



Successful industry/academia cooperation: From simple via complex to lucid solutions

Jakob Stoustrup

Automation & Control, Department of Electronic Systems, Aalborg University, Fr. Bajers Vej 7C, 9220 Aalborg, Denmark

ARTICLE INFO

Article history:

Received 27 March 2013

Accepted 2 May 2013

Recommended by Alessandro Astolfi

Available online 7 June 2013

Keywords:

Industrial/academic cooperation

Industrial control

Control project management

Advanced control

Impact of control technology

ABSTRACT

The control literature is rich on impressive applications of advanced control, and within almost any industrial sector there are numerous examples of successful advanced control applications. Nevertheless, there is a widespread belief that there is still a wide potential for increased cooperation between academia and industry within this area.

In this position paper, it is advocated that one of the enablers for successful cooperation between industry and academia within the control area is a proper framework for cooperation projects between companies and universities. Some suggestions for such a framework based on elaborate experience are proposed. In particular, the paper points to the importance of designing projects with project time explicitly scheduled for a phase based on reflection of complex academic solutions. In this phase, the objective is to mimic the behavior of the complex controllers by less complex but industrially feasible solutions. The proposed approach is illustrated by three case studies of successful industrial/academic cooperation.

© 2013 European Control Association. Published by Elsevier Ltd. All rights reserved.

1. Introduction

The scientific area of control theory and control engineering can in many ways be said to be a mature area in the sense that it has a very long history, it has gone through several development phases and it is a technology that has become pervasive in our societies. An intuitive understanding of the behavior of closed loop systems can with some right be said to have been associated with the design of adobe houses for the right thermal behavior or with the design of fire places and smoke openings for maximal ventilation at the beginning of civilization. The conscious design of feedback mechanisms definitely predates Watt's steam engine, although it is difficult to put a precise date to the first real application of feedback design. An early example often mentioned is the water clock of Ktesibios from the third century BC, see Landels [12].

Ever since the antics, the history has been full of successful applications of control methods. Control technology is ubiquitous, and it has become virtually impossible to count the number of explicitly designed control loops in our surroundings, embedded in cars, indoor climate systems, refrigerators, cell phones, etc. Nonetheless, there is a persistent concern within the scientific control community that there are vast unexploited potentials of

control theory within industry, and there is considerable frustration caused by these unexploited potentials. The analysis is probably correct to some extent, as several reports have analyzed the return-of-investment for advanced control projects within several industries and documented that compared to many other technologies, investments in control technology seems to be considerably under-emphasized, see e.g. Samad and Annaswamy [16] and Brisk [4].

To the extent where the potential for industrial/academic cooperation is not fully exploited, this can have several reasons. In many cases, such an endeavor is never pursued, either because such cooperation has never been discussed between the company and a potential academic partner, or because the parties could not agree on a contract. In other cases, an attempt has been made, but it has not been successful in the sense that no results of the cooperation has been implemented in a product (or in the production process, if that has been the scope).

The classical paper [5] challenges frequent statements given in the academic literature, leaving "[...] *the unmistakable impression that those who conceive the theory are in some sense leagues ahead of those who would use it*". Instead, it proceeds to support the thesis that theory "[...] *has some rugged terrain to traverse before it meets the needs of those who would apply it*". At the end, the paper concludes: "*The gap is present indeed, but contrary to the views of many, it is the theoretician who must close it*".

A very recent paper [17], which deserves significant attention, challenges the control community to consider how a novel control

E-mail address: jakob@es.aau.dk

URL: <http://es.aau.dk/staff/jakob>

technology fits into an industrial value chain. In particular, it argues that the customer and the end-user are usually distinct, and whereas the customer might be able to appreciate the virtues of a new control technology, the end-user might not demand the added value sufficiently to justify the added cost for the customer. If a new technology does not provide sufficient added value to each link in the value chain, it will not be implemented.

In this position paper, we shall discuss how a cooperation project between an industrial and an academic partner can be designed in such a way that the probability for a successful outcome is improved. There is clearly a high number of conditions that must be satisfied for such a successful cooperation, starting with agreeing on a contract, assigning the right team, etc., and this paper does not in any sense have the ambition to encompass all those complicated aspects. It just addresses a few aspects regarding the project phase plan that is possibly sometimes overlooked. The approach suggested is based on the experience of the author and his co-workers and a large number (100+) of industrial partners, and is probably applicable to a class of cooperation projects between industry and academia, but definitely not to all.

2. Solutions to industrial problems: from simple via complex to lucid

No prudent industry is prepared to take risks that are unnecessary in the sense that they are not matched by extraordinarily high promises of benefits. Some of the risks associated with applications of advanced control can to some extent be captured by the notion of *complexity*. This does not mean that industry in general will shun complexity. In fact, most industrial systems have in themselves a huge complexity and sometimes the competitive edge is based on mastering complexity. In the advanced control context described here, however, the somewhat ambiguous notion of complexity will be understood in a more narrow sense, e.g. a control solution characterized by having one or more of the following properties:

A crucial dependency of a high fidelity model. Such a model is often difficult and/or expensive to obtain and to maintain. Moreover, as many systems are subject to frequent engineering or operational changes this dependency compromises robustness of the solution.

A large number of parameters in the algorithm. This is challenging in particular if the parameters need to be tailored to the application at hand. In addition to the challenges mentioned under the previous property, this can make the controller difficult to tune and expensive to adjust to new product versions.

A 'monolithic' or highly centralized structure of the algorithm. This makes the controller difficult to integrate with other SW components, e.g. exception handling or monitoring. It also makes debugging more challenging.

A large extent of SW code required for its implementation. This is not always relevant, but crucial if the controller is to be implemented on a strictly resource limited embedded platform.

In designing new control architectures and algorithms, a key (and time-consuming) aspect concerns the analysis of functional and operational requirements. As it is with software engineering, a successful requirements analysis is a key factor in a good controller design. Below, we shall argue that it is crucial that also elements of complexity as indicated by the examples in the above list are taken into consideration.

Using the existing actuators and sensors only is often a requirement for an advanced control project. However, investing in measuring additional physical entities in general may lead to a

complete redesign of the control system providing, at the same time, a simplification and a performance improvement. In the example section of this paper, Section 3, we shall give two examples, where it was crucial to restrict the control solution to existing instrumentation, and one example where the key to the breakthrough was adding one additional low-cost sensor, which in that case was uncomplicated.

In the sequel, we shall discuss various approaches to projects targeted at improving the performance of industrial control loops. There are probably as many types of such projects as there projects themselves. However, in order to facilitate the discussion, we shall introduce a coarse taxonomy based on stereotypical project types, including the approach proposed in this paper. The exposition is deliberately kept at an abstract and non-comprehensive level as the purpose of this section is exclusively to facilitate the discussion of the specific project design aspects addressed in this paper.

2.1. Consultancy type control projects

There is no doubt that a huge unexploited potential across many industries throughout the world is constituted by improper tuning of millions of simple control loops. In Åström [1] it was argued that the largest economical potential for control engineering was constituted by tuning existing PID loops. One example is within Heating Ventilation & Air-Conditioning (HVAC) systems, where Komareji [8] summarizes a technological study based on a survey of a large number of actual HVAC installations in commercial buildings. Although such systems always are delivered with a detailed installation manual, the study mentioned nevertheless documents that a majority of actual HVAC systems are left untuned, i.e. with default settings.

The large potential of making fairly simple changes to existing control loops is mainly due to the large number of such loops worldwide. In many cases the improvement will be significant but perhaps not radical. For the discussion in this paper, the development in a 'consultancy' project is illustrated in Fig. 1 in a stereotypical embodiment. In this stereotype, a consultancy project is characterized by a moderate increase in complexity in the proposed solutions relative to the existing/original solution. In fact, if only tuning is performed, no complexity is added at all. If the control software is modified, a common type of change in this type of project could be to introduce handling of certain special types of operating conditions by conditional software clauses ('if' statements in the source code). In the stereotype consultancy project type illustrated in Fig. 1, the solution complexity grows slowly with time. As the solution complexity grows, also the performance grows, but at the same moderate pace. Typically,

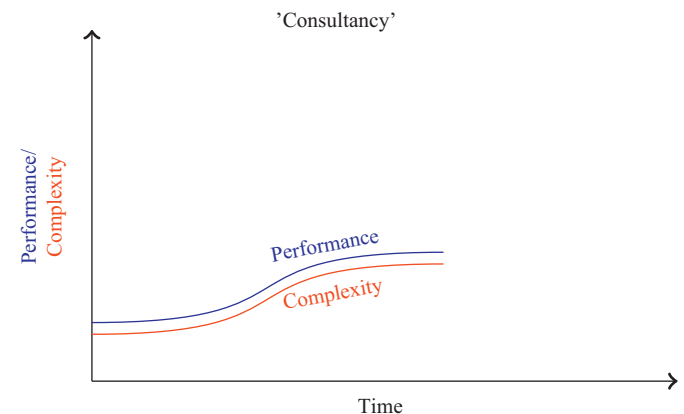


Fig. 1. Project of 'Consultancy' type. The solution only involves a moderate increase in complexity, but also has only a moderate increase in the performance.

the total effort for a project of this type is typically very limited, such that even if the gain in performance is only moderate, the return-of-investment can still be very high.

2.2. Academic type control projects

One of the true virtues of control theory is the provided insight into complex systems. As control theory has matured through its long history, it has also produced solutions itself of an increasing complexity. Therefore, it is a typical attribute of a cooperation project between academia and industry that the academic partner contributes to the solution of a given industrial problem by proposing controllers of a high complexity, compared to existing or state-of-the-art industrial control solutions.

Compared to projects of the ‘Consultancy’ type, typically ‘Academic’ projects (Fig. 2) are much more effort demanding. On the other hand, they also typically have a relatively higher potential for a much higher gain in performance, as it has been argued in Samad and Annaswamy [16]. This suggests a higher potential return-of-investment for this project type.

A major challenge, however, is that academic/industrial cooperation projects often involve a disproportionately high risk, which eventually can compromise the expected return-of-investment. This is not only true for control projects but for almost any kind of academic/industrial cooperation. Throughout the history of the European Union, it has on one hand been documented that the research projects carried out under European Framework Programmes have had high scientific impacts, see e.g. Arnold et al. [3]. On the other hand, the industrial impact in comparison has been much more modest, although every Framework Programme has tightened the criteria for project exploitation plans and taken other initiatives to facilitate industrial exploitation of results obtained.

There are of course numerous reasons, why a seemingly successful research project with significant research results carried out in cooperation between leading academic groups and competent industrial partners fails to exploit solutions obtained. One of the reasons often mentioned, however, is that lack of practicality and/or an excessive level of complexity, prevents the industrial application of the academic results. According to Arnold et al. [3] only around 30% of all projects in FP5 fulfilled important goals concerning development of demos and prototypes.

A somewhat consequential risk for academic/industrial cooperation is caused by skepticism by companies. In part based on the experience mentioned above, some companies will not endeavor into joint academic/industrial project, due to the anticipated risk

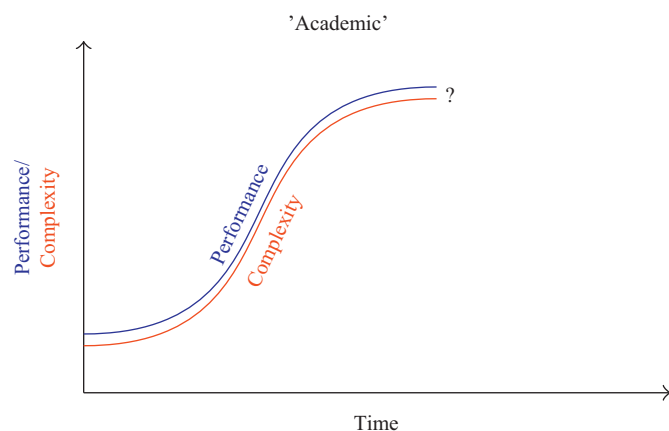


Fig. 2. Project of ‘Academic’ type. The solution provides a significant improvement of performance, but also involves a dramatic increase in complexity.

of failure. This implies that major revenue potentials could be left unexploited.

2.3. Industrial/academic control cooperation: from simple via complex to lucid

Given the premise that a high complexity of solutions to control problems can be preventive for their industrial application, a challenging question is how to design academic/industrial cooperation projects, where on one hand, state-of-the-art (or beyond) control theory is employed and on the other hand, the obtained solutions are still industrially feasible. This is clearly a complex question with no simple answer. It is the position of this paper, however, that in a number of cases, one component in the answer to this complex question is a project design that has handling of the complexity issue built in. The idea is illustrated in Fig. 3, which depicts projects carried out in accordance with, what in this paper is referred to as the *from-simple-via-complex-to-lucid* approach. Projects designed in accordance with this approach have two overall phases.

The starting point for such a project is often an industrial solution that from a control theoretic point of view consists of simple components, e.g. based on a number of single loop PI controllers with no or elementary set-point adaptation. In the first phase of the project, the development of complexity and performance is similar to the ‘Academic’ project type. The solutions proposed will be based on state-of-the-art control theory and the objective is to find ‘optimal’ solutions, i.e. to determine the limits of performance for the process or system at hand by a possibly hypothetical ultimate control solution.

The crucial component in the *from-simple-via-complex-to-lucid* approach, however, is a subsequent project phase initiated with a scheduled conscious reflection on the results from the first phase of the project. The objective of this reflection is not so much to consider *what* was successfully obtained as to *why* this was obtained. It is the position of this paper that a dramatic increase of a highly complex control structure with a very advanced algorithm, often to a large extent, is due to a small number of fairly simple mechanisms that have implicitly been invoked by the advanced and complex control solution. It is, however, also the experience from a significant number of successful advanced control projects carried out as joint academic/industrial projects by colleagues of the author that the importance of these mechanisms would not have been realized without developing the highly complex solutions.

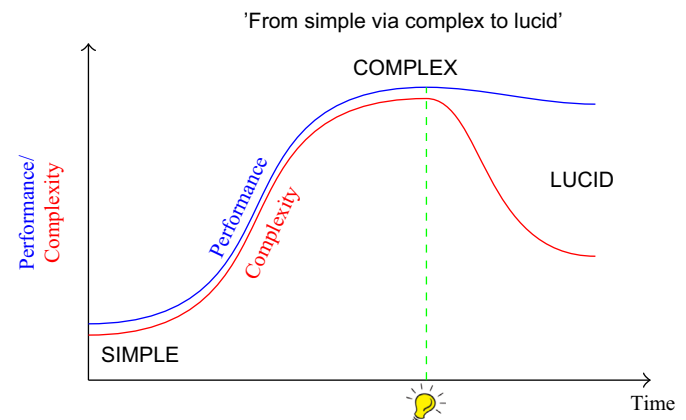


Fig. 3. Project of proposed ‘Simple–Complex–Lucid’ type. In the first phase of the project, the development of complexity and performance is similar to the ‘Academic’ project type. In the second phase, however, the complexity is reduced significantly, with almost the same high performance as for the solution of full complexity.

In the second phase of a project with the *from-simple-via-complex-to-lucid* approach the complexity is reduced significantly, with almost the same high performance as for the solution of full complexity. The main instrument to this is to pursue non-complex ways to invoke the central mechanisms, resulting from the reflection, for realizing the potential displayed in the first phase. In a typical end-result, the actual performance obtained is not quite as high as for the full-blown complex solution. However, almost the same performance can often be obtained by significantly less complex means. The central issue at stake is that these solutions might be only marginally more complex than the original starting point and thereby industrially feasible. In the next section, we shall present a number of case studies that illustrates this point.

A key aspect in the transition from complex to lucid solutions is often the evaluation of the constraints imposed by the automation/computation technological platform (PLC, DSP, etc.). As pointed out in [17], such constraints might be very severe in some applications, especially in embedded control solutions, whereas in others, they might be virtually non-existent. Examples of the latter can be found in the process control industry, where existing solutions already involves solving online optimization problems with millions of decision variables on very power platforms, see [17]. In either case, however, it should be feasible to integrate and maintain a new solution seamlessly with the existing process management system, which is a highly nontrivial constraint.

In many cases, having derived the complex solution, in itself creates value by providing a benchmark for other solutions. Even if the complex solution does not itself lead directly to a 'lucid' solution, it can still serve an important purpose in terms of quantifying how suboptimal more realistic solutions perform.

3. Case studies

In the following, three case studies are presented, which illustrate the principles outlined above. As it happens, the case studies are all (dominantly) thermodynamical systems. However, similar experience has been obtained with a wide range of other types of systems and we strongly believe that the proposed conclusions apply to a broad range of industrial applications.

The described case studies were all carried out in cooperation between academic and industrial partners, and the solutions have all been applied in actual products (the industrial details are left out for competitive reasons).

3.1. Control of Heating, Ventilation & Air-Conditioning systems

The case study described in this subsection is based on Komareji [8]. In this work, the control of Heating, Ventilation & Air-Conditioning (HVAC) systems was studied. The objective for an HVAC system is to deliver conditioned air to maintain thermal comfort and indoor air quality. Such systems account for around 11% of the total electricity consumption in Europe (according to figures from the EC Joint Research Centre). In the present study, the objective was to reduce the electrical power consumption of such systems. The potential for such savings is related to the highly nonlinear dependency of power consumption on operating conditions. Both for fans and pumps, the electrical power consumption has an approximately cubic growth with air and water flow, respectively.

The Heating, Ventilation & Air-Conditioning (HVAC) system, which will be considered here, is shown in Fig. 4. The system is a lab setup, which mimics a typical HVAC system composed of two heat exchangers: a heat recovery part and a water-to-air heat exchanger (an air coil).

In general the heat recovery part has the mission of transferring heat from the exhausted room air to the fresh sucked air. Throughout this process there can be either mixing between the exhausted and fresh air or no mixing between them at all. Here, a heat recovery wheel is applied as a heat recovery part. As can be seen in Fig. 4, there are two separate ducts for conducting the exhausted room air and the fresh outdoor air. An aluminum wheel rotates between two ducts and recovers the thermal energy from exhausted air. It should be noted that there is no mixing between the two air streams. The temperature of the fresh air that leaves the heat recovery part is controllable through the wheel rotation speed manipulation.

After preliminary warming up of the fresh air, it goes through the water-to-air heat exchanger for the final heating. The main task of the water-to-air heat exchanger is to transfer thermal energy from hot water to the fresh air through the coil. The coil is connected to a hydronic circuit which supplies the hot water. The hydronic circuit for the system considered is a so-called primary-tertiary system.

As a typical HVAC system operates in quasi-steady state operation a majority of the time, the main concern for the control system is to ensure that the operating condition is always optimal from an energy consumption point of view. To that end, a system-wide nonlinear constrained optimization problem was formulated and solved, see Komareji et al. [10,11]. The main result was that with a standardized load profile based on a Northern European climate, the electrical power could be reduced by 76% in average!

Based on this result a systemwide nonlinear MPC scheme was developed, which could realize this reduction. This controller was verified in simulation and validated on the lab system. The result was a highly complex and computationally demanding controller that required substantial tuning to reproduce the simulation results.

A closer study of the results obtained showed, however, that the main energy savings were due to the modified pump operation of the water-to-air heat exchanger subsystem.

The main revelation in studying the results obtained was that for any operating point the bypass, i.e. the horizontal pipe segment shown in the upper right of Fig. 4, always had a zero flow! In hindsight, the reason for this is almost trivial. For a non-zero bypass flow, there will always be a mixing of water of different temperatures either to the left (positive flow) or to the right (negative flow) of the bypass valve. This causes an exergy loss without contributing to the energy transfer, which is clearly suboptimal. This in turn implies that any optimal solution must have an operating point with zero bypass flow.

Based on this realization, a control system structure along with a zero bypass flow compensator could be developed as shown in Fig. 5. In this control structure the main controller is an MPC controller and a compensator, which is slower than the MPC controller deals with the bypass flow problem. A simple PI controller can be applied as the bypass flow compensator.

A careful study of the operation of the MPC control action enabled a further simplification, see Komareji et al. [9]. Fig. 6 shows the simplified optimal control structure. One PI controller (C_1) determines the primary water flow (q_{ws}) through the information from the inlet temperature feedback. Tertiary water flow (q_{wt}) is controlled by a PI controller (C_2) which tries to keep the tertiary water flow close to the primary water flow. Actually C_2 is the same as the bypass compensator, but C_2 is a fast controller here. The variable speed pump acts as an actuator to control the tertiary water flow. Because the variable speed pump is much faster than the primary valve, which acts as an actuator to control the primary water flow, the two controllers are decoupled in time domain.

Existing state-of-the-art commercial HVAC control systems is mainly composed of an architecture of PI controllers. In comparison,




Fig. 4. The HVAC system.

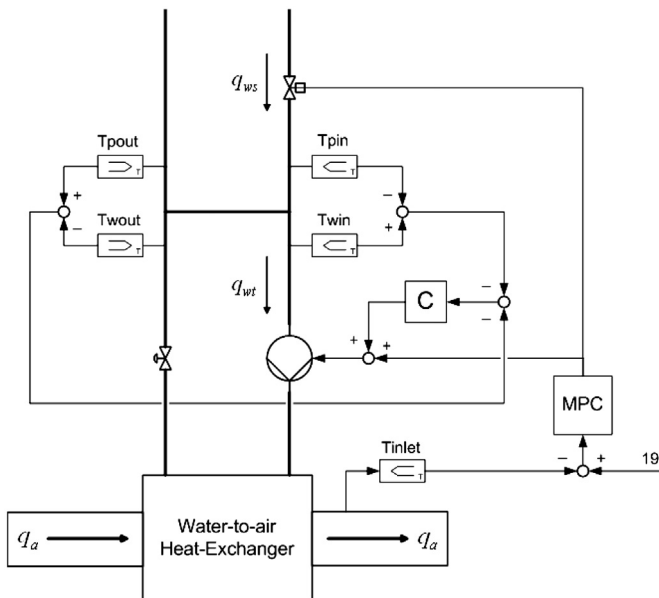


Fig. 5. The control system along with the bypass compensator.

the simplified optimal scheme required one additional temperature sensor (with a cost of approx. 1€) and one additional PI controller, with an easy tuning procedure. This convinced one of the involved industrial partners to implement the solution in a new product, which has been very successful.

The main lesson learned was that a substantial energy reduction was obtained due to a simple observation that can be explained in a first year thermodynamics course. In spite of the trivial nature of the observation, it has never previously been applied in any commercial product by the numerous vendors of HVAC systems. The resulting simplified solution, achieving close to the original 76% electrical energy reduction by the ‘complex’ solution mentioned above, was an essentially trivial extension of existing commercial controllers. Nevertheless, it seems unlikely that it would have been developed without the previous development of the extremely complex global optimization based controller, i.e. that it required the *simple-complex-lucid* approach to succeed.

3.2. Radiator control in pump based systems

The case study described in this subsection is based on Tahersima [18]. This work addresses in part radiator based central

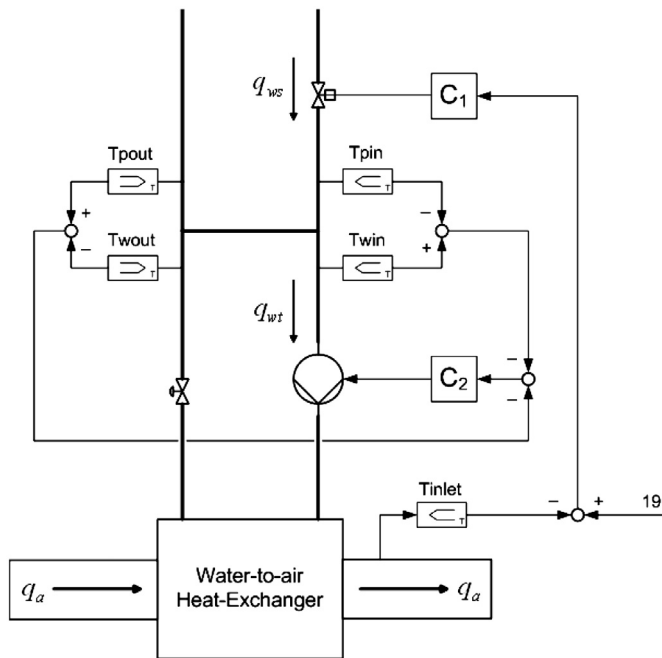


Fig. 6. Simplified optimal control scheme.

heating systems using pump control, which will be described below. As a side comment, another part of the above reference addresses geothermal heat pump based domestic heating control, which also lead to a commercial implementation following the simple–complex–lucid scheme, but this will not be reported here due to space limitation; the interested reader is referred to Tahersima [18].

In countries, where for climatic reasons heating is required, a widespread solution is hydronic floor heating and hydronic radiators controlled by thermostatic radiator valves (TRV) combined with a central heat source through a pump. Due to strong nonlinear dependency on operating conditions, designing controllers for such systems is highly challenging. In fact, the performance of almost all traditional TRV based systems are unsatisfactory seen from a control perspective, as the systems become unresponsive at high heat demands (typically during winter) and often oscillatory at low heat demands (late spring or early fall, i.e. at the endpoints of the heating season). Sometimes the oscillations are too small to be noticed by the users, but a vast majority of actual installations have smaller or larger oscillations in low heat demand situations.

This problem was studied as a dilemma of stability/performance. The dilemma arises when a TRV is regulated by a linear controller with time invariant parameters. A high loop gain and large time constant are the main reasons of this phenomenon. On the other hand, selecting a smaller controller gain to handle the instability situation will result in a poor performance of the radiator when the heat demand is high, see Figs. 7 and 8.

A widespread ‘folklore’ in the industry pointed to the highly nonlinear behavior of the TRVs as the main course for the oscillations widely observed across countries and vendors. In Anderson et al. [2], a validated model was able to reproduce observed oscillations reliably. However, for this simulation model, substituting the TRV with a fictitious component with a completely linear behavior did not remove the oscillations. This study, thus, effectively proved the common hypothesis to be wrong. This strongly motivated further studies.

The system considered in Tahersima [18] is a room heated by a hydronic radiator equipped with a thermostatic valve. A central heating system with multiple radiator circuits equipped with TRVs

is demonstrated in Fig. 9. Disturbances that excite the system are the ambient temperature and solar radiation, see Tahersima et al. [19].

The dilemma between stability and performance arises when the TRV is controlled by a fixed linear controller, see Tahersima et al. [19]. To deal with the dilemma, a gain scheduled controller is proposed which is designed based on a proposed Linear Parameter Varying (LPV) model of the radiator. Parameters of the LPV model are developed in a closed form as a function of the system operating point and the radiator specific dimensioning characteristics. The block diagram of the system, shown in Fig. 10, illustrates the closed loop control system. The chosen values of all parameters are in accordance with the typical experimental and standard values.

This proposed LPV controller was indeed able to handle the full operational envelope. The controller was implemented in a test house. In each operating point it performed just as well as individually tuned PI controllers in that operating point, see Tahersima et al. [20].

A closer study of the validated LPV models that were used to derive this controller revealed that indeed the TRV was not responsible for the high gain variation of linearized models in different operating points. Rather, the gain variations were caused by the radiator dynamics themselves. The somewhat complex models are described in Tahersima [18]. However, at the core of these models are lumped equations for the radiator with temperatures T_i of the individual segments as states. These models contain several nonlinear terms, but are dominated by bilinear terms of the form:

$$\frac{dT_i}{dt} = C_v \cdot q \cdot (T_{i-1} - T_i) + C_r$$

where q is the (indirectly) controlled flow. For some operating point $(\bar{T}_i, \bar{T}_{i-1}, \bar{q})$, the small signal gain from flow to temperature is then

$$G_{Tq}(s) = \frac{C_v(\bar{T}_{i-1} - \bar{T}_i)}{s + C_v\bar{q}}$$

The DC-gain of this transfer function, i.e. $G_{Tq}(0) = C_v(\bar{T}_{i-1} - \bar{T}_i)/C_v\bar{q}$ clearly tends to infinity as \bar{q} tends to zero, i.e. as the heat demand tends to zero. This explains why no fixed LTI controller can stabilize the system for all operating conditions (as the system has more than 180° phase contribution).

It is interesting to note that this phenomenon will occur for any system dominated by a bilinear term between control signal and a controlled state, which indeed is true for a large number of thermal and fluid dynamical systems.

Based on this understanding, it was possible to design a simple gain-scheduled controller that performed comparable to the LPV controller. The architecture for this controller is shown in Fig. 11. The idea is to use an adaptive transformation to remove the flow dependency as shown in the inner loop. After this an LTI lead–lag type controller can ensure that the transformed system always behaves similar to the high demand situation. By choosing high demand as the desired situation, the closed loop system is given the incentive to place the dominant poles as far as possible from the imaginary axis, and as a result as fast as possible.

The resulting behavior is shown in Figs. 12 and 13. It can be seen that the simple gain-scheduled controller performs well in the high heat demand situation without compromising stability in low heat demand situations. Further, the control structure includes very few tuning parameters, which are easy to determine. This significantly contributed for making the solution industrially feasible.




Fig. 7. Performance of a controller that is designed to suit the high demand condition is shown in both low and high heat demand weather conditions.




Fig. 8. Performance of a controller that is designed to suit the low demand condition is shown in both low and high heat demand weather conditions.




Fig. 9. Schematic of the radiator hydraulic circuit.

Existing state-of-the-art systems have controllers, which are of PI type. In comparison, the simplified resulting controller is a classical cascaded controller with a single scheduling parameter.

Although, the final resulting controller has a classical structure, it is doubtful whether it would have been possible to come up with this simple structure without the preceding in-depth analysis and the obtained understanding of the challenge. Further, it would not have been possible to verify that this simple structure has a performance, which is very close to optimal.

3.3. Control of the steam-water cycle in fossil fired power plants

The case study described in this subsection is based on Mølbak [14], Mortensen [13] and Hangstrup [6]. The background for this line of work was a desire in the late 1980s to review the potential for improving the load-following capability of Danish power plants, as the country saw a rapidly increasing penetration of wind power in the Danish power system. This penetration became the world's largest, which it still is (35.3% in 2012). At the same




Fig. 10. Closed loop control system of the room and radiator.




Fig. 11. Block diagram of the closed loop system with linear transformation.




Fig. 12. (Top) ambient temperature and (bottom) room temperature for three controllers. The results of simulation with flow adaptive controller together with two fixed PI controllers are shown. The PI controller designed for the high demand situation encounters instability in the low heat demand condition.




Fig. 13. (Top) ambient temperature and (bottom) room temperature for three controllers. The results of the simulation with flow adaptive controller together with two fixed PI controllers are shown. The PI controller designed for the low demand condition is very slow for the high demand situation.

time as experiencing the dramatically increased intermittency of power production from wind turbines, the Danish power production also was subject to a politically determined decentralization. As an end result the requirement for load-following capability of the Danish power plants that remained under central dispatch increased drastically.

A unique opportunity for industrial/academic cooperation on this issue appeared when it was decided to take one of the blocks (Block #2) of the Danish power plant Skærbækværket off-grid (Fig. 14). At the same time, it was decided that this block could be used for a period of time by the R&D department, in part to address the issues mentioned above. The R&D department chose to involve a university partner and go for an ambitious high-gain/high-risk solution.

Except for the wind farms, the major Danish power plants are all fossil/bio fired power plants. In Fig. 15, the steam-water cycle for Skærbækværket Block #2 is shown.

In the series of works mentioned above, various control loops in the system shown in Fig. 15 were considered, all with the objective of improving load-following without compromising plant lifetime by disproportionate wear of components.

At the end, a solution was found, which is summarized in Mortensen et al. [15], for which the authors received the *Control Engineering Practice Paper Prize* in 1999. The solution is based on a 2DOF LQ solution, which was tested for various operating points for Skærbækværket Block #2. Later, gain-scheduling was studied for this solution, see Hangstrup et al. [7]. The results were convincing with an improvement of load-following capabilities up to almost 100% in the best cases. Based on this, the involved company decided to pursue the research in this direction.

The focus from this point on was to integrate the found solution with the existing power plant control system. The challenges were numerous. First of all, there is a host of legal and safety related issues why algorithms in a commissioned system cannot simply be replaced. Second, it was crucial that an easy fall-back strategy would work in a fail-safe manner. The final solution found was simple but effective. One element in the solution was to implement the controller as a piggy-back component to the existing system. Thus, the 2DOF LQ solution was recomputed, but this time

as a controller intended to give additive control signals to an existing closed-loop system. Another element was to investigate frequency by frequency, which authority was needed by the controller in order to obtain the indicated performance. By doing this, it was possible to map the action of the rather high order LQ solution to a small number of extremely low order linear filters. These filters were implemented and were found to have a performance close to the full-blown model based 2DOF LQ solution.

Due to the simplicity and fail-safe features of the final solution obtained, it was decided to implement this solution on four Danish power plants. It is difficult to provide exact numbers for the average gain in load-following capability obtained, since the power system changed dramatically during the many years it took to complete the project series described above. It is estimated, however, to be in the range of 30% for average load conditions. This solution has subsequently been one of the enablers (along with a huge number of other enablers) for the high penetration of wind power in the Danish power grid.

In this case, the solution was built on top of a highly complex SCADA based control system. The required add-on, however, was just a couple of low order filters that provided control signal adjustments in a piggy-back architecture.

Although the resulting solution was extremely simple (which it had to be), none of the researchers involved in this long progress could envisage that the solution obtained could have been produced in any other way than it was, i.e. starting from a *simple* solution (except that a power plant SCADA system is anywhere close to 'simple') and via a *complex* to the final highly simplified feasible solution, which became *lucid* after lengthy and careful interpretations of the found complex solution.

4. Conclusions

There is a large number of potential challenges that can prevent an industrial/academic cooperation in achieving its full potential or even cause it to fail altogether. Sometimes such a cooperation ends even before it has started if the parties (or their legal departments) are not able to agree on a contract.

This paper has focused on just one of these many potential obstacles for a fruitful industrial/academic cooperation. Nevertheless, it is our belief that this aspect deserves attention. Advanced control solutions as they appear in the control theoretical literature are often not industrially feasible. Although this observation is not very subtle and even fairly widely accepted, nevertheless in our experience, very few research proposals with industrial and academic partners have an explicitly formulated project phase, dedicated to the transition from an advanced control solution of high complexity to a solution having the low complexity required for a vast number of industries.

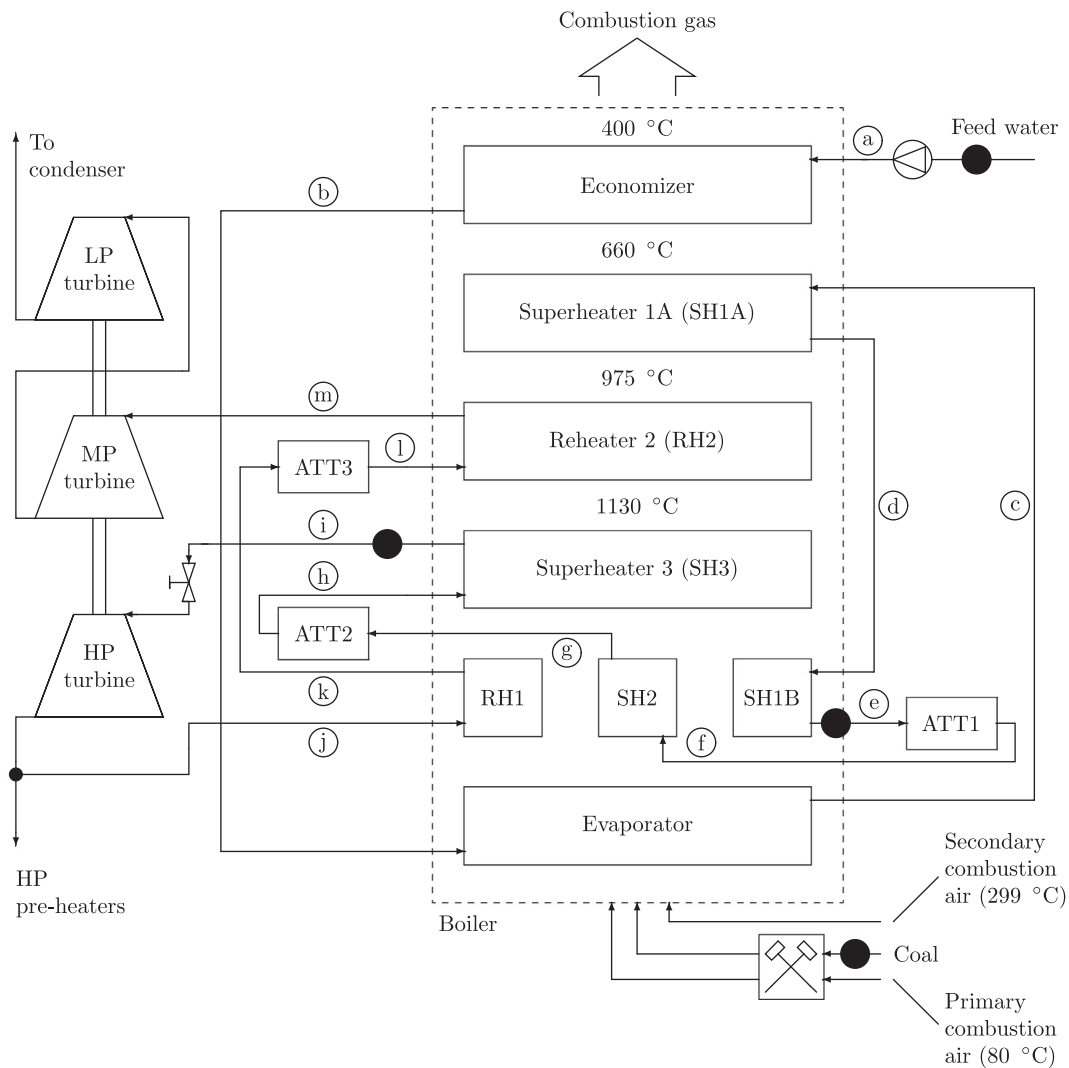
In our experience, a conscious project design following the *Simple* → *Complex* → *Lucid* formula can increase the probability of a successful industrial application of advanced control technology significantly.

It is our hope that the thoughts presented in this paper can stimulate a number of fruitful discussions in the control community. The ideas presented here are clearly not self-contained and could highly benefit from the experience of other academic groups and industries, possibly combining them with other aspects of successful project designs. This could eventually lead the way into unleashing more of the unexploited potential for industrial/academic cooperation within the control systems area.

Finally, it is our hope that the proposed approach along with a number of supporting success stories can be part of a story-telling that can convince some skeptics to endeavor into academic/



Fig. 14. Danish power plant Skærbækværket.



- | | |
|---|--|
| (a) Temperature: 263 °C , Pressure: 216 bar | (h) Temperature: 458 °C , Pressure: 187 bar |
| (b) Temperature: 360 °C , Pressure: 212 bar | (i) Temperature: 545 °C , Pressure: 185 bar |
| (c) Temperature: 375 °C , Pressure: 198 bar | (j) Temperature: 354 °C , Pressure: 42.4 bar |
| (d) Temperature: 437 °C , Pressure: 194 bar | (k) Temperature: 450 °C , Pressure: 41.2 bar |
| (e) Temperature: 462 °C , Pressure: 192 bar | (l) Temperature: 417 °C , Pressure: 40.8 bar |
| (f) Temperature: 445 °C , Pressure: 191 bar | (m) Temperature: 545 °C , Pressure: 40 bar |
| (g) Temperature: 468 °C , Pressure: 188 bar | |

Fig. 15. The steam flow and combustion air/gas flow. The temperatures and pressures are steady state values, valid at 100% load. (ATT) is a spray attenuator.

industrial cooperation, perhaps by transforming some high-gain/high-risk projects to high-gain/less-risk projects.

Acknowledgments

The author would like to thank Editor-in-Chief A. Astolfi for inviting the paper. Also, the author is largely indebted to the anonymous reviewer comments, which provided altogether new perspectives to the presentation, and the final version of the paper benefited significantly from these.

Last, but not least, the author would like to thank several colleagues for stimulating discussions, inspiring the points made

above. They are too many to mention, but the author would like explicitly to thank Professor Andrew G. Alleyne who both participated in the early discussions and also gave invaluable recommendations for this paper that strongly improved it in the final version.

References

[1] K. Åström, *Tuning and adaptation*, Plenary at IFAC World Congress, 1996.
 [2] P. Andersen, T. Pedersen, J. Stoustrup, N. Bidstrup, J. Svendsen, B. Lovmand, Elimination of oscillations in a central heating system using pump control, in: Proceedings of the 2000 American Control Conference, Chicago, IL, 2000, p. 135–139. <http://dx.doi.org/10.1109/ACC.2000.878790>.

- [3] E. Arnold, B. Mahieu, J. Stroyan, D. Campbell, M. Carlberg, F. Giaracca, et al., Understanding the Long Term Impact of the Framework Programme, Technical Report, European Policy Evaluation Consortium (EPEC), 2011.
- [4] M.L. Brisk, Process control: potential benefits and wasted opportunities, in: Fifth Asian Control Conference, vol. 1, 2004, p. 10–16.
- [5] A.S. Foss, Critique of chemical process control theory, *IEEE Transactions on Automatic Control* 18 (6) (1973) 646–652.
- [6] M.E. Hangstrup, Strategies for Industrial Multivariable Control: With Application to Power Plant Control, Ph.D. Thesis, Aalborg University, Aalborg, Denmark, 1998.
- [7] M. Hangstrup, J. Stoustrup, P. Andersen, T. Pedersen, Gain-scheduled control of a fossil-fired power plant boiler, in: Proceedings of the IEEE International Conference on Control Applications, HI, USA, 1999, p. 905–909. <http://dx.doi.org/10.1109/CCA.1999.800913>.
- [8] M. Komareji, Efficient Water Supply in HVAC Systems, Ph.D. Thesis, Aalborg University, Aalborg, Denmark, 2009.
- [9] M. Komareji, J. Stoustrup, H. Rasmussen, N. Bidstrup, F. Nielsen, P. Svendsen, Simplified optimal control in HVAC systems, in: Proceedings of the 2009 IEEE Multi-conference on Systems and Control, St. Petersburg, Russia, 2009, pp. 1033–1038. <http://dx.doi.org/10.1109/CCA.2009.5280724>.
- [10] M. Komareji, J. Stoustrup, H. Rasmussen, N. Bidstrup, P. Svendsen, F. Nielsen, Optimal model-based control in HVAC systems, in: Proceedings of the 2008 American Control Conference, Seattle, WA, USA, 2008, pp. 1443–1448. <http://dx.doi.org/10.1109/ACC.2008.4586695>.
- [11] M. Komareji, J. Stoustrup, H. Rasmussen, N. Bidstrup, P. Svendsen, F. Nielsen, Optimal set-point synthesis in HVAC systems, in: Proceedings of the 2007 American Control Conference, New York City, USA, 2007, pp. 5076–5081. <http://dx.doi.org/10.1109/ACC.2007.4282452>.
- [12] J. Landels, *Engineering in the Ancient World*, University of California Press, Berkeley, 1978, ISBN 0-520-03429-5.
- [13] J.H. Mortensen, Control Strategy for Unleashing the Load-Following Capability of Power Plants, Ph.D. Thesis, Aalborg University, Aalborg, Denmark, 1997.
- [14] T. Mølbak, Optimization of Boiler Control, Ph.D. Thesis, Aalborg University, Aalborg, Denmark, 1990.
- [15] J.H. Mortensen, T. Moelbak, P. Andersen, T.S. Pedersen, Optimization of boiler control to improve the load-following capability of power-plant units, *Control Engineering Practice* 6 (12) (1998) 1531–1539.
- [16] T. Samad, A. Annaswamy, The Impact of Control Technology, IEEE Control Systems Society, 2011.
- [17] T. Samad, G. Stewart, Perspectives on innovation in control systems technology: compatibility with industry practices, *IEEE Transactions on Control Systems Technology* 21 (2) (2013) 284–288.
- [18] F. Tahersima, An Integrated Control System for Heating and Indoor Climate Applications, Ph.D. Thesis, Aalborg University, Aalborg, Denmark, 2012.
- [19] F. Tahersima, J. Stoustrup, H. Rasmussen, An analytical solution for stability-performance dilemma of TRV controlled hydronic radiators, *Energy and Buildings*, vol. 64, 2013, p. 439–446.
- [20] F. Tahersima, J. Stoustrup, H. Rasmussen, Eliminating oscillations in TRV controlled hydronic radiators, in: Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference, Orlando, FL, IEEE, 2011, pp. 4407–4412. <http://dx.doi.org/10.1109/CDC.2011.6161216>.