Published in final edited form as:

IEEE Trans Neural Netw Learn Syst. 2012 September ; 23(9): . doi:10.1109/TNNLS.2012.2202400.

Convergence and Rate Analysis of Neural Networks for Sparse Approximation

Aurèle Balavoine [Student Member, IEEE], Justin Romberg [Senior Member, IEEE], and Christopher J. Rozell [Senior Member, IEEE]

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA

Abstract

We present an analysis of the Locally Competitive Algorithm (LCA), which is a Hopfield-style neural network that efficiently solves sparse approximation problems (e.g., approximating a vector from a dictionary using just a few nonzero coefficients). This class of problems plays a significant role in both theories of neural coding and applications in signal processing. However, the LCA lacks analysis of its convergence properties, and previous results on neural networks for nonsmooth optimization do not apply to the specifics of the LCA architecture. We show that the LCA has desirable convergence properties, such as stability and global convergence to the optimum of the objective function when it is unique. Under some mild conditions, the support of the solution is also proven to be reached in finite time. Furthermore, some restrictions on the problem specifics allow us to characterize the convergence rate of the system by showing that the LCA converges exponentially fast with an analytically bounded convergence rate. We support our analysis with several illustrative simulations.

Keywords

Exponential convergence; global stability; locally competitive algorithm; Lyapunov function; nonsmooth objective; sparse approximation

I. Introduction

Sparse approximation has generated substantial interest in a wide range of research communities over the last two decades, including signal processing, machine learning, statistics, and computational neuroscience (see [2], [3] and references therein).¹ Specifically, sparse approximation involves solving an optimization problem to represent a signal using just a few atoms from some specified (possibly overcomplete) dictionary. In addition to describing compelling models of neural coding for sensory information [4], [5], this approach has led to state-of-the-art results in many types of inverse problems. One example of a regime that has leveraged this type of signal model is the theory of compressed sensing (CS) [6], [7]. In this domain, highly undersampled signals are recovered by solving a sparse approximation problem, thereby shifting the burden of data acquisition from the front-end sensor to a computationally intensive back end.

Because of the increasing interest in the sparse approximation paradigm, significant effort has been made to design efficient algorithms for solving (or approximately solving) least squares optimization problems that are regularized with a sparsity-inducing penalty (e.g., [8]–[10]). However, all of these algorithms are developed to run on digital computers in discrete time, which are both implausible for neural systems and suffer from several drawbacks as engineering approaches. In particular, the computational time required by these algorithms presents a barrier to real-time signal processing applications with high-dimensional signals at significant data rates. Specifically, digital algorithms tend to have storage requirements and convergence times that scale unfavorably with the dimensionality of the signals being approximated. Additionally, the power consumption of digital solutions can be prohibitive for many applications. Given these considerations, a fast and low-power method for solving sparse approximation problems in a parallel architecture would be valuable both for engineering systems and viable models of neural coding.

Analog neural networks have long been proposed for solving optimization problems [11], with an early example being Hopfield's pioneering results [12] using networks to solve linear optimization problems. Analog neural networks offer several potential advantages over comparable digital algorithms, including their ability to be implemented in analog architectures that are highly parallel, fast, and power-efficient. Recent advances in very-large-scale integration (VLSI) reconfigurable analog chips [13] make the design of such systems more feasible and affordable than has often been true in the past. A recent neural network architecture, called a Locally Competitive Algorithm (LCA) [14], has been proposed to solve the types of nonsmooth optimization problems that come up in sparse approximation. These Hopfield-style networks appear to efficiently solve a whole family of sparse approximation problems by incorporating ideal thresholding nonlinearities in the network dynamics. The results in [14] provide encouraging evidence that biological systems and engineering applications can use neural networks to solve these important sparse approximation problems.

However, despite encouraging evidence of its performance, the LCA lacks strong convergence guarantees and estimates of the convergence rate. Furthermore, the specifics of the LCA architecture violate many of the assumptions used in previous analysis, making the extensive literature on neural network convergence inapplicable for this system. For example, in previous work the activation function is often assumed to be linear [15], piecewise linear [16], or nonlinear but increasing and bounded [17]–[20]. In contrast, the LCA activation function is nonlinear and unbounded. Other analyses assume an interconnection matrix which is positive definite [16], [21], [22] or nonsingular [23], [24], while the LCA interconnection matrix may have negative eigenvalues as well as a nontrivial nullspace (because of the approximation dictionary being overcomplete). Other analyses consider a nonsmooth objective function with constraints which shows convergence only when the constraints are nonzero and convex [25], [26]. Finally, other relevant analyses assume that the objective function is convex [27] or piecewise linear [28], while the work in [29] focuses on the controllability of the network path.

The main contributions of this paper are to present a formal analysis of the LCA network architecture and its convergence properties, despite: 1) an activation function that is nonsmooth and not necessarily bounded or increasing and 2) a potentially singular interconnection matrix. Section III contains our first main result, which states that the fixed points of the LCA network correspond to critical points of the objective function. In the special case, where the objective is convex, this set coincides with the global minima of the objective. In addition, we show that the network is globally stable and that the outputs are quasi-convergent, in the sense that they get infinitely close to a set of fixed points. Finally, in the case where the objective function has isolated minima, we show that the LCA

converges to the solution of the sparse approximation problem for any initial point (i.e., a much stronger condition than just the nonincreasing property of the energy function that was shown in previous work [14]). This section also shows that the LCA is well behaved in that it converges in a finite number of switches (i.e., nodes crossing above or below threshold). Section IV expands on these general results to show that, under additional mild conditions on the problem specifics, the LCA actually converges *exponentially fast* to the solution. Furthermore, we give an analytic expression for this convergence rate that depends on the properties of the detailed approximation problem. Finally, Section V presents simulation results showing the correspondence of our analytic results with empirical observations of the network behavior.

II. Background

Before presenting our main results, in this section we briefly give a precise statement of the sparse approximation problems of interest, a description of the LCA architecture, and some preliminary observations on the LCA network dynamics that will be useful in the subsequent analysis.

A. Sparse Approximation

As mentioned above, sparse approximation is an optimization program that seeks to find the approximation coefficients of a signal on a prescribed dictionary, using as few nonzero elements as possible. To fix notation, we denote the input signal by $y \in \mathbb{R}^M$, the unit-norm dictionary elements by the columns of the $M \times N$ matrix $\Phi = [\phi_1, \dots, \phi_N]$, and the by $a \in \mathbb{R}^N$. Generally, $M \ll N$ (i.e., the approximation dictionary is overcomplete), and the problem of recovering a from y is underdetermined. While similar in spirit to the well-known winner-take-all problem [30], sparse approximation problems are generally formulated as the solution to an optimization program because this approach can often yield strong performance guarantees in specific applications (e.g., recovery in a CS problem). In the most generic form, the objective function is the sum of a quadratic data fidelity term (i.e., mean squared error) and a regularization term that uses a sparsity-inducing cost function $C(\cdot)$

$$\min_a V(a(t)) = \frac{1}{2} \|y - \Phi a\|_2^2 + \lambda \sum_{n=1}^N C(a_n) \quad (1)$$

where the parameter λ is a tradeoff between the two terms in the objective function. The “ideal” sparse approximation problem has a cost function that simply counts the number of nonzero elements, resulting in a nonconvex objective function that has many local minima [31].

One of the most widely used programs from this family is known as Basis Pursuit DeNoising (BPDN) [32], which is given by the objective function

$$\min_a \frac{1}{2} \|y - \Phi a\|_2^2 + \lambda \|a\|_1. \quad (2)$$

In BPDN, the ℓ_1 -norm ($\|a\|_1 = \sum_n |a_n|$) is used as a convex surrogate the idealized counting norm. This program has gained in popularity as researchers have shown that, in many cases of interest, substituting the ℓ_1 norm yields the same solution as using an idealized (and generally intractable) counting norm [33]. However, BPDN illustrates the canonical challenge of sparse approximation problems. Despite being convex, the BPDN

objective function contains a nonsmooth nonlinearity which makes it considerably more difficult than a classic least-squares problem.

In the context of computational neuroscience, sparse approximation has been proposed as a neural coding scheme for sensory information. In one interpretation, programs such as BPDN can be viewed as Bayesian inference in a linear generative model with Gaussian noise and a prior with high kurtosis to encourage sparsity (e.g., the Laplacian prior in the case of BPDN) [5]. Given the prevalence of probabilistic inference as a successful description of human perception [34] and the theoretical benefits of sparse representations [3], it has long been conjectured that sensory systems may encode stimuli via sparse approximation. In fact, in classic results, sparse approximation applied to the statistics of natural stimuli in an unsupervised learning experiment has been shown to be sufficient to qualitatively and quantitatively explain the receptive field properties of simple cells in the primary visual cortex [4], [35] as well as the auditory nerve fibers [36]. Only recently have there been proposals of efficient neural networks that could efficiently solve the necessary optimization problems to implement this type of encoding [14], [35], [37], [38].

B. LCA Structure

Our primary interest will be the LCA [14], which is an analog continuous-time dynamical system that is a type of Hopfield-style network. In particular, each node in the LCA network is characterized by the evolution of a set of internal state variables, $u_n(t)$ for $n = 1, \dots, N$, and uses a nonlinear activation function $T(\cdot)$ to produce output variables $a_n(t)$ for $n = 1, \dots, N$. The activation function is typically a nonsmooth nonlinear function (such as a thresholding function) to induce sparsity in the outputs. The dynamics of the internal state variables are governed by a set of coupled nonlinear ordinary differential equations (ODEs)

$$\begin{aligned} \tau \dot{u}(t) &= -u(t) - (\Phi^T \Phi - I) a(t) + \Phi^T y \\ a(t) &= T_\lambda(u(t)). \end{aligned} \quad (3)$$

In this network, each node is associated with a single dictionary element ϕ_n , $n = 1, \dots, N$, and the node outputs will be shown to be the solution to the optimization problem of interest. The architecture of a typical LCA is shown in Fig. 1.

The inputs to the LCA network are feedforward connections computing the vector of driving synaptic inputs $\Phi^T y$, reflecting how well the signal y matches each dictionary element. The network also has recurrent inhibitory or excitatory connections between the nodes, modulated by weights corresponding to the interconnection matrix $G = \Phi^T \Phi - I$ (i.e., a modified Grammian matrix for the dictionary). The more overlap there is between a pair of nodes (characterized by the inner product between their dictionary elements), the stronger will be the potential inhibition between those nodes. While the modulating weight is symmetric between any pair of nodes, the total inhibition is not, because it is also modulated by the activity of each individual node. Moreover, the matrix G potentially has both negative and positive eigenvalues, as well as a nontrivial nullspace. This inhibition structure ensures that nodes that carry the same information will not become active at the same time, thus meeting the goals of sparse approximation. The time constant τ is dependent on the physical properties of the solver implementing the ODEs. For our purposes, we will often assume $\tau = 1$ without loss of generality.

It was shown in [14] that the objective function in (1) is nonincreasing along the LCA trajectory with the following relationship ($\forall a_n \in \mathbb{R}$ such that $a_n \geq 0$) between the cost penalty term $C(\cdot)$ and the activation function $T(\cdot)$

$$\lambda \frac{dC(a_n)}{da_n} = u_n - a_n = u_n - T_\lambda(u_n). \quad (4)$$

In the case of BPDN, the cost penalty is simply $C(x) = |x|$ and the activation function obtained from (4) is the soft-thresholding function (Fig. 2) defined by

$$a_n(t) = T_\lambda(u_n(t)) = \begin{cases} 0, & |u_n(t)| \leq \lambda \\ u_n(t) - \lambda \operatorname{sign}(u_n(t)), & |u_n(t)| > \lambda. \end{cases}$$

This function often arises in connection to algorithms for minimizing the absolute value of the coefficients (see [39]). Generalizing from the soft threshold, we focus on “thresholding” activation functions $T(\cdot)$ of the form

$$a_n(t) = T_\lambda(u_n(t)) = \begin{cases} 0, & |u_n(t)| \leq \lambda \\ f(u_n(t)), & |u_n(t)| > \lambda \end{cases} \quad (5)$$

where the function $f(\cdot)$ is real-valued defined and continuous on the domain $\mathcal{D} = (-\infty, -\lambda] \cup [\lambda, +\infty)$, differentiable on the interior of $\mathcal{D}$, and satisfying the conditions

$$f(-u_n) = -f(u_n) \quad \forall u_n \in \mathcal{D} \quad \text{and} \quad f(\lambda) = 0 \quad (6a)$$

$$f'(u_n) > 0 \quad \forall u_n \in \mathcal{D} \quad (6b)$$

$$f(u_n) \leq u_n \quad \forall u_n \in \mathcal{D} \quad \text{s.t.} \quad u_n \geq 0. \quad (6c)$$

Remark 1: Equation (6a) ensures that $T(\cdot)$ is continuous on all $\mathbb{R}$ and, consequently, that $u(t)$ and $a(t)$ are continuous with respect to time. This ensures that the cost function $C(\cdot)$ is differentiable everywhere except at the origin. Equation (6b) makes $f(\cdot)$ a bijection on $\mathcal{D}$. Finally, (6a) and (6c) ensure, by a simple computation, that $C(\cdot)$ is positive and nondecreasing with the absolute value of the coefficient. Notice that the ℓ_1 -minimization objective function satisfies the three criteria in (6), in addition to being convex (which is not required, but guarantees that the system will find a global minima). In Fig. 3, we plot a more generic stylized activation function that satisfies the conditions (6).

As shown in Fig. 3, the activation function (5) is composed of two operating regions. The first region corresponds to the case where the internal state u_n is below the threshold λ , in which case the output a_n is zero. We call the nodes in this region *inactive nodes*. The second region corresponds to the case where the internal state u_n is above threshold, in which case (6b) guarantees that the output a_n is strictly increasing with u_n . We call the nodes that are above threshold *active nodes*, and we denote by $\mathcal{A}$ the *active set* (i.e., the set of indices corresponding to those nodes, $\mathcal{A} = \{n \in [1, M] : |u_n(t)| > \lambda\}$). On the contrary, we denote the set of indices corresponding to nodes that are below threshold by $\mathcal{I}$ and call them the *inactive set*. While these two sets do change with time as the network evolves, for the sake of readability we omit the dependence on time in the notation.

Owing to the nonsmooth nature of the objective (1), we require the generalized notion of a subgradient (see [40] for more details). Note that $V(a)$ is differentiable everywhere except at points a that contain zeros, due to the nondifferentiability of $C(\cdot)$ at the origin noted above

for the cases of interest in sparse approximation. The subgradient $\partial V(\cdot)$ extends the classic notion of a gradient $\nabla V(\cdot)$ at those points of discontinuity. The properties in (6) allow us to use the following simple definition:

$$\partial V(a) = \text{co} \left\{ \lim_{i \rightarrow \infty} \nabla V(a_i) : a_i \rightarrow a, a_i \notin \mathcal{S}, a_i \notin \Omega_V \right\}$$

where co is the convex hull, Ω_V is the set of points where V fails to be differentiable, and $\mathcal{S}$ is any set of Lebesgue measure 0 in $\mathbb{R}^N$. In other words, it is the smallest convex set containing the gradients of the function as we approach the discontinuity from any direction. Note that, when a is not a point of discontinuity, this simply reduces to $\nabla V(a)$.

C. LCA Node Dynamics

Because nodes can cross threshold and go from the inactive set to the active set (and vice versa), the LCA can be thought of as a switched system [41] where the ODEs change at each switching time (i.e., when a node crosses above or below threshold). In between two switching times, the active set and inactive set Γ^c are fixed. Therefore, the dynamics on the two sets can be considered separately to facilitate analysis. In the following, $\Phi_{\mathcal{T}}$ is the matrix composed of the columns of Φ indexed by the set $\mathcal{T}$. Similarly, $u_{\mathcal{T}}$ and $a_{\mathcal{T}}$ refer to the elements in the original vectors indexed by $\mathcal{T}$. We also denote by $\{t_k\}_{k \in \mathbb{N}}$ the sequence of times for which the system switches from the set of active nodes Γ_{k-1} to Γ_k . The notation $T(u(t))$ refers to the vector $[T(u_1(t)), \dots, T(u_N(t))]^T$, and we denote by $F(t)$ the associated Jacobian matrix with respect to the internal state variables u . Because the activation function has the form in (5), the matrix $F(t)$ is diagonal with diagonal elements equal to zero for indices in the inactive set Γ^c and equal to $f'(u_n)$ for indices n in the active set Γ .

Using the chain rule and (5), we can calculate the derivative with respect to time of the outputs on the active set as

$$\dot{a}_{\Gamma}(t) = F'_{\Gamma}(t) \dot{u}_{\Gamma}(t) = \text{diag}\left(f'(u_n(t))\right)_{n \in \Gamma} \dot{u}_{\Gamma}(t). \quad (7)$$

As a consequence, the ODE (3) can be rewritten for nodes in the active set as follows:

$$\dot{a}_{\Gamma}(t) = F'_{\Gamma}(t) \left[-u_{\Gamma}(t) + a_{\Gamma}(t) - \Phi_{\Gamma}^T \Phi_{\Gamma} a_{\Gamma}(t) + \Phi_{\Gamma}^T y \right]. \quad (8)$$

The active nodes follow this ODE until the next switch occurs, changing the sets of active and inactive nodes. Similarly, we can rewrite the set of differential equations acting on the inactive nodes. Since the output of the activation function for nodes in Γ^c is zero, the ODE (3) on the inactive set becomes

$$\dot{u}_{\Gamma^c}(t) = -u_{\Gamma^c}(t) - \Phi_{\Gamma^c}^T \Phi_{\Gamma} a_{\Gamma}(t) + \Phi_{\Gamma^c}^T y. \quad (9)$$

Separating the system dynamics into two separate characterizations for the active nodes and inactive nodes Γ^c yields two sets of differential equations that are partially decoupled, and this will be crucial for our subsequent analysis. This partial decoupling is possible because only the active nodes in the system produce inhibitory signals, coupling their dynamics to the dynamics of each node in the inactive set via the interconnection matrix.

However, because inactive nodes do not inhibit active nodes, the dynamics on the active set are independent of the inactive set.

III. Convergence Results

In [14], the authors show that the LCA network trajectory is nonincreasing on the energy surface corresponding to the desired objective function. While this is a necessary behavior for a network that solves an optimization problem, it is not sufficient to actually show that the network converges to a fixed point (or set of points), let alone reaches a minimum of that desired objective function. Both of these are important guarantees to make before relying on this system as a viable model of neural processing or as an implementation in engineering applications.

In Section III-A, we analyze the general convergence properties of the LCA. In particular, in this section the interconnection matrix $T - I$ may have many negative or zero eigenvalues, which complicates the analysis. The main results of this first section show that the fixed points of the LCA correspond to critical points of the objective (1) and the outputs of the network converge from any initial state to the set of fixed points. In the case of isolated critical points of (1), we show that the LCA is globally convergent in the sense that starting from any initial state, the system converges to a fixed point. Elaborating on these results, Section III-B shows that the LCA (even under very general assumptions) is very well behaved, converging in a finite number of switches (i.e., nodes crossing above or below threshold) and therefore recovering the solution support (i.e., the set of nonzero elements) in finite time. Owing to the difficulties described above, our approach relies on classic results from nonsmooth analysis, such as subgradients and a generalized chain rule (see Appendix A).

A. Global Asymptotic Convergence

To begin, we recall some useful definitions of stability and concepts from Lyapunov theory (see [42] for more background). There exist several notions of stability associated with dynamical systems that describe the evolution of the nodes or their outputs both locally and globally. For a neural network described by a differential equation of the form

$$\dot{x} = G(x), \quad x \in \mathbb{R}^N \quad (10)$$

with outputs $z = T(x)$ defined as in (5), we say that a constant vector $x^* \in \mathbb{R}^N$ is a *fixed point* of (10) if and only if

$$G(x^*) = 0. \quad (11)$$

On the other hand, the outputs of a system reach a *critical point* a^c of the objective function $V(\cdot)$ when they satisfy the inclusion

$$0 \in \partial V(a^c). \quad (12)$$

Note that, if $V(\cdot)$ is convex, then the critical points correspond exactly to the global minima of V . We say that (10) is (Lyapunov) *stable* at x^* if, for each $\epsilon > 0$, there exists $R > 0$ such that for all x_0 with $\|x_0 - x^*\| < R$, and all the solutions $x(\cdot)$ with initial state x_0

$$\|x(t) - x^*\| < \epsilon \quad \forall t > 0. \quad (13)$$

This is clearly a local notion of stability, guaranteeing that, once the trajectory is close to a fixed point, it remains nearby. But this type of stability is insufficient to guarantee global

convergence, which means that trajectories approach a fixed point as time goes to infinity. The following notion of stability is slightly stronger, guaranteeing that trajectories approach at least a set of fixed points regardless of the initial state.

Definition 1: We say that the outputs of (10) are *globally quasi-convergent* if there exists a set $\mathcal{E} = \{z^* \in \mathbb{R}^N : \dot{z}^* = 0\}$ such that for all $x_0 \in \mathbb{R}^N$, the outputs $z(\cdot) = T(x(\cdot))$ with initial state x_0 satisfy $\lim_{t \rightarrow +\infty} z(t) \in \mathcal{E}$.

Finally, the strongest and most desirable form of convergence which guarantees that the nodes converge to a single fixed point is stated in the next definition.

Definition 2: We say that (10) is *globally convergent*, or equivalently *globally asymptotically stable*, if there exists a fixed point x^* at which the system is stable, and if for all initial states $x_0 \in \mathbb{R}^N$ the solutions $x(\cdot)$ satisfy $\lim_{t \rightarrow +\infty} x(t) = x^*$.

Typically, global convergence is established through the use of a so-called Lyapunov function V . The notation $\dot{V}$ refers to the derivative with respect to time, i.e., dV/dt .

Definition 3: A function $V: \mathbb{R}^N \mapsto \mathbb{R}$ is a *weak Lyapunov function* on $\mathbb{R}^N$ if:

1. $V(x) > 0 \forall x \neq 0$;
2. V is continuous on $\mathbb{R}^N$;
3. $V(x) \rightarrow 0 \forall x \in \mathbb{R}^N$;
4. V is radially unbounded

$$\lim_{\|x\| \rightarrow +\infty} V(x) = +\infty.$$

Similarly, a function is called a *strict Lyapunov function* if it meets the above conditions, with the exception of having a strict inequality in condition 3) (i.e., $V(x) < 0, \forall x \in \mathbb{R}^N$).

Remark 2: When it is possible to find a weak Lyapunov function for a given dynamical system, the first theorem of Lyapunov [42] guarantees that any fixed point of the system is stable in the sense of (13). However, to show global convergence of a system, the second theorem of Lyapunov requires a strict Lyapunov function.

One can check that the objective function in (1) is a weak Lyapunov function for (3), thus guaranteeing that the LCA is Lyapunov stable (which improves on the stability result obtained in [14]). However, it is not a strict Lyapunov function. Indeed, (1) only depends on the active nodes, meaning that it could stop decreasing while subthreshold nodes are still evolving. Continued evolution of the inactive nodes could cause a node in $\mathcal{C}$ to become active, thereby causing the objective function to start decreasing again. As a consequence, condition 3) in Definition 3 is not satisfied with strict inequality, and the standard approach of using the objective function as a Lyapunov function is not sufficient to show global convergence of the LCA. To show global convergence of the system to a fixed point, it is necessary to account for the dynamics on both active and inactive nodes.

Our main convergence theorem guarantees the global quasi-convergence of the LCA toward a set of critical points of the objective function. In the case of isolated critical points, this implies global convergence.

Theorem 1: The LCA system defined in (3), with an activation function of the form (5) satisfying conditions (6):

1. has fixed points that are critical points of the objective function defined in (1);
2. has globally quasi-convergent outputs;
3. is globally convergent, provided that the critical points of (1) are isolated.

Note that part 2 of the above theorem only relates to the output variables, while part 3 is a much stronger condition on the entire dynamical system (including subthreshold states). In the highly relevant case of convex objective functions with unique minima (e.g., BPDN), part 3 of the above theorem applies directly and we have the strongest possible notion of convergence. In fact, for most dictionaries (e.g., random Gaussian), the minimum of the ℓ_1 -minimization objective function (2) is unique [43]. Recent results on subgradient dynamical systems [26], [44] lead us to believe that part 3 could still hold in the case where the fixed points of the system are not isolated and there exists a subspace of solutions to (1), but this conjecture is beyond the scope of this paper. Section V-A illustrates the convergence behavior of the LCA in simulation. The proof of this theorem is in Appendix A, and relies on generalized notions of a subgradient due to the nonsmooth nature of the objective.

B. Convergence in a Finite Number of Switches

In the theorem below, we strengthen the previous result to establish that, under some mild conditions, the support of the final solution is reached with a finite number of switches. To prove this result, it is sufficient to assume that no node in the solution lies exactly on the threshold λ . This assumption precludes unwanted infinite oscillation behavior on the boundaries, known as Zeno behavior. In other words, we assume that there exists a margin ($r > 0$) above and below the threshold which contains no node in u^* . One would expect this condition to hold with near-certainty for any signal that was not pathologically constructed.

Theorem 2: If (3) converges to a fixed point u^* such that there exists $r > 0$

$$\begin{aligned} \forall n \in \Gamma_*, \quad & \left| u_n^* \right| \geq \lambda + r \\ \forall n \in \Gamma_*^c, \quad & \left| u_n^* \right| \leq \lambda - r \end{aligned}$$

then the system converges after a finite number of switches. This implies that the neural network recovers the support of the solution u^* in finite time. Section V-A explores the number of switches during the convergence of the LCA in simulation.

Proof: Let $\mathcal{A}$ be the set of active nodes in u^* . By contradiction, assume that the sequence of switching times $\{t_k\}_{k \in \mathbb{N}}$ is infinite. Since the LCA converges to u^* , we have

$$u(t_k) \xrightarrow[k \rightarrow +\infty]{} u^*.$$

As a consequence, for $r > 0$, there exists $K \in \mathbb{N}$ such that $\forall k \geq K, \|u(t_k) - u^*\|_2 < r$. To begin, we show that for all $k \geq K$, the state variables $u(t_k)$ are in the subsystem $\mathcal{A}$. For all $k \geq K$, we have two cases.

1. Nodes that are above threshold in u^* are above threshold in $u(t_k)$. Indeed, $\forall n \in \Gamma_*$, we have

$$\begin{aligned} r > |u_n(t_k) - u_n^*| &\geq |u_n^*| - |u_n(t_k)| \geq \lambda + r - |u_n(t_k)| \\ \Rightarrow |u_n(t_k)| &> \lambda. \end{aligned}$$

Moreover, nodes are active with the correct sign; otherwise, we would have

$$\begin{aligned} r > |u_n(t_k) - u_n^*| &= |u_n(t_k)| + |u_n^*| > \lambda + \lambda + r \\ \Rightarrow 0 &> \lambda \end{aligned}$$

which is a contradiction.

2. Nodes that are below threshold in u^* are below threshold in $u(t_k)$. Indeed, $\forall n \in \Gamma_{*}^c$, we have

$$\begin{aligned} |u_n(t_k)| - \lambda &\leq |u_n(t_k)| - |u_n^*| - r \leq |u_n(t_k) - u_n^*| - r < 0 \\ \Rightarrow |u_n(t_k)| &< \lambda. \end{aligned}$$

As a consequence, for all $k \leq K$, $k \neq *$. However, k and $k+1$ must be different to define the switching time t_{k+1} and we reach a contradiction. This proves that after a finite number of switches K , there cannot be any switching out of subsystem $*$.

IV. Exponential Convergence Rate

With Theorem 1 showing that the LCA is globally convergent, the most pressing issue remaining is to determine the convergence rate of the network. Such a bound will be especially important for implementations, which must guarantee solution times. In this section, we show that, under some additional conditions on the problem's specifics [given in (16)], the LCA network converges exponentially fast to a unique fixed point u^* . We also give an analytic characterization of the convergence speed.²

To begin this section, we recall the definition of exponential convergence.

Definition 4: The dynamical system in (10) is *exponentially convergent* to the solution x^* if there exists a constant $c > 0$ such that for any initial point $x(0)$, there exists a constant $\kappa_0 > 0$ [which may depend on $x(0)$] for which the trajectory $x(t)$ of the system satisfies

$$\|x(t) - x^*\| \leq \kappa_0 e^{-ct} \quad \forall t \geq 0.$$

The constant c is referred to as *convergence speed* of the system.

In order to state the main theorem of this section, we define the two following quantities. The first constant is denoted by α and provides a bound on the derivative of $f(\cdot)$ in (5)

$$\forall t \geq 0, \forall n=1, \dots, N \quad \left| f'(u_n(t)) \right| \leq \alpha. \quad (14)$$

Note that the constant α is always well defined since the trajectories $u_n(t)$ are bounded and the function $f(\cdot)$ is continuous. The second constant, denoted by β , is the smallest positive

²We reintroduce the time constant τ in this discussion of the LCA, since it appears in the expression for the convergence speed.

constant such that for any active set Γ visited by the algorithm and any vector x in $\mathbb{R}^N$ with active set $\tilde{\Gamma} = \Gamma \cup \Gamma_*$, where Γ_* is the active set of the solution to (1); we have

$$(1 - \delta) \|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1 + \delta) \|x\|_2^2. \quad (15)$$

The constant δ depends on the singular values of the matrix $\Phi_{\tilde{\Gamma}}$ and on the sequence of active sets visited by the system. It may not be well defined for any matrix Φ or any input y . However, in many interesting cases in CS, the constant δ is close to 0 and the dictionary elements are almost orthogonal for any small enough active set [45]. The following theorem shows that this constant directly relates to the convergence speed of the neural network.

Theorem 3: Under conditions (6) on the activation function in (5), and provided that the constants α and β , defined in (14) and (15), respectively, satisfy

$$\delta < \min \left(1, \frac{1}{(2\alpha - 1)^2} \right) \quad (16)$$

the LCA system defined in (3) is globally exponentially convergent to a unique equilibrium, with convergence speed $(1 - \delta)/2$.

If condition (16) is satisfied, the expression given for the convergence speed is positive and thus meaningful. It depends on the eigenvalues of the matrix $\Phi_{\tilde{\Gamma}}^T \Phi_{\tilde{\Gamma}}$, which vary with the active set $\tilde{\Gamma}$. A careful analysis of the sequence of active sets visited by the network is required to obtain a good estimate of δ . Since such a study is application dependent, we do not address this question here. Note that in the very interesting case of the soft-threshold function, $\beta = 1$ and so condition (16) reduces to $\alpha < 1$. The time constant τ of the physical solver implementing the LCA neural network appears in the expression of the speed of convergence. Lowering this time constant means the system will converge faster. Analog systems can have smaller time constants than their digital counterparts, which scale better with the problem size. Section V-B explores the convergence rate bounds for the LCA in simulation.

To establish the expression of the convergence speed, we use the following Lure-Postnikov-type Lyapunov function:

$$E(t) = \frac{1}{2} \|\tilde{u}(t)\|_2^2 + \sum_{n=1}^N \int_0^{\tilde{u}_n(t)} g_n(s) ds \quad (17)$$

where we again redefine the output and state variables in terms of the distance from any arbitrary fixed point u^* as

$$\begin{aligned} \tilde{u}_n(t) &= u_n(t) - u_n^* \\ \tilde{a}_n(t) &= a_n(t) - a_n^* = T_\lambda(\tilde{u}_n(t) + u_n^*) - T_\lambda(u_n^*) \end{aligned} \quad (18)$$

and the function $g_n(s) = T_\lambda(s + u_n^*) - T_\lambda(u_n^*)$, which captures how a perturbation of the internal state affects the output. Because the second term in (17) is nonnegative, we have trivially that this energy function bounds the mean-squared error of the current state values to their steady-state values, i.e., $(1/2) \|\tilde{u}(t)\|_2^2 \leq E(t)$. The theorem will prove that u^* is unique. The properties presented in the following Lemma (see Appendix B for proof) are useful to prove the main result of this section.

Lemma 1: Assume that the activation function (5) satisfies the conditions (6). Then, the set of variables $\tilde{u}$ and $\tilde{a}$ defined in (18) satisfy the following properties:

1. $\text{sign}(\tilde{a}_n) = \text{sign}(\tilde{u}_n)$;
2. $|\tilde{a}_n| \leq \alpha |\tilde{u}_n|$
3. $\tilde{a}_{\mathcal{T}}^T \tilde{a}_{\mathcal{T}} \leq \alpha \tilde{u}_{\mathcal{T}}^T \tilde{a}_{\mathcal{T}} \leq \alpha^2 \tilde{u}_{\mathcal{T}}^T \tilde{u}_{\mathcal{T}}$ for any $\mathcal{T}$ (in particular for $\mathcal{T} = \tilde{\Gamma}$);
4. $\sum_{n=1}^N \int_0^{\tilde{u}_n(t)} g_n(s) ds \leq \tilde{u}^T \tilde{a}$.

Armed with this lemma, we can now give the proof of the main theorem.

Proof of Theorem 3: We begin by noting that since u^* and a^* are the internal state variables and the outputs at a fixed point of (3); we can rewrite the dynamics in terms of the new variables

$$\tau \dot{\tilde{u}}(t) = -\tilde{u}(t) - (\Phi^T \Phi - I) \tilde{a}(t). \quad (19)$$

Using the chain rule, we find the expression for the derivative with respect to time of the energy function (17) as

$$\tau \dot{E}(t) = \tau \frac{dE(\tilde{u}(t))^T}{d\tilde{u}(t)} \frac{d\tilde{u}(t)}{dt} = (\tilde{u}(t) + \tilde{a}(t))^T \tau \dot{\tilde{u}}(t).$$

Using (19), this becomes

$$\tau \dot{E}(t) = -(\tilde{u}(t) + \tilde{a}(t))^T (\tilde{u}(t) + (\Phi^T \Phi - I) \tilde{a}(t)).$$

Using the properties in Lemma 1 leads to the desired expression (and removing the term t to increase readability)

$$\begin{aligned} \tau \dot{E}(t) &= -(\tilde{u} + \tilde{a})^T (\tilde{u} + (\Phi^T \Phi - I) \tilde{a}) \\ &= (\tilde{u} + \tilde{a})^T (-\tilde{u} + \tilde{a} - \Phi^T \Phi \tilde{a}) \\ &= -\tilde{u}^T \tilde{u} + \tilde{a}^T \tilde{a} - \tilde{a}^T \Phi^T \Phi \tilde{a} - \tilde{u}^T \Phi^T \Phi \tilde{a}. \end{aligned}$$

Adding and subtracting the term $(1/2) \tilde{u} \Phi^T \Phi \tilde{u}$ results in

$$\begin{aligned} \tau \dot{E}(t) &= -\tilde{u}^T \tilde{u} + \tilde{a}^T \tilde{a} - \tilde{a}^T \Phi^T \Phi \tilde{a} + \frac{1}{2} (\tilde{u} - \tilde{a})^T \Phi^T \Phi (\tilde{u} - \tilde{a}) - \frac{1}{2} \tilde{u} \Phi^T \Phi \tilde{u} - \frac{1}{2} \tilde{a}^T \Phi^T \Phi \tilde{a} \\ &= -\tilde{u}^T \tilde{u} + \tilde{a}^T \tilde{a} - \frac{3}{2} \|\Phi \tilde{a}\|_2^2 + \frac{1}{2} \|\Phi (\tilde{u} - \tilde{a})\|_2^2 - \frac{1}{2} \|\Phi \tilde{u}\|_2^2. \end{aligned}$$

We separate the vectors into their components onto the active set $\tilde{\Gamma} = \Gamma \cup \Gamma^*$ of $\tilde{a}$ and its complement $\tilde{\Gamma}^c$ as

$$\begin{aligned}
\frac{1}{2} \|\Phi(\tilde{u} - \tilde{a})\|_2^2 - \frac{1}{2} \|\Phi\tilde{u}\|_2^2 &= \frac{1}{2} \left\| \Phi_{\tilde{\Gamma}} \left(\tilde{u}_{\tilde{\Gamma}} - \tilde{a}_{\tilde{\Gamma}} \right) \right\|_2^2 + \frac{1}{2} \left\| \Phi_{\tilde{\Gamma}^c} \tilde{u}_{\tilde{\Gamma}^c} \right\|_2^2 - \frac{1}{2} \|\Phi\tilde{u}\|_2^2 \\
&= \frac{1}{2} \left\| \Phi_{\tilde{\Gamma}} \left(\tilde{u}_{\tilde{\Gamma}} - \tilde{a} \right) \right\|_2^2 - \frac{1}{2} \left\| \Phi_{\tilde{\Gamma}} \tilde{u}_{\tilde{\Gamma}} \right\|_2^2.
\end{aligned}$$

We can now use the definition of in (15)

$$\begin{aligned}
\tau \dot{E}(t) &\leq -\tilde{u}^T \tilde{u} + \tilde{a}^T \tilde{a} - \frac{3}{2} (1 - \delta) \|\tilde{u}\|_2^2 + \frac{1}{2} (1 + \delta) \left\| \tilde{u}_{\tilde{\Gamma}} - \tilde{a} \right\|_2^2 - \frac{1}{2} (1 - \delta) \left\| \tilde{u}_{\tilde{\Gamma}} \right\|_2^2 \\
&= -\frac{1}{2} (1 - \delta) \tilde{u}^T \tilde{u} - (1 - \delta) \tilde{u}^T \tilde{a} - \frac{1}{2} (1 - \delta) \tilde{u}_{\tilde{\Gamma}}^T \tilde{u}_{\tilde{\Gamma}} - 2\delta \tilde{u}_{\tilde{\Gamma}}^T \tilde{a} + 2\delta \tilde{a}^T \tilde{a} - \frac{1}{2} (1 + \delta) \left\| \tilde{u}_{\tilde{\Gamma}^c} \right\|_2^2 \\
&\leq -\frac{1}{2} (1 - \delta) \tilde{u}^T \tilde{u} - (1 - \delta) \tilde{u}^T \tilde{a} - \frac{1}{2} (1 - \delta) \tilde{u}_{\tilde{\Gamma}}^T \tilde{u}_{\tilde{\Gamma}} - 2\delta \tilde{u}_{\tilde{\Gamma}}^T \tilde{a} + 2\delta \tilde{a}^T \tilde{a}.
\end{aligned}$$

Now, using property 3) of Lemma 1, the expression becomes

$$\begin{aligned}
\tau \dot{E}(t) &\leq -\frac{1}{2} (1 - \delta) \|\tilde{u}\|_2^2 - (1 - \delta) \tilde{u}^T \tilde{a} + \left[-\frac{1}{2\alpha} (1 - \delta) - 2\delta + 2\delta\alpha \right] \tilde{u}_{\tilde{\Gamma}}^T \tilde{a} \\
&= -\frac{1}{2} (1 - \delta) \|\tilde{u}\|_2^2 - (1 - \delta) \tilde{u}^T \tilde{a} - \frac{1}{2\alpha} \left[1 - \delta(2\alpha - 1)^2 \right] \tilde{u}_{\tilde{\Gamma}}^T \tilde{a}.
\end{aligned}$$

From (16), we know that $1 - (2 - 1)^2 \geq 0$, and so

$$\tau \dot{E}(t) \leq -\frac{1}{2} (1 - \delta) \|\tilde{u}\|_2^2 - (1 - \delta) \tilde{u}^T \tilde{a}.$$

Finally, using property 4) of Lemma 1 and the fact that $\alpha \geq 1$, we arrive at

$$\begin{aligned}
\tau \dot{E}(t) &\leq -\frac{1}{2} (1 - \delta) \|\tilde{u}\|_2^2 - (1 - \delta) \sum_{n=1}^N \int_0^{\tilde{u}_n(t)} g_n(s) ds \\
&= - (1 - \delta) E(t).
\end{aligned}$$

Knowing already that $E(t)$ is an upper bound on the quantity of interest, we can integrate both sides of the above inequality (i.e., apply Gronwall's inequality) to get a bound on $E(t)$

$$\frac{1}{2} \|u(t) - u^*\|_2^2 \leq E(t) \leq e^{-(1-\delta)t/\tau} E(0) \quad (20)$$

proving that the LCA is exponentially convergent to a unique fixed point u^* . Taking the square root, this expression also shows that the convergence speed in the exponential is $(1 - \delta)/(2\tau)$, as stated in the theorem.

V. Simulations

In this section, we illustrate the previous theoretical results in simulation. Each plot is based on the following canonical sparse approximation problem. We generate a “true” sparse vector $a_0 \in \mathbb{R}^N$, with $N = 512$ and $s = 5$ nonzero entries. We select the locations of the nonzeros uniformly at random and draw their amplitudes from a normal Gaussian distribution (then normalizing to have unit norm). We choose a dictionary Φ as a union of the canonical basis and a sinusoidal basis having dimensions $M \times N$ with $M = 256$. The vector of measurements is $y = \Phi a_0 + \eta \in \mathbb{R}^M$, where η is Gaussian random noise with standard deviation 0.0062. We use an LCA with a soft-threshold activation function, with

the threshold set to $\epsilon = 0.025$ and $u(0) = 0$. We simulate the LCA dynamical equations (3) through a discrete approximation in `MATLAB` with a step size of 0.001 and a solver time constant chosen to be equal to $\tau = 0.01$.

A. Convergence Results

From Theorem 1, the LCA should converge and recover the solution to the sparse approximation problem (1), which has a unique minimizer. Since the signal a_0 to recover is sparse, the outputs of the neural network are expected to converge to a solution close to the initial signal a_0 . Fig. 4 shows the evolution of a few nodes $u_n(t)$ selected at random. We see that both active and inactive nodes converge relatively quickly. Fig. 5 shows the fixed point reached by the LCA system and compares it to the initial signal a_0 and to the solution obtained using a standard digital solver for the same optimization program [8]. The solution reached by the network possesses exactly five nonzero entries which correspond to the nonzero entries in a_0 . The recovered amplitudes are very close to the initial amplitudes (it cannot be exact due to the added measurement noise) and to the ones produced by the reference digital solver. However, the LCA produces a sparse vector while the digital solver returns many small but nonzero entries that would have to be removed by postprocessing.

To illustrate the global convergence behavior, we also ran the LCA for 30 randomly generated initial points. We selected two nodes from the final active set and plotted the trajectories in the space defined by those two nodes. Fig. 6 clearly shows that the solution is attractive for any of those initial points.

To illustrate the number of switches used by the system (see Theorem 2), we generate 1000 sparse vectors a_0 and measurements y and simulate the LCA dynamics. Fig. 7 shows a histogram of the number of switches needed for the system to converge. The figure illustrates that the number of switches before convergence of the neural network is finite and of the order of the dictionary size. This illustrates that this solver takes an efficient path toward the solution.

B. Convergence Rate Results

To illustrate the convergence rate result in Theorem 3, it is necessary to find an expression for the convergence speed $(1 - \epsilon)/\epsilon$ that appears in the exponential term in (20). This term bounds the error squared

$$\frac{1}{2} \|u(t) - u^*\|_2^2$$

which is normalized to have initial value at $t = 0$ of 1 and plotted using a log-scale on Fig. 8. Note that the constant ϵ defined in (15) depends on the sequence of active sets $\mathcal{K}_k$ visited by the system. However, it is very difficult to predict for a given input signal what sequence of active sets the algorithm is going to visit. To estimate this upper bound, we compute the constant ϵ using the matrix $\Phi_{\mathcal{K}^*}$ composed of the dictionary elements that are active in the final solution. The corresponding upper bound on the decay $e^{-(1 - \epsilon)t/\tau}$ is plotted in Fig. 8. During the transient phase, the number of active nodes is actually larger and therefore we expect $1 - \epsilon$ to be smaller than this estimate and the nodes to converge more slowly at some point during the transient phase. As a consequence, we keep track of the largest support visited by the network and compute the corresponding ϵ . This second upper bound $e^{-(1 - \epsilon_{\max})t/\tau}$ on the convergence rate is plotted on the same figure. As expected, the theoretical decay computed with the maximum support visited is an upper bound for the convergence speed. However, it can be seen that this estimate is very pessimistic and that the bound computed with the final active set is a better estimate. This simulation illustrates

that the theoretical exponential convergence appears to capture the essential system behavior.

VI. Conclusion

This paper presented a mathematical analysis of the convergence properties and convergence rate of the LCA, a neural network designed specifically for the challenging sparse approximation problem. Despite a nonsmooth activation function and possibly singular interconnection matrix that prevent the application of existing analysis approaches, we have shown that the system is globally convergent to the optimal solution. In addition, under some mild assumptions on the solution, we have shown that the trajectories follow a reasonable path and reach the final active set in finite time. Finally, under slightly stronger assumptions on the problem specifics (applicable at least in CS recovery problems), we established that the LCA is exponentially convergent and with a convergence rate that depends on problem-specific parameters.

This collection of results and analysis leads us to conclude that performance guarantees can be made for the LCA system that make it plausible for implementation in engineering applications and as a model of biological information processing. Indeed, providing such guarantees makes it easier to justify the expense associated with developing analog VLSI implementations, which could eventually result in significant improvements to the speed and power consumption necessary for real-time signal processing applications. Our future work will concentrate on finding reasonable estimates of the theoretical convergence speed (especially in well-studied special cases such as CS recovery) and further characterizations of the LCA system dynamics that may open up new applications of this system for time-varying input signals.

Acknowledgments

This work was supported in part by the National Science Foundation under Grant CCF-0905346 and the National Institutes of Health under Grant R01-EY019965.

Appendix A

Proof of Theorem 1

Building on the earlier definition of a subgradient to give a notion that nondifferentiable functions can still be well behaved, a function $g: X \rightarrow \mathbb{R}$ (where X is a Banach space) is said to be *regular* [40, Def. 2.3.4] at x in X if:

1. for all X , the usual one-sided directional derivative

$$g'(x;v) = \lim_{t \downarrow 0} \frac{g(x+tv) - g(x)}{t} \quad \text{exists;}$$

2. for all X , $g(x; \cdot) = g^\circ(x; \cdot)$, where $g^\circ(x; \cdot)$ is the generalized directional derivative

$$g^\circ(x;v) = \lim_{\substack{y \rightarrow x \\ t \downarrow 0}} \sup \frac{g(y+tv) - g(y)}{t}.$$

It can easily be seen that, since the function $C(\cdot)$ is defined from $\mathbb{R}$ to $\mathbb{R}$ and is differentiable on $\mathbb{R} \setminus \{0\}$, and the function $T(\cdot)$ is continuous on all $\mathbb{R}$, then $C(\cdot)$ admits left and right derivatives and is clearly regular on $\mathbb{R}$. This implies that $V(\cdot)$ is regular on $\mathbb{R}^N$, and by [40, Prop. 2.3.3] we have that

$$\partial V(a(t)) = -\Phi^T y + \Phi^T \Phi a(t) + \lambda \partial C(a(t)) \quad (21)$$

where $C(a(t)) = [C(a_1(t)), \dots, C(a_N(t))]^T$. We also recall the following result [40, Th. 2.3.9 (iii)], which is a generalization of the chain rule for regular functions.

Lemma 2: Suppose that $V(a): \mathbb{R}^N \rightarrow \mathbb{R}$ is regular in $\mathbb{R}^N$ and that $a(t): [0, +\infty) \rightarrow \mathbb{R}^N$ is strictly differentiable on $[0, +\infty)$. Then, $V(a(t))$ is regular on $\mathbb{R}^N$, and we have

$$\frac{d}{dt} V(a(t)) = \zeta^T \dot{a}(t) \quad \forall \zeta \in \partial V(a(t)). \quad (22)$$

Note that, from this theorem, since $V(a(t))$ is regular, we can choose any element in $\partial V(a(t))$ to compute the time derivative of $V(\cdot)$ along the trajectories of the neural network. Armed with these tools, we proceed with the proof of Theorem 1.

Proof of Theorem 1: Beginning with part 1 of the theorem, we first show that any fixed point of system (3) is a critical point of the objective in (1). From (11), any fixed point u^* of (3) satisfies the relationship $u(t) = 0$. Let Γ^* be the active set at the fixed point u^* . Equations (8) and (9) yield

$$-u_{\Gamma^*}^* + f(u_{\Gamma^*}^*) - \Phi_{\Gamma^*}^T \Phi_{\Gamma^*} a_{\Gamma^*}^* + \Phi_{\Gamma^*}^T y = 0 \quad (23)$$

$$-u_{\Gamma^c}^* - \Phi_{\Gamma^c}^T \Phi_{\Gamma^*} a_{\Gamma^*}^* + \Phi_{\Gamma^c}^T y = 0. \quad (24)$$

According to (12), and using (21), a point a^* is a critical point of $V(\cdot)$ if and only if

$$\Phi^T y - \Phi^T \Phi a^* \in \lambda \partial C(a^*). \quad (25)$$

For the nodes in the active set $n \in \Gamma^*$, $C(a_n)$ admits a usual gradient, as defined in (4), and thus (25) yields

$$\Phi_{\Gamma^*}^T y - \Phi_{\Gamma^*}^T \Phi_{\Gamma^*} a_{\Gamma^*}^* = u_{\Gamma^*}^* - f(u_{\Gamma^*}^*).$$

This is the same as condition (23). For nodes in the inactive set Γ^c , we need to determine the subgradient of $C(\cdot)$ at zero. To do so, note that using (4) and the continuity of the function $f(\cdot)$ at 0 [condition (6a)], we have

$$\lim_{a \downarrow 0} \lambda \nabla C(a) = \lim_{u \downarrow \lambda} u - f(u) = \lambda.$$

As a consequence, we find that

$$\lim_{a \downarrow 0} \nabla C(a) = 1 \quad \text{and} \quad \lim_{a \uparrow 0} \nabla C(a) = -1.$$

From the definition of the subgradient, these limits show that $C(0) = [-1, 1]$. Using this, the condition in (25) restricted to the set of inactive nodes is equivalent to

$$\Phi_{\Gamma_*}^T y - \Phi_{\Gamma_*}^T \Phi_{\Gamma_*} a_{\Gamma_*}^* \in [-\lambda, \lambda].$$

We immediately see that this condition is the same as (24), since by definition of the inactive nodes $u_{\Gamma_*}^* \in [-\lambda, \lambda]$. This shows that the fixed points a^* coincide with the critical points of the objective function (1).

Moving on to establishing the convergence result in part 2 of the theorem, we first note that, from our analysis above, the same reasoning leads to the conclusion that

$$-\dot{u}(t) = u(t) - f(u(t)) + \Phi^T \Phi a(t) - \Phi^T y \in \partial V(a(t)). \quad (26)$$

Using the chain rule (22) with $\dot{a}(t) = -u(t) - V(a(t))$ from (26) and using the expression in (7), we get the relationship

$$\begin{aligned} \frac{dV(a(t))}{dt} &= -\dot{u}(t)^T \dot{a}(t) \\ &= -\dot{u}(t)^T F'(t) \dot{u}(t) \\ &= -\sum_{n \in \Gamma} f'(u_n(t)) |\dot{u}_n(t)|^2. \end{aligned} \quad (27)$$

This expression is valid for all time $t \geq 0$. In addition, because the function $f(\cdot)$ satisfies (6b), $f(u_n(t)) > 0$, and thus

$$\begin{aligned} \frac{dV(a(t))}{dt} &\leq 0 \quad \text{for all } t \geq 0, \quad \text{and} \\ \frac{dV(a(t))}{dt} &< 0 \quad \text{for all } t \geq 0 \quad \text{s.t.} \quad \|\dot{u}_\Gamma(t)\|_2 \neq 0. \end{aligned}$$

This means that $V(a(t))$ is nonincreasing for all $t \geq 0$. Since $V(a(t))$ is continuous, bounded below by zero, and nonincreasing, $V(a(t))$ converges to a constant value V^* , and its time derivative $\dot{V}(a(t))$ tends to zero as $t \rightarrow \infty$. Using (27) and (6b), we conclude that

$\lim_{t \rightarrow \infty} \|\dot{u}_\Gamma(t)\|_2 = 0$. As a consequence, we also have $\lim_{t \rightarrow \infty} \|\dot{a}(t)\|_2 = 0$, so the outputs converge to the set $E = \{a : \text{s.t. } \dot{a}(t) = 0\}$, and the LCA outputs are quasi-convergent.

Moving on to part 3 of the theorem, we assume that the critical points of (1) are isolated. We need to show that both active and inactive nodes converge to a single fixed point. From part 2, we know that, since the nodes converge to E , after some time, the active nodes will be within a ball of radius R around one element $a^* \in E$. However, since critical points of (1) are isolated, there exists a ball $\mathcal{B}_\epsilon(a^*)$ of radius $\epsilon > 0$ around a^* that does not contain any critical point

$$a^* \in E \quad \text{and} \quad \forall a \in \mathcal{B}_\epsilon(a^*), a \neq a^* \Rightarrow a \notin E.$$

Since the system is stable, we know that once the trajectory gets close enough (within a ball of radius R) to one element in E , it cannot leave a ball of radius around this fixed point [see (13)]. As a consequence, the outputs remain within the ball $\mathcal{B}_\epsilon(a^*)$, which contains only the fixed point a^* . This proves that the active nodes converge to the point a^* in E .

$$\lim_{t \rightarrow +\infty} a(t) = a^*. \quad (28)$$

This implies that the ODE (3) can be written in terms of the distance $\tilde{a}(t) = a(t) - a^*$ of the outputs from the solution

$$\dot{u}(t) = -u(t) - \Phi^T \Phi a^* + \Phi^T y + a^* - \Phi^T \Phi \tilde{a}(t) + \tilde{a}(t).$$

We let

$$u^* = -\Phi^T \Phi a^* + \Phi^T y + a^*$$

and rewrite the ODE (3) as

$$\dot{u}(t) = -u(t) + u^* - (\Phi^T \Phi - I) \tilde{a}(t).$$

Solving this ODE for all $t \geq 0$ yields

$$u(t) = u^* + e^{-t} (u(0) - u^*) + e^{-t} \int_0^t e^s (\Phi^T \Phi - I) \tilde{a}(s) ds.$$

While it is difficult to say anything directly about the trajectory of the system, it is helpful to consider a surrogate trajectory that is a straight line in the state-space $u^* + e^{-t} (u(0) - u^*)$. This linear path obviously converges to the fixed point u^* , and if we are able to show that the actual trajectory $u(t)$ asymptotically approaches this idealized linear path, we will have established that the system converges to u^* . To take this approach, we examine the quantity

$$h(t) = u(t) - u^* - e^{-t} (u(0) - u^*)$$

which is the deviation from the linear path. Consider the norm of this deviation

$$\begin{aligned} \|h(t)\|_2 &= \|u(t) - u^* - e^{-t} (u(0) - u^*)\|_2 \\ &= \left\| e^{-t} \int_0^t e^s (\Phi^T \Phi - I) \tilde{a}(s) ds \right\|_2 \\ &\leq e^{-t} \left\| \Phi^T \Phi - I \right\|_2 \int_0^t e^s \|\tilde{a}(s)\|_2 ds. \end{aligned}$$

To show convergence to zero, we split the integral into two parts. Since $\tilde{a}(t) \xrightarrow{t \rightarrow +\infty} 0$, then for any $\tilde{\epsilon} > 0$, there exists a time $t_c \geq 0$ such that $\forall t \geq t_c, \|\tilde{a}(t)\|_2 \leq \tilde{\epsilon}$. Moreover, since $\|\tilde{a}(t)\|_2$ is continuous and goes zero as t goes to infinity, it admits a maximum $\mu, \forall t \in \mathbb{R}$. This yields, for all $t \geq 2t_c$

$$\begin{aligned} \|h(t)\|_2 &\leq e^{-t} \left\| \Phi^T \Phi - I \right\|_2 \mu \int_0^{t_c} e^s ds + e^{-t} \left\| \Phi^T \Phi - I \right\|_2 \tilde{\epsilon} \int_{t_c}^t e^s ds \\ &\leq \left\| \Phi^T \Phi - I \right\|_2 \mu [e^{-t+t_c} - e^{-t}] + \left\| \Phi^T \Phi - I \right\|_2 \tilde{\epsilon} [1 - e^{-t+t_c}] \\ &\leq \left\| \Phi^T \Phi - I \right\|_2 \mu [e^{-t/2} - e^{-t}] + \left\| \Phi^T \Phi - I \right\|_2 \tilde{\epsilon}. \end{aligned}$$

Since the left term converges to 0 and $\tilde{\epsilon}$ can be chosen arbitrarily small, this shows that the trajectory $u(t)$ converges to the trajectory $u^* + e^{-t}(u(0) - u^*)$ as t goes to infinity, and thus we can conclude that $u(t) \xrightarrow{t \rightarrow +\infty} u^*$.

Because we have shown separately that both the active and inactive nodes converge for any initial state, it concludes our proof that the system is globally convergent.

Appendix B

Proof of Lemma 1

Proof: Each of the four cases will be treated separately.

1. For any $\tilde{u}_n \in \mathbb{R}$, let $z_n = \text{sign}(\tilde{u}_n)$. Since the thresholding function is nondecreasing (6b), and $T(-u_n) = -T(u_n)$ from (6a), we have

$$\begin{aligned} z_n = \text{sign}(\tilde{u}_n) &\Rightarrow 0 \leq z_n \tilde{u}_n \\ &\Rightarrow z_n u_n^* \leq z_n \tilde{u}_n + z_n u_n^* \\ &\Rightarrow T_\lambda(z_n u_n^*) \leq T_\lambda(z_n \tilde{u}_n + z_n u_n^*) \\ &\Rightarrow z_n T_\lambda(u_n^*) \leq z_n T_\lambda(\tilde{u}_n + u_n^*) \\ &\Rightarrow 0 \leq z_n [T_\lambda(\tilde{u}_n + u_n^*) - T_\lambda(u_n^*)] \\ &\Rightarrow 0 \leq z_n \tilde{a}_n \\ &\Rightarrow \text{sign}(\tilde{a}_n) = z_n = \text{sign}(\tilde{u}_n). \end{aligned}$$

2. We can separate this proof into four cases. If $|\tilde{u}_n + u_n^*| \leq \lambda$ and $|u_n^*| \leq \lambda$, we have $|\tilde{a}_n| = |T_\lambda(\tilde{u}_n + u_n^*) - T_\lambda(u_n^*)| = 0 \leq \alpha |\tilde{u}_n|$. If $|\tilde{u}_n + u_n^*| \leq \lambda$ and $|u_n^*| > \lambda$, according to the mean value theorem, since the function $f(\cdot)$ is continuous on $[\text{sign}(u_n^*)\lambda, u_n^*]$, differentiable on $(\text{sign}(u_n^*)\lambda, u_n^*)$ and $f(\text{sign}(u_n^*)\lambda) = 0$, there exist $c \in (\text{sign}(u_n^*)\lambda, u_n^*)$ such that

$$\begin{aligned} |\tilde{a}_n| &= |f(u_n^*)| = |f'(c)(u_n^* - \text{sign}(u_n^*)\lambda)| \\ &= |f'(c)|(|u_n^*| - \lambda) \quad (\text{since } |u_n^*| > \lambda) \\ &\leq \alpha(|u_n^*| - \lambda) \\ &\leq \alpha|\tilde{u}_n| \quad (\text{since } |u_n^*| - |\tilde{u}_n| \leq |\tilde{u}_n + u_n^*| \leq \lambda). \end{aligned}$$

If $|\tilde{u}_n + u_n^*| > \lambda$ and $|u_n^*| \leq \lambda$, according to the mean value theorem, there exists $c \in (\text{sign}(\tilde{u}_n + u_n^*)\lambda, \tilde{u}_n + u_n^*)$ such that

$$\begin{aligned} |\tilde{a}_n| &= |f(\tilde{u}_n + u_n^*)| \\ &= |f'(c)(\tilde{u}_n + u_n^* - \text{sign}(\tilde{u}_n + u_n^*)\lambda)| \\ &= |f'(c)|(|\tilde{u}_n + u_n^*| - \lambda) \\ &\leq \alpha(|\tilde{u}_n| + |u_n^*| - \lambda) \leq \alpha|\tilde{u}_n|. \end{aligned}$$

Finally, if $|\tilde{u}_n + u_n^*| > \lambda$ and $|u_n^*| > \lambda$, according to the mean value theorem, there exists $c \in (u_n^*, \tilde{u}_n + u_n^*)$ such that $|\tilde{a}_n| = |f(\tilde{u}_n + u_n^*) - f(u_n^*)| = |f'(c)\tilde{u}_n| \leq \alpha|\tilde{u}_n|$.

3. Using properties 1) and 2), we have

$$\begin{aligned}
\tilde{a}_{\mathcal{T}}^T \tilde{a}_{\mathcal{T}} &= \sum_{n \in \mathcal{T}} \tilde{a}_n \tilde{a}_n = \sum_{n \in \mathcal{T}} |\tilde{a}_n| |\tilde{a}_n| \\
&\leq \sum_{n \in \mathcal{T}} \alpha |\tilde{u}_n| |\tilde{a}_n| = \alpha \sum_{n \in \mathcal{T}} \tilde{u}_n \tilde{a}_n = \alpha \tilde{u}_{\mathcal{T}}^T \tilde{a}_{\mathcal{T}} \\
&\leq \sum_{n \in \mathcal{T}} \alpha |\tilde{u}_n| \alpha |\tilde{u}_n| = \alpha^2 \sum_{n \in \mathcal{T}} \tilde{u}_n \tilde{u}_n = \alpha^2 \tilde{u}_{\mathcal{T}}^T \tilde{u}_{\mathcal{T}}.
\end{aligned}$$

4. Since the soft-threshold function is nondecreasing, then the function $g_n(s) = T_{\lambda}(s + u_n^*) - T_{\lambda}(u_n^*)$ is also nondecreasing. Moreover, $g_n(0) = 0$ and $g_n(\tilde{u}_n(t)) = \tilde{a}_n(t)$. As a consequence, we can bound the integral by

$$\int_0^{\tilde{u}_n(t)} g_n(s) ds \leq (g_n(\tilde{u}_n(t)) - g_n(0)) \tilde{u}_n(t) = \tilde{a}_n \tilde{u}_n.$$

As a consequence

$$\sum_{n=1}^N \int_0^{\tilde{u}_n(t)} g_n(s) ds \leq \sum_{n=1}^N \tilde{a}_n \tilde{u}_n = \tilde{a}^T \tilde{u}.$$

Biography

Aurèle Balavoine (S'09) received the B.S. degree from Ecole Supérieure d'Electricité, Gif-sur-Yvette, France, in 2008, and the M.S. degree from the Georgia Institute of Technology, Atlanta, in 2009. She is currently pursuing the Ph.D. degree with the School of Electrical and Computer Engineering, Georgia Institute of Technology.

Her current research interests include sparsity-based signal processing, mathematical optimization, dynamical systems, and neural networks.

Justin Romberg (S'98–M'03–SM'12) received the B.S.E.E., M.S., and Ph.D. degrees from Rice University, Houston, TX, in 1997, 1999, and 2004, respectively.

He was a Researcher with Xerox PARC, Palo Alto, CA, in 2000, a Visitor with the Laboratoire Jacques-Louis, Lions, Paris, in 2003, and a Fellow with UCLA's Institute for Pure and Applied Mathematics, Los Angeles, CA, in 2004. He was a Post-Doctoral Scholar with the Department of Applied and Computational Mathematics, California Institute of Technology, Pasadena, from 2003 to 2006. He joined the faculty of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, in 2006, and is currently an Associate Professor and a member of the Center for Signal and Image Processing.

Dr. Romberg was a recipient of ONR Young Investigator Award in 2008 and the PECASE Award and a Packard Fellowship in 2009. He was named a Rice University Outstanding Young Engineering Alumnus in 2010. From 2008 to 2011, he was an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY.

Christopher J. Rozell (S'00–M'09–SM'12) received the B.S.E. degree in computer engineering and the B.F.A. degree in performing arts technology (music technology) in 2000 from the University of Michigan, Ann Arbor, and the M.S. and Ph.D. degrees in electrical engineering from Rice University, Houston, TX, in 2002 and 2007, respectively.

He was a Texas Instruments Distinguished Graduate Fellow with Rice University, and following graduate school was a Post-Doctoral Research Fellow with the Redwood Center

for Theoretical Neurosciences, University of California, Berkeley. He joined the faculty at the Georgia Institute of Technology, Atlanta, in 2008, where he is currently an Assistant Professor with the School of Electrical and Computer Engineering, and a member of the Laboratory for Neuroengineering, and the Center for Signal and Image Processing. His current research interests include constrained sensing systems, sparse representations, statistical signal processing, and computational neuroscience.

REFERENCES

1. Balavoine A, Rozell CJ, Romberg J. Global convergence of the locally competitive algorithm. *Proc. IEEE Digital Signal Process. Workshop.* Jan.2011 :431–436.
2. Elad M, Figueiredo MAT, Ma Y. On the role of sparse and redundant representations in image processing. *Proc. IEEE.* Jun; 2010 98(6):972–982.
3. Olshausen BA, Field DJ. Sparse coding of sensory inputs. *Current Opinion Neurobiol.* Aug; 2004 14(4):481–487.
4. Olshausen BA, Field DJ. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature.* Jun.1996 381(6583):607–609. [PubMed: 8637596]
5. Olshausen BA, Field DJ. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vis. Res.* Dec.1997 37(23):3311–3325. [PubMed: 9425546]
6. Candes EJ, Romberg J, Tao T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory.* Feb.2006 52(2):489–509.
7. Donoho DL. Compressed sensing. *IEEE Trans. Inf. Theory.* Apr.2006 52(4):1289–1306.
8. Kim S-J, Koh K, Lustig M, Boyd S, Gorinevsky D. An interior-point method for large-scale L1-regularized least squares. *IEEE J. Sel. Topics Signal Process.* Dec.2007 1(4):606–617.
9. Figueiredo MAT, Nowak RD, Wright SJ. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE J. Sel. Topics Signal Process.* Dec.2007 1(4):586–598.
10. Daubechies I, Defrise M, De Mol C. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Commun. Pure Appl. Math.* Aug.2004 57(11):1413–1457.
11. Cichocki, A., Unbehauen, R. *Neural Networks for Optimization and Signal Processing.* Wiley; New York: 1993.
12. Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proc. Nat. Acad. Sci.* Apr.1982 79(8):2554–2558. [PubMed: 6953413]
13. Twigg C, Hasler P. Configurable analog signal processing. *Digital Signal Process.* Dec; 2009 19(6):904–922.
14. Rozell CJ, Johnson DH, Baraniuk RG, Olshausen BA. Sparse coding via thresholding and local competition in neural circuits. *Neural Comput.* Oct.2008 20(10):2526–2563. [PubMed: 18439138]
15. Hespanha JP. Uniform stability of switched linear systems: Extensions of LaSalle's invariance principle. *IEEE Trans. Autom. Control.* Apr.2004 49(4):470–482.
16. Liu S, Wang J. A simplified dual neural network for quadratic programming with its KWTa application. *IEEE Trans. Neural Netw.* Nov.2006 17(6):1500–1510. [PubMed: 17131664]
17. Yang H, Dillon TS. Exponential stability and oscillation of Hopfield graded response neural network. *IEEE Trans. Neural Netw.* Sep.1994 5(5):719–729. [PubMed: 18267846]
18. Liang X-B, Wang J. A recurrent neural network for nonlinear optimization with a continuously differentiable objective function and bound constraints. *IEEE Trans. Neural Netw.* Nov.2000 11(6):1251–1262. [PubMed: 18249851]
19. Forti M, Tesi A. Absolute stability of analytic neural networks: An approach based on finite trajectory length. *IEEE Trans. Circuits Syst. I, Reg. Papers.* Dec.2004 51(12):2460–2469.
20. Lu W, Wang J. Convergence analysis of a class of nonsmooth gradient systems. *IEEE Trans. Circuits Syst. I, Reg. Papers.* Dec.2008 55(11):3514–3527.
21. Xia Y, Feng G, Wang J. A recurrent neural network with exponential convergence for solving convex quadratic program and related linear piecewise equations. *Neural Netw.* Sep.2004 17(7): 1003–1015. [PubMed: 15312842]

22. Ferreira LV, Kaszkurewicz E, Bhaya A. Support vector classifiers via gradient systems with discontinuous righthand sides. *Neural Netw.* Dec.2006 19(10):1612–1623. [PubMed: 17011165]
23. Forti M, Tesi A. New conditions for global stability of neural networks with application to linear and quadratic programming problems. *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.* Jul. 1995 42(7):354–366.
24. Li L, Huang L. Dynamical behaviors of a class of recurrent neural networks with discontinuous neuron activations. *Appl. Math. Model.* Dec.2009 33(12):4326–4336.
25. Forti M, Nistri P, Quincampoix M. Generalized neural network for nonsmooth nonlinear programming problems. *IEEE Trans. Circuits Syst. I, Reg. Papers.* Sep.2004 51(9):1741–1754.
26. Forti M, Nistri P, Quincampoix M. Convergence of neural networks for programming problems via a nonsmooth Łojasiewicz inequality. *IEEE Trans. Neural Netw.* Nov.2006 17(6):1471–1486. [PubMed: 17131662]
27. Xue X, Bian W. Subgradient-based neural networks for nonsmooth convex optimization problems. *IEEE Trans. Circuits Syst. I, Reg. Papers.* Sep.2008 55(8):2378–2391.
28. Liu QS, Wang J. Finite-time convergent recurrent neural network with a hard-limiting activation function for constrained optimization with piecewise-linear objective functions. *IEEE Trans. Neural Netw.* Apr.2011 22(4):601–613. [PubMed: 21402513]
29. Lin H, Antsaklis PJ. Stability and stabilizability of switched linear systems: A survey of recent results. *IEEE Trans. Autom. Control.* Feb.2009 54(2):308–322.
30. Maass W. On the computational power of winner-take-all. *Neural Comput.* 2000; 12(11):2519–2535. [PubMed: 11110125]
31. Natarajan BK. Sparse approximate solutions to linear systems. *SIAM J. Comput.* 1995; 24(2):227–234.
32. Chen SS, Donoho DL, Saunders MA. Atomic decomposition by basis pursuit. *SIAM Rev.* Mar. 2001 43(1):129–159.
33. Donoho DL, Tanner J. Sparse nonnegative solutions of under-determined linear equations by linear programming. *Proc. Nat. Acad. Sci. United States Amer.* 2005; 102(27):9446–9451.
34. Knill DC, Pouget R. The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends Neurosci.* 2004; 27(12):712–719. [PubMed: 15541511]
35. Rehn M, Sommer FT. A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. *J. Comput. Neurosci.* Apr.2007 22(2):135–146. [PubMed: 17053994]
36. Smith EC, Lewicki MS. Efficient auditory coding. *Nature.* Feb.2006 439(7079):978–982. [PubMed: 16495999]
37. Fischer S, Redondo R, Perrinet L, Cristóbal G. Sparse approximation of images inspired from the functional architecture of the primary visual areas. *EURASIP J. Adv. Signal Process.* 2007; 2007(1):1–17.
38. Perrinet L. Feature detection using spikes: The greedy approach. *J. Physiol. (Paris).* Jul.2004 98(4–6):530–539. [PubMed: 16310348]
39. Bioucas-Dias JM, Figueiredo MAT. A new TwIST: Two-step iterative shrinkage/thresholding algorithms for image restoration. *IEEE Trans. Image Process.* Dec.2007 16(12):2992–3004. [PubMed: 18092598]
40. Clarke, FH. *Optimization and Nonsmooth Analysis.* SIAM; Philadelphia, PA: Jan.. 1987
41. DeCarlo RA, Branicky MS, Pettersson S, Lennartson B. Perspectives and results on the stability and stabilizability of hybrid systems. *Proc. IEEE.* Jul.2000 88(7):1069–1082.
42. Bacciotti, A.; Rosier, L. *Liapunov Functions and Stability in Control Theory.* Springer-Verlag; New York: 2005.
43. Fuchs J-J. On sparse representations in arbitrary redundant bases. *IEEE Trans. Inf. Theory.* Jun. 2004 50(6):1341–1344.
44. Bolte J, Daniilidis A, Lewis A. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM J. Optim.* 2007; 17(4):1205–1223.
45. Candes EJ, Romberg JK, Tao T. Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pure Appl. Math.* Aug.2006 59(8):1207–1223.

Fig. 1. Block diagram of the LCA [14], which is a Hopfield-style network architecture for solving sparse approximation problems.

Fig. 2. Soft-thresholding activation function. When this is used for the thresholding function $T(\cdot)$, the LCA solves the popular BPDN optimization problem used in many sparse approximation applications.

Fig. 3. Generic activation function satisfying conditions (6). The area in gray represents where the activation function must lie in order to satisfy conditions (6b) and (6c).

Fig. 4.

Plot of the evolution with respect to time of several LCA nodes $u_k(t)$. The plain lines correspond to nodes that are active in the final solution and the dashed lines correspond to nodes that are inactive in the final solution.

Fig. 5.

Output a^* of the LCA after convergence. Only nonzero elements are plotted. The fixed point reached by the system is very close to the initial sparse vector used to create the measurement vector (it cannot be exact due to noise). The solution is also very close to a standard digital solver [8] run using the same inputs. Note that the LCA produces many coefficients that are exactly zero (therefore not plotted).

Fig. 6.
Trajectories of $u(t)$ in the plane defined by two nodes chosen randomly from the active set.
Trajectories are shown for 30 random initial states.

Fig. 7.
Histogram (in percentage) of the number of switches the LCA requires before convergence over 1000 trials.

Fig. 8.

Convergence behavior of $(1/2) \|u(t) - u^*\|_2^2$. The dashed line shows the theoretical decay in (20) with $*$ computed by using the final solution support. The dash-dot line shows the theoretical decay with $_{\max}$ computed on the largest support visited. While both estimates are showing theoretically correct behavior, the estimate of the rate based on $*$ is more empirically accurate than the conservative estimate based on $_{\max}$.