Plantwide Control

Recent Developments and Applications

Edited by

GADE PANDU RANGAIAH
Department of Chemical and Biomolecular Engineering, National University of Singapore, Singapore

VINAY KARIWALA
School of Chemical and Biomedical Engineering, Nanyang Technological University, Singapore

ISBN 978-0-470-98014-9

WILEY
A John Wiley & Sons, Ltd., Publication

11

Economic Plantwide Control

Sigurd Skogestad
Department of Chemical Engineering, Norwegian University of Science and Technology (NTNU), Trondheim, Norway

11.1 Introduction

A chemical plant may have thousands of measurements and control loops. The term plantwide control does not imply the tuning and behavior of each of these loops, but rather the control philosophy of the overall plant with emphasis on the structural decisions. In practice, the control system is usually divided into several layers, separated by timescale (see Figure 11.1).

In this introduction, I will temporarily switch to a less formal tone and try to tell the reader how I personally got involved in the plantwide control field. My interest dates back to 1983 when I started my PhD work at Caltech working with Professor Manfred Morari as my supervisor. The main theme of my PhD work was robust control, but as an application I worked on distillation column control which is an excellent example of a plantwide control problem. I was inspired by Greg Shinskey's book on Distillation Control, for which the second edition came out in 1984 (Shinskey, 1984). In particular, I liked his systematic procedure which involved computing the steady-state relative gain array (RGA) for 12 different control structures ('configurations'); the DV-configuration, LV-configuration, ratio configuration, and so on. However, when I looked in more detail at the procedure, I discovered that its theoretical basis was weak.

First, it did not actually include all structures and it even eliminated the DB-configuration (where D and B remain as degrees of freedom after using reflux L and boil-up V for level control) as 'impossible' even though it is workable in practice (Finco *et al*. 1989). Second, controllability theory tells us that the steady-state RGA is actually not useful itself,

Plantwide Control: Recent Developments and Applications, First Edition. Edited by Gade Pandu Rangaiah and Vinay Kariwala.
© 2012 John Wiley & Sons, Ltd. Published 2012 by John Wiley & Sons, Ltd.

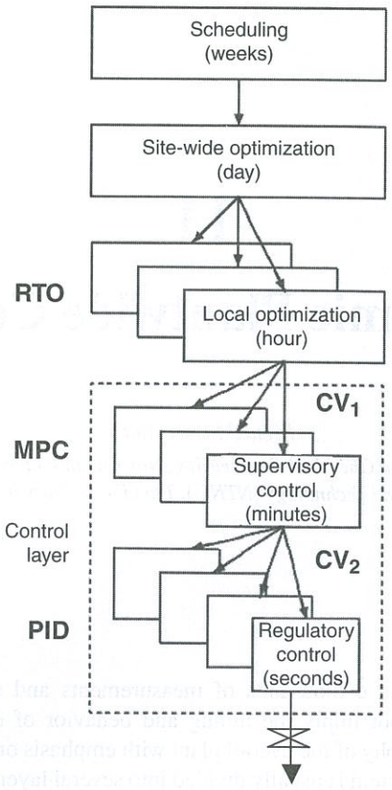

Figure 11.1 *Typical control hierarchy in a chemical plant.*

except that we should avoid pairing on negative gains (Skogestad and Postlethwaite, 2005). Third, the procedure focused on dual-composition control, while in practice we usually use single-end control (e.g., because it may be optimal economically to use maximum heating to maximize the recovery of the valuable product; Skogestad, 2007).

Furthermore, when I studied the distillation column control problem in more detail, I discovered that there were several control objectives which were often conflicting. First, there was the issue of 'stabilizing control' which involved closing the level and pressure loops and possibly also a temperature loop, so that the column did not drift and could be controlled manually without too much effort. Second, there was the issue of 'economic control' (advanced and supervisory control) which involves keeping the column close to its economically optimal operation. In many cases, economic control was the same as dual-composition control, but not always.

Depending on market conditions and disturbances, the best economic mode of operation actually changes. For a distillation column, it is always optimal to control the valuable product at its specification to avoid product 'give-away'. However, for the low-value product it is often optimal to over-purify in order to minimize the loss of valuable product. If product

prices are sufficiently high (compared to energy prices), then it is optimal to use maximum energy (boil-up) to get maximum over-purification. The important conclusion from this is that the optimal configuration will change depending on market conditions, so there is no single 'best' control configuration even for a given column (Skogestad, 2007).

Enough about distillation! Another influence on my work was the famous critique article on process control by Foss (1973). He writes:

> The central issue to be resolved ... is the determination of control system structure. Which variables should be measured, which inputs should be manipulated and which links should be made between the two sets? There is more than a suspicion that the work of a genius is needed here, for without it the control configuration problem will likely remain in a primitive, hazily stated and wholly unmanageable form. The gap is present indeed, but contrary to the views of many, it is the theoretician who must close it.

Foss states that determination of control system structure which, for process control is the same as what I call plantwide control, is "the central issue to be resolved in control". This statement should inspire people to work on the plantwide control. Then, he adds that this most likely will require "the work of a genius". I am not sure if this addition is entirely correct, and may not have helped to inspire people to work on plantwide control. Nevertheless, it did inspire me and I have worked on the problem since then. After 25 years, I am finally approaching a situation where I have a reasonably clear picture on how to approach the problem.

This chapter summarizes, extends and discusses the plantwide control procedure of Skogestad (2004). An important feature of this procedure is to start with the optimal *economic* operation of the plant, and then attempt to design a control structure that implements optimal operation while also considering the more basic requirements of robustness and stability. The procedure is split into a top-down part, based on plant economics, and a bottom-up part. The bottom-up parts aims to find a simple and robust 'stabilizing' or 'regulatory' control structure, which can be used under most economic conditions. A key step in the procedure is the selection of appropriate controlled variables (CVs) both for the upper supervisory (economic) control layer (CV_1) and the lower regulatory layer (CV_2).

11.2 Control Layers and Timescale Separation

The term 'plantwide control' can be defined as 'control structure design applied to chemical plants'. Here, control structure design is not the tuning and behavior of each control loop, but rather the control philosophy of the overall plant with emphasis on the structural decisions:

- selection of controlled variables (CVs), that is, outputs
- selection of manipulated variables (MVs), that is, inputs
- selection of (extra) measurements
- selection of control configuration (structure of overall controller that interconnects the controlled, manipulated and measured variables)
- selection of controller type (proportional-integral-derivative or PID, decoupler, model predictive control or MPC, optimal control (LQG), ratio, etc.).

Control structure design (i.e., plantwide control) thus involves all the decisions necessary to make a block diagram (used by control engineers) or process and instrumentation diagram (used by process engineers) for the entire plant. This includes the structural decisions, but it does not involve the actual design of each individual controller block.

In a mathematical sense, the plantwide control problem is a formidable and almost hopeless combinatorial problem involving a large number of discrete decision variables; this is probably why progress in the area has been relatively slow. In addition, the problem has been poorly defined in terms of its objective. Usually in control, the objective is that the CV (output) should remain close to its setpoint. However, what should we control? What are the CVs? The answer lies in considering the overall plant objective, which is to minimize cost (= maximize profit) while satisfying operational constraints imposed by the equipment, market demands, product quality, safety, environment and so on. We will return to this.

The overall mathematical problem is, in principle, not so difficult to formulate. With today's computing power, it may even be solvable for some restrictive cases. It would involve obtaining a detailed dynamic and steady-state model of the complete plant, defining all the operational constraints, defining all available measurements and manipulations, defining all expected disturbances, defining expected, allowed or desirable ranges for all variables and then designing a non-linear controller that keeps all the controlled variables close to the setpoints or constraints while using the possible remaining degrees of freedom (DOFs) to minimize the cost. This would involve a single centralized controller, which at each time-step collects all the information and computes the optimal changes in the MVs.

Although such a single centralized solution is foreseeable on some very simple processes, it seems safe to assume that it will never be applied to any normal-sized chemical plant. There are many reasons for this; one is that in most cases acceptable control can be achieved with simple structures where each controller block only involves a few variables. Such simple control systems can be designed and tuned with much less effort, especially when it comes to the modeling and tuning effort. After all, most real plants operate well with simple control structures. A related example is control of biological systems. These are extremely complex and there is little computing power available (the brain has many good features but extensive computations is not one of its strong features) to do the optimal centralized control task. We therefore have to rely on very simple, but still effective, control strategies.

So how are real systems controlled in practice? The main simplification is to decompose the overall control problem into many simple control problems. This decomposition involves two main principles:

1. Decentralized (local) control. This 'horizontal decomposition' of the control layer is mainly based on spatial separation, for example, by using local control of individual process units.
2. Hierarchical control. This 'vertical decomposition' is mainly based on timescale separation. In a typical process we have the following layers (see Figure 11.1):
   - scheduling (weeks)
   - site-wide optimization (days)
   - local optimization (hours)
   - supervisory (predictive, advanced) control (minutes)
   - regulatory control (seconds).

We generally have more multivariable coordination as we move upwards in the hierarchy. Such a hierarchical (cascade) decomposition with layers operating on different timescales is used in the control of all real (complex) systems including biological systems and airplanes, so the issues raised in this section are of general interest and not limited to process control.

The upper three layers in Figure 11.1 deal explicitly with economic optimization, and are not considered in this chapter. We are concerned with the two lower control layers, the supervisory ('advanced') and regulatory control layers, where the main objective is to track the setpoints given by the layer above. A very important structural decision, probably more important than the controller design itself, is then the choice of the CVs, for which the setpoints are given. In the upper supervisory control layer, we want to select CVs that are favorable from an economic point of view (CV_1).

Typically, PID controllers are used in the lower regulatory control layer where 'stabilization' of the plant is the main issue. In the upper supervisory ('advanced') control layer, traditionally single-loop PID control has been used, complemented by 'advanced' elements such as static decouplers, feedforward elements, selectors, split-range controller and various logic elements. Over the last 25 years, MPC has gradually taken over as a unifying tool to replace PID control and most of these elements. In the above (local) optimization layer, decisions are usually made manually although real-time optimization (RTO) is used for a few applications (especially in the refining and petrochemical industries).

No matter what procedure we choose to use, the following decisions must be made when designing a plantwide control strategy:

- **Decision 1.** Select 'economic' (primary) controlled variables (CV_1) for the supervisory control layer. Their setpoints CV_{1s} link the optimization layer with the control layers.
- **Decision 2.** Select 'stabilizing' (secondary) controlled variables (CV_2) for the regulatory control layer. Their setpoints CV_{2s} link the two control layers.
- **Decision 3**. Locate the throughput manipulator (TPM).
- **Decision 4.** Select pairings for the stabilizing layer, that is, pair inputs (valves) and controlled variables (CV_2). Here, 'valves' refers to the original manipulated variables.

Decisions 1 and 2 are illustrated in Figure 11.2 where the matrices $CV_1 = H\,y$ and $CV_2 = H_2\,y$ and H and H_2 represent a selection, or in some cases a combination, of the available measurements y. This chapter deals with continuous operation of chemical processes, although many of the arguments also hold for batch processes.

11.3 Plantwide Control Procedure

Over the years, going back to the early work of Buckley (1964) from the DuPont company, several approaches have been proposed for dealing with plantwide control issues. Nevertheless, taking into account the practical importance of the problem, the literature is relatively scarce. Larsson and Skogestad (2000) provide a good review and divide the available approaches into two main approaches.

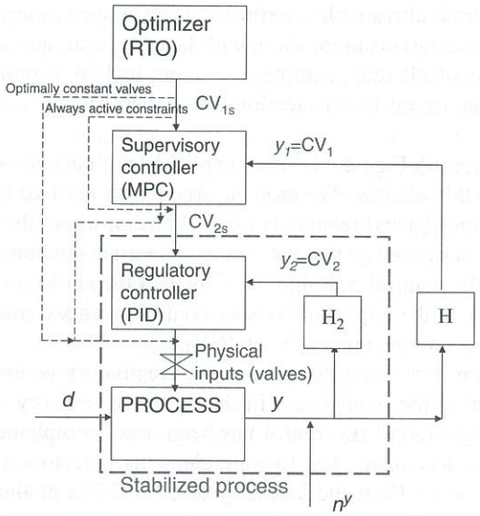

Figure 11.2 *Block diagram of control hierarchy illustrating the selection of controlled variables (H and H_2) for optimal operation ($CV_1 = Hy$) and stabilization ($CV_2 = H_2 y$), respectively.*

First, there is the process-oriented (engineering or simulation-based) approaches of Buckley (1964), Shinskey (1984), Douglas (1988), Downs (1992), Luyben *et al.* (1997, 1998) and Konda *et al.* (2005). One problem here is the lack of a really systematic procedure, and another is little consideration of economics. Second, there is the optimization or mathematically oriented (academic) approaches of Narraway and Perkins (1993), Hansen *et al.* (1998), Kookos and Perkins (2002), Chen and McAvoy (2003) and Engell (2007). The problem here is that the resulting optimization problems are intractable for a plantwide application. A hybrid between the two approaches is therefore more promising (Zheng *et al.*, 1999; Larsson and Skogestad, 2000; Vasbinder and Hoo, 2003; Skogestad, 2004; Ward *et al.*, 2006).

The stepwise plantwide control procedure of Luyben *et al.* (1997, 1998) has been applied in a number of simulation studies. In this chapter, we mainly discuss the seven-step plantwide control procedure of Skogestad (Larsson and Skogestad, 2000; Skogestad 2004). It was inspired by Luyben's procedure, but it has clearly been divided into a top-down part (mainly concerned with steady-state economics) and a bottom-up part (mainly concerned with stabilization and pairing of loops).

Skogestad's procedure consists of the following steps:

I. Top-down part (focus on steady-state optimal operation)
   Step S1. Define operational objectives (economic cost function J and constraints)
   Step S2. Identify steady-state degrees of freedom u and determine the optimal steady-state operation conditions, including active constraints
   Step S3. Identify candidate measurements y and select primary controlled variables $CV_1 = Hy$ (Decision 1)
   Step S4. Select the location of TPM (Decision 3)

II. Bottom-up part (focus on the control layer structure).
   Step S5. Select the structure of regulatory (stabilizing) control layer
   - Select 'stabilizing' controlled variables $CV_2 = H_2 y$ (Decision 2)
   - Select inputs (valves) and 'pairings' for controlling CV_2 (Decision 4)

   Step S6. Select the structure of supervisory control layer
   Step S7. Select structure of (or assess need for) optimization layer (RTO).

The top-down part (steps S1–S4) is mainly concerned with economics, and steady-state considerations are often sufficient. Dynamic considerations are more important for steps S5–S7, although steady-state considerations are also important here. This means that it is important in plantwide control to involve engineers with a good steady-state understanding of the plant. A detailed analysis in steps S2 and S3 requires that we have a steady-state model available and that optimization is performed for the given plant design ('rating mode') for various disturbances.

11.4 Degrees of Freedom for Operation

The issue of degrees of freedom for operation is often confusing and not as simple as would be expected. First, note that we are talking about operation and so the equipment is assumed to be fixed. Second, note that the degrees of freedom (in control often referred to as MVs or inputs u) change depending on where we are in the control hierarchy. This is illustrated in Figures 11.1 and 11.2, where we see that the degrees of freedom in the optimization and supervisory control layers are not the physical degrees of freedom (valves), but rather setpoints for the controlled variables in the layer below. We have:

- DOFs for optimization (steady-state DOFs, u), $MV_{opt} = CV_{1s}$
- DOFs for supervisory control, MV_1=CV_{2s} + unused valves
- (Physical) DOFs for stabilizing control, MV_2 = valves (dynamic process inputs).

In process control applications, the economics of the plant are primarily determined by the (pseudo) steady-state behavior (Morari *et al.*, 1980), and so the steady-state DOFs are usually the same as the economic DOFs.

We refer to the physical DOFs (dynamic process inputs) as 'valves', because this is usually what they are in process control. The stabilizing control system may not use all of the valves. Some valves may not be used dynamically because they are optimally constant, for example, a bypass valve may always be closed (and should then be included in the set CV_1). Other valves may not be needed for stabilizing control, and these 'unused' valves can be used by the supervisory control system to improve dynamic control performance (see Figure 11.2).

11.5 Steady-state DOFs

Identifying the physical DOFs (valves) is relatively straightforward, but it is more difficult to identify the steady-state DOFs (referred to as u in the following discussion). Actually,

it is the *number* of economic (steady-state) degrees of freedom n_u and not the variables themselves which is most important to establish, because it gives the number of controlled variables (CV_1) needed to select in step S3. Of course, as we start to formulate the model and optimize the process, the number of DOFs n_u will be determined; it is however very useful to have an independent and simpler method for finding n_u. First, it is useful for checking. Second, there are cases where we do not have a good mathematical model or do not want to spend the time on optimizing the process. Two approaches are described next: valve counting and potential degrees of freedom.

11.5.1 Valve Counting

One approach is to first identify all the physical (*dynamic*) degrees of freedom (valves). However, because the economics usually depend mainly on the steady state, we should not include variables that have no or a negligible effect on the economics (steady state) such as inputs with only a dynamic effect. We then have:

$$\#\text{ steady-state degrees of freedom } n_u = \#\text{ valves} - \#\text{ variables with no steady-state effect}$$

For example, even though a heat exchanger may have three valves (one valve on the cooling water and bypass valves on both the hot and cold sides), it usually has only one degree of freedom at steady-state (i.e., the amount of heat transferred); two of these three valves only have a dynamic effect from a control point of view. In addition, we need to exclude valves that are used to control variables with no steady-state effect (usually, liquid levels). For example, liquid levels in the reboiler and condenser in a distillation column need to be controlled but their actual values have no steady-state effect. On the other hand, the liquid levels in reactors do have a steady-state effect unless the reactions are in equilibrium.

Example 11.1. DOFs by valve-counting. A simple distillation column has 6 dynamic degrees of freedom (valves) as seen in Figure 11.3: feed F, bottom product B, distillate product D, cooling, reflux L and heat input (which indirectly determines the boil-up V). However, two degrees of freedom (e.g., B and D) must be used to control the condenser and reboiler levels (M_1 and M_2) which have no steady-state effect. This leaves 4 degrees of freedom at steady-state. For the common case with a given feed flow and given column pressure, only 2 steady-state degrees of freedom remain. In step S3 we need to identify controlled variables (CV_1) associated with these, for example, top and bottom composition (x_D and x_B).

11.5.2 Potential Steady-state DOFs

An alternative simple method is to find the potential number of degrees of freedom from the flowsheet as given in Table 11.1. The reason for the word 'potential' is that there may not always be a valve to actually adjust the degree of freedom (e.g., a process-process heat exchanger has 1 potential DOF, but if there is no bypass then it cannot be used in practices).

Example 11.2. Potential DOFs from flowsheet. Consider again the distillation column in Figure 11.3. According to Table 11.1, the column shell itself has 0 steady-state DOFs,

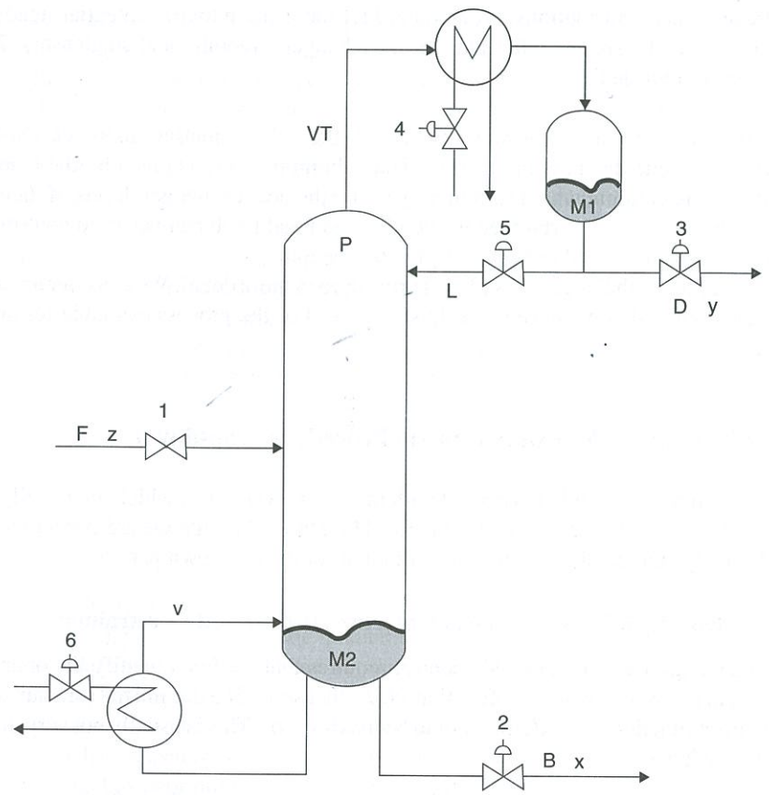

Figure 11.3 *The six dynamic degrees of freedom (valves) for a typical distillation column.*

Table 11.1 *Potential number of steady-state degrees of freedom (DOFs) for some units. Data from Skogestad (2004) and Araujo* et al. *(2007a, 2007b); for extension to closed cycles see Jensen and Skogestad (2009).*

- each external feed stream: 1 (feed rate)
- splitter: n–1 (split fractions) where n is the number of exit streams
- mixer: 0
- compressor, turbine, pump: 1 (work/speed)
- adiabatic flash tank: 0[a]
- liquid phase reactor: 1 (holdup reactant)
- gas phase reactor: 0[a]
- heat exchanger: 1 (bypass or flow)
- column (e.g., distillation) excluding heat exchangers: 0[a] + no. of sidestreams
- pressure[a] : add 1 DOF at each extra place pressure is set (using an *extra* valve, compressor or pump), for example, in adiabatic flash tank, gas phase reactor or column

[a]Pressure is normally assumed to be given by the surrounding process and is not a DOF.

but there are other contributions. From Table 11.1 there are in total 4 potential steady-state DOFs: feed rate (1); spitter (reflux) (1); heat exchangers (reboiler and condenser) (2); and distillation column shell (0).

This is the same number as was found with the valve counting approach. Note that the 4 DOFs include the column pressure. The column pressure is given by the amount of vapor inside the column, which is indirectly set by the heat exchanger duties. If there were non-condensable gases (inerts) then we would need to add a 'bleed valve' to control insert build-up, which would add one more degree of freedom.

We now go through Skogestad's (2004) procedure in more detail. We consider an existing plant, and assume that we have a steady-state model of the process available for analysis purposes.

11.6 Skogestad's Plantwide Control Procedure: Top-down

The top-down part is mainly concerned with the plant economics, which are usually determined primarily by the steady-state behavior. Therefore, although we are concerned about control, steady-state models are usually sufficient for the top-down part.

11.6.1 Step S1: Define Operational Objectives (Cost J and Constraints)

A systematic approach to plantwide control requires that we first quantify the operational objectives in terms of a scalar cost function J (\$/s) that should be minimized (or equivalently, a scalar profit function $P = -J$, that should be maximized). This is usually not very difficult; typically, we have

$$J = \text{cost feed} + \text{cost utilities (energy)} - \text{value products (\$/s)}$$

Note that fixed costs and capital costs are not included, because they are not affected by plant operation on the timescale we consider (around 1 hour). The goal of operation (and of control) is to minimize the cost J, subject to satisfying the operational constraints ($g \leq 0$) including safety and environmental constraints. Typical operational constraints are minimum and maximum values on flows, pressures, temperatures and compositions. For example, all flows, pressures and compositions must be non-negative.

11.6.2 Step S2: Determine the Steady-state Optimal Operation

What is the optimal way of operating the process? We should answer this question before designing the control system. For example, we may find that a valve (for example a bypass) should always be closed (active constraint). This valve should then not be used for (stabilizing) control unless we are willing to accept the loss implied by 'backing off' from the optimal operating conditions. To determine the steady-state optimal operation, we first need to obtain a steady-state model. We then need to identify degrees of freedom and expected disturbances, and perform optimization for the expected disturbances as follows.

(a) **Identify steady-state degrees of freedom u.** To optimize the process, we first need to identify the steady-state degrees of freedom u as has already been discussed. Actually, it is the number of us which is important, because it does not really matter which variables we include in u as long as they make up an independent set.

(b) **Identify important disturbances d and their expected range.** The most important disturbances are usually related to the feed rate (throughput) and feed composition, and external variables such as temperature and pressure of the surrounding process We should also include as disturbances possible changes in specifications and active constraints (such as purity specifications or capacity constraints) and changes in parameters (such as equilibrium constants, rate constants and efficiencies). Finally, we need to include the expected changes in prices of products, feeds and energy as 'disturbances'.

(c) **Optimize the operation for the expected disturbances.** Here we specify the disturbances d and vary the degrees of freedom optimally $u_{\text{opt}}(d)$ in order to minimize the cost J, while satisfying the constraints. The main objective is to find the *constraints regions* (sets of active constraints) and the optimal nominal setpoints in each region.

Mathematically, the steady-state optimization problem can be formulated as

$$\begin{aligned} &\min_u \ J(u, x, d) \\ &\text{subject to} \quad \text{model equations:} \quad f(u, x, d) = 0 \\ &\qquad\qquad\ \ \text{operational constraints:} \quad g(u, x, d) \leq 0 \end{aligned}$$

where u are the steady-state degrees of freedom, d are the disturbances, x are internal variables (states), $f = 0$ represent the mathematical model equations and possible equality constraints (like a given feed flow) and $g \leq 0$ represents the operational constraints (such as a maximum or non-negative flow or a product composition constraint). The process model $f = 0$ is often represented indirectly in terms of a commercial software package (process simulator) such as Aspen or Hysis/Unisim. This usually results in a large non-linear equation set which often has poor numerical properties for optimization.

A major objective of the optimization is to find the expected regions of active constraints. Together with obtaining the model, the optimization step S2 is often the most time-consuming step in the entire plantwide control procedure. In many cases, the model may not be available or we do not have time to perform the optimization. In such cases, a good engineer can often perform a simplified version of steps S1–S3 by using process insight to identify the expected active constraints and possible 'self-optimizing' controlled variables (CV_1) for the remaining unconstrained degrees of freedom.

An important point is that we cannot expect to find a single control structure that is optimal because the set of active constraints will change depending on disturbances and economic conditions (prices). However, the control system should be prepared for the future by performing offline analysis and optimization to identify expected regions of active constraints. The optimally active constraints will vary depending on disturbances (feed rate and composition, outdoor temperature, product specifications) and market conditions (prices). Note that there are generally two main modes of operation, depending on market conditions (Rijnsdorp, 1991).

11.6.2.1 Mode I: Given Throughput (Buyer's Market)

This is usually the nominal mode for which the control system is originally set up to handle. Usually, it corresponds to a 'maximize efficiency' situation where there is some tradeoff

between utility (energy) consumption and recovery of valuable product, corresponding to an unconstrained optimum.

11.6.2.2 Mode II: Maximum Throughput (Seller's Market)

When the product prices are sufficiently high compared to the prices on raw materials (feeds) and utilities (energy), it is optimal to increase the throughput as much as possible, although the efficiency will usually drop. However, as the feed rate is increased, constraints in various units are usually encountered until we eventually reach the bottleneck where a further increase is infeasible.

11.6.3 Step S3: Select Economic (Primary) Controlled Variables, CV_1 (Decision 1)

Here, we are concerned with implementing the optimal operation points found in Step S2 in a robust and simple manner. To make use of all the economic degrees of freedom (inputs u), we need to identify as many economic controlled variables (CV_1) as we have inputs (u). In short, the issue is: what should we control?

For economic optimal operation, the rules for selecting primary controlled variables (CV_1) are

1. **CV_1-rule 1**: control active constraints
2. **CV_1-rule 2**: for the remaining unconstrained degrees of freedom: Control self-optimizing variables.

We need to find one CV_1 for each steady-state degree of freedom u. Here, self-optimizing variables are defined as variables for which close-to-optimal operation with constant setpoints can be achieved, even when there are disturbances (Skogestad, 2000). Active constraints may be viewed as self-optimizing variables because operation is optimized by keeping their values constant; normally however we refer to the unconstrained self-optimizing variables because they are much harder to find. The two rules are discussed in detail in the following sections.

11.6.3.1 CV1-rule 1

In general, the active constraints should be selected as controlled variables, which may be viewed as the obvious self-optimizing variables to be kept constant. The active constraints come out of the analysis in step S2 or may, in some cases, be identified based on physical insight. The active constraints could be input constraints (in the set u) or output constraints.

Input constraints are usually trivial to implement; we just set the input at its optimal minimum or maximum so no control is really needed. For example, if we are operating a very old car then optimal operation (defined as minimum driving time, $J = T$) may be achieved with the gas pedal at its maximum position.

For output constraints we need a controller; a simple single-loop feedback controller is often sufficient. For example, if we have a better car then the maximum speed limit (say 80 km/h) is likely an active constraint and should be selected as the controlled variable (CV_1). To control this, we may use a cruise controller (automatic control) which adjusts the engine power to keep the car close a given setpoint. In this case, we need to *back off* from the speed limit (say to a setpoint of 72 km/h) to avoid exceeding the speed limit if there

is a steady-state measurement error ($n^y = 5$ km/h) or a dynamic control error (3 km/h). In general, we want to minimize the back-off because any back-off results in a loss (i.e., a larger $J = T$) which can never be recovered.

The back-off is the safety margin from the active constraint and is defined as the difference between the constraint value and the chosen setpoint, that is,

$$\text{Back-off} = |\,\text{Constraint} - \text{Setpoint}\,|$$

In the car driving example, back-off is 8 km/h if the speed limit can never be exceeded (hard constraint), but it can be reduced to 5 km/h if the limit is on average speed (soft constraint). This is because the dynamic control error will average out if the controller has integral action.

The active constraints should always be selected as CVs because the optimum is not 'flat' with respect to these variables. There is therefore often a significant economic penalty if we back off from an active constraint, and so tight control of the active constraints is usually desired. If a constrained optimization method is used for the optimization, we can quantify the loss by using the Lagrange multiplier λ associated with the constraint, that is:

$$\text{Loss} = \lambda \times \text{back-off}$$

We see that variables with a large Lagrange multiplier should have a small back-off to reduce the economic loss.

How large a back-off is required? For input (valve) constraints, we usually need no back-off unless we choose to use the input for stabilization in the lower regulatory (stabilizing) layer (in that case, we need some range to use it for control). For output constraints, we have two cases:

- soft output constraints (only average value matters): back-off = measurement error (bias n^y)
- hard output constraints (must be satisfied at all times): back-off = measurement error (bias n^y) + control error (dynamic).

To reduce the back-off, we need accurate measurements of the output constraints. For hard output constraints, we also need tight control with a small dynamic control error. For hard output constraints, we have the 'squeeze and shift' rule: by squeezing the output variation with improved control, we can shift the setpoint closer to its limit (i.e., reduce the back-off for the car example by using a better cruise controller).

11.6.3.2 CV_1-rule 2

The main steps involved in selecting the unconstrained self-optimizing variables are as follows.

(a) **Identify candidate measurements y.** We must first identify all the candidate measurements y together with their expected static measurement error n^y . In general, we should include the inputs (e.g., flow rates), including those used to control active constraints, in the set y.

(b) **Select primary (economic) controlled variables $\mathbf{CV_1 = H}\,y$ (Decision 1)** (see Figure 11.2). If possible, we want to control single measurements for simplicity and, in this case, H is a selection matrix. More generally, we may control measurement

combinations and, in this case, H is the 'full' matrix. In general, this step must be repeated for each constraint region. To reduce the need to switch between regions, consider using the same CV_1s in several regions; this is non-optimal however and may even lead to unfeasibility.

It is usually simple to identify and control the active constraints. The more difficult question is: what should we use the remaining unconstrained degrees of freedom for? Does it even make a difference what we control? The answer to this is: 'yes'.

As an example, consider optimal operation of a marathon runner where the objective is to adjust the power u to minimize the time $J = T$. This is an unconstrained problem; we cannot simply run at maximum ($u = u_{max}$) as for a sprinter. A simple policy is constant speed ($CV_1 = c_1$ = speed), but it is not optimal if there are disturbances d caused by wind or hilly terrain. A better choice is to run with constant heart rate ($CV_1 = c_2$ = pulse), which is easy to measure with a pulse clock. With a constant heart rate (c_2 = pulse = constant), the speed c_1 will automatically increase when we run downhill, as would be expected for optimal operation. Heart rate c_2 is clearly a better self-optimizing variable than speed c_1. One problem with the feedback is that it also introduces a measurement error (noise) n^y which may also contribute to the loss (see Figure 11.2).

In summary, the problem at hand is to choose the matrix H such that keeping the controlled variables $c = \mathrm{H}\,y$ constant (at a given setpoint c_s) gives close-to-optimal operation despite the presence of disturbances d (which shift the optimum) and measurement errors n^y (which give an offset from the optimum).

11.6.3.3 Qualitative Approaches

The following four requirements (Skogestad, 2000) are useful for identifying a 'good' $c = \mathrm{H}\,y$, mainly for the case where we select to control single measurements (c = selected y).

1. The *optimal* value of c is insensitive to disturbances. This means that $\mathrm{d}c_{opt}/\mathrm{d}d$ should be small. Note that this is not saying that the sensitivity of c to disturbances ($\mathrm{d}c/\mathrm{d}d$) should be small. Of course, $\mathrm{d}c/\mathrm{d}d$ should not be too large, as this would make control of c difficult. On the other hand, $\mathrm{d}c/\mathrm{d}d$ should not be too small (e.g., $\mathrm{d}c/\mathrm{d}d = 0$) because then the disturbance cannot be detected and corrected by controlling the variables c. In summary, we are not really concerned about $\mathrm{d}c/\mathrm{d}d$, but we do want $\mathrm{d}c_{opt}/\mathrm{d}d$ to be small.
2. The variable c is easy to measure and control accurately
3. The value of c is sensitive to changes in the manipulated variable u; that is, the gain, $G = \mathrm{d}c/\mathrm{d}u$ from u to c is large (so that even a large error in controlled variable c results in only a small variation in u). Equivalently, the optimum should be 'flat' with respect to the variable c.
4. For cases with two or more controlled variables c, the selected variables should not be closely correlated.

All four requirements should be satisfied. For example, for the operation of 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).

11.6.3.4 Quantitative Approaches

There are two main approaches for finding the matrix H, that is, to identify self-optimizing CV_1s associated with the unconstrained DOFs, described in the following.

1. **'Brute force' approach.** Given a set of controlled variables $CV_1 = c = \mathrm{H}\,y$, we computes the cost $J(c, d)$ when we keep c constant ($c = c_s + \mathrm{H}\,n^y$) for various disturbances d and measurement errors n^y. In practice, this is done by running a large number of steady-state simulations to try to cover the expected future operation. Typically, expected extreme values in the parameter space (for d and n^y) are used to compute the cost for alternative choice for the controlled variables (matrix H). The advantage is that this method is simple to understand and apply, and it also works for non-linear plants and even for changes in active constraint. Only one nominal optimization is required to find the setpoints. The main disadvantage with this method is that the analysis for each H is generally time-consuming and it cannot be guaranteed that all important cases are covered. In addition, there exist an infinite number of choices for H, and it can never be guaranteed that the best cs are found.
2. **'Local' approaches** are based on a quadratic approximation of the cost. This is discussed in more detail in Alstad *et al.* (2009). The main local approaches are as follows.
   (a) **Maximum gain rule**: A quantitative version of Requirements 1 and 3 given above is that 'sensitive' variables should be controlled with a large-scale gain $|\mathrm{G}|$ from the inputs u to $c = \mathrm{H}\,y$. This rule is good for pre-screening and also yields good insight.
   (b) **Nullspace method**: The optimal measurement sensitivity $\mathrm{F} = \mathrm{d}y^{opt}/\mathrm{d}d$ is first obtained and we select H such that $\mathrm{HF} = 0$. This method yields optimal measurement combinations for the case with no noise, $n^y = 0$.

   Each column in F expresses the optimal change in y when the unconstrained independent variable u is adjusted so that the system remains optimal with respect to the disturbance d. Any active constraints are assumed held constant. If we have a model of the process, then it is in principle straightforward to obtain F numerically. Then, assuming that we have at least as many (independent) measurements y as the sum of the number of (independent) inputs u and disturbances d, the optimal is to select $c = \mathrm{H}\,y$ such that $\mathrm{HF} = 0$.

   Note that H is a non-square matrix, and so $\mathrm{HF} = 0$ does not require that $\mathrm{H} = 0$ (which is a trivial uninteresting solution), but rather that H is in the nullspace of F^{T}.
   (c) **Exact local** method (loss method): extends the nullspace method to the case with noise and to any number of measurements; for details see Alstad *et al.* (2009).

The use of these methods is discussed in many papers by Skogestad and co-workers; for example, see Downs and Skogestad (2011) for some practical applications of the nullspace method.

11.6.3.5 Regions and Switching

Note that new self-optimizing variables must be identified (offline) for each region, and that switching of controlled variables is required as a new region (online) is encountered. In practice, it is easy to identify when to switch when a constraint is encountered. It may seem

less obvious when to switch out of a constraint, but it is actually straightforward; simply monitor the value of the unconstrained CVs from the neighbouring regions and switch out of the constraint region when the unconstrained CV reaches its setpoint.

Example 11.3. Consider a recycle process where it is optimal to keep the inert fraction in the purge at 5% using the purge flow as a degree of freedom (unconstrained optimum). However, during operation, there may be a disturbance (e.g., increase in feed rate) so that the recycle compressor reaches its maximum capacity. The recycle compressor was used to control pressure and, since it is still required to control pressure, the purge flow has to take over this task. This means that one has to give up controlling the inert fraction, which will drop below 5%. In summary, we have gone from an unconstrained operating region (I) where we control the inert fraction to a constrained region (II) where the compressor is at maximum load. In region II, we keep the recycle flow at its maximum. How do we know when to switch back from region II to region I? It is simple: we monitor the inert fraction, and when it reaches 5% we switch back to controlling it (region I).

In general, we would like to simplify the control structure and reduce the need for switching. This may require using a suboptimal CV_1 in some regions of active constraints. In this case, the setpoint for CV_1 may not be its nominally optimal value (which is the normal choice), but rather a 'robust setpoint' which reduces the loss when we are outside the nominal constraint region.

11.6.4 Step S4: Select the Location of TPM (Decision 3)

The main purpose of a process plant is to transform feedstocks into more valuable products, and this involves moving mass through the plant. The amount of mass moved through the plant, as expressed by the feed rate or product rate, is determined by specifying one degree of freedom which we refer to as the throughput manipulator (TPM). The location of the TPM is an important decision that links the top-down and bottom-up parts of the procedure. The TPM or process 'gas pedal' is usually a flow but not always, and it is usually set by the operator (manual control).

Definition 11.1. A TPM is a degree of freedom that affects the network flow and which is not directly or indirectly determined by the control of the individual units, including their inventory control (Aske and Skogestad, 2009).

Some plants, for example, with parallel units, may have more than one TPM. The TPM has traditionally been placed at the feed to the plant. One important reason is that most of the control structure decisions are made at the design stage (before the plant is build) where the feed rate is considered fixed; there is usually little thought given to the future operation of the plant where it is likely that we will want to maximize feed (throughput) and moving the TPM may be better.

11.6.4.1 Where should the TPM be Located?

In principle, the TPM may be located anywhere in the plant, although the operators often prefer to have it at the feed and so this will be the default choice. Note that from a purely

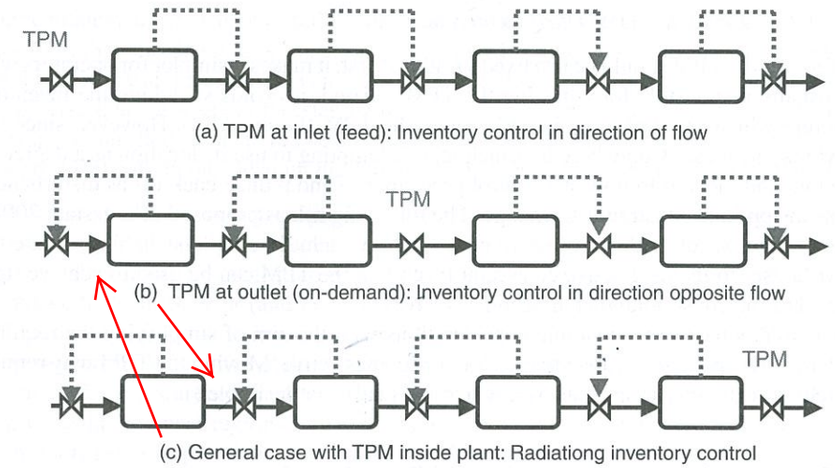

Figure 11.4 *Radiation rule: local-consistency requires a radiating inventory control around a fixed flow (TPM) (Price and Georgakis, 1993; Aske and Skogestad, 2009).*

steady-state point of view the location of the TPM does not matter, but it is important dynamically. There are two main concerns when locating the TPM, as follows.

1. **Economics.** The location has an important effect on economics because of the possible back-off if active constraints are not tightly controlled. In particular, the economic loss may be large for the maximum throughput case if the bottleneck unit is not tightly controlled. In this case, the TPM should be located close to the bottleneck to reduce the back-off from the active constraint that has the largest effect on the production rate.
2. **Structure of regulatory control system**. Because of the radiation rule (Price and Georgakis, 1993), the location of the throughput manipulator has a profound influence on the structure of the regulatory control structure of the entire plant (see Figure 11.4). An underlying assumption for the radiation rule is that we want 'local' consistency of the inventory control system (Aske and Skogestad, 2009). In theory, the radiation rule may not be adhered to by allowing for 'long' inventory loops; this is not common for obvious operational reasons, including the risk of emptying or overfilling tanks, startup, tuning and increased complexity.

Most plants have one 'gas pedal' (TPM), but there may be more than one TPM for plants with parallel units, splits and multiple alternative feeds or products. Note that the feeds usually need to be set in a fixed ratio, so adding a feed does not usually give an additional TPM. For example, for the reaction $A + B \rightarrow C$, we need to have the molar ratio F_A/F_B close to 1 to have good operation with a small loss of reactants. There is only one TPM even if there are two feeds, F_A and F_B.

If we only consider a part of a process plant, then this part may have no TPM. There will instead be a given flow, typically a feed or product, that acts as a disturbance which the control system must be set up to handle. This may also be viewed as having the TPM at a fixed location. For example, the product rate may be given for a utility plant and the feed rate may be given for an effluent treatment plant.

11.6.4.2 Moving the TPM During Operation

Preferably, the TPM should be in a fixed location. First, it makes it simpler for operators who are usually responsible for adjusting the TPM; second, it avoids switching the inventory structure which should be 'radiating' around the TPM (Figure 11.4). However, since the TPM may be located anywhere in principle, it is tempting to use its location as a degree of freedom and move it to improve control performance and reduce back-off as disturbances cause the optimal constraints to change. The following rule is proposed (Skogestad, 2004): to get tight control of the new active constraint and achieve simple switching, locate the TPM 'close' to the next active constraint (such that the TPM can be used to achieve tight control of the constraint when it becomes active).

The rule is based on economic considerations with the aim of simplifying the required switching when the next capacity constraint becomes active. Moving the TPM may require switching regulatory loops however, which is usually not desirable.

11.7 Skogestad's Plantwide Control Procedure: Bottom-up

11.7.1 Step S5: Select the Structure of Regulatory (Stabilizing) Control Layer

The main purpose of the regulatory layer is to 'stabilize' the plant, preferably using a simple control structure with single-loop PID controllers. 'Stabilize' here means that the process does not drift too far away from acceptable operation when there are disturbances. Think about learning to ride a bicycle: before attempting to do more high-level tasks such as following the road from A to B, the bicycle needs to be stabilized.

The regulatory layer is the fastest control layer, and is therefore also used to control variables that require fast and tight control such as economically important active constraints. In addition, the regulatory layer should follow the setpoints given by the supervisory layer.

As discussed in more detail below, the main decisions in Step S5 are to (a) select controlled variables (CV_2) (Decision 2) and (b) to select inputs (valves) and 'pairings' for controlling CV_2 (Decision 4). Interestingly, decision (a) on selecting CV_2 can often be based on steady-state arguments, whereas dynamic issues are the primary concern when selecting inputs (valves) and pairings in decision (b).

Note that we do not use up any degrees of freedom in the regulatory control layer because the setpoints CV_{2s} are left as manipulated variables (MVs) for the supervisory layer (see Figure 11.2). By allowing for cascade loops, the stabilization layer may in theory be designed independent of the supervisory (economic) control layer. However, when closing a stabilizing loop, we do use up some of the time window as given by the closed-loop response time (bandwidth) of the stabilizing loop. In addition, cascade loops add complexity. We would therefore like to simplify and reduce the need for cascade loops.

11.7.1.1 Step S5(a): Select Stabilizing CV_2 (Decision 2).

First of all, we should stabilize the process by controlling drifting variables such as inventories (level and pressure), reactor temperature and temperature profile in distillation column. This simplifies the supervisory (economic) control tasks, by providing for local/fast disturbance rejection and reducing non-linearity in the model. Secondly we should include active constraints (CV_1) that require tight control (typically, hard output constraints) in CV_2 for the regulatory layer. This will reduce the required back-off (recall the 'squeeze and shift' rule). On the other hand, it is usually not necessary for tight control of unconstrained CV_1 variables because the optimum is usually relatively flat.

To systematically select the stabilizing $CV_2 = H_2 y$, the behavior of the stabilized or partially controlled plant should be considered with the variables CV_2 being controlled (see Figure 11.2), taking into account the following two main objectives of the regulatory layer.

- **Local disturbance rejection** (indirect control of primary variables CV_1): with the variables CV_2 controlled, the effect of the disturbances on the primary variables CV_1 should be small. This is to get fast control of the variables CV_1, which may be important to reduce the control error (and thus the back-off) for some variables such as active output constraints.
- **Stabilization** (minimize state drift): more generally, the objective is to minimize the effect of the disturbances on the internal variables (states) x. One reason is to keep the process in the linear region close to the nominal steady-state and avoid the process drifts into a region of operation where it is difficult to recover. The advantage of considering some measure of all the states x is that the regulatory control system is then not tied to a particular control objective (CV_1) which may change with time, depending on disturbances and prices.

When considering disturbance rejection and stabilization in the regulatory layer, it is the behavior at the closed-loop time constant of the above supervisory layer which is of main interest. Since the supervisor layer is usually relatively slow, it is again (as with the selection of CV_1) often sufficient to consider the steady-state behavior when selecting CV_2. However, dynamics are the key issue when selecting the corresponding valves/pairings in Step 5(b).

11.7.1.2 Step S5(b): Select Inputs (Valves) for Controlling CV_2 (Decision 4)

Next, we need to find the inputs (valves) that should be used to control CV_2. Normally, single-loop (decentralized) controllers are used in the regulatory layer and so the objective is to identify pairings. The main rule is to 'pair close' so that the dynamic controllability is good with a small effective delay, in order that the interactions between the loops are small. In addition, the following requirements should also be considered.

1. Local consistency for the inventory control (Aske and Skogestad, 2009). This implies that the inventory control system is radiating around the given flow (Figure 11.4).
2. Tight control of important active constraints (to avoid back-off).
3. Avoid selecting variables that may optimally saturate (steady state) as MVs in the regulatory layer, because this would require either reassignment of regulatory loop (complication penalty) or back-off for the MV variable (economic penalty).
4. Avoid reassignments (logic) in the regulatory layer.

In practice, in order to make the regulatory selection (step S5(b)) more manageable, it can be divided into Step S5(b1): Structure of inventory control layer (closely related to Step S4), followed by Step S5(b2): Structure of remaining regulatory control system.

The regulatory layer should preferably be independent of the economic control objectives (regions of steady-state active constraints), which may change depending on disturbances, prices and market conditions. In principle, choices for CV_1 (Decision 1) and CV_2 (Decision 2) are therefore independent of each other. In order to simplify and reduce the need for

cascade loops, it is however an advantage if we have $CV_1 = CV_2$ (at least for some variables).

11.7.2 Step 6: Select Structure of Supervisory Control Layer

The supervisory or 'advanced' control layer has three main tasks.

1. Control the primary (economic) controlled variables (CV_1) using the setpoints to the regulatory layer plus any remaining (unused) valves (see Figure 11.2) as MVs.

   Interactions may be quite significant at this timescale, and so multivariable control (typically, MPC) should be considered (see below). The supervisory layer may use additional dynamic DOFs, including level setpoints, to improve the dynamic response. At steady state, these extra variables may be 'reset' to their ideal resting values. The supervisory layer may also make use of measured disturbances (feedforward control).

   If the primary controlled variables (CV_1) are not measured, then 'soft sensors' based on other available measurements may be used for their estimation. The 'soft sensors' are usually static, although dynamic state estimators (Kalman filter, moving horizon estimation) may be used to improve the performance. These are not common in process control however, because the supervisory layer is usually rather slow.
2. Supervise the performance of the regulatory layer. The supervisory layer should take action to avoid saturation of MVs used for regulatory control. If an MV in the regulatory layer saturates, then control of the corresponding CV_2 is lost resulting in a large drift away from the desired operating point.
3. Switch controlled variables and control strategies when disturbance or price changes cause the process to enter a new region of active constraints.

There are two main alternatives in terms of the controller used in the supervisory layer.

1. **Advanced single loop control**, that is, PID control with additional 'fixes' such as feedforward (ratio), decouplers, logic, selectors and split range control. In many cases, some of these tasks are moved down to the regulatory layer. With single-loop control, an important decision is to select pairings. Note that the issue of finding the right pairings is more difficult for the supervisory layer because the interactions are usually much stronger at slower timescales; measures such as the RGA may be helpful.
2. **Multivariable control (usually MPC).** Although switching and logic can be reduced when using MPC, it generally cannot be completely avoided. In general, it may be necessary to change the performance objective of the MPC controllers as we switch regions.

11.7.3 Step 7: Structure of Optimization Layer (RTO) (Related to Decision 1)

The task of the RTO layer is to update the setpoints for CV_1 and to detect changes in the active constraint regions that require switching of controlled variables (CV_1). In most cases with a self-optimizing choice for the primary controlled variables, the benefits of the RTO layer are too low to justify the costs of creating and sustaining the detailed steady-state model which is usually required for RTO. In addition, the numerical issues related to optimization are very difficult; even offline optimization is difficult.

11.8 Discussion

The chapter has outlined the plantwide control procedure of Skogestad (2004). It is reasonably systematic, but still far from the level where it can be automated. In general, the involvement of the engineer is key and iteration in the procedure is required. One reason why a completely automated (algorithmic) procedure is not suggested, and is probably not desirable, is that this would require a much more detailed specification of the problem. First, the objectives of the regulatory layer would need to be clearly defined. In practice, such specifications are usually established in an iterative fashion as needed. In addition, an automated procedure would require a detailed steady-state and dynamic model. Furthermore, since 'simplicity of the control layer' is an objective, it would require a formal definition of what is meant by simplicity. This does not mean that we should not continue working towards a more automated or detailed procedure, but just that we should be aware of its 'costs' in terms of the problem definition.

The plantwide control procedure described in this chapter has been applied to many case studies (mostly simulation studies), but they have not been included in this chapter to save space. Some relevant applications include: distillation (Skogestad, 2000); Tennessee Eastman process (Larsson *et al.*, 2001); recycle process (Larsson *et al.*, 2003; Jagtap *et al.*, 2011); refrigeration process (Jensen and Skogestad, 2007); HDA process (Araujo *et al.*, 2007a, b); ammonia plant (Araujo and Skogestad 2008); Eastman extraction plant with recycle, parallel units (Downs and Skogestad, 2011); and CO_2 capture plant (Panahi and Skogestad, 2011).

A discussion of the plantwide control procedure described in this chapter and its comparison with that of Luyben *et al.* (1997; 1998) is given in Skogestad (2011).

11.9 Conclusions

Control structure design deals with the structural decisions of the control system, including what to control and how to pair the variables to form control loops. Although these are very important issues, these decisions are in most cases made in an ad hoc fashion; that is, they are based on experience and engineering insight without considering the details of each problem. In this chapter, a systematic procedure for control structure design for complete chemical plants (plantwide control) is presented. It starts by carefully defining the operational and economic objectives and the degrees of freedom available to fulfill them. The operation is then optimized for expected future disturbances to identify constraint regions. In each region, we should control the active constraints and identify self-optimizing variables for the remaining unconstrained degrees of freedom. Following the decision on where to locate the throughput manipulator, a bottom-up analysis is performed to determine secondary controlled variables and structure of control system (pairing).

References

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

Araujo, A. and Skogestad, S. (2008) Control structure design for the ammonia synthesis process. *Computers and Chemical Engineering*, **32**(12), 2920–2932.
Araujo, A.C.B., Govatsmark, M. and Skogestad, S. (2007a) Application of plantwide control to the HDA process. I Steady-state and self-optimizing control. *Control Engineering Practice*, **15**, 1222–1237.
Araujo, A.C.B., Hori, E.S. and Skogestad, S. (2007b) Application of plantwide control to the HDA process. II Regulatory control. *Industrial & Engineering Chemistry Research*, **46**(15), 5159–5174.
Aske, E.M.B. and Skogestad, S. (2009) Consistent inventory control. *Industrial & Engineering Chemistry Research*, **48**(44), 10892–10902.
Buckley, P.S. (1964) *Techniques of Process Control*, Wiley, New York, Chapter 13.
Chen, R. and McAvoy, T.J. (2003) Plantwide control system design: Methodology and application to a vinyl acetate process. *Industrial & Engineering Chemistry Research*, **42**(20), 4753–4771.
Douglas, J.M. (1988) *Conceptual Design of Chemical Processes*, McGraw-Hill, New York.
Downs, J.J. (1992) Distillation control in a plantwide control environment, in *Practical Distillation Control* (ed. W.L. Luyben), Van Nostrand Reinhold, New York, 413–439
Downs, J.J. and Skogestad, S. (2011) An industrial and academic perspective on plantwide control. *Annual Reviews in Control*, **17**, 99–110.
Engell, S. (2007). Feedback control for optimal process operation. *Journal of Process Control*, **17**, 203–219.
Finco, M.V., Luyben, W.L. and Polleck, R.E. (1989) Control of distillation columns with low relative volatility. *Industrial & Engineering Chemistry Research*, **28**, 76–83.
Foss, A.S. (1973) Critique of chemical process control theory. *American Institute of Chemical Engineering Journal*, **19**(2), 209–214.
Hansen, J.E., Jorgensen, S.B., Heath, J. and Perkins, J. (1998) Control structure selection for energy integrated distillation column. *Journal of Process Control*, **8**, 185–195.
Jagtap, R., Kaistha, N. and Skogestad, S. (2011) Plantwide control for economic operation of a recycle process with side reaction. *Industrial & Engineering Chemistry Research*, **50**, 8571–8584.
Jensen, J.B. and Skogestad, S. (2007) Optimal operation of simple refrigeration cycles. Part II: Selection of controlled variables. *Computers and Chemical Engineering*, **31**, 1590–1601.
Jensen, J.B. and Skogestad, S. (2009) Steady-state operational degrees of freedom with application to refrigeration cycles. *Industrial & Engineering Chemistry Research*, **48**(14), 6652–6659.
Konda, N.V.S.N.M., Rangaiah, G.P. and Krishnaswamy, P.R. (2005) Plantwide control of industrial processes: An integrated framework of simulation and heuristics. *Industrial & Engineering Chemistry Research*, **44**(22), 8300–8313.
Kookos, I.K. and Perkins, J.D. (2002) An Algorithmic method for the selection of multivariable process control structures. *Journal of Process Control*, **12**, 85–99.
Larsson, T. and Skogestad, S. (2000) Plantwide control: A review and a new design procedure. *Modelling, Identification and Control*, **21**, 209–240.
Larsson, T., Hestetun, K., Hovland, E. and Skogestad, S. (2001) Self-optimizing control of a large-scale plant: the Tennessee Eastman process. *Industrial & Engineering Chemistry Research*, **40**(22), 4889–4901.
Larsson, T., Govatsmark, M.S., Skogestad, S. and Yu, C.C. (2003) Control structure selection for reactor, separator and recycle processes. *Industrial & Engineering Chemistry Research*, **42**(6), 1225–1234.
Luyben, M.L., Tyreus, B.D. and Luyben, W.L. (1997) Plantwide Control Design Procedure. *American Institute of Chemical Engineering Journal*, **43**, 3161–3174.
Luyben, W.L., Tyreus, B.D. and Luyben, M.L. (1998) *Plantwide Process Control*, McGraw-Hill, New York.
Morari, M., Arkun, Y. and Stephanopoulos, G. (1980) Studies in the synthesis of control structures for chemical processes. Part I. *American Institute of Chemical Engineering Journal*, **26**, 209–214.
Narraway, L.T. and Perkins, J.D. (1993) Selection of control structure based on economics. *Computers and Chemical Engineering*, **18**, S511–S515.
Panahi, M. and Skogestad, S. (2011) Economically efficient operation of CO2 capturing process, part I: Self-optimizing procedure for selecting the best controlled variables. *Chemical Engineering and Processing*, **50**, 247–253.

Price, R.M. and Georgakis, C. (1993) Plantwide regulatory control design procedure using a tiered framework. *Industrial and Engineering Chemistry Research*, **32**, 2693–2705.
Rijnsdorp, J.E. (1991) *Integrated Process Control and Automation*. Elsevier.
Shinskey, F.G. (1984) *Distillation Control: For Productivity and Eenergy Conservation*, 2nd edition. McGraw-Hill, New York.
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.
Skogestad, S. (2007) The dos and don'ts of distillation columns control. *Chemical Engineering Research and Design (Transactions of Institute of Chemical Engineering, Part A)*, **85**(A1), 13–23.
Skogestad, S. (2011) Plantwide control, Chapter in *Ullmann's Encypledia of Industrial Chemistry*, Volume on Process Systems Engineering. Wiley, New York.
Skogestad, S. and Postlethwaite, I. (2005) *Multivariable Feedback Control*, 2nd edn, Wiley, Chichester, UK.
Vasbinder, E.M. and Hoo, K.A. (2003) Decision-based approach to plantwide control structure synthesis. *Industrial & Engineering Chemistry Research*, **42**, 4586–4598.
Ward, J.D., Mellichamp, D.A. and Doherty, M.F. (2006) Insight from economically optimal steady-state operating policies for dynamic plantwide control. *Industrial & Engineering Chemistry Research*, **45**, 1343.
Zheng, A., Mahajanam, R.V. and Douglas, J.M. (1999) Hierarchical procedure for plantwide control system synthesis. *American Institute of Chemical Engineering Journal*, **45**(6) 1255–1265