Communications in Information and Systems
Vol. 9, No. 1, pp. 41-58, 2009

002

NONLINEAR CONTROL WITH LIMITED INFORMATION*

DANIEL LIBERZON†

Abstract. This paper discusses several recent results by the author and collaborators, which are united by the common goal of making nonlinear control theory more robust to imperfect information. These results are also united by common technical tools, centering around input-to-state stability (ISS), small-gain theorems, Lyapunov functions, and hybrid systems. The goal of this paper is to present an overview of these results which highlights their unifying features and which is more accessible to a general audience than the original technical articles.

1. Introduction. When feedback control theory is taught in undergraduate engineering courses, one draws feedback loops featuring a plant and a controller, assuming perfect instantaneous signal transmission between the two. However, a real feedback loop may contain, for example, a camera that collects measurements, a computer that stores these measurements and calculates a control signal, and a wired or wireless communication medium through which the control signal is sent to the actuators. In such a system, in addition to the plant and controller dynamics, information-processing devices—i.e., devices that collect, encode, transmit, and decode information—play an important role. Moreover, a feedback loop of the above kind may be just a part of a larger interconnected system; in this work, however, we concentrate on a single loop because fundamental problems remain unresolved even for that case.

Besides being an issue of obvious theoretical interest, information flow in a feedback loop is an important consideration in many application-related scenarios:

Coarse sensing The information provided by the sensors may be very limited, in order to keep the sensors cheap or due to the presence of hard-to-reach areas or physical limitations of the sensors;

Limited communication capacity Even though in modern applications a lot of communication bandwidth is usually available, there are also multiple resources competing for this bandwidth (examples include networked control systems where many control loops share a network cable or wireless medium, or microsystems with many sensors and actuators on a small chip);

Cautious information transmission Even if one *can* transmit a lot of information, one may not *want* to do it because of security considerations or when embedded sensors have a short battery life;

Event-driven actuators A seemingly different scenario, but one that is similar and in some sense dual to the first one in this list, is when one utilizes limited

*Dedicated to Roger Brockett on the occasion of his 70th birthday.

†Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, E-mail: liberzon@uiuc.edu. This work was supported by the NSF ECCS-0701676 grant.

actuation (think, for example, of a manual transmission on a car).

The field of control with limited information—of which Roger Brockett is one of the pioneers—has been a very active area of research over the past couple of decades. Some representative references include [1, 2, 3, 4, 5, 6, 7, 8, 9]. This list is by no means complete, and does not reflect the explosion of activity that took place in the last few years. In previous work, deterministic as well as stochastic models have been considered both for system dynamics and for communication channels. Systematic efforts have been directed at merging tools from information theory with those of systems and control theory. On the other hand, with the exception of a handful of recent works (see [10], [11], and the work by the author to be further discussed below), the results have been limited to *linear* plants and controllers.

The goals of our recent work have been to treat *nonlinear* dynamics and to build a *unified framework* for handling several issues arising in control with limited information, namely:

- Quantization;
- Time delays;
- Disturbances.

Our approach towards achieving this goal can be summarized by the following steps:

1. Model the above effects (quantization, delays, disturbances) via *deterministic* additive error signals, which we presently denote by e to facilitate the discussion;
2. Design a nominal control law ignoring these errors, i.e., a control law acting on perfect information, which we for simplicity assume to be a static state feedback $u = k(x)$;
3. ("certainty equivalence") Apply the above control law to the imperfect/corrupted signals, resulting in $u = k(x + e)$, and combine it with an estimation procedure aimed at reducing e to 0 with time.

As natural as it sounds, this approach is too naive and does not work in general as described, especially when the plant is nonlinear. It is well-known—and easy to show by examples—that even if the error e converges to 0 and the controller stabilizes the plant when $e \equiv 0$, in the presence of a nonzero error the state can grow very large, and even blow up in finite time, before the error has a chance to converge. This means that in addition to being stabilizing in the nominal situation where the error is zero, the controller must possess a suitable form of robustness with respect to this error.

Thus at the heart of our approach will be the task of characterizing such robustness of the controller to errors. This characterization will revolve around the concept of *input-to-state stability (ISS)*, introduced by Eduardo Sontag about 20 years ago [12] and by now well-established in nonlinear systems and control theory (see [13] for a recent exposition). Supporting tools for ISS which we will use include *Lyapunov* *functions* and *small-gain theorems*. Since our plant and controller dynamics evolve in continuous time but information typically arrives and gets updated at discrete instants of time, we will see that the theory of *hybrid systems* [14] is also relevant.

The presentation style of this paper is somewhat informal; precise technical assumptions, proofs, and other details can be found in the articles cited.

2. Quantizer. By a *quantizer* we mean a map from a Euclidean space $\mathbb{R}^k$ to a finite subset of $\mathbb{R}^k$. The variable being quantized, $z \in \mathbb{R}^k$, can represent any signal in a feedback loop: a state, a measured output, or a control input. The quantizer maps this continuous signal into a piecewise constant one taking finitely many values.

We speak of quantizers in very general terms here. The quantizer mapping can include an implementation of an encoder, a communication channel, and a decoder, but such a lower-level description will not be used here. Geometrically, a quantizer can be represented as a covering of $\mathbb{R}^k$ by a finite number of regions, with a selected quantization point in each region. However, the exact shape or other geometric properties of these regions will not be important to us.

We will refer to the quantity $e := q(z) - z$ as the *quantization error*. Since the domain of q is unbounded while its set of values is finite, the quantization error cannot be globally bounded. Rather, it can only be bounded on a bounded subset of $\mathbb{R}^k$. Accordingly, we assume that there exist real numbers $M > \Delta > 0$ such that

$$|z| \leq M \quad \Rightarrow \quad |q(z) - z| \leq \Delta. \tag{1}$$

We will refer to M and Δ as the *quantization range* and the *quantization error bound*, respectively, and assume that these values are known to us. Outside the ball of radius M around the origin the quantizer saturates, i.e., the quantization error bound is not valid and we cannot "trust" the quantized measurements. Thus we need to be able to tell from a quantized measurement that z is outside this ball. To this end, it is convenient to assume that

$$|z| > M \quad \Rightarrow \quad |q(z)| > M - \Delta.$$

The above two properties represent all we need to know about the quantizer map.

3. Quantization and ISS. In this section we will introduce input-to-state stability and its role in the context of the quantized feedback stabilization problem. Suppose that we are given a general nonlinear system

$$\dot{x} = f(x, u), \qquad x \in \mathbb{R}^n, \; u \in \mathbb{R}^m \tag{2}$$

and a feedback law $u = k(x)$ that renders the closed-loop system globally asymptotically stable (GAS). Now suppose that the state x is quantized and the actual controller applied to the system is

$$u = k(q(x)) = k(x + e)$$

where $e = q(x) - x$ is the quantization error. The resulting closed-loop system will in general no longer be GAS. First, quantizer saturation will prevent convergence from initial conditions with large magnitude (larger than M, the quantization range). Second, asymptotic convergence requires arbitrarily high precision close to the origin, while the quantizer with finitely many values can only provide accuracy up to Δ, the quantization error bound. Thus, what we may expect in place of GAS is convergence from a larger set (whose size is determined by M) to a smaller set (whose size is determined by Δ). Our task now is to obtain a result establishing such a property.

To this end, assume that there exists a positive definite and radially unbounded $\mathcal{C}^1$ function $V : \mathbb{R}^n \to \mathbb{R}$ that satisfies the following:

$$|x| \geq \rho(|e|) \quad \Rightarrow \quad \frac{\partial V}{\partial x} f(x, k(x+e)) \leq -\alpha(|x|) \tag{3}$$

where ρ and α are some functions of class[1] $\mathcal{K}_\infty$ (actually, it is enough for α to be just positive definite).

We will discuss this assumption in more detail, but for the moment let us proceed. The property (3) allows us to conclude almost immediately that every solution that starts in the largest sub-level set of V contained in the M-ball around 0 converges to the smallest sub-level set of V containing the $\rho(\Delta)$ ball around 0. Indeed, in the annulus between these two level sets V must decrease, because the quantization error e does not exceed Δ by virtue of (1) applied with $z = x$.

To write down specific expressions for the above level sets, it is useful to know that since V is positive definite and radially unbounded, there exist functions α_1, α_2 of class $\mathcal{K}_\infty$ for which

$$\alpha_1(|x|) \leq V(x) \leq \alpha_2(|x|). \tag{4}$$

The precise result is then as follows.

THEOREM 1 ([15]).

Assume that we have

$$\alpha_1(M) > \alpha_2 \circ \rho(\Delta). \tag{5}$$

Then the sets

$$\mathcal{R}_1 := \{x : V(x) \leq \alpha_1(M)\}$$

and

$$\mathcal{R}_2 := \{x : V(x) \leq \alpha_2 \circ \rho(\Delta)\}$$

[1] A function $\alpha : [0,\infty) \to [0,\infty)$ is said to be of *class* $\mathcal{K}$ if it is continuous, strictly increasing, and $\alpha(0) = 0$. If α is also unbounded, then it is said to be of *class* $\mathcal{K}_\infty$. A function $\beta : [0,\infty) \times [0,\infty) \to [0,\infty)$ is said to be of *class* $\mathcal{KL}$ if $\beta(\cdot, t)$ is of class $\mathcal{K}$ for each fixed $t \geq 0$ and $\beta(r,t)$ is decreasing to zero as $t \to \infty$ for each fixed $r \geq 0$.

are invariant regions for the quantized closed-loop system. Moreover, all solutions that start in the set $\mathcal{R}_1$ enter the smaller set $\mathcal{R}_2$ in finite time.

The role of the condition (5) is simply to ensure that $\mathcal{R}_2$ is indeed contained in $\mathcal{R}_1$ and thus the statement is not vacuous. Convergence from $\mathcal{R}_1$ to $\mathcal{R}_2$ is what we have in place of GAS because of quantization.

Now, one may argue that we simply posed an unmotivated (and possibly very strong and hard to check) assumption in order to get the result we want. It turns out, however, that the existence of a Lyapunov function V satisfying (3) is an equivalent characterization of *input-to-state stability (ISS)* of the closed-loop system with respect to e. Defined in the time domain, the ISS property states that

$$|x(t)| \leq \beta(|x(0)|, t) + \gamma(\|e\|_{[0,t]}) \qquad \forall\, t \geq 0 \tag{6}$$

where β is a function of class $\mathcal{KL}$, γ is a function of class $\mathcal{K}_\infty$, and $\|e\|_{[0,t]} := \sup_{0\leq s\leq t} |e(s)|$. As a consequence, if e stays bounded or converges to 0 then so does x. We are not claiming that this assumption is easy to check or is not strong, but at least there are established tools at one's disposal for studying it (more on this below). In our setting, the ISS concept provides a natural way to formulate the type of robustness required from the controller.

EXAMPLE 1. *For a stabilizable linear plant*

$$\dot{x} = Ax + Bu$$

we can find a stabilizing state feedback gain K and a quadratic Lyapunov function $V(x) = x^T Px$ solving the Lyapunov equation

$$(A + BK)^T P + P(A + BK) = -I.$$

Then the closed-loop system

$$\dot{x} = (A + BK)x + BKe$$

is automatically ISS with respect to e because $A + BK$ is Hurwitz. An elementary calculation shows that for ρ we can use the function

$$\rho(\Delta) := 2\|PBK\|\Delta.$$

The ultimate bound on the solutions obtained from Theorem 1 directly depends on the quantization error bound Δ. Thus a closely related question is how to design a quantizer that minimizes Δ over a given region (i.e., for a given value of M) and for a given number of quantization values. This problem is treated in [16].

As we said, quantization can affect the control input instead of (or together with) the plant state. The treatment of input quantization in the ISS framework is quite similar; we will illustrate this when we study the case of time delays.

4. Dynamic quantization. So far we have treated the quantizer as static, i.e., given a priori and impossible to change in real time. However, there are many situations where the quantizer is a part of the design (although subject to certain limitations) and can be modified as the system evolves. One useful example to keep in mind is a digital camera with zooming capability: we can zoom in and out while of course keeping the number of pixels fixed.

We can formalize the above idea by introducing quantized measurements of the form

$$q_\mu(z) := \mu q\Big(\frac{z}{\mu}\Big)$$

where $\mu > 0$. The quantization range for this quantizer is $M\mu$ and the quantization error bound is $\Delta\mu$. We can think of μ as the "zoom" variable: increasing μ corresponds to zooming out and essentially obtaining a new quantizer with larger quantization range and quantization error bound, whereas decreasing μ corresponds to zooming in and obtaining a quantizer with a smaller quantization range but also a smaller quantization error bound. In the camera example, μ corresponds to the inverse of the focal length.

It is often more realistic in practice, and easier for analysis purposes, to vary the zoom variable μ in a piecewise constant fashion rather than continuously. Then the closed-loop system can be viewed as a *hybrid system*, with μ being its discrete state.

With the help of dynamic quantization, we can improve the result of Theorem 1 and even possibly recover GAS. First, if the quantizer is initially saturating, we zoom out fast enough to overcome instability of the open-loop system and eventually capture the state within the quantizer range. Then, once convergence to the attracting set has occurred, we zoom in on this set in such a way that for the new value of μ this set will be inside the larger invariant set $\mathcal{R}_1$, and the result can be applied again. (Note that in the presence of μ, both M in the definition of $\mathcal{R}_1$ and Δ in the definition of $\mathcal{R}_2$ should be multiplied by μ.) Iterating in this way, we can achieve GAS if certain conditions hold on the nonlinear functions involved, although in practice there are limitations on how small or how large the zoom variable can be.

To determine when a zoom-in should take place, we can use either the quantized measurements or an upper bound on the convergence time. For detailed descriptions of such dynamic quantization schemes, see [4, 15]. It was later pointed out to the author by Dragan Nešić that an alternative approach to designing the zooming strategy, and to analyzing the resulting system, can be based on an ISS nonlinear small-gain theorem. Essentially, what one does is make sure that the x-dynamics are ISS with respect to μ, that the (discrete) μ-dynamics are ISS with respect to x in a suitable discrete sense, and that the composition of the two ISS gain functions (i.e., the functions playing the role of γ in (6)) is smaller than the identity function. This approach is quite fruitful as it lends itself to generalizations and points to interesting general questions related to small-gain analysis of hybrid systems; see [17, 18] for further information on this topic.

5. Adding time delays. We now consider the feedback stabilization problem in the presence of both quantization and time delays. There are many results in the literature on time-delay systems giving conditions under which asymptotic stability is preserved under sufficiently small delays. Our objectives here are different: since quantization alone destroys asymptotic stability, we are instead looking for a result similar to Theorem 1, but hopefully for delays not necessarily very small. Such a result is made possible by the work of Andy Teel [19], whose approach is based on the ISS nonlinear small-gain theorem.

In order to illustrate a variety of cases, here we consider a scenario where the *control input* (rather than the plant state as before) is quantized, and the feedback law acts on delayed measurements of the state. (Other possible architectures can be analyzed in very similar ways.) Mathematically, the closed-loop system is represented as

$$\dot{x}(t) = f(x(t), q(k(x(t-\tau))))$$

or, equivalently, as

$$\dot{x} = f(x, k(x) + \theta + e)$$

where

$$\theta(t) := k(x(t-\tau)) - k(x(t))$$

and

$$e(t) := q(k(x(t-\tau))) - k(x(t-\tau)).$$

Here e is again the quantization error, while θ is a new error term reflecting the presence of the time delay τ. We can rewrite θ as

$$\theta(t) = -\int_{t-\tau}^{t} k'(x(s)) f(x(s), k(x(s-\tau)) + e(s)) ds \tag{7}$$

(assuming that the feedback law is differentiable). Noting the range of the time arguments inside this integral, and applying the integral mean value theorem, we can show the existence of some functions γ_1, γ_2 of class $\mathcal{K}_\infty$ for which

$$|\theta(t)| \le \max\left\{\tau\gamma_1(\|x\|_{[t-2\tau,t]}), \tau\gamma_2(\|e\|_{[t-2\tau,t]})\right\}. \tag{8}$$

We now assume as in Section 3 that the feedback law guarantees the ISS property, except that, due to the system architecture considered here, the additive input enters in the control signal rather than in the argument of the control law. In other words, we are now asking for ISS with respect to actuator errors and not measurement errors. The assumption we make is the existence of a $\mathcal{C}^1$ function V satisfying

$$|x| \geq \rho(|v|) \quad \Rightarrow \quad \frac{\partial V}{\partial x} f(x, k(x) + v) \leq -\alpha(|x|)$$

as well as (4), where $\rho, \alpha, \alpha_1, \alpha_2$ are of class $\mathcal{K}_\infty$. We know that, for $v = \theta + e$, this corresponds to an estimate of the form

$$|x(t)| \leq \beta(|x(0)|, t) + \gamma_\theta \left(\|\theta\|_{[0,t]}\right) + \gamma_e \left(\|e\|_{[0,t]}\right)$$

where β is of class $\mathcal{KL}$ and γ_θ, γ_e are of class $\mathcal{K}_\infty$. With the help of (8) and some calculations, this can be brought to

$$|x(t)| \leq \beta(|x(0)|, t) + \alpha_1^{-1} \circ \alpha_2 \circ \rho(2\tau\gamma_1(\|x\|_{[t-2\tau,t]})) + \tilde{\gamma}_e \left(\|e\|_{[t-2\tau,t]}\right)$$

where $\tilde{\gamma}_e(r) := \alpha_1^{-1} \circ \alpha_2 \circ \rho(2 \max\{\tau\gamma_2(r), r\})$.

This is where the small-gain argument comes in: if the function $\alpha_1^{-1} \circ \alpha_2 \circ \rho(2\tau\gamma_1(\cdot))$ is smaller than identity, then the second term on the right-hand side can be eliminated and we recover ISS with respect to e only. After that, we can proceed in the same way as in Section 3. It is important to observe two things. First, the above function depends explicitly on the delay, thus the small-gain condition imposes an upper bound on the delay. Second, a nonlinear function may in general have infinite slope both near 0 and at infinity; thus it can be fit below the identity function only on an interval of the form $[\varepsilon, \Lambda] \subset (0, \infty)$. This implies that we will obtain convergence from a set whose size is determined by Λ (which is to be chosen based on the quantization range M) to a smaller set whose size is determined by ε (which is to be selected based on the quantization error bound Δ).

Take κ to be some class $\mathcal{K}_\infty$ function with the property that

$$\kappa(r) \geq \max_{|x| \leq r} |k(x)| \qquad \forall\, r \geq 0. \tag{9}$$

The precise result is this.

Theorem 2 ([20]). *Assume that*[2]

$$\|x\|_{[-\tau,\tau]} \leq E_0$$

for some known $E_0 > 0$. *Assume that for some* $\Lambda > \varepsilon > 0$ *we have*

$$\max\left\{\alpha_1^{-1} \circ \alpha_2(E_0), \varepsilon, \tilde{\gamma}_e(\Delta\mu)\right\} < \Lambda < \kappa^{-1}(M\mu)$$

[2] In the original statement of this result in [20], the right endpoint of the interval in this inequality was the initial time t_0 (here taken to be 0) while it should be $t_0 + \tau$. If we only have this assumption the way it was written in [20], then we have to use the knowledge of system dynamics to generate a bound on the interval $[t_0, t_0 + \tau]$. The reason is that the formula (7) is valid only when the system already satisfies the differential equation on the whole interval of integration, hence it is only valid starting from $t = t_0 + \tau$. We thank Emilia Fridman for pointing out this issue.

and

$$\alpha_1^{-1} \circ \alpha_2 \circ \rho(2\tau\gamma_1(r)) < r \qquad \forall\, r \in (\varepsilon, \Lambda]. \tag{10}$$

Then the solution of the closed-loop system satisfies the bound

$$\|x\|_{[-\tau,\infty)} \le \max\left\{\alpha_1^{-1} \circ \alpha_2(E_0), \varepsilon, \tilde{\gamma}_e(\Delta\mu)\right\} \tag{11}$$

and the ultimate bound

$$\|x\|_{[-\tau+T,\infty)} \le \max\left\{\varepsilon, \tilde{\gamma}_e(\Delta\mu)\right\} \tag{12}$$

for some $T > 0$.

As in Section 4, a dynamic quantization scheme can be used to enlarge the region of attraction and improve convergence; see [20] for details. The above result demonstrates that within the ISS framework, quantization and time delays can be naturally treated in a unified manner.

6. External disturbances. In order to have a control scheme that can handle realistic situations of the type we mentioned in the introduction, it is also important to be able to attenuate external disturbances entering the plant dynamics. In the recent paper [21] we considered linear plants with quantized state feedback which are affected by a disturbance. We assumed that the quantization was dynamic, as in Section 4. Unlike the earlier works treating quantization and disturbances [22, 7, 23], the disturbance was taken to be *completely unknown*. In other words, neither the values of the disturbance nor any a priori knowledge of its bounds are available to the control designer.

To see the difficulties involved in applying the approach described in Sections 3 and 4 in the presence of a disturbance, suppose that for a given value of μ the state has entered the attracting set $\mathcal{R}_2$ and so we are supposed to zoom in to facilitate further convergence. However, at this point a high value of the disturbance might "hit" the system and force the state to exit the quantization range, invalidating subsequent analysis. To counteract this effect, we need to return to the zooming-out stage and wait until the state is "re-captured". Then we can restart the zooming-in stage, but of course there is no guarantee that the same phenomenon will not occur again.

From the preceding discussion we infer two things. First, we must switch repeatedly between zooming in and zooming out, hence the dynamic quantization strategy will necessarily be more complicated than the one sketched in Section 4. We will not give details of the scheme here, but they can be found in [21]. Second, unless the disturbance converges to 0 we cannot hope to achieve asymptotic stability. This latter point is rather obvious, but what is less obvious is which property we should take as a characterization of desirable disturbance attenuation.

A very natural candidate for formulating the desired property is again furnished by the ISS concept. In particular, ISS with respect to the disturbance would guarantee that the state of the system is bounded if the disturbance is bounded, and converges to 0 if the disturbance converges to 0. The main result of [21] is that we can indeed achieve such a property. While we cannot go into any analysis details here, we mention that the ISS small-gain theorem—which, as we explained, plays a central role in establishing Theorem 2 and also supports the dynamic quantization design in the disturbance-free case—is relevant in this case as well (actually, the special case of a cascade connection is used).

One interesting technical remark is that the ISS gain functions that we obtain grow faster than any linear function both near 0 and at infinity. It turns out that this is not an artifact of our control design, but rather a consequence of a result by Nuno Martins who showed, using techniques from information theory, that it is impossible to achieve ISS with linear gain for any linear system with finite data rate feedback [24]. Thus, in the presence of state quantization it is indeed necessary to formulate the disturbance attenuation problem in terms of nonlinear ISS gains, in other words, quantization makes the problem inherently nonlinear despite the fact that the given open-loop system is linear. This further supports the use of ISS, which now describes the control objective rather than a control design assumption as in the previous sections. A related result is that the same scheme as in [21] also achieves nonlinear-gain l_2 stabilization [25].

It is possible to extend these results to nonlinear plants, under the assumption that the controller renders the closed-loop system ISS with respect to both the quantization error and the disturbance (in the linear case this is true for every stabilizing feedback). This extension is conceptually straightforward in view of what we already discussed, although the control algorithm becomes less constructive.

7. Minimal data rate. One of the central questions in the recent research on control with limited information has been: how much (or how little) information is enough to stabilize the plant? See, e.g., [5, 26, 7, 22] for somewhat different but closely related answers to this question. The amount of information is usually defined by data rate or, more precisely, by a suitable notion of channel capacity.

The results that we have presented so far do not directly address this question, because they do not focus on the minimal data rate. Since we assumed the quantized measurements to be continuously available, the data rate required to implement the above control strategies is actually infinite. This, however, can be easily corrected by introducing time sampling (this was done in [21] where a sampled-data implementation was developed and compared with a continuous-time one). Still, the above results do not look for the smallest possible number of quantization values that will do the job, and instead simply assume that this number is large enough.

If the number of values that the quantizer takes is small, one needs to be more active in extracting information from the system. The idea of dynamically scaling the quantizer, which we described in Section 4, represents one example of how this can be done. In addition to scaling, one can also move the quantizer around in the state space, by making it follow the state estimates which are generated on-line. With this modification, the number of necessary quantization values can be made much smaller. Combining this idea with the present approach based on ISS, we obtain results for nonlinear systems, whereas the earlier references only dealt with linear systems. However, unlike in the linear case, the required condition on the number of quantization regions—see (16) below—is possibly conservative and is not guaranteed to be both necessary and sufficient for stabilization (although we do recover known tight bounds by properly specializing our results to the linear case).

We now briefly describe the basic set-up; see [27] for more details. Consider again the general nonlinear control system (2), where f is locally Lipschitz. We assume again that a nominal feedback law $u = k(x)$ is given which provides ISS with respect to measurement errors in the sense of (6). We assume that measurements are to be received by the controller at discrete times $0, \tau, 2\tau, \dots$, where $\tau > 0$ is a fixed sampling period. At each of these sampling times, the measurement received by the controller must be a number in the set $\{0, 1, \dots, N\}$, where N is a fixed positive integer. Thus the data available to the controller consists of a stream of integers $q_0(x(0)), q_1(x(\tau)), q_2(x(2\tau)), \dots$ where the q_k's are appropriate (dynamically updated) state quantizer functions with N values. It is convenient to assume that the number $N^{\frac{1}{n}}$ is an odd integer, otherwise we can work with a smaller N for which this is true.

Assume that the initial condition satisfies $\|x(0)\|_\infty \le E_0$ for some known $E_0 > 0$, where $\|x\|_\infty := \max\{|x_i| : 1 \le i \le n\}$ is the infinity norm on $\mathbb{R}^n$. Let L be the Lipschitz constant for the function f on the region

$$\{(x, u) : \|x\|_\infty \le D,\ \|u\|_\infty \le \kappa(D)\} \tag{13}$$

where

$$D := \beta(E_0, 0) + \gamma\big(N^{\frac{1}{n}} E_0\big) + N^{\frac{1}{n}} E_0 \tag{14}$$

and κ is a function satisfying (9). Define

$$\Lambda := e^{L\tau} \ge 1. \tag{15}$$

With this notation in place, the result is as follows.

Theorem 3 ([27]). *If*

$$\Lambda < N^{\frac{1}{n}} \tag{16}$$

then there exists a quantized feedback control strategy that makes $x(t)$ converge to 0.

In view of the definition of Λ via the formula (15), the inequality (16) characterizes the trade-off between the amount of information provided by the quantizer at each sampling time and the required sampling frequency. This relationship depends explicitly on the Lipschitz constant L, which can be interpreted as a measure of expansiveness of the open-loop system. In the case of linear systems, in place of L one has the norm of the system's A-matrix.

The above result assumes the knowledge of an upper bound on the initial state. Without such a bound, it is still possible to stabilize the system by applying a zooming-out strategy and generating an upper bound on the state first; see [27]. However, it should be noted that the condition on the data rate depends on the initial state. In this sense, the stabilization result is "semi-global."[3]

8. Discussion and further research directions.

8.1. The ISS assumption. The results expressed by Theorems 1, 2 and 3 relied on the assumption that the nominal feedback law provides ISS with respect to the quantization error. In the case of state quantization, this error appears as the measurement error, while in the case of input quantization it appears as the actuator error. In both situations the ISS requirement is restrictive in general. (For linear systems and linear stabilizing feedback laws, such robustness with respect to additive errors is of course automatic.) Regarding ISS with respect to actuator errors, it was shown in [12] that for globally asymptotically stabilizable systems affine in controls, such a feedback always exists. The assumption of ISS with respect to measurement errors is more difficult to satisfy, as was demonstrated by way of counterexamples in [28] and [29]. This problem has received a lot of attention in the literature, and positive results for some classes of systems have been obtained [30, 31, 32, 29, 33, 34, 35]. Further progress on this problem is obviously important for making our results applicable to larger classes of systems.

The ISS requirement appears to be essentially necessary for obtaining results of the kind we presented. ISS is of course stronger than GAS under zero inputs (0-GAS). However, 0-GAS does imply some weaker forms of ISS. For example, it implies the ISS bound (6) for initial conditions in a given compact region and under sufficiently small inputs. This was the basis for replacing the ISS assumption by just 0-GAS in [36] and in [20]. In particular, in [36] the initial state is assumed to belong to a known region and the data rate is increased to make the quantization error small enough before the state leaves this region. (The stabilization result in that paper is semi-global in the same sense as Theorem 3.) In [20], an alternative result is presented in which ISS is not required but the small-gain condition restricts the quantization error bound

[3] In the original statement of this result in [27], we incorrectly referred to the resulting property of the closed-loop system as GAS. We thank Yoav Sharon for pointing out this inaccuracy.

to be sufficiently small. It is important to realize that one is not really relaxing ISS but rather (explicitly or implicitly) using the fact that for sufficiently small errors on a suitable region ISS follows automatically from 0-GAS. In other words, the ISS property still plays a crucial role.

One aspect that the ISS assumption does not take into account is that the error input with respect to which we want to have ISS is not an arbitrary signal, but is defined by the behavior of the overall closed-loop system. For instance, in the setting of Section 7 it is possible to make this error decay to 0 exponentially fast. For this reason, as noted in [37], it is actually sufficient to have the weaker integral ISS property plus a technical condition on the corresponding gain. The paper [27] also proposes a different approach, in which ISS of the closed-loop system is replaced by an ISS assumption on the impulsive system describing the evolution of the estimated state. The resulting assumption is more constructive to check, and has motivated further research on ISS of impulsive systems reported in [38].

8.2. Output quantization and ISS observers. The problem of quantized *output* feedback was considered, e.g., in [15, 26] but only for linear systems. In the recent paper [39] we described an ISS framework in which this problem can be treated for nonlinear systems (recovering earlier linear results as a special case). Consider the plant

$$\begin{aligned} \dot{x} &= f(x,u) \\ y &= h(x) \end{aligned} \tag{17}$$

where y is the measured output. As the nominal controller we now take a dynamic (observer-based) output feedback law of the form

$$\begin{aligned} \dot{z} &= g(z,u,y) \\ u &= k(z) \end{aligned} \tag{18}$$

Assuming that only quantized measurements $q(y)$ of the plant output y are available to the controller, we consider the "certainty equivalence" quantized output feedback law based on (18), namely,

$$\begin{aligned} \dot{z} &= g(z,u,q(y)) \\ u &= k(z) \end{aligned} \tag{19}$$

Writing $e = q(y) - y$ for the output quantization error, we can represent the closed-loop system that results from interconnecting (17) and (19) as

$$\begin{aligned} \dot{x} &= f(x,k(z)) \\ \dot{z} &= g(z,k(z),h(x)+e) \end{aligned} \tag{20}$$

It is this system that we need to be ISS with respect to e. If this ISS property is satisfied, then the analysis can proceed almost in the same way as in Section 3, except we would now be working in the (x, z)-space; see [39] for details.

One way to achieve ISS of (20) is to ask for the following two properties:

1. The system

$$\dot{x} = f(x, k(z)) = f\big(x, k(x + (z - x))\big)$$

satisfies

$$|x(t)| \leq \beta_1(|x(0)|, t) + \gamma_1\left(\|z - x\|_{[0,t]}\right) \tag{21}$$

for some β_1 of class $\mathcal{KL}$ and γ_1 of class $\mathcal{K}_\infty$;

2. The system (20) satisfies

$$|z(t) - x(t)| \leq \beta_2(|z(0) - x(0)|, t) + \gamma_2\left(\|e\|_{[0,t]}\right) \tag{22}$$

for some β_2 of class $\mathcal{KL}$ and γ_2 of class $\mathcal{K}_\infty$.

Property 1 states that the static state feedback $u = k(x)$ should render the x-subsystem ISS with respect to measurement errors, i.e., the system $\dot{x} = f(x, k(x+d))$ should be ISS with respect to d. In our case d is the observer error, i.e., the difference between the observer state and the plant state, but for the purposes of control design it can be viewed as a general measurement disturbance. This is the same assumption that we encountered in the state quantization case.

Property 2 means that the z-subsystem, acting as an observer for the x-subsystem, yields an observer error satisfying an ISS contraction property with respect to additive errors at the output. Even though it is easy to state and appears to be very natural, this property seems difficult to satisfy except for very special classes of systems, and has not been systematically studied in the observer design literature. A notable exception[4] is the paper [40], which leads to some promising research avenues.

8.3. Disturbances and minimal data rate. An obvious gap in the above developments is that the attenuation of completely unknown external disturbances was not achieved under a minimal data rate: the results of Section 6 do not strive for a minimal data rate, while the results of Section 7 are for the disturbance-free case. Closing this gap requires a nontrivial combination and extension of the ideas mentioned in those two sections. This direction is being pursued in the ongoing work; see [41]. The ISS small-gain theorem (and its special case for cascade connections) again plays a central role in the analysis.

[4]We thank Jung-Su Kim for calling our attention to this paper.

8.4. Modeling uncertainty. When one thinks of control with limited information, another important aspect to consider is lack of information about the dynamical model of the plant itself. While the adaptive control literature provides a large body of results on control of uncertain plants, it does not consider quantization and modeling uncertainty together. The recent paper [42] studies the problem of stabilizing uncertain linear systems with quantization. The plant uncertainty is dealt with by the supervisory adaptive control framework, which employs switching among a finite family of candidate controllers. This enables counterparts of the results presented in Sections 3 and 4 for uncertain systems.

8.5. Control with limited information and without estimation. A somewhat general question prompted by the above results, and by similar developments in the literature, is whether or not it is really necessary to incorporate a state estimation procedure into the control design. Is it perhaps possible to solve interesting control problems by acting directly on the available (possibly very coarse) state measurements, without attempting to reconstruct (even asymptotically) the precise values of the state? The recent paper [43] shows that at least in some cases the answer is "yes." That paper considers a multi-agent system in which each agent moves like a Dubins car and has an extremely limited sensor which reports only the presence of another agent within some sector of its windshield. Using a very simple quantized control law with three values, each agent tracks another agent assigned to it by maintaining that agent within this windshield sector. Lyapunov analysis is used to show that by acting autonomously in this way, the agents will achieve rendezvous if the initial assignment graph is connected. In contrast to the results surveyed here, the approach of [43] does not involve any estimation procedure aimed at reconstructing coordinate information (of course the control objective is also different there).

REFERENCES

[1] D. F. Delchamps. *Stabilizing a linear system with quantized state feedback.* IEEE Trans. Automat. Control, 35(1990), pp. 916–924.

[2] W. S. Wong and R. W. Brockett. *Systems with finite communication bandwidth constraints I: state estimation problems.* IEEE Trans. Automat. Control, 42(1997), pp. 1294–1299.

[3] W. S. Wong and R. W. Brockett. *Systems with finite communication bandwidth constraints II: stabilization with limited information feedback.* IEEE Trans. Automat. Control, 44(1999), pp. 1049–1053.

[4] R. W. Brockett and D. Liberzon. *Quantized feedback stabilization of linear systems.* IEEE Trans. Automat. Control, 45(2000), pp. 1279–1289.

[5] G. N. Nair and R. J. Evans. *Stabilization with data-rate-limited feedback: tightest attainable bounds.* Systems Control Lett., 41(2000), pp. 49–56.

[6] N. Elia and S. K. Mitter. *Stabilization of linear systems with limited information.* IEEE Trans. Automat. Control, 46(2001), pp. 1384–1400.

[7] S. Tatikonda and S. K. Mitter. *Control under communication constraints.* IEEE Trans. Automat. Control, 49(2004), pp. 1056–1068.

[8] A.S. Matveev and A.V. Savkin. *The problem of LQG optimal control via a limited capacity communication channel*. Systems Control Lett., 53(2004), pp. 51–64.

[9] J. Baillieul. *Feedback designs in information-based control*. In B. Pasik-Duncan, editor, Stochastic Theory and Control, volume 280 of Lecture Notes in Control and Information Sciences, pages 35–57. Springer, Berlin, 2002.

[10] G. N. Nair, R. J. Evans, I. M. Y. Mareels, and W. Moran. *Topological feedback entropy and nonlinear stabilization*. IEEE Trans. Automat. Control, 49(2004), pp. 1585–1597.

[11] C. De Persis. *n-bit stabilization of n-dimensional nonlinear feedforward systems*. IEEE Trans. Automat. Control, 30(2005), pp. 299–311.

[12] E. D. Sontag. *Smooth stabilization implies coprime factorization*. IEEE Trans. Automat. Control, 34(1989), pp. 435–443.

[13] E.D. Sontag. *Input to state stability: Basic concepts and results*. In P. Nistri and G. Stefani, editors, Nonlinear and Optimal Control Theory, pages 163–220. Springer-Verlag, Berlin, 2007.

[14] A. van der Schaft and H. Schumacher. *An Introduction to Hybrid Dynamical Systems*. Springer, London, 2000.

[15] D. Liberzon. *Hybrid feedback stabilization of systems with quantized signals*. Automatica, 39(2003), pp. 1543–1554.

[16] F. Bullo and D. Liberzon. *Quantized control via locational optimization*. IEEE Trans. Automat. Control, 51(2006), pp. 2–13.

[17] D. Nešić and D. Liberzon. *A small-gain approach to stability analysis of hybrid systems*. In Proc. 44th IEEE Conf. on Decision and Control, pages 5409–5414, 2005.

[18] D. Liberzon and D. Nešić. *Stability analysis of hybrid systems via small-gain theorems*. In J. Hespanha and A. Tiwari, editors, Proc. 9th Int. Workshop on Hybrid Systems: Computation and Control, volume 3927 of *Lecture Notes in Computer Science*, pages 421–435. Springer, Berlin, 2006.

[19] A. R. Teel. *Connections between Razumikhin-type theorems and the ISS nonlinear small gain theorem*. IEEE Trans. Automat. Control, 43(1998), pp. 960–964.

[20] D. Liberzon. *Quantization, time delays, and nonlinear stabilization*. IEEE Trans. Automat. Control, 51(2006), pp. 1190–1195.

[21] D. Liberzon and D. Nešić. *Input-to-state stabilization of linear systems with quantized state measurements*. IEEE Trans. Automat. Control, 52(2007), pp. 767–781.

[22] J. P. Hespanha, A. Ortega, and L. Vasudevan. *Towards the control of linear systems with minimum bit-rate*. In Proc. 15th Int. Symp. on Mathematical Theory of Networks and Systems (MTNS), 2002.

[23] G. N. Nair and R. J. Evans. *Stabilizability of stochastic linear systems with finite feedback data rates*. SIAM J. Control Optim., 43(2004), pp. 413–436.

[24] N. C. Martins. *Finite gain l_p stabilization requires analog control*. Systems Control Lett., 55(2006), pp. 949–954.

[25] T. Kameneva and D. Nešić. *On l_2 stabilization of linear systems with quantized control*, IEEE Trans. Automat. Control, 53(2008), pp. 399–405.

[26] D. Liberzon. *On stabilization of linear systems with limited information*. IEEE Trans. Automat. Control, 48(2003), pp. 304–307.

[27] D. Liberzon and J. P. Hespanha. *Stabilization of nonlinear systems with limited information feedback*. IEEE Trans. Automat. Control, 50(2005), pp. 910–915.

[28] R. A. Freeman. *Global internal stabilizability does not imply global external stabilizability for small sensor disturbances*. IEEE Trans. Automat. Control, 40(1995), pp. 2119–2122.

[29] N. C. S. Fah. *Input-to-state stability with respect to measurement disturbances for one-dimensional systems*. ESAIM J. Control, Optimization and Calculus of Variations, 4(1999), pp. 99–122.

[30] R. A. Freeman and P. V. Kokotović. *Global robustness of nonlinear systems to state measurement disturbances*. In Proc. 32nd IEEE Conf. on Decision and Control, pages 1507–1512, 1993.

[31] R. A. Freeman and P. V. Kokotović. *Robust Nonlinear Control Design: State-Space and Lyapunov Techniques*. Birkhäuser, Boston, 1996.

[32] R. A. Freeman. *Time-varying feedback for the global stabilization of nonlinear systems with measurement disturbances*. In Proc. 4th European Control Conf., 1997.

[33] Z.-P. Jiang, I. M. Y. Mareels, and D. Hill. *Robust control of uncertain nonlinear systems via measurement feedback*. IEEE Trans. Automat. Control, 44(1999), pp. 807–812.

[34] R. G. Sanfelice and A. R. Teel. *On hybrid controllers that induce input-to-state stability with respect to measurement noise*. In Proc. 44th IEEE Conf. on Decision and Control, pages 4891–4896, 2005.

[35] C. Ebenbauer, T. Raff, and F. Allgöwer. *Certainty-equivalence feedback design with polynomial-type feedbacks which guarantee ISS*. *IEEE Trans. Automat. Control*, 52:716–720, 2007.

[36] C. De Persis and A. Isidori. *Stabilizability by state feedback implies stabilizability by encoded state feedback*. Systems Control Lett., 53(2004), pp. 249–258.

[37] C. De Persis. *Nonlinear stabilizability via encoded feedback: the case of integral ISS systems*. Automatica, 42(2006), pp. 1813–1816.

[38] J. P. Hespanha, D. Liberzon, and A. R. Teel. *Lyapunov conditions for input-to-state stability for impulsive systems*. Automatica, 44(2008), pp. 2735–2744.

[39] D. Liberzon. *Observer-based quantized output feedback control of nonlinear systems*. In Proc. 17th IFAC World Congress, pages 8039–8043, 2008.

[40] H. Shim, J. H. Seo, and A. R. Teel. *Nonlinear observer design via passivation of error dynamics*. Automatica, 39(2003), pp. 885–892.

[41] Y. Sharon and D. Liberzon. *Input-to-state stabilization with minimum number of quantization regions*. In Proc. 46th IEEE Conf. on Decision and Control, pages 20–25, 2007.

[42] L. Vu and D. Liberzon. *Stabilizing uncertain systems with dynamic quantization*. In Proc. 47th IEEE Conf. on Decision and Control, pages 4681–4686, 2008.

[43] J. Yu, S. LaValle, and D. Liberzon. *Rendezvous without coordinates*. In Proc. 47th IEEE Conf. on Decision and Control, pages 1803–1808, 2008.