Self-optimizing control: From key performance indicators to control of biological systems [1]

Sigurd SKOGESTAD

Dept. of Chemical Engineering, Norwegian University of Science and Technology, N–7491 Trondheim, Norway

Abstract. The topic of this paper is how to implement optimal decisions in an uncertain world. A study of how this is done in practical systems - from the nationwide optimization of the economy by the Central Bank to the optimal use of reseources in a single cell - shows that a common approach is to use feedback strategies where selected controlled variables are kept at constant values. For example, the Central Bank may adjust the interest rate (independent input variable) in order to keep the inflation constant (selected controlled variable). The goal of this paper is to present a unified framework for selecting controlled variables based on the idea of self-optimizing control, and to provide a number of examples.

Keywords. Optimal operation, active constraint, controlled variable, constrol structure design

1 Introduction

The national economy, the goverment, companies and businesses, consumers, chemical process plants, biological systems, and so on, are all decision makers that make up a complex hiearchical decision system(Findeisen *et al.* 1980). At each level, there are available degrees of freedom (decision variables or "inputs") that generally are adjusted locally in order to optimize the local behavior. We are here *not* concerned with the optimal coordination of all these decision makers (which is certainly very interesting), but rather on studying how these individual "players" make and more importantly implement their decisions.

A major problem in making the right decision is that the world is changing. These changes, which we can not affect, are here denoted *disturbances* d. They include changes in exogeneous variables (such as the outdoor temperature), as well as parameter variations in the system (e.g., aging of system components). A common strategy in practice is to use a simple feedback strategy where the degrees of freedom u are adjusted to keep selected controlled variables c at constant values c_s ("setpoints"). The idea is to get "self-optimizing control" where "near-optimal operation" is indirectly achieved, without the need for continuously solving the above optimization problem. In this paper we study this is more detail, and provide a number of examples.

We assume that optimal operation of the system can be quantified in terms of a scalar cost function (performance index) J_o which is to be minimized with respect to the available degrees of freedom u_o,

$$\min_{u_o} J_o(x, u_o, d) \tag{1}$$

subject to the constraints

$$g_1(x, u_o, d) = 0; \quad g_2(x, u_o, d) \leq 0 \tag{2}$$

Here d represents the exogenous disturbances that affect the system, including changes in the model (typically represented by changes in the function g_1), changes in the specifications (constraints), and changes in the parameters (prices) that enter in the cost function (and possibly in the constraints). x represents the internal states. We have available measurements $y = f_o(x, u_o, d)$ that give information about the actual system behavior during operation (y also includes the cost function parameters (prices), measured values

[1] 8th Intl. Symp. on Process Systems Engineering (PSE'2003), 22-27 June 2003, Kunming, China.

of other disturbances d, and measured values of the independent variables u_o). For simplicity, we do not in this paper include time as a variable. The equality constraints ($g_1 = 0$) include the model equations, which give the relationship between the independent variables (u_o and d) and the states (x). The system must generally satisfy several inequality constraints ($g_2 \leq 0$), for example, we usually require that selected variables are positive. The cost function J_o is in many cases a simple linear function of the independent variables with prices as parameters. In many cases it is more natural formulate the optimization problem as a maximization of the profit P, which may formulated as a minimization problem by selecting $J_o = -P$.

In most cases some of inequality constraints are active (i.e. $g_{2c} = 0$) at the optimal solution. Implementation to achieve this is usually simple: We adjust the corresponding number of degrees of freedom u_o such that these active constraints are satisfied (the possible errors in enforcing the constraints should be included as disturbances). In some cases this consumes all the available degrees of freedom. For example, if the original problem is linear (linear cost function with linear constraints g_1 and g_2), then it is well known that from Linear Programming theory that there will be no remaining unconstrained variables.

For nonlinear problems (e.g. g_1 is a nonlinear function), the optimal solution may be unconstrained, and such problems are the focus of this paper. The reason is that it is for the remaining unconstrained degrees of freedom (which we henceforth call u) that the selection of controlled variables is an issue.

For simplicitly, let us write the remaining unconstrained problem in reduced space on the form

$$\min_u J(u, d) \tag{3}$$

where u represents the remaining unconstrained degrees of freedom, and where we have eliminated the states $x = x(u, d)$ by making use of the model equations. J is then not a simple function in the variables u and d, but rather a functional. For any value of the disturbances d we can then solve the (remaining) unconstrained optimization problem (3) and obtain $u_{\text{opt}}(d)$ for which

$$\min_u J(u, d) = J(u_{\text{opt}}(d), d) \stackrel{\text{def}}{=} J_{\text{opt}}(d)$$

The solution of such problems has been studied extensively, and is not the issue of this paper. In this paper the concern is implementation, and how to handle variations (known or unknown) in d in a simple manner. In the following we let d^* denote the nominal value of the disturbances.

Let us first assume that the disturbance variables are constant, i.e., $d = d^*$. In this case implementation is simple: We keep u constant at $u_s = u_{\text{opt}}(d^*)$ (here u_s is the "setpoint" or desired value for u), and we will have optimal operation. (Actually, this assumes that we are able to implement $u = u_s$, which may not be possible in practice due to an implementation error $n = u - u_s$ (Skogestad 2000)). But what happens if d changes? In this case $u_{\text{opt}}(d)$ changes and operation is no longer optimal. What value should we select for u_s in this case? Two "obvious" approaches are

1. If we do not have any information on how the system behaves during actual operation, or if it is not possible to adjust u once it has been selected, then the optimal policy is to find the best "average" value u_s for the expected disturbances, which would involve "backing off" from the nominally optimal setpoints by selecting u_s different from $u_{\text{opt}}(d^*)$. The solution to this problem is quite complex, and depends on the expected disturbance scenario. For example, we may use stochastic optimization (Birge and Louveaux 1997). In any case, operation may generally be far from optimal for a given disturbance d.
2. In this paper we assume that the uconstrained degrees of freedom u may be adjusted freely. Then, if we have information (measurements y) about the actual operation, and we have a model of the system, we may use these measurements in to update the disturbances d, and based on this perform a reoptimization to compute a new optimal value $u_{\text{opt}}(d)$, which is subsequently implemented, $u = u_{\text{opt}}(d)$.

Both of these approaches are complex and require a detailed model of the system, and are not likely to be used in practice, except in special cases. Is there any simpler approach that may work?

2 Implementation of optimal operation: Self-optimizing control

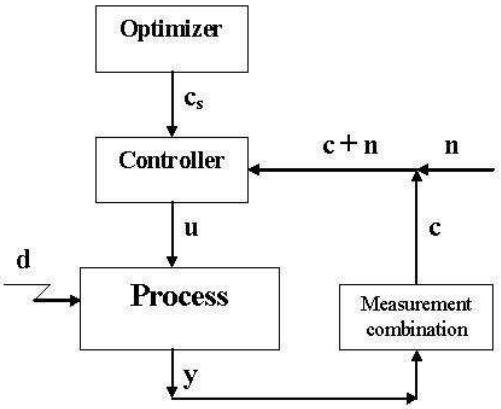

Figure 1: Implementation with separate optimization and control layers. Self-optimizing control is when near-optimal operation is achieved with c_s constant.

We assume in the rest of this paper that we have available measurements $y = f_y(u, d)$ about how the actual operation is proceeding, and that the values of u may be adjusted freely. If we look at how real systems operate, then we see that in most cases a feedback solution is used, whereby the degrees of freedom u are adjusted in order to keep certain controlled variables c at constant values, where c is a selected subset of the available measurements y; see Figure 1.

Example 1. Central Bank. *Consider the role of the Central Bank in a country, which has available one degree of freedom, namely the interest rate (u). The measurements y may in this case include the inflation rate (y_1), the unemployment rate (y_2), the consumer spending (y_3) and the investment rate (y_4). In addition, we also know the chosen interest rate ($y_5 = u$). The simplest policy would be to do nothing, that is, keep the interest rate constant (corresponds to the choice $c = y_5 = u$). A more common policy today is for the Central Bank to adjust the interest rate (u) in an attempt to keep the inflation rate constant (corresponds to the choice $c = y_1$). A typical desired value (setpoint) for the inflation rate is $c_s = 2.5\%$.*

What is the motivation behind attempting to keep c constant at c_s? Obviously, the idea must be that the optimal value of c, denoted $c_{\text{opt}}(d)$, depends only weakly on the disturbances d, such that by keeping c at this value, we indirectly obtain optimal, or at least near-optimal, operation (Morari *et al.* 1980). More precisely, we may define the **loss** L as the difference between the actual value of the cost function obtained with a specific control strategy, e.g. adjusting u to keep $c = c_s$, and the truly optimal value of the cost function, i.e.

$$L(u, d) = J(u, d) - J_{\text{opt}}(d) \tag{4}$$

> **Self-optimizing control** (Skogestad 2000) is when we can achieve an acceptable loss with constant setpoint values for the controlled variables (without the need to reoptimize when disturbances occur)

Let us summarize how the optimal operation may be implemented in practice:

1. A subset of the degrees of freedom u_o are adjusted in order to satisfy the active constraints (as given by the optimization).
2. The remaining unconstrained degrees of freedom (u) are adjusted in order to keep selected constrolled variables c at constant desired values (setpoints) c_s. These variables should be selected to minimize the loss.

Ideally, this results in "self-optimizing control" where no further optimization is required, but in practice some infrequent update of the setpoints c_s may be required. If the set of active constraints changes, then one may have to change the set of controlled variables c, or at least change their setpoints, since the optimal values are expected to change in a discontinuous manner when the set of active constraints change.

We next present some simple examples to illustrate the above ideas.

Example 2. Cake baking. *Let us consider the final process in cake baking, which is to bake it in an oven. Here the are two independent variables, the heat input ($u_1 \stackrel{\text{def}}{=} Q$) and the baking time ($u_2 \stackrel{\text{def}}{=} t$). It is a bit more difficult to define exactly what J is, but it could be quantifed as the average rating of a test panel (where 1 is the best and 10 the worst). One disturbance will be the room temperature. A more important disturbance is probably uncertainty with respect to the actual heat input, for example, due to varying gas pressure for a gas stove, or difficulty in maintaining a constant firing rate for a wooden stove. In practice, this seemingly complex optimization problem, is solved by using a thermostat to keep a constant oven temperature (e.g., keep $c_1 = T_{oven}$ at 200°C). and keeping the cake in the oven for a given time (e.g., choose $c_2 = u_2 = 20$ min). The feedback strategy, based on measuring the oven temperature c_1, gives a self-optimizing solution where the heat input (u_1) is adjusted to correct for disturbances and uncertainty. The optimal value for the controlled variables (c_1 and c_2) are obtained from a cook book, or from experience. An improved strategy may be to measure also the temperature inside the cake, and take out the cake when a given temperature is reached (i.e., u_2 is adjusted to get a given value of $c_2 = T_{cake}$.*

Example 3. Long distance running. *Consider a runner who is participating in a long-distance race, for example a marathon. The cost function to be minimized is the total running time, $J = T$. The independent variable u is the energy input (or something similar). Of course, the runner may perform some "on-line" optimization of his/her body, but this is not easy (especially if the runner is alone), and a constant setpoint policy may probably be more efficient.*

The most common and simplest strategy is to run at the same speed as the other runners (e.g. $c = y_1 =$ distance to best runner, with $c_s = 1$m) , until one is no longer to maintain this speed. However, this does not work if the runner is alone.

Another possible strategy is to keep constant speed ($c = y_2$ speed). However, this policy is not good if the terrain is hilly ($d =$ slope of terrain), where it is clearly optimal to reduce the speed. This policy, as well as the previous one, may also give infeasability, *since the the runner may not able to maintain the desired speed, for example, towards the end of the race.*

A better self-optimizing strategy for a lone runner may be to keep a constant heart rate ($c = y_3 =$ heart rate). In this case, a constant setpoint strategy seems more reasonable, as the speed will be reduced while running uphill.

Example 4. Biology. *Biological systems, for example a single cell, have in place very complex chemical and biochemical reaction newtworks, of which significant parts have the function of a feedback control systems (Savageau 1976) (Doyle and Csete 2002). Indeed, Doyle (lecture, Santa Barbara, Feb. 2002) speculates that many of the supposedly unimportant genes in biological systems are related to control, and compares this with an airplane (or a chemical plant) where the majority of the parts of the system are related to the control system. Biological systems at the cell level are obviously not capable of performing any "on-line" optimization of its overall behavior. Thus, it seems reasonable to assume that biological systems have instead developed self-optimizing control strategies of the kind discussed in this paper. A challenge is to find out how these complex systems work and what the controlled variables are. Biological systems have developed and been optimized over millions of years. If we could identify the controlled variables, then we can also do further "reverse engineering" in an attempt to identify the cost function J_o that nature has been attempting to minimize.*

Example 5. Business systems. *Business systems are very complex with a large number of degrees of freedom (u's), measurements, disturbances and constraints, and may also be regarded as . The overall*

objective of the system is usually to maximize the profit (or more specifically, the net present value of the future profit, $J = -NPV$) (although, businesses are often criticized for using other shorter-term objectives, such as maximizing this years share price, but we will leave that discussion). In any case, it is clear that few managers base their decisions on performing a careful optimization of their overall operation. Instead, managers often make decions about "company policy" which in many cases involved keeping selected controlled variables (c's) at constant values. For example, the common approach of identifying **key perfomance indicators (KPIs)** *for the business, may be viewed as the selection of appropriate controlled variables c. Some examples of KPIs or "value metrics" (Koppel 2001) may be*

- *Time for the business to respond to an order from a customer*
- *Energy consumption per unit produced*
- *Number of accidents per unit produced*
- *Number of employees per unit produced*
- *Degree of automation in the plant*

The optimal value for these variables are typically obtained by analyzing how other succesful businesses perform (benchmarking to find "best business practice").

3 Optimal choice of controlled variables

In most cases the controlled variables c are selected simply as a subset of the measurements y, but more generally we may allow for variable combinations and write $c = h(y)$ where the function $h(y)$ is free to choose. If we only allow for linear variable combinations then we have

$$\Delta c = H\Delta y \tag{5}$$

where the constant matrix H is free to choose. Does there exist a variable combination with zero loss for all disturbances, that is, for which $c_{\rm opt}(d)$ is independent of d? As proved by Alstad and Skogestad (2002) the answer is "yes" for small disturbance changes, provided we have at least as many independent measurements (y's) as there are independent variables (u's and d's). The derivation Alstad and Skogestad (2002) is surprisingly simple: In general, the optimal value of the y's depend on the disturbances, and we may write this dependency as $y_{\rm opt}(d)$. For "small" disturbances the resulting change in the optimal value of $y_{\rm opt}(d)$ depends linearly on d, i.e.

$$\Delta y_{\rm opt}(d) = F\Delta d \tag{6}$$

where the sensitivity $F = dy_{\rm opt}(d)/dd$ is a constant matrix. We would like to find a variable combination $\Delta c = H\Delta y$ such that $\Delta c_{\rm opt} = 0$. We get $\Delta c_{\rm opt} = H\Delta y_{\rm opt} = HF\Delta d = 0$. This should be satisfied for any value of Δd, so we must require that H is selected such that

$$HF = 0 \tag{7}$$

This is always possible provided we have at least as many (independent) measurements y as we have independent variables (u's and d's) (Alstad and Skogestad 2002): First, we need one c (and thus one extra y) for every u, and, second, we need one extra y for every d in order to be able to get $HF = 0$.

Example 1. Central Bank (continued). *For this problem we have u = interst rate and $J = -$ National Product. An important constraint in this problem is that $u \geq 0$ (because a negative interest rate will result in an unstable situation), but in most cases this constraint will not be active, so we have an unconstrained optimization problem with one degree of freedom. The measurements y may include the inflation rate (y_1), the unemployment rate (y_2), the consumer spending (y_3) and the investment rate (y_4). There are many disturbances, for example, d_1 ="the mood" of the consumers, d_2 = global politics, including possible*

wars, d_3 = oil prices, etc. As mentioned earlier, and a common policy is to attempt to keep the inflation rate constant, i.e. $c = y_1$. However, with the large number of disturbances, it is unlikely that this choice is always self-optimizing. Even if we assume that there there is only one major disturbance (e.g. d_1=consumer mood), then from the results presented above we need to combine at least two measurements. This could, for example, be a corrected inflation goal based on using the interest rate, $c = h_1 y_1 + h_2 u$, but more generally we could use additional measurements, $c = h_1 y_1 + h_2 y_2 + h_3 y_3 + h_4 y_4 + h_5 u$. The parameters for such a corrected inflation goal could be obtained by reoptimizing the model for the national economy with alternatives disturbances, using the approach just outlined.

In the above example, the prices were assumed constant. From physical considerations, it is clear that the introduction of price changes may be taken care of by introducing a "price correction" on each controlled variable, but that price changes otherwise will not affect the problem of selecting controlled variables. The reason is that prices appear only in the optimization part of the block diagram in Figure 1, so that it is not prossible to detect price changes in the process itself.

4 Conclusion

The selection of controlled variables for different systems may be unified by making use of the idea of self-optimizing control. The idea is to first define quantitavely the operational objectives through a scalar cost function J to be minimized. The system then needs to be optimized with respect to its degrees of freedom u_o. From this we identify the "active constraints" which are implemented as such. The remaining unconstrained degrees of freedom u are used to controlled selected controlled variables c at constant setpoints. In the paper it is discussed how these variables should be selected. We have in this paper not discussed the implementation error $n = c - c_s$ which may be critical in some applications (Skogestad 2000).

The full version of this paper with an additional detailed example (blending of gasoline) is available at the home page of S. Skogestad, see:
`http://www.nt.ntnu.no/~skoge/publications/2003/selfopt_pse2003/`

References

Alstad, V. and S. Skogestad (2002). Robust operation by controlling the right variable combination. *2002 AIChE Annual Meeting, Indianapolis, USA. (Available from the home page of S. Skogestad).*

Birge, J.R and F. Louveaux (1997). *Introduction to stochastic programming*. Springer.

Doyle, J.C. and M.E. Csete (2002). Reverse engineering of biological complexity. *Science* **295**, 1664–1669.

Findeisen, W, F.N. Bailey, M. Brdys, K. Malinowski, P. Tatjewski and A. Wozniak (1980). *Control and coordination in Hierarchical Systems*. John Wiley & sons.

Koppel, L.B (2001). Business process control: The outer loop. *Proc. Symp. Chemical Process Control (CPC'6), Tuscon, Arizona, Jan. 2001*.

Morari, M., G. Stephanopoulos and Y. Arkun (1980). Studies in the synthesis of control structures for chemical processes. Part I.. *AIChE Journal* **26**(2), 220–232.

Savageau, M.A. (1976). *Biochemical systems analysis*. Addison-Wesley.

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