Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Plug-and-Play Process Control: Improving Control Performance through Sensor Addition and Pre-filtering ⋆

Jan Bendtsen * **Klaus Trangbaek** * **Jakob Stoustrup** *

* *Department of Electronic Systems, Automation and Control, Aalborg University, Denmark (email: {dimon, ktr, jakob}@es.aau.dk)*

Abstract:
An important issue in the area of reconfigurable systems is how to respond correctly if new components are added. We consider the problem of improving control performance for a system where a new set of sensors becomes available. It is assumed that a complete re-design of the control system is undesirable for various reasons. The sensor dynamics are unknown and must be identified via experiments. The paper demonstrates how new sensor information can be fused with existing sensor information and fed to the existing control system, either based on knowledge of the existing plant or in an entirely data-driven fashion. The method is illustrated on a numerical example.

Keywords: Filtering, controller reconfiguration

1. INTRODUCTION

1.1 Plug-and-play process control: an approach towards incremental controller design

The life-time of a controller for an embedded control system might be just as long as the life-time of the embedded system itself, especially if the control system has been designed to handled aging components (e.g. by adaptive control methods) and/or faulty components (e.g. by fault tolerant control methods.

In contrast, the life-time of a high level control system for a complex, industrial process is typically very short, as industrial control processes are often characterized by constant, structural modifications.

The short life-time of high level process control systems is often a limiting factor for companies, when they have to decide whether to invest in advanced control design projects. Obviously, the payback time has to be shorter than the controller life-time, but this precondition might not be satisfied for complicated processes that are subject to frequent, structural changes.

It would thus be highly desirable if new control system hardware could enter into the system in a "plug-and-play" fashion, as known from the PC industry since the 1990's; the system automatically recognises the new device, and dynamics made observable by new sensors are identified, whereupon a controller that utilises new measurements and/or actuation channels can be designed. In a true "plug-and-play" fashion, the hope would then be that the controller be re-designed automatically using efficient numerical tools, thus reducing or even removing the load on the designer.

⋆ This work is supported by The Danish Research Council for Technology and Production Sciences.

The problem here is that a vast majority of control design methodologies are *monolithic* in the sense that they embark from a model of an uncontrolled (open-loop) system and outputs a full, multivariable control system, that does not exploit any knowledge or functionality from previous designs.

This in itself poses the problem that it can be difficult to merge the designed multivariable feedback controller with all the other software that is part of the control and automation system. In fact, in a typical automation system, the feedback control algorithm constitutes at most a few percent of the total number of source code lines.

Thus, there is a need for a novel control design paradigm, where structural changes of a plant can be made by incremental changes of the control system, and where, ultimately, the control system will be able to autonomously re-configure itself in order to accommodate such changes.

This vision can be exemplified by the industrial control scenario described in Knudsen et al. (2008). This case study concerns a district heating system, where the size of the network is increased. If the existing pumping capacity of the network is not sufficient, a number of new pumps and pressure sensors are added. It is non-trivial how to extend the control system in such a way that both the existing and the new part of the network performs adequately, without having to redesign the entire control system, which is not feasible. A common practical solution is to integrate network extensions by separating them from the existing network by heat exchangers. This solution, however, is sub-optimal, and Knudsen et al. (2008) demonstrate that a better solution can be obtained by applying a plug-and-play control concept.

Another industrial case study is described in Michelsen et al. (2008). This paper concerns the problem of tem-

perature control of a domestic heating system. For such a system, new heating elements may be added after the initial installation, i.e., the house might have been built with a floor heating system, and an electrical heater might then be added later. The challenge is now that the built-in controllers of the individual heat emitters might be conflicting, as they are typically all designed to control the temperature individually. In Michelsen et al. (2008) it is described how to design these control systems in a way such that they can be added one by one without compromising the overall system performance.

As a third example, in Trangbaek et al. (2008), a buffer tank example is studied. The fluid level is controlled by a pump and a valve in series. The tank is disturbed by an unmeasured load flow. The only measurement is the fluid level. At the original design, the valve was manually operated, meaning that the control system could only use the pump to control the fluid level, and a single loop controller is designed to that end. After some time, it is found that the performance is not satisfactory and that the strain on the pump is too high. Therefore the manual valve is replaced with an electronically controlled one, and Trangbaek et al. (2008) describes a systematic procedure for attaching an add-on controller in a "piggy-back configuration on top of the existing one while retaining stability at all times.

As a common denominator for these three case studies, a control system is already running before the structural modification, and in each case, it is not acceptable to take the plant out of operation in order to re-commission the entire control system.

The three examples are described here to exemplify the need for a new control design paradigm as outlined above. Such a paradigm is the goal of a research program carried out by a consortium of companies and universities in Italy and Denmark, see Plug & Play Process Control.

1.2 Adding a new sensor

This paper deals with another example of a problem that should be addressed by this new control design paradigm: the situation where a new sensor is plugged into a control system, providing measurements hitherto unavailable to the control system. Typically, when introducing a new sensor in this way, the system designer hopes to obtain better information about specific parts of the plant, e.g., extra temperature or pressure sensors in a power plant, leading to (for instance) higher-bandwidth control or better observability of certain plant states.

On the other hand, as outlined above, the control engineers working at the plant are rarely keen on de-commissioning the existing control system, since the workload involved and the system downtime this incurs, likely poses an unacceptable burden on the plant, economically and otherwise.

So far, not much work has been done in this area. Mostly, the issue of on-line reconfigurable control systems has been done in the field of fault-tolerant control, see e.g., Bao et al. (2003); Blanke et al. (2003); Zhang et al. (2004). However, in fault-tolerant control, the situation is usually the reverse of the one outlined above, namely that a sensor or actuator breaks down, and the control system has to be able to survive the loss. Benitez-Perez and Garcia-Nocetti (2003) treats replacement of "dumb" sensors and actuators in an existing distributed control system with "intelligent" ones, along with the associated communication scheduling problem, but the numbers of in- and outputs remain the same before and after the reconfiguration. Apparently, only Niemann (2006) explicitly considers the problem of expanding an existing control system with new sensors/actuators, but full knowledge of both control and plant dynamics is required, and a design method is not presented.

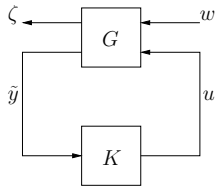

Fig. 1. System setup. The existing system G is stabilised by the existing controller K.

With complete plant and sensor knowledge, one might also consider an approach like Abdelrahman and Kandasamy (2003). As mentioned, however, the control system cannot always be re-designed from the bottom up. Rather than redesigning the control system, this paper instead proposes to use new measurements together with the existing measurements in a data driven, *sensor fusion*-like fashion (cf. e.g. Brooks and Iyengar (1998); Dasarathy (1994) and the references therein). In addition to avoiding the control redesign, the approach chosen here does not require estimation of the new dynamics revealed by the new sensor either; it only requires knowledge of the "old" plant dynamics. To deal with the case where no such knowledge is available, we propose an entirely data-driven method for sensor fusion as an alternative.

The paper proceeds as follows. Section 2 describes the setup, focusing on the existing control loop and the closed loop with an extra sensor added. Then, we present an approach to the aforementioned sensor fusion design based on knowledge of existing plant dynamics and then propose a purely data-driven filter design to fuse the new sensor measurement with the existing measurements, assuming no plant knowledge is available. Section 3 then compares the approaches through numerical examples, and finally Section 4 discusses the results and points out various issues that should be addressed in future work.

2. INCORPORATING A NEW SENSOR

2.1 Setup

We consider the setup sketched in Fig. 1. The plant G has the (known) state space realisation

$$\begin{aligned} x_{k+1} &= Ax_k + B_u u_k + B_w w_k \\ \zeta_k &= C_\zeta x_k + D_\zeta w_k \\ \tilde{y}_k &= \tilde{C} x_k + D_w w_k \end{aligned} \tag{1}$$

which maps disturbance signals $w \in \mathbb{R}^{m_w}$ and control signals $u \in \mathbb{R}^m$ to performance outputs $\zeta \in \mathbb{R}^{p_\zeta}$ and original measurement outputs $\tilde{y} \in \mathbb{R}^p$. $A, B_w, B_u, C_\zeta, D_\zeta, \tilde{C}$ and D_w are constant real matrices of appropriate dimensions,

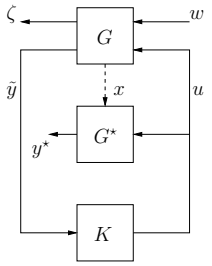

Fig. 2. A new measurement $y^\star$ becomes available, providing measurements from the subsystem $G^\star$.

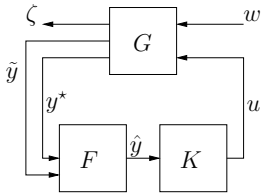

Fig. 3. The filter F is introduced to fuse the sensor measurements, providing an improved output estimate $\hat{y}$ to be used for the exiting feedback control.

with (A, B_u) and $(A, \tilde{C})$ controllable and observable pairs, respectively. The disturbance $w \in \mathbb{R}^{(n+p+p_\zeta)}$ contains reference signals y_r as well as state and measurement noise. The noise is assumed to be sampled from a zero-mean Gaussian sequence with covariance matrix equal to the identity matrix. The actual effects on the states and outputs are modelled via B_w and D_w, while D_ζ models the effect of reference and noise on the performance. G is controlled by a controller $\tilde{K}$ with *fixed* transfer function.

At some point a new sensor is connected to the plant. The new sensor provides access to measurements $y^\star \in \mathbb{R}^{p^\star}$ from a new dynamical subsystem $G^\star$, whose states are affected by the control inputs and system states. The dynamics may be due to the sensor itself, as well as dynamics in the plant. However, we shall assume that the subsystem's states do not affect the original system G; see Fig. 2. Thus, even if the sensor dynamics were known, the new plant dynamics have so far been unobservable. As a consequence, the new subsystem is considered entirely unknown at the point where the new sensor comes on-line.

As outlined earlier, we assume that the controller itself cannot be reconfigured. Thus, we propose to design a sensor fusion filter, denoted F in Fig. 3, which includes the new sensor measurement in the existing measurements, hopefully improving the 'quality' of the signals provided to the controller.

We assume that the controller itself is optimal with respect to some performance objective given the original measurements. Thus, no improvement can be achieved by any causal filtering of these without using the new measurement. What we can hope to achieve is to remove measurement noise by using information from the new sensor.

In the following we will propose various methods for designing such a filter.

2.2 Subspace-based Filtering

In this section we present a method for designing the filter $F(z)$ using a model of G (but not of $G^\star$ or K) and measurements. The design consists of two parts. First we generate a sequence of desirable measurements, in the sense that we attempt to remove measurement noise efficiently using a *non-causal* filter. Then we find a causal filter that will generate this desired result using all measurements.

Note that x_k in (1) is not measured, but can be estimated by various means, since G is assumed to be known. Here we propose to employ *subspace methods* (see e.g., van Overschee and Moor (1996)) to estimate the states, since they are not restricted to open-loop stable systems. In the following, we give a brief review of this approach.

We consider the known system G. For a given number of block rows $i > n$, we define the following block Hankel matrices along with an extended observability matrix Γ_i:

$$U_{0|i-1} = \begin{bmatrix} u_0 & u_1 & u_2 & \dots & u_{j-1} \\ u_1 & u_2 & u_3 & \dots & u_j \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ u_{i-1} & u_i & u_{i+1} & \dots & u_{i+j-2} \end{bmatrix}, \quad \Gamma_i = \begin{bmatrix} \tilde{C} \\ \tilde{C}A \\ \tilde{C}A^2 \\ \vdots \\ \tilde{C}A^{i-1} \end{bmatrix}$$

$$\tilde{Y}_{0|i-1} = \begin{bmatrix} \tilde{y}_0 & \tilde{y}_1 & \tilde{y}_2 & \dots & \tilde{y}_{j-1} \\ \tilde{y}_1 & \tilde{y}_2 & \tilde{y}_3 & \dots & \tilde{y}_j \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \tilde{y}_{i-1} & \tilde{y}_i & \tilde{y}_{i+1} & \dots & \tilde{y}_{i+j-2} \end{bmatrix}$$

Note that, since $(A, \tilde{C})$ are assumed to be observable, the rank of Γ_i is equal to n.

Next, we define the extended reversed controllability matrix Δ_i as

$$\Delta_i = \begin{bmatrix} A^{i-1}B_u & A^{i-2}B_u & \dots & AB_u & B_u \end{bmatrix}$$

and a lower block triangular Toeplitz matrix T_i incorporating all the system matrices as

$$T_i = \begin{bmatrix} D_w & 0 & 0 & \dots & 0 \\ \tilde{C}B_u & D_w & 0 & \dots & 0 \\ \tilde{C}AB_u & \tilde{C}B_u & D_w & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \tilde{C}A^{i-2}B_u & \tilde{C}A^{i-3}B_u & \tilde{C}A^{i-4}B_u & \dots & D_w \end{bmatrix}$$

Then the optimal state sequence

$$X_i = [x_i, x_{i+1}, x_{i+2}, \dots, x_{i+j-1}]$$

is computed as (see van Overschee and Moor (1996)):

$$X_i = \begin{bmatrix} \Delta_i - A^i\Gamma_i^\dagger T_i & A^i\Gamma_i^\dagger \end{bmatrix} \begin{bmatrix} U_i \\ \tilde{Y}_i \end{bmatrix} \tag{2}$$

This state estimate can then be used to generate an optimal filtered existing output sequence $\hat{y}_{opt}$:

$$\hat{y}_{opt,k} = \tilde{C}x_k, \qquad k = i, \dots, i+j-1$$

Note that the filtering is noncausal and as such can be expected to give better results than a Kalman filter. Also, a Kalman filter design would require knowledge of the full system, which we assume is not available here.

Alternatively, we can generate y_opt by a low pass filtering. A causal low pass filter will create a phase shift, and cannot

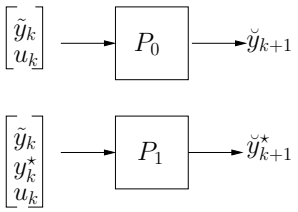

Fig. 4. Setup for data-driven filter design.

be expected to produce anything useful. But by using a non-causal filter, zero phase shift can be achieved. More specifically, the original measurements are first filtered through a standard low pass filter. The filtered sequence is then reversed in time and filtered again through the same filter, and reversed again. Of course, no choice of a causal F can reproduce this filtered sequence if only the original measurements are provided, but in some cases the additional measurements will help. The point is that by a lucky selection of filter poles, mainly measurement noise will be removed, although the best way to choose these poles is still an open question.

In order to compute F in Fig. 3, we parameterise F using some appropriate set of parameters θ (for instance as an ARX model) and solve the following minimisation problem

$$\min_\theta \left\| \hat{y}_{opt,k} - F(\theta, z) \begin{bmatrix} y_k^\star \\ \tilde{y}_k \\ u_k \end{bmatrix} \right\| \tag{3}$$

for $k = 1, \ldots, N$ and some appropriate signal norm $\|\cdot\|$.

2.3 Data-Driven Filtering

If the plant model is *not* known, it is not possible to generate the optimal state sequence, and hence $\hat{y}_{opt}$ in the previous section is not known. To devise an entirely data-driven approach we instead turn our attention to optimal predictors.

Consider the setup in Fig. 4. From measurement data, we can design two predictors for $\tilde{y}$: $P_0 = [P_{0y}(\alpha_0, z)\ P_{0u}(\beta_0, z)]$ uses control signals and only the old output measurements, whereas $P_1 = [P_{1y}(\alpha_1, z)\ P_{1u}(\beta_0, z)]$ also uses the new measurement, while still only predicting the old measurements. In each case α and β represents free parameters. In the design, we enforce $\beta_1 = \beta_0$, which is equivalent to forcing the part concerned with the control signal to be the same for the two filters, i.e., $P_{1u}(\beta_0, z) = P_{0u}(\beta_0, z)$.

Assuming that the new measurement contains useful information, we would now expect P_1 to be a better predictor than P_0. If not, we would like the design method to produce $F = I$, since the controller is assumed to be optimal.

This can be achieved with the setup in Fig. 5. Letting P_1 generate a good prediction of $\tilde{y}$, the inverse of P_0 can then generate a filtered version of the measurement to be fed to the controller. Observe that the control signal may be needed for the filter design, but is not necessary in the real-time implementation.

To sum up, the procedure is as follows:

(1) Identify P_0 based on measurements of $\tilde{y}_k$ and u_k. Extract the mapping $P_{0u}(\beta_0, z)$

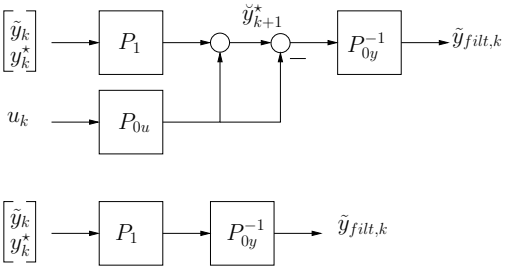

Fig. 5. Construction of $F(z) = P_{0y}(\alpha_0, z)^{-1} P_{1y}(\alpha_1, z)$.

(2) Identify P_1 based on measurements of $\tilde{y}_k, y_k^\star$ and u_k subject to the constraint that $P_{1u}(\beta_0, z) = P_{0u}(\beta_0, z)$
(3) Construct F as $F = P_{0y}(\alpha_0, z)^{-1} P_{1y}(\alpha_1, z)$.

3. NUMERICAL EXAMPLES

This section illustrates the methods proposed in the previous section through two numerical examples. First, a simple single-input, single-output open-loop unstable system described by

$$x_{k+1} = \begin{bmatrix} 1.05 & 0.8 \\ 0 & 0.8 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ 2 \end{bmatrix} u_k + \begin{bmatrix} w_{1,k} \\ w_{2,k} \end{bmatrix}$$
$$\tilde{y}_k = x_{1,k} + w_{3,k}$$

With $E(w_k w_k^T) = I$, a standard LQG controller is designed for this system using standard techniques, see e.g., Zhou et al. (1995). Since the controller is optimal, no performance gain is possible by simply filtering the existing measurements.

A new sensor then comes online, adding the extra output

$$y_k^\star = x_{2,k} + w_{4,k}$$

where the standard deviation of w_4 is one-tenth of the standard deviation of w_3. Utilising this new output, we design three different filters, denoted $F_{opt}(z)$, $F_{zpf}(z)$ and $F_{FIR}(z)$, respectively. F_{opt} is designed assuming full system infomation is available, using optimal controller design techniques; this filter is included purely for comparison purposes. $F_{zpf}(z)$ is designed via non-causal zero phase lag filtering, while $F_{FIR}(z)$ is designed as outlined in Section 2.3 using second-order finite-impulse response filters.

The following table lists the variances of the control and state sequences with and without each of the new filters:

	Original	F_{opt}	F_{FIR}	F_{zpf}
$E(u^T u)$	1.1468	0.6483	0.5618	0.6328
$E(x_1^T x_1)$	9.86365	3.95976	5.3780	9.7832
$E(x_2^T x_2)$	5.4543	2.2210	2.6055	4.1245

From the closed-loop frequency responses shown in Fig. 6 and the table, it is seen that the filters in each case improve the performance by including the new sensor measurement. Furthermore, it is worth noting that the introduction of the filters leads to performance improvements in almost the entire frequency range, except for the channel from w_1 to x_1. From Fig. 7 it is noted that the shapes of the frequency responses of both $F_{zpf}(z)$ and $F_{FIR}(z)$ are similar to the optimal full-information filter $F_{opt}(z)$ in spite of the fact that no model information was used in the design of these filters.

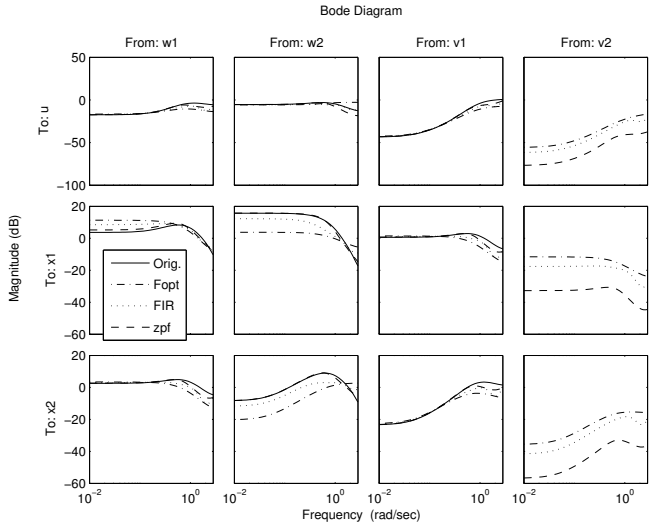

Fig. 6. Magnitude Bode diagrams of the closed loop transfer functions from the four noise channels w_1, W_2, w_3 and w_4 to the control signal u and the states x_1 and x_2.

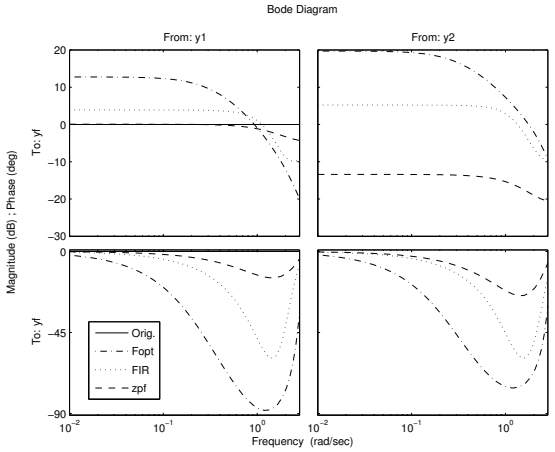

Fig. 7. Frequency response of the designed filters.

Next we consider an open-loop unstable system with two outputs and two controllable inputs of the form (1) with

$$A = \begin{bmatrix} 0.9220 & 0.1592 & -0.1579 & -0.0012 \\ 0.2046 & 0.9654 & 0.1139 & 0.0654 \\ 0.1503 & -0.0304 & 0.8254 & -0.0804 \\ 0.1424 & -0.5030 & 0.0621 & 0.7526 \end{bmatrix}$$

$$B_u = \begin{bmatrix} 0.0006 & 0.0008 \\ 0.0002 & 0.0015 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \tilde{C} = [I\ 0]$$

An LQG controller with integrators for both outputs has been designed for this system. The correct model of G and the actual noise covariances have been used in the design.

The system is affected by reference signals, state and measurement noise. Thus, w is defined as $w_k = \begin{bmatrix} y_{ref,k}^T & \nu_k^T \end{bmatrix}^T$, where $y_{ref,k} \in \mathbb{R}^2$ and ν is a 6-dimensional stationary random process with mean 0 and covariance $E(\nu\nu^T) = I$. The influence of the disturbance on the states and outputs

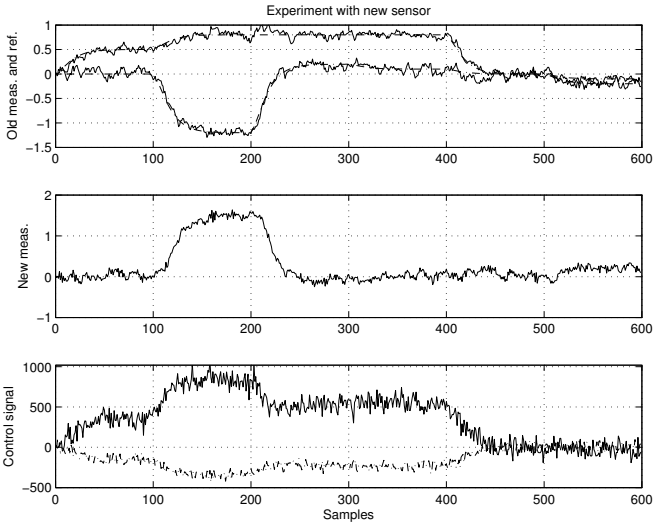

Fig. 8. Initial experiment with new sensor. The top plot shows the measured output $\tilde{y}$ with measurement noise removed (full line) plotted together with the reference y_{ref} (dash-dotted line). The middle plot shows the new sensor output $y^\star$. The bottom plot shows the control signal u.

is given by $B_w = \begin{bmatrix} 0 & 10^{-2}I & 0 \end{bmatrix}$ and $D_w = \begin{bmatrix} -I & 0 & 10^{-1}I \end{bmatrix}$, respectively.

Then, a new sensor is introduced into the system, providing measurements from a new subsystem $G^\star$ (cf. 2) with the following (unknown) dynamics:

$$\begin{aligned} x_{k+1}^\star &= A^\star x_k^\star + \begin{bmatrix} \tilde{A}^\star & B^\star \end{bmatrix} \begin{bmatrix} x_k \\ u_k \end{bmatrix} + \eta_k \\ y_k^\star &= C^\star x_k^\star + \xi_k \end{aligned} \tag{4}$$

with $A^\star, B^\star$ and $C^\star$ represent subsystem dynamics and the in- and output matrices of the subsystem, respectively, and η_k and ξ_k normally distributed independent random samples. $\tilde{A}^\star$ represents the coupling between the states of G and the subsystem states.

$$\tilde{A}^\star = \begin{bmatrix} 0 & 0 & 0.2 & 0.1 \\ 0 & 0 & -0.3 & 0.64 \end{bmatrix}, \quad A^\star = \begin{bmatrix} 0 & 0.8 \\ 0.2 & -0.12 \end{bmatrix}$$

$$B^\star = \begin{bmatrix} 0 & 0 \\ 0.0002 & 0.00165 \end{bmatrix} \quad C^\star = [0\ 0\ 0\ 0\ 1\ 0]$$

An experiment is carried out where the existing measurement $\tilde{y}$ is used for output feedback and the new measurement $y^\star$ is simply recorded. The control signals and measurements are shown in Fig. 8. The new sensor is simulated with measurement noise with one-tenth the standard deviation of the sensors that yield the existing output measurements.

Next, equation (2) is used to generate the optimal filtered state sequence based on the known system $(A, B_u, \tilde{C})$ and the input-output data collected during the experiment. From these filtered states, a desired output sequence $\{\hat{y}\}$ was then computed and used for the design of a sensor fusion filter

$$F(z) = \left[\begin{array}{c|c} A_F & B_F \\ \hline C_F & D_F \end{array}\right]$$

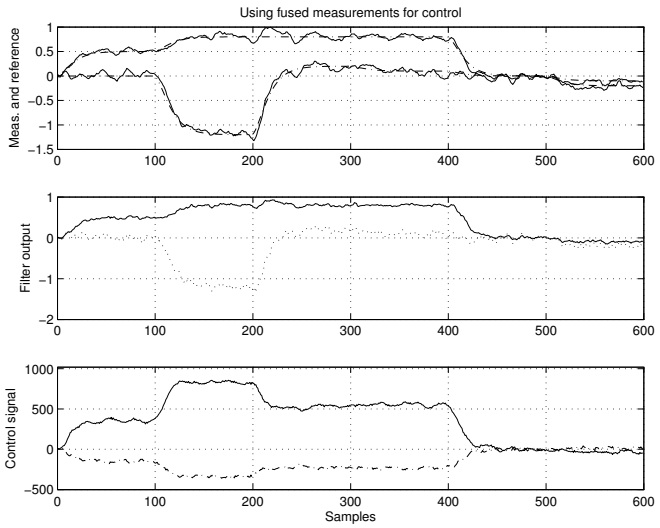

Fig. 9. Experiment where the new sensor measurements are fused with the existing measurements. The control signal is clearly more smooth than in Figure 8.

F was found by approximating the following minimisation:

$$\min_{A_F,B_F,C_F,D_F} \left\| \hat{y} - F(z) \begin{bmatrix} \tilde{y} \\ y^\star \\ u \end{bmatrix} \right\|_2 \tag{5}$$

by a standard subspace identification method, restricting F to be of order two.

The sensor fusion filter was then inserted in closed loop and the experiment was repeated (with the same noise sequence for comparison). The result can be seen in Fig. 9. From this figure, it is observed that although the output performance has only been improved by a small amount, the control is clearly much smoother than before the filter was introduced. Deeper examination reveals, however, that the filter does not actually use the new measurement considerably (i.e., the gain from $y^\star$ to $\tilde{y}$ is very small). Instead, the performance improvement seems to be due to a noise smoothing effect in the filter.

The filter needs access to the control signal in the above implementation. However, by removing u from the optimisation problem (5), we obtain a filter without this need. This has also been tested on the example, again resulting in a performance improvement (albeit not quite as significant).

4. CONCLUSION

This paper identified a need for a novel direction in control engineering, which enables designers to make incremental changes to the control system of a 'living' process. That is, a process that changes with time in a structural fashion, such as when new actuators or sensors are added. It then proposed a novel method for fusing new sensor measurements with existing sensor measurements from an existing plant, so that they may be used in an existing control loop. It is assumed that the controller cannot be changed; this situation is fairly common in large-scale plants where existing controllers are implemented using programmable logic controllers with various safety logics etc.

Two methods for designing the filter were presented, one based on partial plant knowledge being available, for instance acquired during plant operation, and one that does not asssume any plant knowledge at all. Perhaps surprisingly, the method that ignores plant knowledge appears to work best, at least in simulations. This is likely to be an issue of applying open-loop estimation methods in closed-loop.

REFERENCES

M. Abdelrahman and P. Kandasamy. Integration of multiple sensor fusion in controller design. *ISA Transactions*, 42:197205, 2003.

J. Bao, W. Z. Zhang, and P. L. Lee. Decentralized fault-tolerant control system design for unstable processes. *Chemical Engineering Science*, 58:50455054, 2003.

H. Benitez-Perez and F. Garcia-Nocetti. Reconfigurable distributed control using smart peripheral elements. *Control Engineering Practice*, 11:975–988, 2003.

M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki. *Diagnosis and Fault-Tolerant Control*. Springer Berlin-Heidelberg, 2003.

R. R. Brooks and S. S. Iyengar. *Multiple-Sensor Fusion-Fundamentals and Applications with Software*. Prentice Hall, Englewood Cliffs, NJ, 1998.

B. V. Dasarathy. *Decision Fusion*. IEEE Press, New York, 1994.

Torben Knudsen, Klaus Trangbaek, and Carsten Skovmose Kallese. Plug and play process control applied to a district heating system. In *Proceedings of the 2008 IFAC World Congress*, Seoul, Korea, July 2008.

Axel G. Michelsen, Roozbeh Izadi-Zamanabadi, and Jakob Stoustrup. Towards automatic model based controller design for reconfigurable plants. In *Proceedings of the 2008 IFAC World Congress*, Seoul, Korea, July 2008.

H. Niemann. Parameterisation of extended systems. *IEE Proceedings: Control Theory and Applications*, 153:221–227, 2006.

Plug & Play Process Control. Research program of the Danish Research Council for Technology and Production Sciences. http://www.control.aau.dk/plugandplay, 2006.

Klaus Trangbaek, Jakob Stoustrup, and Jan Bendtsen. Stable controller reconfiguration through terminal connections. In *Proceedings of the 2008 IFAC World Congress*, Seoul, Korea, July 2008.

P. van Overschee and B. De Moor. *Subspace Identification for Linear Systems*. Kluwer Academic Publishers, 1996.

X. D. Zhang, T. Parisini, and M. M. Polycarpou. Adaptive fault tolerant control of nonlinear uncertain systems: An informationbased diagnostic approach. *IEEE Trans. Automatic Control*, 49:12591274, 2004.

K. Zhou, J. Doyle, and K. Glover. *Robust And Optimal Control*. Prentice-Hall International, 1995. ISBN 0-13-456567-3.