23rd International Symposium on Mathematical Theory of Networks and Systems
Hong Kong University of Science and Technology, Hong Kong, July 16-20, 2018

Noise-Induced Limitations to the Scalability of Distributed Integral Control

Emma Tegling and Henrik Sandberg

***Abstract*— We study performance limitations of distributed feedback control in large-scale networked dynamical systems. Specifically, we address the question of how the performance of distributed integral control is affected by measurement noise. We consider second-order consensus-like problems modeled over a toric lattice network, and study asymptotic scalings (in network size) of $\mathcal{H}_2$ performance metrics that quantify the variance of nodal state fluctuations. While previous studies have shown that distributed integral control fundamentally improves these performance scalings compared to distributed proportional feedback control, our results show that an explicit inclusion of measurement noise leads to the opposite conclusion. The noise's impact on performance is shown to decrease with an increased inter-nodal alignment of the local integral states. However, even though the controller can be tuned for acceptable performance for any given network size, performance will degrade as the network grows, limiting the *scalability* of any such controller tuning. In particular, the requirement for inter-nodal alignment increases with network size. We show that this in practice implies that large and sparse networks will require any integral control to be centralized, rather than distributed. In this case, the best-achievable performance scaling, which is shown to be that of proportional feedback control, is retrieved.**

Keywords: **Networked Control Systems, Large Scale Systems, Fundamental Limitations.**
AMS subject classification: **93A14, 93A15, 93C05.**

I. PROBLEM FORMULATION

A. System dynamics

Consider a networked dynamical system modeled over the discrete toric lattice $\mathbb{Z}_L^d$, with a total of $N = L^d$ nodes. The local dynamics are of second order, meaning that there are two states x_k and v_k, at each network site $k \in \mathbb{Z}_L^d$. These states can be thought of as, respectively, the position and velocity deviations of the k^{th} agent in a formation control problem, but may also capture, for example, phase and angular frequency in coupled oscillator networks. The system dynamics are modeled as follows:

$$\begin{bmatrix}\dot{x}\\ \dot{v}\end{bmatrix} = \begin{bmatrix}0 & I\\ F & G\end{bmatrix}\begin{bmatrix}x\\ v\end{bmatrix} + \begin{bmatrix}0\\ I\end{bmatrix}u + \begin{bmatrix}0\\ I\end{bmatrix}w, \tag{1}$$

where u is a control input and w is a disturbance. The linear feedback operators F and G define convolutions of the states x and v with the function arrays $f = \{f_k\}$ and $g = \{g_k\}$ over $\mathbb{Z}_L^d$, i.e., $(Fx)_k = \sum_{l\in\mathbb{Z}_L^d} f_{k-l}x_l$. This structure implies that the state feedback is *spatially invariant* with respect to $\mathbb{Z}_L^d$. The feedback in (1) is referred to as *static* if the control input $u = 0$, in which case the feedback is simply proportional to state deviations.

E. Tegling and H. Sandberg are with the School of Electrical Engineering and Computer Science, and the ACCESS Linnaeus Center at KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden (tegling, hsan@kth.se). This work was supported by the Swedish Research Council under Grant 2013-5523 and Grant 2016-00861

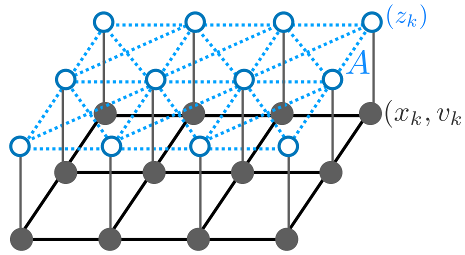

Fig. 1: Example structure of the distributed integral controller. The inter-nodal alignment of integral states z_k takes place over a communication network (dashed lines), while the state feedback interactions take place over the physical network (solid lines). We show that if the controller is subject to noise, the inter-nodal alignment through A becomes increasingly important. For large and sparse physical networks all-to-all communication or centralized integral control will be necessary.

An example of the dynamics in (1) is nearest-neighbor consensus for $d = 1$:

$$\begin{aligned}\ddot{x}_k = \dot{v}_k = f_+(x_{k+1}-x_k) + f_-(x_{k-1}-x_k) + g_+(v_{k+1}-v_k)\\ + g_-(v_{k-1}-v_k) - f_o x_k - g_o v_k + u_k + w_k,\end{aligned} \tag{2}$$

where $f_+, f_-, f_o, g_+, g_-, g_o \geq 0$ are fixed gains. We refer to terms like $(x_{k+1}-x_k)$ as *relative feedback* and to terms like $-f_o x_k$ as *absolute feedback*. Absolute feedback is well-known to be beneficial for control performance in networked dynamical systems, but the corresponding measurements are often not available, see e.g. [1], [2]. Here, we will assume that *only relative* measurements of the (generalized) position state x are available, i.e., $f_o = 0$ in (2). Absolute measurements of the (generalized) velocity are, however, available to each controller.

We remark that the analysis here is not limited to nearest-neighbor feedback, but we assume that measurements are available from a neighborhood of width $2q$, with q fixed. We further assume that feedback interactions are symmetric around each site k.

B. Distributed integral control

Distributed integral control in networked dynamical systems is motivated by a desire to eliminate stationary controller errors which arise through the standard static feedback (as this is essentially just proportional control). We consider

TABLE I: Asymptotic performance scalings for the system (1) with (i) $u = 0$ (static feedback), (ii). u as in (3) with $v^m = v$ (distributed integral control, noiseless) and (iii) u as in (3) with $v^m = v + \varepsilon\eta$ (noisy distributed integral control). The notation $\sim$ defines scalings as follows: $u(N) \sim v(N) \quad \Leftrightarrow \quad \underline{c}v(N) \leq u(N) \leq \bar{c}v(N)$, for any $N > \bar{N}$, with $\bar{N}$ fixed and the constants $\underline{c}, \bar{c} > 0$ independent of N, algorithm parameter $\beta = \max\{||f||_\infty, ||g||_\infty\}$ and relative noise intensity ε.

	Local error	**Global error**
(i) Static feedback	$V_N \sim \frac{1}{\beta}$ for any d	$V_N \sim \frac{1}{\beta}\begin{cases} N & d=1 \\ \log N & d=2 \\ 1 & d \geq 3 \end{cases}$
(ii) Distributed integral control (noiseless)	$V_N \sim \frac{1}{\beta}$ for any d	$V_N \sim \frac{1}{\beta}$ for any d
(iii) Noisy distributed integral control	$V_N \sim \frac{\varepsilon^2}{\beta}\begin{cases} N & d=1 \\ \log N & d=2 \\ 1 & d \geq 3 \end{cases}$	$V_N \sim \frac{\varepsilon^2}{\beta}\begin{cases} N^3 & d=1 \\ N & d=2 \\ N^{1/3} & d=3 \\ \log N & d=4 \\ 1 & d \geq 5 \end{cases}$

the control input u to be a *distributed integral controller* on the form:

$$u = z$$
$$\dot{z} = -c_o v^m + Az, \tag{3}$$

where v^m is the velocity measured by the controller, $c_o > 0$ is a fixed (integral) gain and A is a feedback operator subject to the same assumptions as F in (1). An example of the control law (3) is:

$$\dot{u}_k = \dot{z}_k = a_+(z_{k+1} - z_k) + a_-(z_{k-1} - z_k) - c_o v_k^m, \tag{4}$$

where $a_+, a_- > 0$ are fixed gains. This controller integrates the absolute velocity measurements, but also aligns the integral state z over the network through the consensus or *distributed averaging filter* represented by the operator A. It is useful to think of the information exchange through A as taking place over a communication network layer, separate from the physical network as in Fig. 1.

The controller (3) have been proposed in the context of frequency control in power networks [3], [4]. Its main advantage is that it can be implemented in a distributed fashion.

C. Performance evaluation

We are concerned with the performance of the system (1), and in particular, with how well the performance of the control laws scale as the network size $N \to \infty$. In line with related work [2], [5]–[9], we characterize performance through the steady state variance of nodal state fluctuations, when the system is driven by a white noise disturbance input w. This variance is measured through the squared $\mathcal{H}_2$ norm from w to a performance output y:

$$\mathbf{V}_N := \sum_{k \in \mathbb{Z}_L^d} \lim_{t \to \infty} \mathbb{E}\{y_k^*(t) y_k(t)\}. \tag{5}$$

We consider the following performance measurements:

Definition 1 (Global error):

$$y_k = x_k - \frac{1}{N} \sum_{l \in \mathbb{Z}_L^d} x_l \tag{6}$$

Definition 2 (Local error):

$$y_k = x_k - x_{k-1} \tag{7}$$

Throughout, we consider the *per-site variance*, which due to the system's spatial invariance is independent of the site k:

Definition 3 (Per-site variance):

$$V_N = \mathbb{E}\{y_k^*(t) y_k(t)\} = \frac{\mathbf{V}_N}{N}. \tag{8}$$

It is the *scaling* of V_N with the system size N as it grows asymptotically that is of interest. A fast scaling in N implies a lack of network coherence and a limited scalability of the control law.

The following results, of which (i) appeared in [2, Corollary 3.2] and (ii) follows from [9, Corollary 1] are the main motivation for this work.

Result 1 (Performance scalings): Consider the system (1) and assume that the velocity measurements are noiseless, that is, $v^m = v$. Then, Table I lists the asymptotic scaling of the per-site variance V_N with

(i) Static feedback, i.e., where the secondary control input $u = 0$, and
(ii) Distributed integral control with u as in (3).

D. Objectives

Result 1 implies that distributed integral control on the form (3) *fundamentally improves* performance in terms of global error compared to static feedback, if the velocity measurements are noiseless. In terms of local error, the variance is bounded in N for both controllers, but can be shown to decrease in absolute terms through integral control. The objective of the present work is to determine to which extent this result is robust to measurement noise in the controller.

The apparent reason for the performance improvement through integral control is namely that integration of the absolute velocity measurements emulates absolute position feedback [9]. Any noise and bias in the velocity measurements is prevented from causing destabilizing drifts in this position feedback by the distributed averaging filter A in (3). Yet, we show here that noise in the velocity measurements may still have a significant effect on performance. We also demonstrate how the performance under noisy measurements depends on the design of the distributed averaging filter A.

II. Limitations of Distributed Integral Control under Noisy Measurements

Result 1 is indeed sensitive to the accuracy of the absolute velocity measurements, and may change radically if they are subject to noise. Here, we model additive measurement noise and let the velocity measurement in (3) be

$$v^m = v + \varepsilon\eta,$$

where the vector η contains uncorrelated white noise and ε is a scaling factor defined through $\mathbb{E}\{\eta(\tau)\eta^T(t)\} = \varepsilon\mathbb{E}\{w(\tau)w^T(t)\}$.

The technical framework presented in [10] allows us to analyze how the measurement noise affects the scaling of local and global performance. It uses a block-diagonalization of the system through discrete Fourier transforms of the function arays associated with the feedback operators. This allows a derivation of closed-form expressions of the systems' $\mathcal{H}_2$ norms, whose scalings in N can then be analyzed.

This paper's main result is presented in the following proposition.

Proposition 2: Consider the system (1) with control input (3) and assume that the velocity measurements are noisy, that is, $v^m = v + \varepsilon\eta$. Then, row (iii) of Table I lists the asymptotic scaling of the per site variance V_N.

Table 1 reveals that the measurement noise η leads to an unfavorable scaling of both local and global error variance – even worse than for distributed static feedback. This may not be an issue for small networks, as the variance is proportional to the factor ε^2, which can be very small. However, it limits the overall scalability of distributed integral control to large networks.

A. Achieving scalable integral control

The error variance of the system under noisy integral control can be partitioned into two terms; one due to the process disturbances w and one due to measurement noise η. It is the scaling of the latter term that causes the unfavorable performance scaling reported in Proposition 2.

For any given system with a fixed network size, it is possible to trade off these terms by tuning the controller and optimize the overall performance. This can, for example, be done by tuning the distributed averaging filter A. However, no such tuning will be scalable to large networks. The situation is illustrated in Fig. 2. In fact, we show that the best-achievable performance scaling is that of distributed static feedback:

Proposition 3: The best-achievable performance scaling for the system (1) under noisy integral control (3) is that of distributed static feedback in Table I.

To retrieve the best-achievable performance scaling according to Proposition 3, the impact of the measurement noise η must be limited. We show that this can only be done in the following ways:

1) Reducing the integral gain c_o: Reducing the gain c_o in (3) reduces (or eliminates) the impact of the measurement noise η. In this case, c_o must be *decreased as* $1/L^2$ *(i.e.* $1/N^{2/d}$*)*. In practice, this implies $c_o \to 0$ and the integral action is eliminated. In this case, the control input u is meaningless.

2) Increasing the distributed averaging gain – centralized averaging: The distributed averaging gain $\bar{a}$ can be defined as $\bar{a} = ||a||_\infty$, where $a = \{a_k\}$ is the function array associated with the operator A in (3). In the example (4) $\bar{a} = a_+ = a_-$.

To prevent the unfavorable performance scaling due to the noise η, the gain $\bar{a}$ must be *increased as* L^2 *(i.e.* $N^{2/d}$*)*. In practice, this means that $\bar{a} \to \infty$ when the lattice size L grows.

While an infinite gain in distributed averaging is not feasible, the same result can be realized as *centralized* averaging integral control, where a central controller has instantaneous access to the integral states at all nodes. The control signal u_k is then the same for all $k \in \mathbb{Z}_L^d$:

$$u_k = z; \quad z = \frac{1}{N}\sum_{k\in\mathbb{Z}_L^d} v_k^m. \tag{9}$$

It is not difficult to show that this controller has the same performance with respect to the errors (6) and (7) as static feedback.

3) Increasing communication network connectivity: Let q_A be the "communication window" for the operator A, i.e. the size of the neighborhood in each coordinate direction within which each controller aligns the state z. That is, each node is allowed to communicate with its $(2q_A)^d$ nearest neighbors. Formally, we can define $q_A := \max_{a_k \neq 0}|k|$.

To prevent the unfavorable performance scaling due to the noise η, we must require $q_A \sim L = N^{1/d}$. This implies that the communication window must scale with the lattice size. For large networks, this in principle leads to all-to-all communication. This can be practically challenging and a centralized approach as in (9) is likely preferable.

III. Conclusions

We have characterized limitations of distributed integral control in terms of the *scaling* of $\mathcal{H}_2$ performance to large networks. We showed that such limitations arise due to noisy measurements, and can only be alleviated by increasing the amount of inter-nodal alignment between controllers (either through the number of connections, or their weights).

This is in contrast to previous results reported in [9], [11], [12] which showed that "little" inter-nodal alignment of integral states (i.e., small gains $\bar{a}$ and few interconnections in communication network) is optimal for performance in the absence of measurement noise. It is intuitively clear that the inter-nodal alignment becomes increasingly important if noise is considered explicitly. Our results are, however, surprising in that it is not enough to increase this alignment in proportion to the noise intensity. Instead, the weights must be

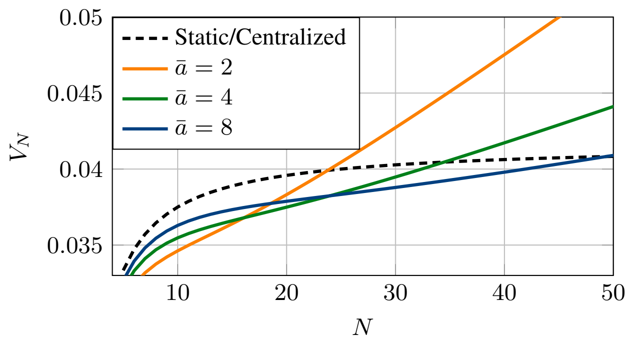

Fig. 2: Scaling of local error variance with static feedback vs. noisy distributed integral control in 1D lattice. For a given network size N, it is possible to set the gain $\bar{a}$ in the distributed averaging filter A so that the integral controller performs better than the static controller. Yet, no such controller scales well to larger networks. A centralized integral controller on the form (9), corresponding to $\bar{a} \to \infty$, will however have the same performance as static feedback.

increased as $\bar{a} \sim L^2 = N^{2/d}$ or the communication window $q_A \sim L = N^{1/d}$.

Naturally, any real-world application will have a finite number of nodes. The distributed integral controller can therefore always be tuned for acceptable performance. Our results imply, however, that such a tuning cannot be done independently of the network size. Therefore, even though the controller is implemented in a distributed fashion, its tuning requires global knowledge.

The results presented in this paper have been derived for a particular distributed integral controller and under the assumption of a spatially invariant. It is fairly straightforward to show that the key results hold also for more general integral controllers on the form $u = z$, $\dot{z} = Az + Bx + Cv^m$. Under more general assumptions on the network topology, it is possible to derive certain bounds on the forms of the expressions in Table 1. However, it is an open and interesting research question whether network heterogeneities, which are present in any real-world application, can be exploited to improve system performance.

Acknowledgements

We wish to thank Hendrik Flamme and Bassam Bamieh for many insightful discussions related to this work.

References

[1] P. Barooah and J. P. Hespanha, "Estimation on graphs from relative measurements," *IEEE Control Systems*, vol. 27, no. 4, pp. 57–74, Aug 2007.

[2] B. Bamieh, M. R. Jovanović, P. Mitra, and S. Patterson, "Coherence in large-scale networks: Dimension-dependent limitations of local feedback," *IEEE Trans. Autom. Control*, vol. 57, no. 9, pp. 2235 – 2249, Sept. 2012.

[3] M. Andreasson, D. Dimarogonas, H. Sandberg, and K. Johansson, "Distributed PI-control with applications to power systems frequency control," in *American Control Conf.*, June 2014, pp. 3183–3188.

[4] J. W. Simpson-Porco, F. Dörfler, and F. Bullo, "Synchronization and power sharing for droop-controlled inverters in islanded microgrids," *Automatica*, vol. 49, no. 9, pp. 2603 – 2611, 2013.

[5] F. Lin, M. Fardad, and M. Jovanović, "Optimal control of vehicular formations with nearest neighbor interactions," *IEEE Trans. Autom. Control*, vol. 57, no. 9, pp. 2203–2218, Sept 2012.

[6] S. Patterson and B. Bamieh, "Consensus and coherence in fractal networks," *IEEE Trans. Control Netw. Syst.*, vol. 1, no. 4, pp. 338–348, Dec 2014.

[7] M. Siami and N. Motee, "Fundamental limits and tradeoffs on disturbance propagation in large-scale dynamical networks," *IEEE Trans. Autom. Control*, vol. 61, no. 12, pp. 4055–4062, 2016.

[8] T. W. Grunberg and D. F. Gayme, "Performance measures for linear oscillator networks over arbitrary graphs," *IEEE Trans. Control Netw. Syst.*, 2017, to appear.

[9] E. Tegling and H. Sandberg, "On the coherence of large-scale networks with distributed PI and PD control," *IEEE Control Systems Letters*, vol. 1, no. 1, pp. 170–175, July 2017.

[10] E. Tegling, P. Mitra, H. Sandberg, and B. Bamieh, "On fundamental limitations of dynamic feedback control in large-scale networks," 2017, arXiv preprint: https://arxiv.org/abs/1710.02880.

[11] X. Wu, F. Dörfler, and M. R. Jovanović, "Topology identification and design of distributed integral action in power networks," in *American Control Conf.*, July 2016, pp. 5921–5926.

[12] M. Andreasson, E. Tegling, H. Sandberg, and K. H. Johansson, "Coherence in synchronizing power networks with distributed integral control," in *56th IEEE Conf. on Decision and Control*, Dec 2017, pp. 6327–6333.