


A decoupled feedback structure for covertly appropriating networked control systems

Conference Paper

Author(s):

Smith, Roy 

Publication date:

2011-01

Permanent link:

<https://doi.org/10.3929/ethz-b-000044492>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

IFAC Proceedings Volumes 44(1), <https://doi.org/10.3182/20110828-6-IT-1002.01721>

A Decoupled Feedback Structure for Covertly Appropriating Networked Control Systems

Roy S. Smith

*ETH-Zurich, Automatic Control Institute, Physikstrasse 3, 8092
Zürich, Switzerland (email: rsmith@control.ee.ethz.ch).*

Abstract: The use of communication networks for the transmission of control system signals (measurements and actuation) is becoming more widespread. If the communication (or the sensors and actuators themselves) can be modified by a malicious agent then the control of the physical plant can be covertly appropriated. A parameterised decoupling structure is introduced which allows the covert agent a wide range of control actions on the physical plant while remaining undetectable from the point of view of the original networked controller. The designer of the covert agent need only have a model of the physical plant; knowledge of the networked controller is not required. A MIMO process control example (based on the control of irrigation canals) is used to illustrate the concepts.

1. INTRODUCTION

The ready availability of components for both wireless and wired network communications has accelerated the adoption of feedback control systems that operate over network communication links. Such networks have been used for monitoring and supervisory control of geographically widespread or separated processes; the SCADA architecture is such an example.

Such systems also have potential security problems. A malicious agent can more readily gain access to the signals and information within the control loop and use these to disrupt or compromise the controlled plant. Such attacks can take place entirely within the communication network or they may include physical interference with the actuators and sensors. The problem arises because the controller receives all of its information about the operation of the physical plant via potentially modifiable information channels.

If the objective of the malicious agent is simply the disruption of the controlled system then there are well known means of attack via the network, one example being denial-of-service attacks. In this paper we are more interested in stealth; the malicious agent does not want to reveal to the controller that the system is compromised. The term covert agent is a more fitting description.

In our scenario we assume that the covert agent can modify the sensing and actuation signals. This may be accomplished from within the network, and/or by modification or augmentation of the physical sensors. If the plant is linear, time-invariant and known to the covert agent then that agent can use a parameterised feedback based structure to gain control of the plant in a manner that cannot be detected by the controller. From the controller's point of view, the effect of every noise and disturbance entering the physical system is identical to the uncompromised case. Perfect characterizations of the noise/disturbance

characteristics and/or probing signals are unable to detect the presence of the covert agent. In this sense, the actions of the covert agent are invisible to the controller. In this paper we present a covert agent architecture that gives these properties.

As networked control architectures becomes more prevalent the security of these systems is a matter of growing concern. There is a significant research effort on the vulnerability of SCADA systems from a computer security point of view (Igre et al. [2006]). See also Amin et al. [2010], Chabukswar et al. [2010] and Sandberg et al. [2010]. Several aspects of this prior work are relevant here. Amin et al. [2010] uses a detailed model and design procedure for an irrigation channel in order to compromise its operation. They make the observation that the attack will be harder to detect if the downstream measurements are also modified. Sandberg et al. [2010] studies the detection of attacks in large-scale MIMO systems and observes that (from at least a measurement point of view) attacks which are consistent with the dynamics will not be detected. The parameterisation given in the current work modifies both the actuation and sensing signals in a way that is consistent with the dynamics from the point of view of the controller and drives the physical plant to an operating point of the covert agent's choosing.

2. COVERT CONTROLLER PARAMETERISATION

The nominally controlled system is illustrated in Figure 1. This is a standard feedback control system; the network serves to emphasize that the only information received by the controller is the measured signal, y_m , transmitted via the network.

In order to be able to write closed-loop transfer functions we assume that the controller is linear and time-invariant. The results given below still hold for a nonlinear or time-varying controller. Consider the controller to given by,




Fig. 1. Nominally connected networked control system. The controller, C , is assumed to receive all information about the physical plant, P , via potentially modifiable communication links.

$$u_c = [C_y \ C_r] \begin{bmatrix} y_m \\ y_{\text{ref}} \end{bmatrix},$$

where r is a reference signal. The measurement, y_m , is comprised of the plant output, additive noise, n , and the effects of a disturbance, w .

$$y_m = P_u u + P_w w + n. \quad (1)$$

In the uncompromised case the actuation is faithfully reproduced giving,

$$u = u_c.$$

This leads to closed-loop transfer functions given by,

$$y_m = (I - P_u C_y)^{-1} P_u C_r y_{\text{ref}} + (I - P_u C_y)^{-1} n + (I - P_u C_y)^{-1} P_w w. \quad (2)$$

We assume that C_y and C_r have been chosen to give the appropriate trade-off between closed-loop stable noise, disturbance and reference tracking responses. It is assumed that any additional delays or communication noise introduced in the network is accounted for in the design of C_y and C_r .

Note that in (2) we have expressed the transfer functions in terms of y_m . The purpose of doing this is to show them in terms of signals *that are available to the controller during operation*. In doing the design one would instead work with transfer functions in terms of y in order to distinguish between the effects of sensor noise and output disturbances.

2.1 Model of the Covert Agent

We now consider the case where a covert agent can both measure and corrupt the signals transmitted via the network, y_m and u_c . This not only applies to the case where the signals are intercepted on the network; it also considers the situation where either the actuators or sensors are compromised. Because the parameterisation given here is in a feedback form we will also refer to the covert agent as a covert controller.

In the following we assume:

1. The covert agent has a model of the plant's control to output mapping, denoted by Π_u ;

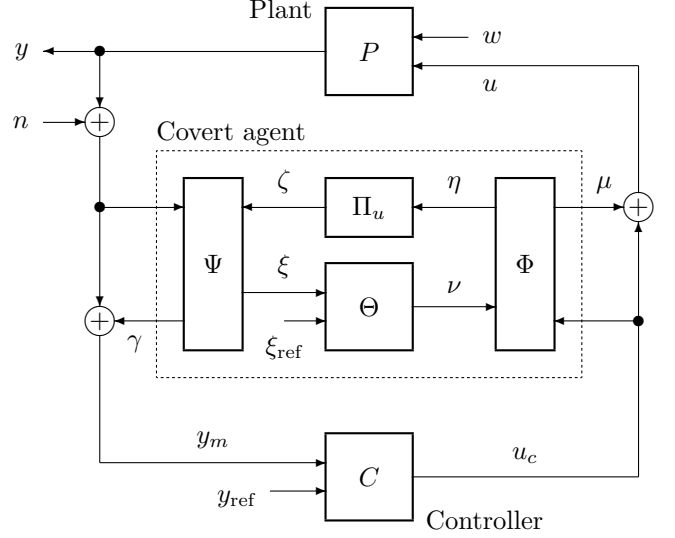


Fig. 2. Interconnection diagram illustrating the covert agent's parameterisation. The use of decoupling and feedback makes the covert agent's actions, specified by Θ and ξ_{ref} , undetectable by the controller, C .

2. The covert agent can only measure and add to existing control or measurement signals; and,
3. The plant is linear and time-invariant.

The analysis given below will assume that the covert agent's model of the plant is exact; $P_u = \Pi_u$. In practice it is sufficient that any error between P_u and Π_u be smaller than the allowable robustness margins for which the controller was designed.

The assumption that the covert agent can only add to the signal is not particularly restrictive but makes it clear that this configuration also addresses the case where the plant actuation is compromised by connecting another actuator in parallel.

The covert agent model is illustrated in Figure 2. The agent measures u_c and y_m and adds to these the signals μ and γ respectively. The internal structure of the agent is given in this form in order to derive its effect in terms of a decoupled control action.

The Φ decoupling block is given by,

$$\begin{aligned} \begin{bmatrix} \gamma \\ \xi \end{bmatrix} &= \Psi \begin{bmatrix} y + n \\ \zeta \end{bmatrix} \\ &= \left(\begin{bmatrix} 0 & 1 \\ (1-\lambda)/\lambda & -1 \end{bmatrix} \otimes I_{ny} \right) \begin{bmatrix} y + n \\ \zeta \end{bmatrix}, \end{aligned}$$

where \otimes denotes the Kronecker product and I_{ny} is an identity matrix of the same dimension as y . Similarly, define I_{nu} as an identity matrix of dimension equal to that of u . The Φ decoupling block is then given by,

$$\begin{bmatrix} \mu \\ \eta \end{bmatrix} = \Phi \begin{bmatrix} u_c \\ \nu \end{bmatrix} = \left(\begin{bmatrix} (\lambda-1) & \lambda \\ (1-\lambda) & -\lambda \end{bmatrix} \otimes I_{nu} \right) \begin{bmatrix} u_c \\ \nu \end{bmatrix},$$

Both Φ and Ψ are functions of a parameter $\lambda \neq 0$ that can be chosen by the covert agent. It will become clear that $1 - \lambda$ determines the amount of disturbance rejection control handled by the covert controller. The choice of λ also has consequences on the range of decoupled control

actions available to the covert controller. These issues will be discussed later and it will be seen that for practical purposes $\lambda = 1$ is a good choice.

To see the effect of this decoupling consider the system in terms of the measurements and actuator for the controller, C . The measurement vector is now given by,

$$\begin{aligned} y_m &= y + n + \gamma \\ &= P_u u + \Pi_u(1 - \lambda)u_c + \Pi_u(-\lambda)\nu + n + P_w w, \\ &= \Pi_u u_c + \lambda(P_u - \Pi_u)(u_c + \nu) + n + P_w w, \end{aligned}$$

and if $P_u = \Pi_u$, then,

$$y_m = P_u u_c + P_w w + n. \quad (3)$$

Note that this is exactly the case for the uncompromised closed-loop given in (1). Using,

$$u_c = C_y y_m + C_r y_{\text{ref}},$$

gives the closed-loop response,

$$\begin{aligned} y_m &= (I - P_u C_y)^{-1} P_u C_r y_{\text{ref}} \\ &\quad + (I - P_u C_y)^{-1} P_w w + (I - P_u C_y)^{-1} n. \end{aligned}$$

As expected, from the point of view of the controller, this is identical to uncompromised closed-loop given in (2). However, as we will subsequently see the action on the true plant output, y , is very different.

Note that the derivation of y_m in (3) used only the linearity of P_u and Π_u and the static equations given by the decoupling blocks Ψ and Φ . No assumptions were made about the controller.

2.2 Design Parameterisation of the Covert Agent

In the structure illustrated in Figure 2 the degrees of freedom available to the covert agent are the value of λ , the choice of operator Θ , and the signal ξ_{ref} . We will interpret the last two as a controller and reference signal respectively.

For simplicity in the following, the Θ operator will be taken to be a linear time-invariant operator of the form,

$$\nu = [\Theta_\xi \ \Theta_r] \begin{bmatrix} \xi \\ \xi_{\text{ref}} \end{bmatrix}.$$

From the definition of Ψ we have,

$$\begin{aligned} \xi &= \left(\frac{1 - \lambda}{\lambda} \right) (y + n) - \zeta \\ &= \left(\frac{1 - \lambda}{\lambda} \right) (P_u u + P_w w + n) - \Pi_u \eta. \end{aligned}$$

Substituting,

$$u = u_c + \mu = \lambda u_c + \lambda \nu,$$

and

$$\eta = (\lambda - 1)u_c - \lambda \zeta,$$

leads to,

$$\begin{aligned} \xi &= (P_u + \lambda(\Pi_u - P_u))\nu + (1 - \lambda)(P_u - \Pi_u)u_c \\ &\quad + \frac{1 - \lambda}{\lambda}(n + P_w w). \end{aligned}$$

We again use the assumption that $P_u = \Pi_u$ to get,

$$\xi = \Pi_u \nu + \left(\frac{1 - \lambda}{\lambda} \right) (n + P_w w).$$

Note that ξ and ν are the input and output (respectively) of Θ . We can therefore interpret the function of Θ as closing a feedback loop around Π_u with the objective of tracking ξ_{ref} . Using this interpretation gives as a closed-loop transfer function,

$$\begin{aligned} \xi &= (I - \Pi_u \Theta_\xi)^{-1} \Pi_u \Theta_r \xi_{\text{ref}} \\ &\quad + \frac{1 - \lambda}{\lambda} (I - \Pi_u \Theta_\xi)^{-1} (P_w w + n). \end{aligned} \quad (4)$$

The only constraint that is placed on Θ is that the closed-loop transfer functions in (4) are stable. We can express the ‘‘design’’ of the covert agent in terms of the design of the closed-loop transfer functions above. Section 2.3 discusses the actions available to the covert agent in more detail.

From the above we can see that the actuation signals, u_c and ν , and measurement signals, y_m and ξ , give a decoupled representation of the compromised system.

$$\begin{bmatrix} y_m \\ \xi \end{bmatrix} = \begin{bmatrix} P_u & 0 \\ 0 & \Pi_u \end{bmatrix} \begin{bmatrix} u_c \\ \nu \end{bmatrix} + \begin{bmatrix} P_w & 1 \\ P_w(1 - \lambda)/\lambda & (1 - \lambda)/\lambda \end{bmatrix} \begin{bmatrix} w \\ n \end{bmatrix}. \quad (5)$$

The decoupled representation allows us to design the covert controller, Θ , without being required to know the networked controller, C .

2.3 Covert Agent Control Options

The only constraint that we have on design of the covert controller is that Θ stabilizes Π_u . This allows a wide range of control actions for the covert controller.

However the covert controller’s objectives are stated in terms of the variable ξ which is not the same as the physical output of the plant, y . From the definition of Ψ and the interconnection shown in Figure 2 we have,

$$y = \lambda \xi + \lambda y_m - n. \quad (6)$$

If the covert controller knows the networked controller’s reference, y_{ref} , then this information can be used to specify a ξ reference value, ξ_{ref} , that gives a desired value for y . For simplicity assume that w and n are both zero mean and that the networked controller, C , has been designed such that,

$$\lim_{t \rightarrow \infty} E\{y\} = y_{\text{ref}},$$

where $E\{\}$ denotes the expected value. The design of the networked controller, C , assumes that,

$$\lim_{t \rightarrow \infty} E\{y\} = \lim_{t \rightarrow \infty} E\{y_m\}$$

but in fact C controls only y_m giving instead,

$$\lim_{t \rightarrow \infty} E\{y_m\} = y_{\text{ref}}.$$

In the absence of the covert controller this would be sufficient to ensure the desired performance of plant in closed-loop with the networked controller.

Suppose now that the covert controller’s objective is to instead set the output y to a covert reference value of y_s . Then implementing Θ as a reference tracking controller such that,

$$\lim_{t \rightarrow \infty} E\{\xi\} = \xi_{\text{ref}},$$

gives (via Equation 6),

$$\lim_{t \rightarrow \infty} E\{y\} = \lim_{t \rightarrow \infty} E\{\lambda \xi + \lambda y_m - n\} = \lambda \xi_{\text{ref}} + \lambda y_{\text{ref}}.$$




Fig. 3. Illustration of the irrigation canal system. The canals are narrow (2.5 m.) compared to their length.

Now by choosing, $\xi_{\text{ref}} = (y_s - \lambda y_{\text{ref}})/\lambda$, we have,

$$\lim_{t \rightarrow \infty} E\{y\} = y_s.$$

In this manner the covert controller can effectively reset the reference tracking set-point to any desired value.

We note here that the choice of $\lambda = 1$ in the decoupling parameterisation makes the selection of ξ_{ref} simpler. In this case

$$\lim_{t \rightarrow \infty} E\{y\} = y_{\text{ref}} + \xi_{\text{ref}}.$$

This allows the covert control to specify its reference plant output, y_s , with respect to the controller's reference value, y_{ref} . In the application example we will use this choice to give a covert controller that drives the plant output to a fixed value less than the reference controller.

3. AN APPLICATION EXAMPLE

The covert controller parameterisation is illustrated on an irrigation canal control example described by Sánchez-Peña et al. [2009]. This class of problem has also been studied from a network security point of view by Amin et al. [2010]. Because the sensing and actuation is geographically separated applying closed-loop control to irrigation canals often requires networked control systems.

The system described by Sánchez-Peña et al. [2009] is briefly outlined here for context, and is illustrated in Figure 3. A reservoir is assumed to be at a fixed height (3.5 m.) and its outlet flow is controlled by a sluice gate with the gate height, u_1 , proportional to the flow. The second actuator is a sluice gate controlling the flow from canal 1 to canal 2. The actuation variable is the gate height, u_2 . Both sluice gates are constrained to a maximum opening of 0.9 metres. The second canal is terminated by a fixed height spillway. The two measured variables of interest for control purposes are the heights at the downstream ends of each canal; h_1 and h_2 respectively.

The full model is described by a pair of partial differential equations known as the Saint-Venant equations. Sánchez-Peña et al. [2009] give a simplified model derived using the approach described by Litrico and Fromion [2004]. The linearized plant is described by,

$$y = \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} = P_u u$$

$$= \begin{bmatrix} 1 & 0 \\ 0 & e^{-t_2 s} \end{bmatrix} \begin{bmatrix} \frac{4.87}{1800s + 1} & \frac{-4.35}{2100s + 1} \\ \frac{1.20}{1900s + 1} & \frac{1.40}{1500s + 1} \end{bmatrix} \begin{bmatrix} e^{-t_1 s} & 0 \\ 0 & 1 \end{bmatrix} u,$$

where $t_1 = 7$ minutes and $t_2 = 15$ minutes.




Fig. 4. Canal water levels under closed-loop control by the nominal networked controller. A step reference change in h_1 (from 0.64 to 0.94) is commanded at 300 minutes. No change is commanded for h_2 . A flow disturbance on at gate 1 (shown in Figure 5) begins at 600 minutes.

The networked reference tracking controller for this plant is based on a robust MIMO Smith-predictor method described by Sánchez-Peña et al. [2009]. The controller incorporates an integral term for zero steady-state tracking error of h_1 and h_2 . In the simulations presented here an Internal Model Control (IMC) implementation is used (see Morari and Zafiriou [1989] for details). The networked controller is designed to be robust with respect to variations in the time delays within the plant. This feature would also be useful for handling the delay uncertainty within the communication network.

The closed-loop simulation given in Figures 4 and 5 establishes a baseline for the performance of the networked control system. A reference step change in h_1 (from an initial value of $h_1 = 0.64$ m. to a value of $h_1 = 0.94$ m.) occurs at 300 minutes. The second canal level is held constant at $h_2 = 0.96$ m. No sensor noise is included in these simulations so that the networked controller's responses can be more clearly seen.

The closed-loop system exhibits a well damped response and zero steady-state tracking errors. There is relatively little cross-coupling between the h_1 and h_2 channels. The flow disturbance on u_1 (from 600 to 1000 minutes) is equivalent to raising sluice gate 1 by 0.01 metres. It gives rise to a disturbance in h_1 of approximately 0.02 metres which is then rejected by the networked controller.

The above simulation scenario is now repeated with a covert controller actively manipulating the actuation and measurement signals. The covert controller structure is that illustrated in Figure 2. The decoupling blocks, Φ and Ψ , use the parameter $\lambda = 1.0$. This has several advantages in this application. The first is that the covert controller's objective can be specified with respect to the networked controller's objective. In this case the covert controller's objective will be to maintain canal 1 height at of 0.1 metres below the networked controller's h_1 reference value.




Fig. 5. Gate position actuation signals under closed-loop control by the nominal networked controller. A step reference change in h_1 is commanded at 300 minutes. At 600 minutes a flow disturbance of 400 minutes duration is introduced.




Fig. 6. Canal water levels in the presence of the covert controller. The command and disturbances signals are identical to those shown in Figures 4 and 5. In this case the covert controller (using $\lambda = 1$) introduces a controlled reduction in the level of h_1 . The commanded reduction is 0.1 metres and begins at 50 minutes.

The second advantage is that the covert controller need not know the linearization biases (steady-state values of u and y at the linearization point) used by the networked controller, C . If $\lambda \neq 1$ these biases must be subtracted from the u_c input to Φ and added to the γ output of Ψ . A corollary of this observation is that Π_u can be a linearized model of the plant P .

The results of the simulation with the covert controller are shown in Figures 6 and 7. The covert controller's reference objective is,




Fig. 7. Gate position actuation signals in the presence of the covert controller. Shown here are the actual gate level commands (u_1 and u_2) as well as the commands generated by the networked controller.

$$\xi_{\text{ref}} = \begin{bmatrix} -0.1 \\ 0 \end{bmatrix},$$

and this step is commanded at the 50 minute time point. As described above this has the effect of reducing the actual canal 1 height, h_1 , to 0.1 metres below the h_1 reference being maintained by the networked controller, C . The canal 2 height, h_2 , is held at the level specified by the networked controller.

Figure 6 shows that these objectives have been accomplished by the covert controller. At the 50 minute time point h_1 (for the plant) drops from $h_1 = 0.64$ to $h_1 = 0.54$. The offset of 0.1 metres on h_1 is maintained when the networked controller raises its h_1 reference value to 0.94 at 300 minutes.

The heights measured by the networked controller are also shown on Figure 6. The largest difference between the covert controller scenario and the baseline scenario in Figure 4 is a 10 minute deviation of 0.03 metres in the measurement of h_2 at the 50 minute point. This is due to the transient coupling between the height channels when the covert controller turns on. This deviation is significantly smaller than that due to the flow disturbance at the 600 minute point and would be undetectable in presence of noise and small flow disturbances.

The sluice gate actuation signals, u , (shown in Figure 7) differ from the commanded actuation signals, u_c , from the 50 minute point onwards. The difference between commanded and actuals signals is due to the action of the covert controller.

Figures 8 and 9 illustrate the operation of the covert controller by showing its internal signals ξ and ν respectively. Using the decoupled representation in (5) we have,

$$\xi = \Pi_u \nu + \frac{1-\lambda}{\lambda} (P_w w + n),$$

and with $\lambda = 1$ this becomes,

$$\xi = \Pi_u \nu. \quad (7)$$




Fig. 8. Measurement and reference signals (ξ and ξ_{ref}) for the covert controller. In this implementation $\lambda = 1$ and the signal ξ is the difference between the covert controller’s reference plant output and the networked controller’s reference plant output.




Fig. 9. Actuation signal, ν , for the covert controller, Θ . With $\lambda = 1$ this actuation signal is added to the networked controllers actuation commands and the covert controller does not react to the flow disturbance.

The design of the covert controller is accomplished by designing Θ as a reference tracking controller for the linearized model Π_u in (7). In this example a linearized version of the networked controller, C , was used for Θ . It is important to note that this need not be the case; Θ can be any controller designed for operation with Π_u .

The ability of the (Π_u, Θ) closed-loop to track a $\xi_1 = -0.1$ step command is illustrated here. Note that with $\lambda = 1$ the covert controller does not respond to the flow disturbance at the 600 minute point. This is evident in the Θ output signals, ν .

The simulation illustrates another important point about the covert controller parameterisation. Although the de-

coupling is exact when $\Pi_u = P_u$ we do not assume that the state of covert controller’s model, Π_u , matches the state of P_u . This is demonstrated by the effect of the flow disturbance at 600 minutes. The physical plant states clearly respond to the disturbance, yet the states of the covert model, Π_u , clearly do not. The reason for this is the linearity and time-invariance assumptions; the Π_u model compensates only for the effect of the covert controller’s actuation on the plant. These effects can be added to those of the networked controller, C , irrespective of the time.

4. DISCUSSION AND FUTURE WORK

Covertly appropriating a networked control system requires only access to both the sensing and actuation signals, and a model of the plant being controlled. The plant model need not be identical to that used for the design of the networked controller but the difference between P_u and Π_u should be less than the robustness margins used in the design of the networked controller, C . A more detailed quantification of this constraint is an area for future work.

This work illustrates the importance of secure communication links and the physical security of the actuators and sensors in a networked control system. A malicious agent need only use the formulation given here to covertly take control of a physical plant. Intrusion detection based on the measurement, actuation or probing signals will not detect such agents.

5. ACKNOWLEDGMENTS

The author would like to thank Bruno Sinópoli (CMU) and Henrik Sandberg (KTH, Stockholm) for helpful discussions. Thanks are also due to Ricardo Sánchez-Peña (Inst. Tecnológico, Buenos Aires) for providing the simulation code for the process control example.

REFERENCES

- Saurabh Amin, Xavier Litrico, S. Shankar Sastry, and Alexandre M. Bayen. Stealthy deception attacks on water SCADA systems. In *13th ACM Int. Conf. Hybrid Systems*, pages 161–170, 2010.
- Rohan Chabukswar, Bruno Sinópoli, Gabor Karsai, Anarita Giani, Himanshu Neema, and Andrew Davis. Simulation of network attacks on SCADA systems. In *Proc. 1st. Workshop on Secure Control Systems*, pages 1–8, April 2010.
- Vinay M. Ijure, Sean A. Laughter, and Ronald D. Williams. Security issues in SCADA networks. *Computers & Security*, 25:498–506, 2006.
- X. Litrico and V. Fromion. Simplified modeling of irrigation canals for controller design. *J. Irrigation & Drainage Engr.*, pages 373–383, 2004.
- Manfred Morari and Evangelhos Zafiriou. *Robust Process Control*. Prentice-Hall, New Jersey, 1989.
- Ricardo S. Sánchez-Peña, Yolanda Bolea, and Vicenç Puig. MIMO Smith predictor: Global and structured robust performance analysis. *J. Process Control*, 19:163–177, 2009.
- Henrik Sandberg, André Teixeira, and Karl H. Johansson. On security indices for state estimators in power networks. In *Proc. 1st. Workshop on Secure Control Systems*, pages 1–6, April 2010.