the feed rate to each refining train is adjusted to match reactor production with refining system production. *The allocation of load to each system is adjusted to achieve equal marginal refining costs.*

Example 5-2 – 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 Figure 5-1. 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.

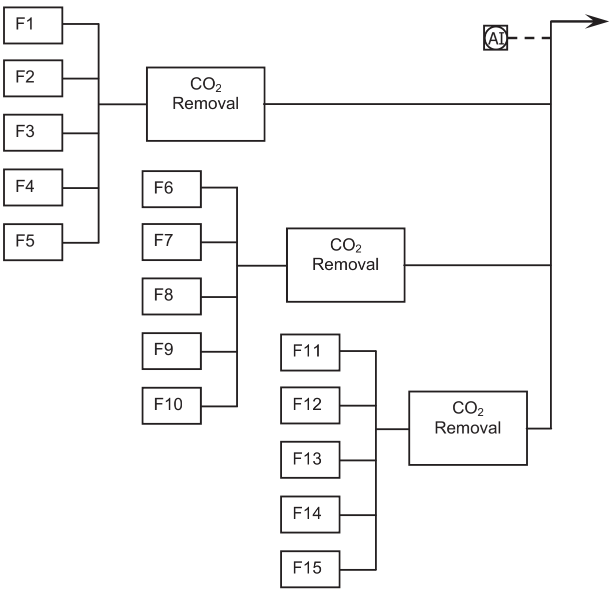

Figure 5-1. Syngas process with fifteen furnaces and three CO_2 removal systems.

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.

Normally with units in parallel, an expected "self-optimizing" strategy is to operate with the same outlet conditions (temperatures or compositions) of all parallel units. This would have been a good strategy if the reactors were identical, but, 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 Figure 5-2. 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

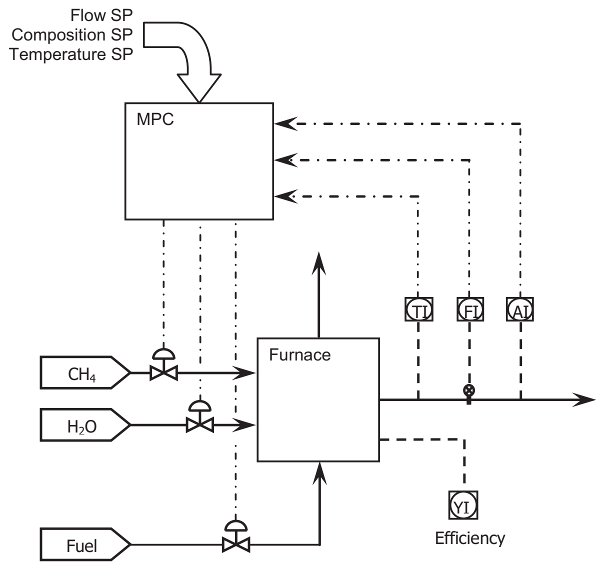

Figure 5-2. Individual syngas furnace control with three manipulated variables and three controlled variables.

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 Figure 5-3.

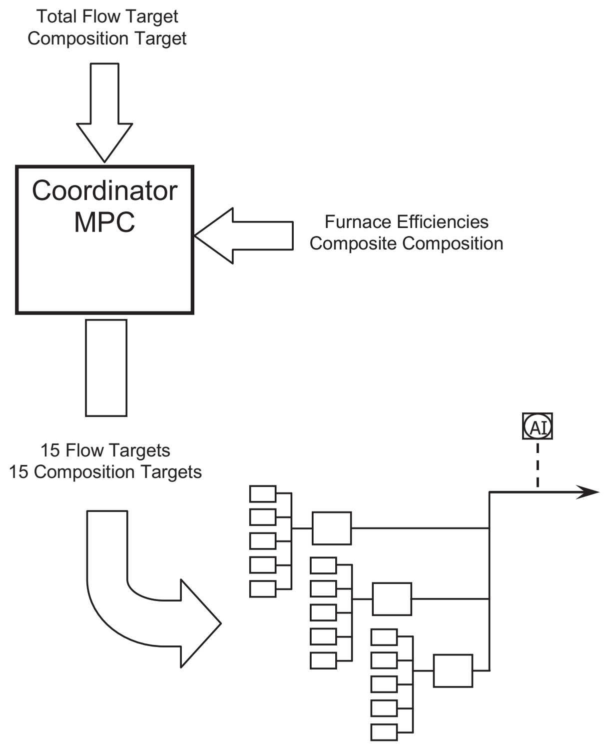

Figure 5-3. Coordination MPC supervising fifteen local furnace MPC controllers.

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 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 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 examples illustrate that the inventory control strategy design not only affects the dynamics between manipulated and controlled variables used for optimization, but also can change the gain as well. 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 is applied to the industrial examples. 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

V. Alstad and S. Skogestad, “Null Space Method for Selecting Optimal Measurement Combinations as Controlled Variables”, *Ind. Eng. Chem. Res.*, **46** (3), 846-853 (2007).

V. Alstad, S. Skogestad and E.S. Hori, “Optimal measurement combinations as controlled variables”, *J. Proc. Control*, **19**, 138-148 (2009).

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

E.M.B. Aske, S. Strand and S. Skogestad, “Coordinator MPC for maximizing plant throughput”, *Computers and Chem. Eng.,* **32**, 195-204 (2008).

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

P.S. Buckley, *Techniques of Process Control*, Wiley, New York (1964), Chapter 13.

R. Chen and T.J. McAvoy, “Plantwide control system design: Methodology and application to a vinyl acetate process”, *Ind. Eng. Chem. Res.*, **42** (20), 4753-4771 (2003).

J.M. Douglas, *Conceptual Design of Chemical Processes*, McGraw-Hill, New York (1988).

J.J. Downs, “Distillation Control in a Plantwide Control Environment”, In: *Practical Distillation Control*, W.L. Luyben (ed.), Van Nostrand Reinhold, 413-439 (1992)

J.J. Downs and E.F. Vogel, “A Plant-Wide Industrial Process Control Problem”, *Computers and Chem. Eng.,* **17**, 245-255 (1993)

I.J. Halvorsen and S. Skogestad, “Optimizing control of Petlyuk distillation: Understanding the steady-state behavior”, *Computers and Chem. Eng.,* **21**, Suppl., S249-S254 (1997).

I.J. Halvorsen, S. Skogestad, J.C. Morud and V. Alstad, “Optimal selection of controlled variables”, *Ind. Eng. Chem. Res.*, **42** (14), 3273-3284 (2003).

E.S. Hori, S. Skogestad and V. Alstad, “Perfect steady-state indirect control”, *Ind. Eng. Chem. Res.*, **44** (4), 863-867 (2005).

NVSNM Konda, G.P. Rangaiah and P.R. Krishnaswamy, “Plantwide control of industrial processes: An integrated framework of simulation and heuristics”, *Ind. Eng. Chem. Res.* , **44** (22), 8300-8313 (2005).

I.K. Kookos and Perkins, J.D., “An Algorithmic method for the selection of multivariable process control structures”, *J. Proc. Control*, 12, 85-99 (2002).

W.L. Luyben, B.D. Tyreus and M.L. Luyben, *Plantwide Process Control*, McGraw-Hill, New York (1998).

W.L. Luyben, “Snowball Effects in Reactor Separator Processes with Recycle”, *Ind. Eng. Chem. Res.* , **33**, 299-305 (1994).

L.T. Narraway, J.D. Perkins, “Selection of control structure based on economics”, *Computers and Chem. Eng.,* **18**, S511–S515 (1993).

S. Skogestad “Plantwide control: the search for the self-optimizing control structure”, *J. Proc. Control*, **10**, 487-507 (2000).

S. Skogestad, “Control structure design for complete chemical plants”, *Computers and Chem. Eng.,* **28** (1-2), 219-234 (2004).

B. Srinivasan, D. Bonvin and E. Visser, “Dynamic optimization of batch processes – III. Role of measurements in handling uncertainty”, *Computers and Chem. Eng.,* **27** (1) 27–44 (2003).

E.M. Vasbinder and K.A. Hoo, “Decision-based approach to plantwide control structure synthesis”, *Ind. Eng. Chem. Res*., **42**, 4586-4598 (2003).

J.D. Ward, D.A. Mellichamp, and M.F. Doherty, “Importance of Process Chemistry in Selecting the Operating Policy for Plants with Recycle”, *Ind. Eng. Chem. Res*., **43**, 3957 (2004).

J.D. Ward, D.A. Mellichamp, and M.F. Doherty, “Insight from Economically Optimal Steady-State Operating Policies for Dynamic Plantwide Control”, *Ind. Eng. Chem. Res*., **45**, 1343 (2006).

C.H. Urbanczyk and R.A. Wattenberger, “Optimization of well rates under gas coning conditions”, *SPE Advanced Technology Series*, **2**, (2), 61-68 (1994).

A. Zheng, R.V. Mahajanam, J.M. Douglas, “Hierarchical procedure for plantwide control system synthesis”, *AIChE Journal*., **45**, (6) 1255-1265 (1999)