The logarithmic quantiser is not optimal for LQ control

Jean-Charles Delvenne

Abstract— **We seek to stabilise a scalar linear system through a finite-capacity communication channel, while minimising a quadratic cost. We show that the logarithmic quantiser strategy is not optimal for the quadratic cost in the limit of low capacities.**

Control with a communication constraint

Imagine a plant that communicates with a physically distant controller through communication channels. The limited capacity of communication lines makes the design of controller able to stabilise the plant a more difficult, and sometimes impossible, task. This family of problems, acutely useful with the generalisation of digital communication technologies, leads to a dialogue between control theory and information theory which has essentially started with [2], and has been pursued actively in the last decade (see, e.g., [1], [6], [4], [5], [8]).

In this paper, we consider the simplest such problem: stabilising the discrete-time scalar linear system $x_{t+1} = ax_t + u_t$, with $|a| > 1$. We know that the initial state is distributed uniformly in the interval $[-1, 1]$. The controller receives the state x_t through a communication channel, with a capacity limited to R bits par channel use (see Fig. 1). An energy function on the state space $x \mapsto E(x)$, typically $E(x) = x^2$, is given. Our task is to design a coding/decoding strategy for the communication channel and a controller in order to make the closed-loop system stable and minimise the cost defined as follows. The cost of such a control strategy is defined as $\mathbb{E} \sum_{t \geq 0} E(x_t)$, where $\mathbb{E}$ denotes the expected value over all initial states. For $E(x) = x^2$, we therefore solve an LQ optimal control problem. This problem is a generalisation of the one considered in [5], [3]. Note that $\mathbb{E} \sum_{t \geq 0} E(x_t)$ can be written as $\mathbb{E}(E(x_0))\tau$, where τ is a time, loosely interpreted as the average time taken by a piece of energy to leave the system.

Usually control problems involve a cost function of the form $x_t^2 + u_t^2$. In the standard framework (in the absence of communication constraint) the term u_t^2 is meant to limit the speed of convergence to the solution and avoid a trivial deadbeat solution. In this case however, the limited capacity already imposes a limited speed of convergence, therefore the penalty on the input is not required.

This paper presents research results of the Belgian Programme on Interuniversity Attraction Poles, initiated by the Belgian Federal Science Policy Office. The scientific responsibility rests with its author.

JCD is with FUNDP (University of Namur), Department of Mathematics, 8, Rempart de la Vierge, 5000 Namur, Belgium, jean-charles.delvenne@math.fundp.ac.be

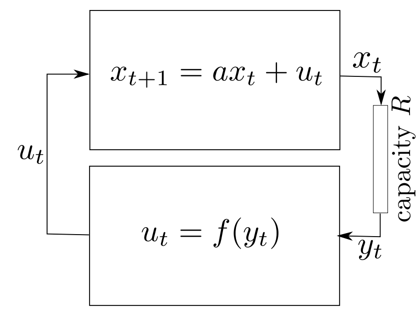

Fig. 1. A channel of capacity R restricts the communication between the plant and the controller. We look for the controller that will minimise a quadratic cost. The logarithmic strategy, which discretises x_t with a logarithmic quantiser, is shown to be suboptimal in the low capacity limit.

The logarithmic strategy

Let us describe the logarithmic strategy [4]. For a parameter $0 < \rho < 1$, the state space $[-1, 1]$ is partitionned into countably many intervals $I_{\pm i} = [\pm\rho^i, \pm\rho^{i+1}[$. At every step, the state x_t is measured and found to be in some interval $I_{\pm i}$. The index $\pm i$ is then transmitted through the communication channel. The controller, receiving $\pm i$, applies the input $u_t = -aq(x_t)$, where the quantisation function $q(x_t)$ sends x_t to the centre of the interval $I_{\pm i}$ to which it belongs.

For simplicity of analysis, we assume that ρ is chosen such that $|a|(1-\rho)$ is exactly equal to ρ^k, for some integer $k \geq 1$. We assumed that x_0 is contained in $[-1, 1]$, and uniformly distributed in this interval. This implies that the support of the distribution of x_1 is exactly $[-\rho^k, \rho^k]$, the support of x_2 is exactly $[-\rho^{2k}, \rho^{2k}]$, and so on. The distribution of $x_2, x_3, \ldots$, however, are no longer uniform on their supports. The number of bits required to encode an interval index is, according to information theory, $\sum_{\pm i} \mathrm{Prob}(x_t \in I_{\pm i}) \log \mathrm{Prob}(x_t \in I_{\pm i})$. Although the distribution of x_t is different for every t, one may compute that the number of bits required is always $R = \log 2 + \frac{H(\rho)}{1-\rho}$, where $H(\rho) = -\rho \log \rho - (1-\rho) \log(1-\rho)$. However, at every stage a different encoding strategy must be used.

Note that the optimal encoding at any time uses variable-length codewords: the number of bits required to encode an interval with small probability is higher. This may be a problem if the time to receive and decode a long codeword is higher, for the unstable dynamics may then be abandoned to itself. The impact of delays in decoding for control problems is investigated in [7]. A variable length code is however acceptable if the channel has actually a much larger capacity,

ISBN 978-963-311-370-7

but is shared with other applications. In this case, the coder may send more than R bits par channel use, as long as on average it respects its quota of R bits.

Elia and Mitter [4] proved that the logarithmic quantiser is the coarsest (in a certain sense) that ensures $E(x_t)$ to be a decreasing function of time. We complement this result by showing that with the linear-quadratic criterion given above, the logarithmic strategy is not always optimal, and we exhibit a better strategy.

A lower bound

We first state a lower bound that every control strategy must satisfy. Although in this paper we are only interested in memoryless, time-invariant controllers, it is possible to prove that for any controller, even time-varying or with memory, that stabilises the system $x_{t+1} = ax_t + u_t$, the cost is bounded from below as follows:

$$\mathbb{E} \sum_{t \geq 0} E(x_t) \geq 2/(\pi e) \frac{1}{1 - 2^{-2(R - \log|a|)}} \tag{1}$$

We therefore must have $R > \log|a|$, as already proved in various more general contexts [8], [3], [6], [9]. Furthermore, we interpret the extra rate $R - \log|a|$ has the rate necessary to reduce the cost. This quantifies the trade-off between rate and cost. Adding more levels to the quantiser will help reduce the cost, i.e., make the state converge more quickly to zero, but will increase the required capacity R.

Actual and ideal cost for the logarithmic strategy

The actual cost of the logarithmic strategy is $\frac{\mathbb{E}E(x_0)}{1-a^2(1-\rho)^3/4(1-\rho^3)}$. As we know that the required rate is $R = \log 2 + \frac{H(\rho)}{1-\rho}$, the lower bound above implies that the ideal strategy would achieve $\mathbb{E}\sum_{t\geq 0} E(x_t) \geq \frac{\mathbb{E}E(x_0)}{1-a^2 2^{-2(1+\frac{H(\rho)}{1-\rho})}}$. By 'ideal strategy', we mean the controller that achieves the general bound (1) with equality, assuming it exists. In Fig. 2, we plot the ratio actual cost/ ideal cost as a function of ρ, for $a = 5$.

We observe a treshold value $\rho = 0.477...$ under which the logarithmic strategy fails to stabilise the system. Just after the treshold, the ratio is unbounded, which means that the logarithmic strategy performs arbitrarily bad as compared to the ideal strategy for low rates R (i.e., low ρ). In the limit $\rho \to 1$, the rate R becomes infinite, which means that the communication constraint tends to disappear. Without the communication constraint, the optimal strategy is trivially the deadbeat, with cost $\mathbb{E}E(x_0)$. With higher and higher rates, the logarithmic strategy becomes faster and faster, with cost converving to $\mathbb{E}E(x_0)$. The non optimality for low rates means that there exists another strategy which, requesting the same rate, will perform with a lower cost. Of course, this is deduced with the assumption the the lower bound (1) shows no conservativeness. This is essentially the case as discussed below.

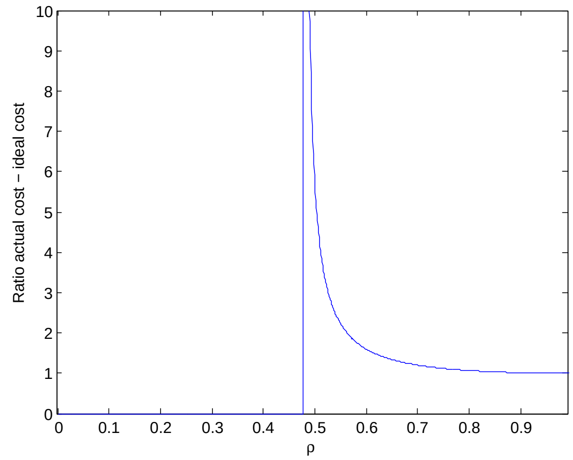

Fig. 2. For $a = 5$, we plot the ratio between the actual cost $\mathbb{E}\sum_{t\geq 0} E(x_t)$ of the logarithmic strategy and an ideal strategy that would achieve the best possible cost with the same rate $R = 1 + \frac{H(\rho)}{1-\rho}$ according to the lower bound (1). The cost is finite only for $\rho > 0.477\ldots$ For $\rho \to 1$ (high rate limit), the ratio trivially converges to one as both actual and ideal converge to the cost of the deadbeat strategy $\mathbb{E}E(x_0)$.

The optimal strategy

We show how to fix the logarithmic strategy for low capacities, using a variant of the so called nested uniform quantiser strategy [5]. It consists in subdividing every interval $[\pm\rho^i, \pm\rho^{i+1}[$ into N intervals of equal size. The controller assigns the input $u_t = -aq(x_t)$, where the quantisation function $q(x_t)$ is the midpoint of the interval to which x_t belongs. We can prove that for any system, by an appropriate choice of N and ρ, we reach a cost close to the ideal cost up to a constant, small factor. The details will be published elsewhere.

Conclusions

We have shown how to study the trade-off cost/rate in a problem of LQ control with communication constraints. We have illustrated on the simple example of the logarithmic quantiser on a scalar linear system. In a further paper we will provide more examples, as well as a more detailed study of the logarithmic quantiser, including proofs and computations.

Acknowledgements

I thank the reviewers for their useful comments.

REFERENCES

[1] R. Brockett and D. Liberzon. Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control*, AC-45:1279–1289, 2000.

[2] D. Delchamps. Stabilizing a linear system with quantized state feedback. *IEEE Transactions on Automatic Control*, AC-35:916–924, 1990.

[3] J.-Ch. Delvenne. An optimal quantized feedback strategy for scalar linear systems. *IEEE Transactions on Automatic Control*, 51(2):298–303, 2006.

[4] N. Elia and S.K. Mitter. Stabilization of linear systems with limited information. *IEEE Transactions on Automatic Control*, AC-46:1384–1400, 2001.

[5] F. Fagnani and S. Zampieri. Stability analysis and synthesis for scalar linear systems with a quantized feedback. *IEEE Transactions on automatic control*, 48(9):1569–1584, 2003.

[6] G. N. Nair, R. J. Evans, I. M. Y. Mareels, and W. Moran. Topological feedback entropy and nonlinear stabilization. *IEEE Transactions on Automatic Control (special issue on Networked Control Systems)*, 49(9):1585–1597, 2004.

[7] A. Sahai and S. Mitter. The necessity and sufficiency of anytime capacity for control over a noisy communication link: Part i. *IT Transactions*, pages 3369–3395, 2006.

[8] S. Tatikonda and S. K. Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004.

[9] H. Touchette and S. Lloyd. Information-theoretic approach to the study of control systems. *Physica A*, 331:140–172, 2004.