

Adding a single state memory optimally accelerates symmetric linear maps

Alain Sarlette

► **To cite this version:**

Alain Sarlette. Adding a single state memory optimally accelerates symmetric linear maps. IEEE Transactions on Automatic Control, Institute of Electrical and Electronics Engineers, 2016, 61 (11), pp.3533. <hal-01093907v2>

HAL Id: hal-01093907

<https://hal.inria.fr/hal-01093907v2>

Submitted on 28 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adding a single state memory optimally accelerates symmetric linear maps

A. Sarlette — alain.sarlette@inria.fr
QUANTIC project team, Paris Sciences Lettres University,
INRIA Rocquencourt, Paris, France
& SYSTeMS, Ghent University, Belgium.

June 2015

Abstract

Previous papers have proposed to add memory registers to the dynamics of discrete-time linear systems in order to accelerate their convergence. In particular, it has been proved that adding one memory slot per agent allows faster convergence towards average consensus. We here prove that this situation *cannot* be improved by adding more memory slots, when the knowledge about the self-adjoint linear map to be accelerated reduces to bounds on its extreme eigenvalues.

1 Main result & related work

1.1 Motivating question: consensus

A consensus algorithm in discrete time [1] writes, for a column vector $x = [x_1; x_2; \dots; x_N] \in \mathbb{R}^N$:

$$x(t+1) = x(t) - \alpha L x(t). \quad (1)$$

Here $x_k \in \mathbb{R}$ is the state of agent k , $\alpha > 0$ is a gain, and L is the Laplacian matrix characterizing interactions among agents: component j of Lx equals $\sum_{k=1}^N w_{j,k} (x_j - x_k)$ with weights $w_{j,k} \geq 0$. Usually L contains many zeros, as each agent interacts with few fellows. For the algorithm to preserve the average of the initial values $x_k(0)$, we assume $w_{j,k} = w_{k,j}$ for all j, k . The Laplacian L is then symmetric nonnegative definite, and if the interactions form a connected graph it has a single eigenvalue $\lambda_1 = 0$ with eigenvector $v_1 = [1; 1; \dots; 1]$; we call v_1 a consensus situation. When $(I_N - \alpha L)$ has nonnegative entries, with I_N the $N \times N$ identity matrix, it can equivalently be viewed as the transition matrix of a Markov chain; symmetric L implies that the limiting distribution is uniform.

For time-invariant L , the convergence of (1) is dictated by the largest eigenvalue, in modulus, of $(I_N - \alpha L)$ — excluding the trivial eigenvalue 1 associated to $\lambda_1 = 0$ and eigenvector v_1 , which spans the target subspace. In the orthonormal basis corresponding to the eigenvectors of L (so-called “modes”), the system decouples into

$$\tilde{x}_k(t+1) = (1 - \alpha \lambda_k) \tilde{x}_k(t), \quad k = 1, 2, \dots, N, \quad (2)$$

with \tilde{x}_k the coefficient of mode k and λ_k the associated eigenvalue. If we only know that the eigenvalues λ_k of L belong to an interval $[\underline{\lambda}, \bar{\lambda}] \subset \mathbb{R}_{>0}$ for $k = 2, 3, \dots, N$, then the convergence speed — in terms of bounds on eigenvalues of (2) guaranteed over all $\lambda_k \in [\underline{\lambda}, \bar{\lambda}]$ — is optimal when α is selected to satisfy $(1 - \alpha\underline{\lambda}) = -(1 - \alpha\bar{\lambda})$. This gives $\alpha = \frac{2}{\bar{\lambda} + \underline{\lambda}}$ and worst eigenvalue $\mu = \frac{\bar{\lambda} - \underline{\lambda}}{\bar{\lambda} + \underline{\lambda}}$.

In practice, the time step of discrete-time consensus is mostly limited by communication speed, much more rarely by local computation power. Hence [2] and later [3, 4, 5] propose to improve convergence speed by adding local dynamics to each agent, namely a memory slot with associated gain $\beta_1 \in \mathbb{R}$:

$$x_k(t+1) = x_k(t) - \alpha \sum_{j=1}^N w_{k,j} (x_k(t) - x_j(t)) - \beta_1 (x_k(t) - x_k(t-1)). \quad (3)$$

The papers analyze in detail only this case of one memory slot and prove that with $\beta_1 < 0$ it allows faster convergence. This spurs the natural question:

How much can be gained by adding more memory slots?

The answer is **strictly nothing**, as we prove next in a more general setting.

1.2 Formal result

Consider a general linear iteration

$$x(t+1) = x(t) + \alpha(b - Ax(t)) \quad (4)$$

where $x = [x_1; x_2; \dots; x_N] \in \mathbb{R}^N$ is a dynamic variable, $\alpha \in \mathbb{R}$ is a constant gain to be chosen, $b \in \mathbb{R}^N$ is a given constant bias and A is a given constant, self-adjoint positive semidefinite linear map on \mathbb{R}^N . We are interested in applications where (4) is applied iteratively in order to converge towards the set of fixed points $\mathcal{S} = \{s \in \mathbb{R}^N : As = b\}$. For consensus, with $A = L$ and $b = 0$, we have $\mathcal{S} = \{\lambda v_1 : \lambda \in \mathbb{R}\}$; in other computational applications \mathcal{S} might reduce to a single point; we always assume \mathcal{S} to be nonempty.

Motivated by the consensus setting, we consider a “memory-accelerated” version of (4):

$$x(t+1) = x(t) + \alpha(b - Ax(t)) + \sum_{m=0}^{M-1} \beta_m x(t-m),$$

with chosen gains $\beta_m \in \mathbb{R}$ for $m = 0, 1, 2, \dots, M-1$. This yields stationary solutions x satisfying $\alpha(b - Ax) + \left(\sum_{m=0}^{M-1} \beta_m\right)x = 0$. Hence we require $\sum_{m=0}^{M-1} \beta_m = 0$ to ensure that \mathcal{S} remains a set of stationary solutions under the “memory-accelerated” dynamics. The latter can then be rewritten as:

$$x(t+1) = x(t) + \alpha(b - Ax(t)) + \sum_{m=1}^{M-1} \beta_m (x(t-m) - x(t)), \quad (5)$$

with freely chosen gains $\beta_m \in \mathbb{R}$. For analysis, we can rewrite (5) in the basis of eigenvectors of A and this just leads to a set of N independent scalar equations

$$\tilde{x}_k(t+1) = (1 - \alpha\lambda_k)\tilde{x}_k(t) + \alpha\tilde{b}_k + \sum_{m=1}^{M-1} \beta_m (\tilde{x}_k(t-m) - \tilde{x}_k(t)) \quad (6)$$

where λ_k are the associated eigenvalues of A , for $k = 1, 2, \dots, N$. Hence it is clear that convergence of (5) is governed by the separate convergence of each modal component \tilde{x}_k associated to its eigenvalue λ_k in (6). Note that for \mathcal{S} to be nonempty, we must have $\tilde{b}_k = 0$ for every k for which $\lambda_k = 0$. Then the modes with $\lambda_k = 0$ of both (4) and (5) (initialized with $\tilde{x}_k(t) = \tilde{x}_k(0)$ for $t < 0$) become trivial, such that convergence towards \mathcal{S} is governed by the *nonzero* eigenvalues of A only. We consider the setting where the choice of $\alpha, \beta_1, \dots, \beta_{M-1}$ is allowed to exploit the following knowledge. Let $|z|$ and $\angle(z)$ respectively denote the modulus and the phase of $z \in \mathbb{C}$.

Assumption 1: *It is known a priori that the nonzero eigenvalues λ_k of A take values in the interval $[\underline{\lambda}, \bar{\lambda}] \subset (0, +\infty)$.*

Definition 2: *Denote $\gamma_m(\lambda)$, $m = 1, 2, \dots, M$ the roots of the characteristic polynomial of the linear difference equation (6) with $\lambda_k = \lambda$. Then we define the convergence speed guarantee ν of (5) for fixed $\alpha, \beta_1, \dots, \beta_{M-1}$ by*

$$\nu := \max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} \max_{m \in \{1, 2, \dots, M\}} |\gamma_m(\lambda)|,$$

i.e. ν denotes the worst possible eigenvalue of (6) over all λ_k satisfying Assumption 1.

We then prove the following result.

Theorem 3: *Consider dynamics (5) with A self-adjoint and satisfying Assumption 1. Then the convergence speed guarantee is optimized, i.e. $\nu \in (0, 1)$ is minimized over all $\alpha, \beta_1, \beta_2, \dots, \beta_{M-1} \in \mathbb{R}$, when using just a single memory slot optimally tuned as in [2], i.e. taking:*

$$\begin{aligned} \beta_m &= 0 \text{ for all } m > 1, \\ \beta_1 &= \beta_{1*} := - \left(\frac{1}{\mu} - \sqrt{\frac{1}{\mu^2} - 1} \right)^2, \\ \alpha &= \alpha_* := \frac{2(1 - \beta_{1*})}{\bar{\lambda} + \underline{\lambda}}, \end{aligned}$$

where $\mu = \frac{\bar{\lambda} - \underline{\lambda}}{\bar{\lambda} + \underline{\lambda}}$. *The corresponding convergence speed guarantee is*

$$\nu_* = \frac{1}{\mu} - \sqrt{\frac{1}{\mu^2} - 1} = \sqrt{-\beta_{1*}}.$$

This ν_* increases with μ , which was the optimal convergence speed guarantee for $M = 1$ i.e. no added memory slot. For $\mu = 1 - \varepsilon$, with ε frequently called the *spectral gap*, we have $\nu_* < 1 - \sqrt{\varepsilon}$ and the bound gets tight as $\varepsilon \rightarrow 1$.

Remark 4: [2] and related work mention that $\beta_1 < 0$ is necessary to get acceleration over the memoryless case. As an auxiliary result, it is not difficult to show for (5) that if we are restricted to $\beta_m \geq 0$ for all m then no acceleration with respect to the memoryless case can be achieved.

1.3 Related acceleration approaches

Before proving Theorem 3 on the basis of complex polynomial analysis, we mention a few related acceleration schemes in the literature. We hope that in the future these viewpoints might lead to a more elegant explanation and intuitive understanding of the limitation to one memory slot.

Optimization techniques The scheme studied here is not unlike techniques proposed earlier by the optimization community, where acceleration methods are a major focus point. For instance the Nesterov method [6] accelerates the gradient descent optimization of a convex function f , i.e. it accelerates

$$x(t+1) = x(t) - \alpha \nabla_x f(x(t))$$

by applying, in its most basic version:

$$\begin{aligned} m(t) &= x(t) - \alpha \nabla_x f(x(t)) \\ x(t+1) &= m(t) + g(t)(m(t) - m(t-1)). \end{aligned}$$

For $f = \frac{1}{2}x^T Ax - b^T x$, this corresponds exactly to the acceleration (5) of (4) with $M = 2$, except that the Nesterov method involves a time-varying step size $g(t)$, whose details go beyond the scope of the present paper. Our result seems to suggest that just adding more memory slots would not further accelerate the Nesterov method.

Robust control The problem setting (6) with $\tilde{b}_k = 0$ can be viewed as an ensemble of closed-loop systems resulting from proportional feedback with different gains λ_k , i.e.:

$$\begin{aligned} y(z) &= H(z)u(z), \quad u(z) = -\lambda_k y(z) \quad \text{with} \\ H(z) &= \frac{\alpha}{z - 1 + \sum_{m=1}^{M-1} \beta_m (1 - z^{-m})} \end{aligned}$$

where $y(z)$ and $u(z)$ are the output and input respectively and $H(z)$ is the plant transfer function. We then want to design the plant $\alpha, \beta_1, \dots, \beta_{M-1}$ to get the fastest possible worst-case performance over the ensemble of feedback gains $\lambda_k \in [\underline{\lambda}, \bar{\lambda}]$. In view of Theorem 3, a plant with $M = 2$ would be best in terms of convergence speed. As a major caveat though, the present paper proves optimality over all $H(z)$ with no zeros and with at least one pole at $z = 1$; the possible practical relevance of this setting in control applications would have to be checked.

Various reformulations of where the uncertainty sits can of course be envisioned, e.g. the uncertain plant could be modeled by A whereas $\alpha, \beta_1, \dots, \beta_{M-1}$ would characterize the controller. From that perspective, (6) can be viewed as an optimal control problem for a so-called (non-symmetric) interval matrix or linear parametric uncertainty system, along the lines of Kharitonov's theorem. Unfortunately the latter does not generalize easily to discrete-time systems [7, 8]. An interesting result when there is a single uncertain parameter is proposed in [9]. Nevertheless, the author found no simple way to use this towards proving Theorem 3.

Accelerated consensus As mentioned in the introduction, our initial motivation was the study of consensus algorithms accelerated through ‘‘local’’ memories; see [2, 3, 4, 5] which propose the scheme (similar to) (5) and analyze the case $M = 2$. Very recently [10] has proposed a scheme with ‘‘extrapolating’’ step, similar to the case $M = 2$ (and to the Nesterov method, see above) but where the limited information known about the network is a bound on the total number of nodes, instead of on the extremal eigenvalues of L .

Researchers have also considered accelerations based on different resources than local memory, for instance allowing *periodically time-varying* gain α in (4). Optimal tuning rules for $\alpha(t)$ can then build on *polynomial-based filtering* [12, 13]. The optimal acceleration strategy

of [11] can also be viewed as an instance of this framework, provided one limits the algorithm to a finite number of Chebyshev polynomials which are then applied periodically. In this latter context, when λ_k can take all values in $[\underline{\lambda}, \bar{\lambda}]$, it is easy to check that the memory-based strategy (6) is faster than the optimal periodically time-varying one. However if more is known about the λ_k , then polynomial-based filtering is straightforward to adapt towards better speed-up. In the extreme case where L has M different eigenvalues whose values are exactly known, it is possible to construct a polynomial of order M that achieves deadbeat convergence in M steps. Similar “finite-time consensus” strategies have been proposed and studied in e.g. [14, 15], with an interpretation of the resulting algorithms as gathering initial state information at a central node, computing its average and redistributing the average through the network.

It seems not straightforward to compare the resources “local memory” and “time-varying gain” outside an applicational context. In particular a standard time-varying implementation requires a stronger type of synchronization among all the nodes, namely they must not only agree on the frequency of updates but also on which gain to apply at each step. This may or may not be restrictive. The Nesterov method combines both memory and time-dependent step lengths. These approaches, along with accelerated Markov chains (see next paragraph), offer alternatives towards further acceleration but apparently *only when more is known than* $\lambda_k \in [\underline{\lambda}, \bar{\lambda}]$.

Accelerated Markov chains The optimal acceleration of the dual of consensus, namely Markov chains, has also been extensively studied in the literature, see e.g. the “lifting” method of [16] and an even faster “pseudo-lifting” in [17]. Those approaches exploit knowledge of the particular network, hence they apply to the case where more is known than only $\lambda_k \in [\underline{\lambda}, \bar{\lambda}]$ (see previous paragraph). In fact [17] proposes an accelerated consensus version which directly connects with the “gather-and-distribute” strategy of [14, 15], although with *a priori* more complex communication requirements: it requires each node to communicate a vector of several values at each time step.

Yet lifted Markov chains cannot be viewed straightforwardly as the dual of accelerated consensus.¹ Whether the result of [2] and/or Theorem 3 can be used to investigate Markov chain acceleration, possibly without requiring detailed network knowledge, remains to be studied. The lifted Markov chain [16] *using* detailed network knowledge is limited to quadratic speed-up.

Information theory Consider the consensus application from the viewpoint of an individual node k . The messages received by node k when *symmetrically* exchanging data with a network, reflect in part the other nodes’ initial values $\{x_j(0) : j \neq k\}$ and in part the influence of $x_k(0)$ on the other nodes in the network. It seems not unreasonable to suspect that local memory can help disentangle these influences. Yet Theorem 3 says that if the network is poorly characterized, in the precise sense $\lambda \in [\underline{\lambda}, \bar{\lambda}]$, then having more than one additional memory is not helpful. This somehow seems to say: taking into account the direct feedback loop from k through its neighbors j and directly back to k does allow improved convergence; but speculating about longer feedback loops does not pay off.

¹In fact, although this is a minor point, even in the unaccelerated case, Markov chains unlike consensus are restricted to systems with positive coefficients in the matrix $I - \alpha L$.

Finally, let us emphasize that Theorem 3 starts with a positive semidefinite self-adjoint matrix A , which can be viewed as a diffusive operator, or in consensus as an undirected communication matrix. In contrast, the state matrix characterizing the accelerated system (5) is not symmetric. Hence it seems to introduce transport, or hidden directed communication (effects), which is known to improve convergence speed in several contexts [16, 18]. This also implies that the acceleration of [2], exactly like Markov chain liftings, cannot be applied iteratively. I.e. the new state matrix, characterizing the system that results from an acceleration according to [2], is not symmetric which precludes applying the technique of [2] once more.

In the following we first prove Theorem 3, before discussing an example in Section 4.

2 Proof, first part: reformulation

We consider the dynamics with memory slots in the form (6). The convergence speed for each mode is governed by the roots of

$$P(z; \lambda_k) = z^M - (1 - \alpha\lambda_k)z^{M-1} + \sum_{m=0}^{M-2} \beta_{M-m-1} (z^{M-1} - z^m), \quad (7)$$

viewed as a function of z parameterized by λ_k . In accordance with Definition 2, for the sequel we replace the set of λ_k by a generic λ which runs through $[\underline{\lambda}, \bar{\lambda}]$.

2.1 Optimal solution with one memory slot

To make the paper self-contained, we outline a proof of optimal tuning for the case $M = 2$, which is covered in [2].

Proposition 5 [adapted from [2]]: *For $M = 2$ the tuning $\beta_1 = \beta_{1*}$, $\alpha = \alpha_*$, as proposed in Theorem 3, minimizes the value of the convergence speed guarantee ν . Moreover, the roots of*

$$P_*(z; \lambda) = z^2 - (1 - \alpha_*\lambda)z + \beta_{1*}(z - 1)$$

take the values $z_{\pm*}(\lambda) = \nu e^{\pm i\theta_\lambda}$ where the map $\lambda \mapsto \theta_\lambda$ is a continuous bijection from $[\underline{\lambda}, \bar{\lambda}]$ to $[0, \pi]$.

Proof: First note that $\beta_1 \geq 1$ always leads to an unstable eigenvalue. Then the roots of $P(z; \lambda)$ are equal to the roots $z_{\pm}(\tilde{\lambda})$ of

$$f(z; \tilde{\lambda}) = z^2 - (1 - \beta_1)\tilde{\lambda}z - \beta_1$$

with $\tilde{\lambda} = (1 - \alpha\lambda - \beta_1)/(1 - \beta_1)$.

We first fix a value of β_1 and study the roots of $f(z; \tilde{\lambda})$ as a function of $\tilde{\lambda}$, which translates into an optimal value for α . Standard function analysis yields:

Property 5: *The worst root $\max\{|z_+(\tilde{\lambda})|, |z_-(\tilde{\lambda})|\}$ of $f(z; \tilde{\lambda})$ is a monotonically (non-strictly) increasing function of $|\tilde{\lambda}|$.*

Hence for given β_1 the choice of α should strive to minimize $\max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} (|\tilde{\lambda}|)$. This is obtained by choosing α such that $(1 - \alpha\underline{\lambda} - \beta_1) = -(1 - \alpha\bar{\lambda} - \beta_1)$, since $|\tilde{\lambda}|$ takes its extremal value(s) when λ is at the boundary of the interval $[\underline{\lambda}, \bar{\lambda}]$. We hence get the condition $\alpha = \frac{2(1-\beta_1)}{\lambda+\bar{\lambda}}$.

Moreover, this implies that $\tilde{\lambda}$ runs through the interval $[-\mu, \mu]$ when λ spans $[\underline{\lambda}, \bar{\lambda}]$, with $\mu = \frac{\bar{\lambda}-\underline{\lambda}}{\bar{\lambda}+\underline{\lambda}}$ independent of β_1 .

It remains to optimize β_1 , assuming that it is associated to its respective optimal α ; by Property 5 and the previous sentence, this reduces to:

$$\min_{\beta_1 \in \mathbb{R}} \max_{\tilde{\lambda} \in \pm\mu} \max \left(|z_{\pm}(\tilde{\lambda})| : f(z_{\pm}(\tilde{\lambda}); \tilde{\lambda}) = 0 \right).$$

A standard function analysis shows that it is optimal to take β_1 such that $z_+(\mu) = z_-(\mu) = -z_+(-\mu) = -z_-(-\mu)$. This yields the expressions of α_* , β_{1*} and ν given in Theorem 3.

The just mentioned optimality condition requires that the discriminant $\Delta_* = (1 - \beta_{1*})^2 \tilde{\lambda}^2 + 4\beta_{1*}$ of $f(z; \tilde{\lambda}) = 0$ takes its zeros for $\tilde{\lambda} = \pm\mu$. Then for all $\tilde{\lambda} \in [-\mu, \mu]$, i.e. all $\lambda \in [\underline{\lambda}, \bar{\lambda}]$, we have $\Delta_* \leq 0$. Therefore $4|z_{\pm*}(\lambda)|^2 = (1 - \beta_{1*})^2 \tilde{\lambda}^2 - \Delta_*$ (sum of squared real and imaginary parts) = $-4\beta_1$ independently of $\lambda \in [\underline{\lambda}, \bar{\lambda}]$. Continuity of the roots of $f(z)$ between $\tilde{\lambda} = \mu$ and $\tilde{\lambda} = -\mu$ yields the last part of the Proposition. \square

2.2 General polynomial property

We now reformulate $P(z; \lambda)$ using Proposition 5.

Claim 6 [to be proved]: *For any $\nu \in (0, 1)$ and $\tilde{P}_{M-1} = \sum_{k=0}^{M-1} a_k y^k$ with $a_1, a_2, \dots, a_{M-1} \in \mathbb{R}$ and $a_{M-1} \neq -1$, there exists $\theta \in [0, \pi]$ for which*

$$\tilde{P}(y; \theta) = (y - e^{i\theta})(y - e^{-i\theta})y^{M-2} + (y - \frac{1}{\nu})\tilde{P}_{M-1}(y) \quad (8)$$

has a root of modulus ≥ 1 .

Proposition 7: *Theorem 3 is true if Claim 6 holds.*

Proof [of Proposition 7]: Let

$$\hat{P}(z; \theta) = (z - \nu e^{i\theta})(z - \nu e^{-i\theta})z^{M-2} + (z - 1) \sum_{m=0}^{M-1} \tilde{\beta}_m z^m.$$

Using $z^2 - (1 - \alpha_* \lambda)z + \beta_{1*}(z - 1) = (z - \nu e^{i\theta\lambda})(z - \nu e^{-i\theta\lambda})$ from Proposition 5, it is straightforward to check that $\hat{P}(z; \theta_\lambda) = \frac{\alpha_*}{\alpha} P(z; \lambda)$ provided

$$\begin{aligned} \tilde{\beta}_k &= \frac{\alpha_*}{\alpha} \sum_{m=0}^k \beta_{M-m-1} \quad \text{for } k \leq M-3, \\ \tilde{\beta}_{M-2} &= \frac{\alpha_*}{\alpha} (\sum_{m=0}^{M-2} \beta_{M-m-1}) - \beta_{1*}, \\ \tilde{\beta}_{M-1} &= \frac{\alpha_*}{\alpha} - 1. \end{aligned}$$

Thus by taking appropriate $\tilde{\beta}_0, \tilde{\beta}_1, \dots, \tilde{\beta}_{M-1} \in \mathbb{R}$, the roots of $\hat{P}(z; \theta_\lambda)$ can be made equal to the roots of $P(z; \lambda)$ from Theorem 3 for any $\alpha, \beta_1, \dots, \beta_{M-1} \in \mathbb{R}$, except for $\alpha = 0$. The latter case is trivially uninteresting since for $\alpha = 0$ any $x \in \mathbb{R}^N$ would be a stationary point of (5), hence no convergence can be obtained. Moreover, the case $\tilde{\beta}_{M-1} = -1$ will never appear since this would require infinite α . The tuning proposed in Theorem 3, with roots $\nu e^{\pm i\theta\lambda}$ and 0, corresponds to $\tilde{\beta}_0 = \tilde{\beta}_1 = \dots = \tilde{\beta}_{M-1} = 0$. The parameter $\lambda \in [\underline{\lambda}, \bar{\lambda}]$ in $P(z; \lambda)$ is transferred to $\theta_\lambda \in [0, \pi]$ in $\hat{P}(z; \theta_\lambda)$, according to Proposition 5.

Hence a sufficient condition to prove Theorem 3 is to show that for any $\tilde{\beta}_0, \tilde{\beta}_1, \dots, \tilde{\beta}_{M-1} \in \mathbb{R}$, with $\tilde{\beta}_{M-1} \neq -1$, there exists some $\theta \in [0, \pi]$ such that $\hat{P}(z; \theta)$ has at least one root with modulus $\geq \nu$. The statement of the Proposition is obtained by defining $y = z/\nu$ and $a_m = \tilde{\beta}_m \nu^{m-M+1}$ for $m = 1, 2, \dots, M-1$. \square

Claim 6 looks like a discrete-time robust (in)stability property. Surprisingly, we know no standard result that would straightforwardly establish it. Examples can be constructed where the relevant roots are never real, or appear only for intermediate values of θ (see Section 4), which seems to rule out simple variants of polynomial roots properties. Hence the remainder of this paper explicitly analyzes $\tilde{P}(y; \theta)$ in the complex plane. Note that for any polynomial $P(z)$, the modulus $|P(z)|$ is continuous for all $z \in \mathbb{C}$, while $e^{i\angle(P(z))}$ is continuous for all $z \in \mathbb{C} \setminus \{z : P(z) = 0\}$.

3 Proof, second part: analyzing $\tilde{P}(y; \theta)$

We split up $\tilde{P}(y; \theta)$ into

$$\begin{aligned} P_1(y; \theta) &:= (y - e^{i\theta})(y - e^{-i\theta})y^{M-2}, \\ P_2(y) &:= -(y - \frac{1}{\nu})\tilde{P}_{M-1}(y). \end{aligned}$$

Any y at which $P_1(y; \theta) = P_2(y)$ is a root of $\tilde{P}(y; \theta)$. We first dispose of two special cases.

Proposition 8:

(a) If \tilde{P}_{M-1} in Claim 6 has a root on the unit circle, then $\tilde{P}(y; \theta)$ has a root y_1 with $|y_1| \geq 1$ for some $\theta \in [0, \pi]$.

(b) If $a_{M-1} < -1$ in Claim 6, then $\tilde{P}(y; \theta)$ has a root y_1 with $|y_1| \geq 1$ for all $\theta \in [0, \pi]$.

Proof: (a) Let y_* the root of \tilde{P}_{M-1} on the unit circle and take $e^{i\theta_*} = y_*$ if $\text{Imaginary}(y_*) \geq 0$, else $e^{-i\theta_*} = y_*$. Then $\tilde{P}(y_*, \theta_*) = 0$ with $|y_*| \geq 1$.

(b) We separately analyze (i) the sign and (ii) the magnitude of both P_1 and P_2 for y belonging to the real positive axis.

(i) Let $y_* \geq 1/\nu > 1$ the largest real root of $P_2(y)$. Both $P_2(y)$ and $P_1(y; \theta)$ are positive for $y \in (y_*, +\infty)$, for any θ .

(ii) $|P_2(y)| < |P_1(y; \theta)|$ for y close to y_* , while for very large $|y|$ we have $|P_2(y)| \simeq |a_{M-1}| |y|^M > |P_1(y; \theta)| \simeq |y|^M$.

Hence for any θ , there exists some $y_1 \in (y_*, +\infty)$ where $|P_1(y_1; \theta)| = |P_2(y_1)|$, which with (i) implies $P_1(y_1; \theta) = P_2(y_1)$. This y_1 is a root of $\tilde{P}(y; \theta)$ with $|y_1| > 1$. \square

From Prop. 8(a), we can reduce our investigation to the case where the roots of P_2 are disjoint from the roots of P_1 . Indeed, we have excluded roots of P_2 on the unit circle, whereas a situation with P_2 and P_1 having the common root $y = 0$ can be reduced to an equivalent algorithm with a lower value of M .

The analysis of $\tilde{P}(y; \theta)$ towards Claim 6 for the remaining cases is essentially a generalization of the proof of Prop. 8(b), from the real line towards the complex plane.

3.1 Partitioning the complex plane with P_1 and P_2

For given $P_1(\cdot; \theta)$ and $P_2(\cdot)$, we can partition the complex plane into a collection of open connected sets where $|P_1| > |P_2|$ (we call these type 1 sets), a collection of open connected sets where $|P_2| > |P_1|$ (we call these type 2 sets), separated by sets where $|P_2| = |P_1|$. Note that, as the roots of P_1 and P_2 can be assumed disjoint (see Prop.8(a)), any root of P_2 unambiguously belongs to a type 1 set and any root of P_1 belongs to a type 2 set.




Figure 1: Illustrating the definition of the type 1 sets (shaded), type 2 sets (plain white) and $\Gamma_R(y_*)$ (darker shading). Small blue crosses represent roots of P_2 , small blue circles the roots of P_1 . The boundary $\partial\Gamma_R(y_*)$ (thick plain and dotted curves) and its subset $\partial\Gamma_{\bar{R}}(y_*)$ (thick plain curves) are also represented. The complete picture is meant to be symmetric around the real axis. The sets drawn here are purely schematic, they are not meant to correspond to any actual polynomials P_1 , P_2 and may possibly be more general than what polynomial properties would allow.

For any root y_* of P_2 with $|y_*| > 1$, let $\Gamma(y_*)$ the type 1 set containing y_* . Then for any $R > |y_*|$, let

$$\tilde{\Gamma}_R(y_*) := \Gamma(y_*) \cap \{y \in \mathbb{C} : 1 < |y| < R\}$$

and define $\Gamma_R(y_*)$ to be the connected component of $\tilde{\Gamma}_R(y_*)$ such that $y_* \in \Gamma_R(y_*)$ (see Figure 1). There is at least one such set corresponding to $y_* = 1/\nu$. We denote the boundary of $\Gamma_R(y_*)$ by $\partial\Gamma_R(y_*)$. Note that for any $y \in \partial\Gamma_R(y_*)$ we have either $|P_2(y)| = |P_1(y; \theta)|$, or $|P_2(y)| < |P_1(y; \theta)|$; the latter case involves points where either $|y| = 1$, or $|y| = R$. We finally define $\partial\Gamma_{\bar{R}}(y_*) = \partial\Gamma_R(y_*) \cap \{y : |y| < R\}$ and reduce the proof of Claim 6 to the following.

Claim 9 [to be proved]: *For any P_2 — excluding the cases already handled by Proposition 8 — there exists $\tilde{\theta} \in [0, \pi]$ and y_* a root of P_2 such that $e^{i\angle(P_2(\tilde{y}))} = e^{i\angle(P_1(\tilde{y}; \tilde{\theta}))}$ for some $\tilde{y} \in \partial\Gamma_{\bar{R}}(y_*)$.*

Proposition 10: *Claim 6 is true if Claim 9 holds.*

Proof: We must show that for any situation satisfying Claim 9, we can find $y_1 \in \mathbb{C}$ with $|y_1| \geq 1$ and $\theta \in [0, \pi]$, such that $P_1(y_1; \theta) = P_2(y_1)$. For given P_2 , excluding the cases covered by Proposition 8, take $\tilde{y}, \tilde{\theta}$ according to Claim 9.

If \tilde{y} belongs to the part of $\partial\Gamma_{\bar{R}}(y_*)$ where $|P_2(y)| = |P_1(y; \tilde{\theta})|$, we have $P_1(\tilde{y}; \tilde{\theta}) = P_2(\tilde{y})$ with $|\tilde{y}| \geq 1$ by construction. Hence the case is closed with $y_1 = \tilde{y}$, $\theta = \tilde{\theta}$.

If $|P_2(\tilde{y})| \neq |P_1(\tilde{y}; \tilde{\theta})|$, then necessarily $|P_2(\tilde{y})| < |P_1(\tilde{y}; \tilde{\theta})|$ and $|\tilde{y}| = 1$. Let $\tilde{y} = e^{i\phi}$ and assume $\phi \in [0, \pi]$; the case $\phi \in [-\pi, 0]$ is the same modulo a few notation changes. As for Prop. 8(b), we separately investigate (i) the phase and (ii) the modulus of P_1 , P_2 .

(i) By computing

$$(\tilde{y} - e^{i\theta})(\tilde{y} - e^{-i\theta})\tilde{y}^{M-2} e^{-i(M-1)\phi} = 2(\cos(\phi) - \cos(\theta)).$$

we see that $e^{i\angle(P_1(\tilde{y}; \theta))} = e^{i\angle(P_1(\tilde{y}; \tilde{\theta}))} = e^{i\angle(P_2(\tilde{y}))}$ for all θ for which $(\cos(\phi) - \cos(\theta))$ takes the

same sign as $(\cos(\phi) - \cos(\tilde{\theta}))$, i.e. for all $\theta \in [\tilde{\theta}, \phi] =: \mathcal{I}_1$ or all $\theta \in (\phi, \tilde{\theta}] =: \mathcal{I}_1$, depending on the ordering of $\tilde{\theta}$ and ϕ .

(ii) $f(\theta) := |P_2(\tilde{y})| - |P_1(\tilde{y}; \theta)|$ is a continuous function of θ , with $f(\tilde{\theta}) < 0$ and $f(\phi) > 0$ by construction. Hence there exists $\theta_1 \in \mathcal{I}_1$ such that $f(\theta_1) = 0$.

Combining (i) and (ii), we have $P_1(y_1; \theta) = P_2(y_1)$ with $y_1 = \tilde{y}$ and $\theta = \theta_1$. \square

3.2 Proving Claim 9

It remains to prove that Claim 9 is true, which in fact is a consequence of *Cauchy's argument principle* in complex analysis. Adapted to the current case, the argument principle states the following fact.

Property 11 [Cauchy's Argument Principle]: *Let $\mathcal{D} \subset \mathbb{C}$ a bounded, simply connected open set whose boundary $\partial\mathcal{D}$ is the image of a closed simple curve $\gamma : [0, 1] \rightarrow \mathbb{C}$. Assume that \mathcal{D} contains n roots of $P_2(y)$ and p roots of $P_1(y; \theta)$ for some fixed θ , while $\partial\mathcal{D}$ contains no roots. Then there exists a continuous function $f(t) : [0, 1] \rightarrow \mathbb{R}$ such that*

$$\exp\left(i\left(\angle(P_1(\gamma(t); \theta)) - \angle(P_2(\gamma(t)))\right)\right) = \exp(if(t)),$$

and any such $f(t)$ satisfies $|f(1) - f(0)| = 2\pi|p - n|$.

We will apply this principle with $\partial\mathcal{D}$ defined by elements of $\partial\Gamma_R(y_*)$. Note that the boundary of any $\Gamma_R(y_*)$ is indeed sufficiently regular since the locus in the complex plane where $|P_1|^2 - |P_2|^2 = 0$ is by definition a planar algebraic curve. We will also use the following related property.

Property 12: *Consider the setting of Property 11 with $|p - n| =: m > 0$. Then there are at least m points $y = \gamma(t)$ on $\partial\mathcal{D}$ where $e^{i\angle(P_2(y))} = e^{i\angle(P_1(y; \theta))}$.*

Proof: The interval $(f(0), f(0) + 2\pi m] \subset \mathbb{R}$ contains m different multiples of 2π and the function $f(t)$ must take each of these values at least once as it continuously evolves from $f(0)$ to $f(1) = f(0) + 2\pi m$. The $\gamma(t)$ associated to these values of $f(t)$ satisfy $e^{i\angle(P_2(y))} = e^{i\angle(P_1(y; \theta))}$. \square

The remainder of the proof relies on the following observations (see Fig. 2). For some fixed θ , let $\mathcal{D}_R(y_*)$ the smallest simply connected set containing $\Gamma_R(y_*)$, where the latter contains $m \geq 1$ roots of P_2 . Note that $\partial\mathcal{D}_R(y_*) \subseteq \partial\Gamma_R(y_*)$.

Case A: If $\mathcal{D}_R(y_*)$ does not contain the open unit disc (see Fig. 2 left), then it contains at least m roots of P_2 but no root of P_1 . Hence applying Properties 11 and 12 with $\mathcal{C} = \mathcal{D}_R(y_*)$, there must be m points on $\partial\mathcal{D}_R(y_*) \subseteq \partial\Gamma_R(y_*)$ where $e^{i\angle(P_2(y))} = e^{i\angle(P_1(y; \theta))}$.

Case B: If $\mathcal{D}_R(y_*)$ contains the open unit disc, then $\partial\Gamma_{\bar{R}}(y_*)$ contains a closed curve γ_0 that separates $\Gamma_R(y_*)$ from the open unit disc and we can define $\mathcal{D}_0(y_*) \subset (\mathcal{D}_R(y_*) \setminus \Gamma_R(y_*))$ the simply connected set whose boundary is γ_0 (see Fig. 2 right). This set contains the M roots of P_1 but at most $M - m$ roots of P_2 . Hence by Properties 11 and 12 there must be m points on $\gamma_0 \subset \partial\Gamma_{\bar{R}}(y_*)$ where $e^{i\angle(P_2(y))} = e^{i\angle(P_1(y; \theta))}$.

Claim 9 is readily true for Case B. For Case A, we need to ensure that at least one of the m points in $\partial\Gamma_R(y_*)$ also belongs to $\partial\Gamma_{\bar{R}}(y_*)$. Towards this, we investigate $e^{i\angle(P_1(y; \theta)) - i\angle(P_2(y))}$ when $|y| = R$ with R very large.




Figure 2: Schematic illustration of Case A (left) and Case B (right) to which we apply Properties 11 and 12. The sets and roots are represented similarly to Fig. 1. For Case A the set in red is $\mathcal{D}_R(y_*)$. For Case B, the striped set is $\mathcal{D}_R(y_*)$, the thick curve is γ_0 and the set enclosed by γ_0 is $\mathcal{D}_0(y_*)$.

Lemma 13: *If in P_2 the coefficient $a_{M-1} > 0$ then there exists $R_1 > 0$ such that $e^{i\angle(P_1(y;\theta)) - i\angle(P_2(y))} \neq 1$ for all y with $|y| > R_1$ and for all $\theta \in [0, \pi]$.*

Proof: Note that $e^{i\angle(P_1(y;\theta)) - i\angle(P_2(y))} = \exp\left(i\angle\left(\frac{P_1(y;\theta)}{P_2(y)}\right)\right)$. Let us evaluate the phase of the ratio of two general polynomials $P_g(y) = \sum_{k=0}^M g_k y^k$ and $P_h(y) = \sum_{k=0}^M h_k y^k$ with real coefficients g_k, h_k and assuming $g_M h_M < 0$. Applying the formula $\frac{a+bi}{c+di} = \frac{1}{c^2+d^2}((ac+bd) + (bc-ad)i)$, with a, b, c, d real and i the imaginary unit, we get with $y = Re^{i\phi}$:

$$\frac{P_g(Re^{i\phi})}{P_h(Re^{i\phi})} = \frac{\sum_{k=0}^M \frac{g_{M-k}}{R^k} (\cos(k\phi) - i \sin(k\phi))}{\sum_{k=0}^M \frac{h_{M-k}}{R^k} (\cos(k\phi) - i \sin(k\phi))} = \eta (g_M h_M + O(1/R) + iO(1/R))$$

for some $\eta > 0$. Thus there exists $R_1 > 0$ such that

$$\text{Real}\left(\frac{P_g(Re^{i\phi})}{P_h(Re^{i\phi})}\right) = g_M h_M + O(1/R) < 0$$

for all $R > R_1$. The Lemma follows by applying the above result with $P_g(y) = P_1(y; \theta)$ and $P_h(y) = P_2(y)$. \square

We are now ready to prove Claim 9 in a few steps.

Proof of Claim 9 for $a_{M-1} \leq 0$: From Prop.8(c) we only need to consider the case $a_{M-1} \in (-1, 0]$. Then there exists $R_1 > 1$ such that $|y|^M \simeq |P_1(y; \theta)| > |P_2(y)| \simeq |a_{M-1}| |y|^M$ for all y for which $|y| > R_1$. Take any root y_* of P_2 with $|y_*| > 1$. Then for $R > R_1$, the points y with $R_1 < |y| < R$ cannot belong to $\partial\Gamma_R(y_*)$, since this would require $|P_1(y; \theta)| = |P_2(y)|$. This implies that $\partial\Gamma_R(y_*)$ contains either all points or no point of the circle $\mathcal{C}_R := \{y \in \mathbb{C} : |y| = R\}$. In case $\partial\Gamma_R(y_*) \cap \mathcal{C}_R = \emptyset$, we have $\partial\Gamma_R(y_*) = \partial\Gamma_{\bar{R}}(y_*)$ and the observations after Proposition 12 allow to directly conclude. In case $\partial\Gamma_R(y_*) = \mathcal{C}_R$, we are necessarily in Case B of the observations after Proposition 12 hence the conclusion is also immediate. \square

Proof of Claim 9 for $a_{M-1} > 0$: Take $R > R_1$ to satisfy Lemma 13. The set $\partial\Gamma_R(y_*) \setminus \partial\Gamma_{\bar{R}}(y_*)$ is a subset of $\{y \in \mathbb{C} : |y| = R\}$, and Lemma 13 implies $e^{i\angle(P_2(y))} \neq e^{i\angle(P_1(y;\theta))}$ for all y for which $|y| = R$. Hence all the points identified in the observations after Proposition 12, where $e^{i\angle(P_2(y))} = e^{i\angle(P_1(y;\theta))}$ on $\Gamma_R(y_*)$, must belong to $\Gamma_{\bar{R}}(y_*)$. \square

With this we have covered all situations and hence concluded the proof of Claim 9, which proves Theorem 3.

4 Example

Consider a linear map A with nonzero eigenvalues $\lambda_k \in [0.0122, 0.9878]$. The optimal acceleration without added memory slot i.e. with $M = 1$ yields $\mu = 0.9756$ and a spectral gap $1 - \mu = 0.0244$. With $M - 1 = 1$ added memory slot, an improved convergence speed guarantee $\nu = 0.8000$ is obtained, with $1 - \nu = 0.2 < \sqrt{0.0244} = 0.1562$, using the optimal parameters $\alpha_* = 3.2800$ and $\beta_{1*} = -0.6400$.

If more is known about the eigenvalues of A then additional memory does help accelerate. For instance consider A with nonzero eigenvalues $\lambda_k \in \Lambda = [0.0122, 0.0182] \cup \{0.9878\}$ i.e. the largest eigenvalue is in fact isolated far away from the others. This does not allow to improve convergence speed with $M = 2$, because for $M = 2$ the bound $\nu \geq 0.8000$ holds as soon as A features both eigenvalues 0.0122 and 0.9878 (see the proof of Prop. 5).

In contrast, $M = 4$ and parameter values $\alpha = 3.6908$, $\beta_1 = -0.9083$, $\beta_2 = 0.006662$, $\beta_3 = 0.06785$ do yield an improved convergence speed guarantee $\tilde{\nu} = 0.7560$ over all $\lambda_k \in \Lambda$ (with $\tilde{\nu}$ defined by straightforward extension of Def. 2). Note that these parameters have been obtained numerically by local search around $\alpha = \alpha_*$, $\beta_1 = \beta_{1*}$ and $\beta_k = 0$ for $k > 1$; we do not exclude the existence of better parameter values. Figure 3 shows $|z_*|$ the corresponding largest root in modulus of (7) as a function of λ_k . It highlights how the improved $|z_*|$ for $\lambda_k \in \Lambda$ comes to the detriment of (much) worse $|z_*|$ for $\lambda_k \in [0.0291, 0.9788]$.

Finally, let us illustrate the proof of Claim 9 and Prop.10 with those values, i.e. $M = 4$ and $\alpha = 3.6908$, $\beta_1 = -0.9083$, $\beta_2 = 0.006662$, $\beta_3 = 0.06785$ for $[\underline{\lambda}, \bar{\lambda}] = [0.0122, 0.9878]$ and hence $\nu = 0.8000$. For three values of $\theta \in \{\pi/40, \pi/10, \pi/4\}$, Figure 4 shows in white the set in \mathbb{C} where $|P_1(y; \theta)| < |P_2(y)|$ and in grey the set where $|P_1(y; \theta)| > |P_2(y)|$. Small circles are the roots of P_1 and small crosses the roots of P_2 . In the middle is the unit circle and the colored lines denote the locus where $e^{i\angle(P_1)} = e^{i\angle(P_2)}$. For all values of θ , we are in Case B of Section 3.2. Correspondingly, we have highlighted in black the curve γ_0 delimiting the simply connected set $\mathcal{D}_0(y_*)$. For all θ , there are two points along γ_0 where $e^{i\angle(P_1)} = e^{i\angle(P_2)}$. For $\theta = \pi/40$ the crossing occurs at some y_1 on the unit circle (red dot on Fig. 4), in a region where $|P_1| > |P_2|$. For $e^{i\theta}$ closer to y_1 , e.g. $\theta = \pi/10$, we still have $e^{i\angle(P_1(y_1;\theta))} = e^{i\angle(P_2(y_1))}$ but y_1 now belongs to a region where $|P_2| > |P_1|$. Hence somewhere in between there must be some $\tilde{\theta}$ for which $|P_1(y_1; \tilde{\theta})| = |P_2(y_1)|$. This precisely corresponds to the point where the plot on Fig. 3 crosses $|z_*| = \nu = 0.8000$, and suffices to prove Theorem 3. We pursue this example towards further insight on bad or good θ . Moving $e^{i\theta}$ beyond y_1 finally changes $\angle(P_1(y_1; \theta))$ by π and the points along the unit circle where $e^{i\angle(P_1)} = e^{i\angle(P_2)}$ start moving; for a large interval of θ values, the points on γ_0 where $e^{i\angle(P_1)} = e^{i\angle(P_2)}$ also satisfy $|P_2| = |P_1|$, as on the third graph, hence the polynomial is worse than the optimal one with $M = 2$. As




Figure 3: Largest root in modulus $|z_*|$ of (7) with $M = 4$, $\alpha = 3.6908$, $\beta_1 = -0.9083$, $\beta_2 = 0.006662$, $\beta_3 = 0.06785$, as a function of $\lambda_k \in [0.0122, 0.9878]$. Reducing the relevant eigenvalue set to $\lambda_k \in [0.0122, 0.0182] \cup \{0.9878\}$ (vertical lines) yields an improved convergence speed guarantee $\tilde{\nu} = 0.7560$ (horizontal full line) with respect to the optimal situation $\nu = 0.8000$ with $M = 2$ (horizontal dashed line). However this happens to the detriment of $|z_*| > 0.8000$ for all $\lambda_k \in [0.0291, 0.9788]$.

θ approaches π a behavior symmetric to the neighborhood of $\theta = 0$ takes place (not shown).

5 Conclusion

We have proved that the maximal achievable acceleration of a linear iterative map by adding $M - 1$ memory slots to each state vector component, is already achieved with $M - 1 = 1$. More precisely, this holds for the performance guarantee of a constant map about which we only know bounds $[\underline{\lambda}, \bar{\lambda}] \subset (0, +\infty)$ on the nonzero eigenvalues of the self-adjoint state update matrix A . The fact that the nonzero eigenvalues of A can take *any* values in $[\underline{\lambda}, \bar{\lambda}]$ is important for the property to hold. Better accelerations can be devised if more is known about A , see e.g. Section 4.

A direct extension of our result would allow each subsystem to follow general linear dy-

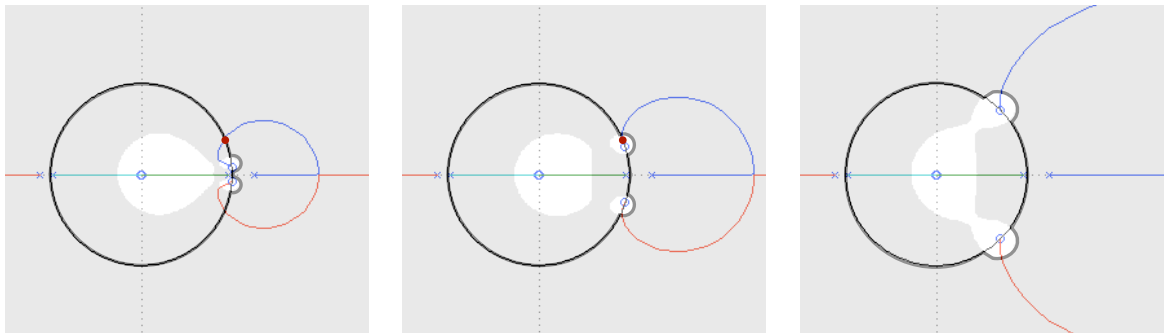


Figure 4: Characteristics of $P_1(y; \theta)$ and $P_2(y)$ for three values of $\theta \in \{\pi/40, \pi/10, \pi/4\}$, illustrating the proof of Claim 9 and Prop.10 on our example polynomial. See main text for details.

namics; thinking of the consensus application with rational input-output transfer function at each node, this would mean characteristic polynomials of the form

$$(z - 1)P(z) + \alpha\lambda_k Q(z)$$

with $P(z)$ and $Q(z)$ two polynomials. The memory slot setting (5) is restricted to $Q(z) = z^{M-1}$, i.e. the consensus algorithm uses only the latest information coming from the network. It is not clear if convergence can speed up with other $Q(z)$.

The basic consensus algorithm (1) has a proven robustness to network incidents². For the case (5) with optimally tuned memory slot, one can construct examples where packet drops lead to instability. Thus the benefit of more memory slots might have to be reevaluated towards robustness.

Acknowledgment

This research has been carried out within the Interuniversity Attraction Poles network DYSCO. The author wants to thank F.Ticozzi and S.Zampieri for stimulating discussions.

References

- [1] J. N. Tsitsiklis, “Problems in decentralized decision making and computation,” *PhD Thesis, MIT*, 1984.
- [2] S. Muthukrishnan, B. Ghosh, and M. H. Schultz, “First- and second-order diffusive methods for rapid, coarse, distributed load balancing,” *Theory Comput. Systems*, vol. 31, pp. 331–354, 1998.
- [3] B. Johansson and M. Johansson, “Faster linear iterations for distributed averaging,” *Proceedings of the 17th IFAC World Congress*, pp. 2861–2866, 2008.
- [4] E. Ghadimi, I. Shames, and M. Johansson, “Accelerated gradient methods for networked optimization,” *arxiv*, vol. 1211.2132, 2012.
- [5] J.Liu, B. Anderson, M.Cao, and S. Morse, “Analysis of accelerated gossip algorithms,” *Automatica*, vol. 49, pp. 873–883, 2013.
- [6] Y. Nesterov, “A method of solving a convex programming problem with convergence rate $O(1/k^2)$,” in *Soviet Mathematics Doklady*, vol. 27, no. 2, 1983, pp. 372–376.
- [7] C. V. Hollot and A. C. Bartlett, “Some discrete-time counterparts to Kharitonov’s stability criterion for uncertain systems,” *IEEE Trans. Automatic Control*, vol. 31, no. 4, pp. 355–356, 1986.

²Stability under switching networks holds: for any directed graphs, as long as $I - \alpha L$ has all components positive, thanks to the common Lyapunov function $V = \max_i x_i - \min_i x_i$ [1]; for undirected graphs, as long as $I - \alpha L$ has all eigenvalues inside the unit circle, thanks to the common Lyapunov function $\sum_i (x_i)^2$. The conditions on $I - \alpha L$ always hold if we start from a stable nominal network and “packet drops” occur, i.e. failures can at times modify L by setting to zero the gain associated to some links.

- [8] P. Vaidyanathan, “Derivation of new and existing discrete-time Kharitonov theorems based on discrete-time reactances,” *IEEE Trans. Acoustics, Speech and Signal Processing*, vol. 38, no. 2, pp. 277–285, 1990.
- [9] B. R. Barmish, “A generalization of kharitonov’s four-polynomial concept for robust stability problems with linearly dependent coefficient perturbations,” *Automatic Control, IEEE Transactions on*, vol. 34, no. 2, pp. 157–165, 1989.
- [10] A. Olshevsky, “Linear time average consensus on fixed graphs and implications for decentralized optimization and multi-agent control,” <http://arxiv.org/pdf/1411.4186v5.pdf>, 2015.
- [11] E. Montijano, J. Montijano, and C. Saúés, “Chebyshev polynomials in distributed consensus applications,” *IEEE Trans. Signal Proc.*, vol. 61, no. 3, pp. 693–706, 2013.
- [12] E.Kokiopoulou and P. Frossard, “Polynomial filtering for fast convergence in distributed consensus,” *arxiv*, vol. 0802.3992, 2008.
- [13] S. Sundaram and C. Hadjicostis, “Finite-time distributed consensus in graphs with time-invariant topologies,” *Proc. American Control Conference*, pp. 711–716, 2007.
- [14] J. Hendrickx, R. Jungers, A. Olshevsky, and G. Vankeerberghen, “Graph diameter, eigenvalues, and minimum-time consensus,” *Automatica*, vol. 50, no. 2, pp. 635–640, 2014.
- [15] L. Georgopoulos, “Definitive consensus for distributed data inference,” *PhD thesis, EPFLausanne*, 2011.
- [16] F. Chen, L. Lovasz, and I. Pak, “Lifting Markov chains to speed up mixing,” *Proc.STOC’99*, 1999.
- [17] K. Jung, D. Shah, and J. Shin, “Distributed averaging via lifted markov chains,” *Information Theory, IEEE Transactions on*, vol. 56, no. 1, pp. 634–647, 2010.
- [18] P. Barooah, P. Mehta, and J. Hespanha, “Mistuning-based control design to improve closed-loop stability margin of vehicular platoons,” *IEEE Trans. Automatic Control*, vol. 54, no. 9, pp. 2100–2113, 2009.