# Convergence Behavior of Belief Propagation: Estimating Regions of Attraction via Lyapunov Functions

**Harald Leisenberger**[1] **Christian Knoll**[1] **Richard Seeber**[2] **Franz Pernkopf**[1]

[1]Signal Processing and Speech Communication Laboratory, Graz University of Technology, Austria
[2]Christian Doppler Laboratory for Model Based Control of Complex Test Bed Systems, Institute of Automation and Control, Graz University of Technology, Austria



## Abstract

In this work, we estimate the regions of attraction for belief propagation. This extends existing stability analysis and provides initial message values for which belief propagation is guaranteed to converge. Our approach utilizes the theory of Lyapunov functions that, however, does not readily yield useful regions of attraction. Therefore, we utilize polynomial sum-of-squares relaxations and provide an algorithm that computes valid Lyapunov functions. This admits a novel way of studying the solution space of belief propagation. Finally, we apply our approach to small-scale models and discuss the effect of the potentials on the regions of attraction.

## 1 INTRODUCTION

Belief propagation (BP) estimates the marginals of an arbitrary joint distribution represented by a graphical model. For tree-structured graphs, BP efficiently computes the exact marginals (Pearl, 1988). While BP often provides accurate estimates on loopy graphs as well, it is only guaranteed to converge on graphs with a single loop (Weiss, 2000).[1]

On graphs with multiple loops, BP may even fail to converge or exhibit multiple fixed points. Sufficient conditions for convergence of BP, however, are limited to models with a unique fixed point (Ihler et al., 2005; Mooij and Kappen, 2007). Alternatively, several approaches exist that successfully improve the convergence behavior of BP; for example by stabilizing BP (Murphy et al., 1999) or by using flexible scheduling schemes (Elidan et al., 2006).

Recently it has been shown that the accuracy may vary considerably between different fixed points (Knoll and Pernkopf, 2019). As the initialization largely determines if and to which fixed point BP converges, sophisticated initialization strategies can improve the performance of BP; for certain model classes, it is even possible to enforce convergence to the best possible fixed point (Knoll et al., 2018; Koehler, 2019). In general, however, the overall impact of the message initialization is not well understood.

In this paper, we provide the theoretical framework for analyzing the role of the initialization. More specifically, we estimate sets of initial messages for which BP is guaranteed to converge. We refer to these sets as *regions of attraction* (ROAs). We emphasize that our approach is not limited to models with a unique fixed point but readily applies to models with multiple fixed points as well.[2] The shape of the ROAs, however, may be complex and intractable to be determined in general. Thus, we are content with computing analytically well-behaved inner bounds of the ROAs.

To estimate valid ROAs, we utilize Lyapunov functions (LFs). LFs are frequently used in control theory for studying stability properties of both linear and nonlinear systems (Tan and Packard, 2004). We consider BP as a nonlinear system and reformulate it to allow for an estimation of the ROAs via LFs. Still, it remains challenging to construct good LFs; we thus rely on the sum-of-squares (SOS) method (Papachristodoulou and Prajna, 2005) for computing the LFs by solving a semidefinite program (SDP). Although this provides ROA estimates reliably, these are often too conservative. Therefore, we explain how to optimize the ROA estimates for BP in an iterative manner.

After establishing the theoretical framework, we estimate the ROAs for a range of models. We focus only on well-understood models that often exhibit multiple fixed points. This allows us to provide accessible experiments for which it is evident how to visualize and interpret the results.

[1]Sometimes, the term loopy belief propagation emphasizes the existence of loops. We will make no such distinction.

[2]One requirement of our method is the knowledge of one or more fixed points. This implies that it cannot readily solve inference tasks in practice by, e.g., selecting good initial values for the messages. It can therefore be utilized for analyzing the stability properties of known fixed points only.

*Accepted for the 37th Conference on Uncertainty in Artificial Intelligence* (UAI 2021).

This paper is structured as follows: Sec. 2 introduces probabilistic graphical models and BP. In Sec. 3 we explain how to estimate ROAs with LFs, before we adapt the SOS method to BP in Sec. 4. In Sec. 5 we present our experimental results on the ROAs of BP. We conclude the paper in Sec. 6.

## 2 BACKGROUND

In this section, we introduce probabilistic graphical models and specify the models to be considered in this work. We then introduce the BP algorithm and summarize important properties regarding its convergence behavior.

### 2.1 PROBABILISTIC GRAPHICAL MODELS

Let $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ be an undirected graph, with $\mathbf{X} = \{1, \dots, N\}$ being the set of nodes and $\mathbf{E}$ being the set of undirected edges. Two nodes $i$ and $j$ are joined by an edge if $(i, j) \in \mathbf{E}$. We account for each edge only once, i.e., $(i, j) = (j, i)$. We further use $N(i)$ to denote the set of neighbors of $i$, i.e., the set of nodes $j$ that are joined to $i$ by an edge $(i, j) \in \mathbf{E}$. Then, the degree $d_i$ of a node is the number of its neighbors $|N(i)|$.

A probabilistic graphical model $\mathcal{U} = (\mathcal{G}, \Psi)$ associates a random variable (RV) $X_i$ to each node $i \in \mathbf{X}$ and is further defined by a set of potentials $\{\Phi_1(\mathbf{x}_1), \dots, \Phi_K(\mathbf{x}_K)\}$. With a slight abuse of notation, we will simply refer to $X_i$ as $i$ because of the inherent connection between nodes and RVs. In this work, we consider pairwise models where the potentials are defined over either one variable (i.e., singleton potentials $\Phi_i(x_i)$) or two variables (i.e., pairwise potentials $\Phi_{i,j}(x_i, x_j)$). Consequently, the joint probability distribution $P_{\mathbf{X}}(\mathbf{x})$ factorizes according to

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \prod_{(i,j)\in\mathbf{E}} \Phi_{i,j}(x_i, x_j) \prod_{i=1}^{N} \Phi_i(x_i), \qquad (1)$$

where $\mathcal{Z}$ is the normalization constant (or partition function). We consider the computation of marginal distributions

$$P_{\mathbf{Y}}(\mathbf{y}) = \sum_{x_i:\, i\in\mathbf{X}\setminus\mathbf{Y}} P_{\mathbf{X}}(\mathbf{x}), \qquad (2)$$

where $\mathbf{Y} \subseteq \mathbf{X}$ may be any set of RVs. Note, however, that we are primarily interested in the singleton marginals $P_i(x_i)$. In general, computing the marginals is an intractable problem (Cooper, 1990). We will discuss how BP allows one to approximate the marginals nonetheless in Sec. 2.3.

### 2.2 MODEL SPECIFICATIONS

In this work, we consider binary pairwise models where the random variables $X_i$ take values in $\mathcal{X} = \{-1, 1\}$ (i.e., Ising models). If $P_{\mathbf{X}}(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathcal{X}^N$, we can express the joint distribution in exponential form by writing (1) as

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \exp\Big(\sum_{(i,j)\in\mathbf{E}} J_{ij} x_i x_j + \sum_{i=1}^{N} \theta_i x_i\Big), \qquad (3)$$

where $J_{ij} \in \mathbb{R}$ is the strength (or coupling) of the edge $(i, j)$ and $\theta_i \in \mathbb{R}$ is the local strength (or field) of the variable $i$. The potentials are thus specified by $\Phi_{i,j}(x_i, x_j) = \exp(J_{ij} x_i x_j)$ and $\Phi_i(x_i) = \exp(\theta_i)$.

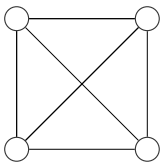

Figure 1: Complete 3-regular graph of size $N = 4$

In our experiments (Sec. 5), we consider models that satisfy the following two assumptions: first, we consider *homogeneous* models where all pairwise potentials are the same (i.e., $J_{ij} = J$) and all local potentials are the same (i.e., $\theta_i = \theta$); second, we consider *d-regular* graphs where all nodes have an equal degree of $d$. We specifically focus on complete graphs of size $N = 4$ (see Fig. 1).[3] These models are relatively well understood and simplify our theoretical considerations. This will be beneficial for interpreting and visualizing our ROA estimates.

### 2.3 BELIEF PROPAGATION (BP)

Belief propagation (BP) is an iterative method to estimate the marginals. Therefore, one computes messages $\mu_{ij}^{(n)}(x_j)$ along all edges, where $ij$ indicates that the message is sent from node $i$ to node $j$, and $(n)$ indicates the current iteration. The messages are updated according to

$$\mu_{ij}^{(n+1)}(x_j) \propto \sum_{x_i\in\mathcal{X}} \Phi_{i,j}(x_i, x_j)\Phi_i(x_i) \prod_{k\in N(i)\setminus j} \mu_{ki}^{(n)}(x_i). \qquad (4)$$

We will discuss how to normalize the messages in Sec. 3.3.

Now, let $\boldsymbol{\mu}^{(n)} = \{\mu_{ij}^{(n)}(x_j), \mu_{ji}^{(n)}(x_i) : (i, j) \in \mathbf{E},\ x_i, x_j \in \mathcal{X}\}$ be the set of all messages at iteration $(n)$ stacked in some arbitrary (but fixed) order and let $\boldsymbol{\mu}^{(n+1)} = \mathrm{BP}(\boldsymbol{\mu}^{(n)})$ be the mapping induced by (4). Here $Z_i \in \mathbb{R}_+$ ensures that all probabilities sum to one. Then, the singleton marginals

[3] Note that this is the smallest possible graph that contains multiple loops. Moreover, this graph behaves the same as the infinite size 3-regular graph and thus qualitatively describes the behavior of grid graphs (Knoll and Pernkopf, 2017).

at iteration $(n)$ can be estimated according to

$$\tilde{P}_i^{(n)}(x_i) = \frac{1}{Z_i}\Phi_i(x_i)\prod_{k\in N(i)}\mu_{ki}^{(n)}(x_i). \tag{5}$$

When updating the messages iteratively, they will – ideally – converge to a fixed point $\boldsymbol{\mu}^\circ$ for which $\boldsymbol{\mu}^\circ = \mathrm{BP}(\boldsymbol{\mu}^\circ)$. In practice, it is often useful to update the messages asynchronously;[4] for our analysis, however, we will only consider plain BP with synchronous message updates.

### 2.4 RELATED WORK

Much of the research on BP's behavior has been driven by its inherent relationship to concepts from statistical physics. In particular, it has been observed that fixed points of BP are in a one-to-one correspondence with stationary points of the so-called Bethe free energy (see Yedidia et al. (2005)).[5] This relationship puts BP on a solid theoretical foundation and provides important insights into its dynamics. In particular, it follows that every stable fixed point is a minimum of the Bethe free energy; note, however, that minima may be unstable as well (Heskes et al., 2003). Moreover, it explains why the presence of loops containing edges with large values $J_{ij}$ is detrimental for BP's performance as this is the main reason for the Bethe free energy being non-convex (see Heskes (2004); Weller et al. (2014)).

Although the Bethe free energy fails to reveal whether a given fixed point is stable and under which conditions BP will converge to it, various approaches aim at minimizing the Bethe free energy directly; the minimization itself may be slow, however, and often leads to suboptimal stationary points (Welling and Teh, 2003; Yuille and Rangarajan, 2003). One can approximate the Bethe free energy with convex functions (Globerson and Jaakkola, 2007; Hazan and Shashua, 2008; Meltzer et al., 2009; Meshi et al., 2009); this, however, often comes at the cost of reduced accuracy. Furthermore, polynomial-time approximation schemes of local minima (Shin, 2012) or the global minimum in attractive models (Weller and Jebara, 2014) exist.

Moreover, it is often necessary to draw from diverse fields of research for understanding the behavior of BP. The stability of BP's fixed points can be analyzed with the Ihara zeta function from graph theory (Watanabe and Fukumizu, 2009). Alternatively, it has been proven beneficial to interpret BP as a *dynamical system*, i.e, as a discrete-time mapping between sets of messages. In this context, it is in principle straightforward to compute all fixed points; yet special considerations are required in practice due to the dramatic growth of the system (Knoll et al., 2017). Moreover, this interpretation allows for analyzing the convergence properties via contraction properties (Mooij and Kappen, 2007) and suggests to analyze the local stability of fixed points (Mooij and Kappen, 2005; Knoll and Pernkopf, 2017).

[4] Utilizing some form of scheduling improves the convergence properties. There is a range of variants available (e.g., Wainwright et al. (2003); Elidan et al. (2006); Sutton and McCallum (2007); Knoll et al. (2015)) with different trade-offs regarding improvement and complexity (see Aksenov et al. (2020) for a comparison).

[5] The Bethe free energy results from approximating the joint distribution via the variational free energy principle (see Mezard and Montanari (2009, Ch.33)).

## 3 LYAPUNOV FUNCTIONS (LFS) AND REGIONS OF ATTRACTION (ROAS)

If one considers BP as a dynamical system, the local stability analysis results from linearizing the system in a fixed point. The behavior of any nonlinear system, however, is far too complex to be reduced to a first-order approximation only.

In this section, we introduce Lyapunov functions (LFs) that enable elaborate stability and convergence analyses. We summarize important properties of LFs and show how to estimate the regions of attraction (ROAs) of BP via LFs. First, we briefly review the fundamentals of stability theory; further details are provided in (Khalil, 2002; Teschl, 2004).

### 3.1 NOTIONS OF STABILITY

Let $F(\boldsymbol{z}) : D \to D$ be a – possibly non-linear – discrete-time system defined on a finite dimensional real domain $D \subseteq \mathbb{R}^m$, inducing an iterative update $\boldsymbol{z}^{(n+1)} = F(\boldsymbol{z}^{(n)})$. Given an initial value $\boldsymbol{z}^{(0)}$, recursive application of $F$ yields a unique sequence $(\boldsymbol{z}^{(0)}, \boldsymbol{z}^{(1)}, \boldsymbol{z}^{(2)}, \ldots)$; we call this sequence the trajectory of $\boldsymbol{z}^{(0)}$. Furthermore, if $F(\boldsymbol{z}^\circ) = \boldsymbol{z}^\circ$, we say that $\boldsymbol{z}^\circ$ is a fixed point (FP) of the system. Our specific interest lies in stable FPs: a FP $\boldsymbol{z}^\circ$ is *locally* stable if a neighborhood $U(\boldsymbol{z}^\circ)$ exists such that the trajectories of initial values inside $U(\boldsymbol{z}^\circ)$ converge to that FP, i.e., if $\lim_{n\to\infty} F^n(\boldsymbol{z}^{(0)}) = \boldsymbol{z}^\circ$. In other words, the trajectory of $\boldsymbol{z}^{(0)}$ is guaranteed to converge if initialized close enough to $\boldsymbol{z}^\circ$. If the trajectory of any starting point in the domain $D$ converges to a FP, then it is referred to as *globally* stable. If no neighborhood exists for which trajectories converge, the FP is unstable.[6]

The simplest approach to assess the stability of a given FP – provided that $F$ is differentiable in $\boldsymbol{z}^\circ$ – is to compute the eigenvalues $\lambda_i$ of the Jacobian $F'(\boldsymbol{z}^\circ)$ whose entries are given by the partial derivatives $(F'(\boldsymbol{z}^\circ))_{u,v} = \frac{\partial F_u(\boldsymbol{z}^\circ)}{\partial z_v}$. If $|\lambda_i| < 1$ for all eigenvalues, the FP is stable. If an eigenvalue exists with $|\lambda_i| > 1$, the FP is unstable. If the largest absolute eigenvalue $\lambda_{\max}$ equals 1, it is impossible to reason about the stability based on the linearization (Teschl, 2004).

[6] In systems theory, a finer distinction between different types of stability is usual, e.g., asymptotic attractivity, etc. For us, it is sufficient to distinguish between global stability, local stability, and instability.

### 3.2 LYAPUNOV FUNCTIONS (LFS)

Computing the eigenvalues of the Jacobian provides a simple and efficient test for stability. This, however, is often not sufficient; one is rather interested in determining which initial values $\boldsymbol{z}^{(0)}$ will converge to $\boldsymbol{z}^\circ$. This convergent region is usually referred to as region of attraction (ROA). In general, ROAs may have complex geometric shapes and cannot be determined computationally. Instead, one can estimate the ROAs using Lyapunov functions (LFs).

The basic idea of LFs is the following: assume we are able to construct an analytical function that (i) attains a global minimum in a FP and (ii) decreases along all trajectories around a FP. Then this implies that all trajectories will continue moving towards the FP and consequently converge to it.[7]

Let $V : D \to \mathbb{R}$ be a continuously differentiable function that is zero at the FP, i.e. $V(\boldsymbol{z}^\circ) = 0$, and let $\Delta V(\boldsymbol{z}) = V(F(\boldsymbol{z})) - V(\boldsymbol{z})$ be the rate of change. Then $V$ is a Lyapunov function (LF) for $\boldsymbol{z}^\circ$ if

$$V(\boldsymbol{z}) > 0 \quad \text{and} \tag{6}$$

$$\Delta V(\boldsymbol{z}) < 0 \tag{7}$$

for all $\boldsymbol{z} \neq \boldsymbol{z}^\circ$ on a neighborhood $E(\boldsymbol{z}^\circ)$.

The existence of a LF (which is, in general, not unique) implies local stability of a FP $\boldsymbol{z}^\circ$ as well (Khalil, 2002). If $E(\boldsymbol{z}^\circ) = D$, then $\boldsymbol{z}^\circ$ is globally stable.

Given a LF, for an arbitrary constant $c > 0$, we define the $c$-sublevel set of $V$ as $\Omega_{V,c} = \{\boldsymbol{z} \in D \,|\, V(\boldsymbol{z}) \leq c\}$ (see Fig. 2a). Whenever there is no ambiguity about the LF, we will omit $V$ in the notation and simply write $\Omega_c$. Sublevel sets of LFs are particularly interesting, as they provide estimates for the ROA. One must be careful, however, in selecting an appropriate sublevel set. Specifically, it has to be ensured that (7) holds on the complete sublevel set $\Omega_c$. This can be illustrated by the following: assume that a trajectory starts in $E(\boldsymbol{z}^\circ) \cap \Omega_{c_2}$, i.e., where (7) is satisfied, but $\Omega_{c_2}$ has been chosen to large to be fully contained in $E(\boldsymbol{z}^\circ)$. Then, in the next time step, the trajectory will jump to another sublevel set $\Omega_{c_1}$ with $c_1 < c_2$ but may leave the set $E(\boldsymbol{z}^\circ)$ and subsequently may never return to $E(\boldsymbol{z}^\circ)$ again (see Fig. 2c). If, however, a trajectory starts in $\Omega_c$ such that $\Omega_c \subseteq E(\boldsymbol{z}^\circ)$, then it will continue iterating to sublevel sets with $c \to 0$ whose infinite intersection is precisely the set $\{\boldsymbol{z}^\circ\}$. Then, $\Omega_c$ is a valid ROA (see Fig. 2b).

As one is usually interested in selecting ROAs as large as possible, one naturally aims for obtaining the largest sublevel set $\Omega_c$ fully contained in $E(\boldsymbol{z}^\circ)$. The size of estimated

[7] The concept of monotonically decreasing functions along system trajectories has its origins in the analysis of physical systems, admitting the interpretation of energy that is dissipated until a stable equilibrium of the system has been established.

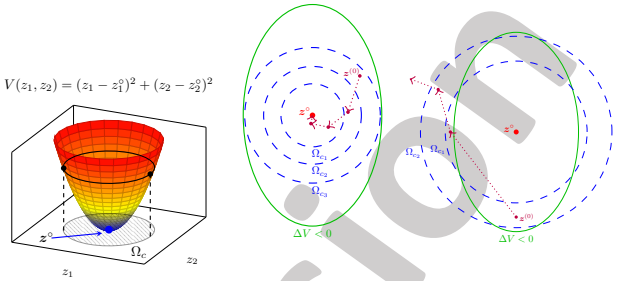


(a) (b) (c)

Figure 2: (a): A Lyapunov function candidate $V$ for FP $\boldsymbol{z}^\circ$ with sublevel set $\Omega_c$. (b): The sublevel set $\Omega_{c_3}$ is selected small enough such that $\Omega_{c_3} \subseteq E(\boldsymbol{z}^\circ)$ and is therefore a valid ROA estimate. (c): The sublevel set $\Omega_{c_2}$ is selected too large and is therefore no valid ROA estimate.

ROAs largely depends on the choice of the corresponding LF. Therefore, one would prefer LFs that are valid on a relatively large neighborhood of the FP. Constructing 'good' LFs is all but a trivial task and has been an intense subject of research (Papachristodoulou and Prajna, 2005). We will address this issue in great detail in Sec. 4.

For the remainder of this section, we focus on LFs that stem from the concept of linearization and that are particularly easy to compute. Note that we will omit the dependence on the FP and refer to the Jacobian $F'(\boldsymbol{z}^\circ)$ as $F'$ for better readability. Given any positive semidefinite (psd.) matrix $Q$, the so-called Lyapunov equation (Khalil, 2002)

$$(F')^T P F' - P = -Q \tag{8}$$

has a unique symmetric psd. solution $P$, whenever $\boldsymbol{z}^\circ$ is a stable FP of $F$. Then the quadratic function $V(\boldsymbol{z}) = (\boldsymbol{z} - \boldsymbol{z}^\circ)^T P (\boldsymbol{z} - \boldsymbol{z}^\circ)$ is a LF for $\boldsymbol{z}^\circ$. Specifically, for $c$ sufficiently small, $\Omega_c$ is a ROA for $\boldsymbol{z}^\circ$.

### 3.3 ROAS OF BELIEF PROPAGATION

Before computing LFs for BP, additional considerations are required. Recall from Sec. 2.3 that BP operates as a discrete-time mapping on the message space. Note that we implicitly assume a synchronous message updating schedule.

In order to enable BP to converge to a FP, it is necessary that all messages remain in a compact subset of the message space. This can be achieved by introducing a normalizing constant $\alpha_{ij}$ (Murphy et al., 1999; Martin et al., 2011). Usually $\alpha_{ij}$ is chosen so that $\sum_{x_j \in \mathfrak{X}} \mu_{ij}^{(n)}(x_j) = 1$. Initializing all messages in $(0, 1)$, it follows that all messages during the whole iteration process take values in $(0, 1)$.

Then, for binary models, the number of independent messages along each direction of an edge reduces to one, since

$\mu_{ij}^{(n)}(+1) = 1 - \mu_{ij}^{(n)}(-1)$. This simple reparametrization is essential for our approach to estimate ROAs in 4.2, which – for computational reasons – requires a reduced number of parameters.[8] Specifically, the dimension of the message space reduces from $4|\mathbf{E}|$ to $2|\mathbf{E}|$.

Putting all our obervations together and identifying the message space $\mathcal{M} = \{\boldsymbol{\mu} \in \mathbb{R}^{2|\mathbf{E}|} \,|\, 0 < \mu_{ij} < 1, (i,j) \in \mathbf{E}\}$ with the unit cube in $\mathbb{R}^{2|\mathbf{E}|}$, we can rewrite the $2|\mathbf{E}|$ update equations (4) induced by the mapping $\mathrm{BP} : \mathcal{M} \to \mathcal{M}$ as

$$\mu_{ij}^{(n+1)} = \alpha_{ij}^{(n)} \big( e^{J_{ij}+\theta_i} \prod_{k \in N(i)\setminus j} \mu_{ki}^{(n)} + e^{-J_{ij}-\theta_i} \prod_{k \in N(i)\setminus j} (1 - \mu_{ki}^{(n)}) \big), \tag{9}$$

$$\alpha_{ij}^{(n)} = \frac{1}{\big(e^{J_{ij}} + e^{-J_{ij}}\big)\big(e^{\theta_i} \prod\limits_{k \in N(i)\setminus j} \mu_{ki}^{(n)} + e^{-\theta_i} \prod\limits_{k \in N(i)\setminus j} (1 - \mu_{ki}^{(n)})\big)}, \tag{10}$$

where $\alpha_{ij}^{(n)}$ normalizes the messages. Note that we have inserted the corresponding expressions for $\Phi_{i,j}(x_i, x_j)$ and $\Phi_i(x_i)$ according to the exponential form (3).

Next, we compute the Jacobian of BP and construct a valid LF according to (8). Taking first order derivatives yields

$$(\mathrm{BP}')_{(i,j),(k,l)} = \tanh(J_{ij}) \frac{\prod\limits_{r \in N(i)\setminus\{j,k\}} \mu_{ri}^{\circ} \prod\limits_{r \in N(i)\setminus\{j,k\}} (1 - \mu_{ri}^{\circ})}{\big(e^{\theta_i} \prod\limits_{r \in N(i)\setminus j} \mu_{ri}^{\circ} + e^{-\theta_i} \prod\limits_{r \in N(i)\setminus j} (1 - \mu_{ri}^{\circ})\big)^2}, \tag{11}$$

if $i = l$ and $k \in N(i)\setminus j$. For any other pairs of edges $(i,j),(k,l)$ the entries $(\mathrm{BP}')_{(i,j),(k,l)}$ are equal to 0.

We state the Lyapunov equation $(\mathrm{BP}')^T P\, \mathrm{BP}' - P = -I$ for BP, where we have chosen $Q = I$ in (8). Due to the symmetry of $P$, the induced equation system only involves the upper triangular submatrix of $P$ and thus consists of $(2|\mathbf{E}|+1)\frac{2|\mathbf{E}|}{2} = 2|\mathbf{E}|^2 + |\mathbf{E}|$ equations and variables. A direct solution by Gaussian elimination requires $\mathcal{O}(|\mathbf{E}|^6)$ operations. Fortunately, there exists an efficient algorithm (Kitagawa, 1977) to solving discrete-time Lyapunov equations that reduces the complexity to $\mathcal{O}(|\mathbf{E}|^3)$ operations. Recall that the solution of (8) specifies a quadratic LF.

**Theorem 1.** *Let $\boldsymbol{\mu}^\circ$ be a stable FP of BP, let BP$'$ be the Jacobian and let $P$ be the solution to* (8)*. Then $V(\boldsymbol{\mu}) = (\boldsymbol{\mu} - \boldsymbol{\mu}^\circ)^T P (\boldsymbol{\mu} - \boldsymbol{\mu}^\circ)$ is a LF for $\boldsymbol{\mu}^\circ$ and can be computed with the computational effort $\mathcal{O}(|\mathbf{E}|^3)$. Moreover, there exists a constant $c > 0$, such that BP converges to $\boldsymbol{\mu}^\circ$ for all initial messages $\boldsymbol{\mu}^{(0)}$ chosen in $\Omega_c = \{\boldsymbol{\mu} \in \mathcal{M} \,|\, V(\boldsymbol{\mu}) \le c\}$.*

[8] A similar reparametrization leads to the so-called cavity update scheme that enjoys wide popularity in statistical physics (Mezard and Montanari, 2009).

Theorem 1 facilitates the construction of a LF for arbitrary probabilistic graphical models. Specifically, there exists a sublevel set $\Omega_c$ such that all messages inside $\Omega_c$ converge to the corresponding FP. Practically, however, we are confronted with two concerns: (i) on the one hand, Theorem 1 does not provide any constructive information on how to choose the parameter $c$, such that $\Omega_c$ is fully contained in $\{\boldsymbol{\mu} \in \mathcal{M} \,|\, \Delta V(\boldsymbol{\mu}) < 0\}$ (see Sec. 3.2 for a detailed discussion); (ii) on the other hand, ROA estimates based on the LF constructed by Theorem 1 may be too conservative and other LFs may be capable to yield larger ROA estimates.

In the next section, we address both issues by resorting to the sum-of-squares technique that automatically constructs (bounded-degree) polynomial LFs and implicitly deals with selecting the optimum sublevel set.

# 4 SUM-OF-SQUARES (SOS) METHOD

In this section, we introduce the Sum-of-Squares (SOS) method that provides the computational tools for automatically computing LFs. We then show how to use SOS to estimate ROAs for BP.

## 4.1 BACKGROUND OF SOS

Sum-of-squares relaxations have gained great popularity for the past 20 years with a variety of applications in control theory (Seeber et al., 2018) as well as in machine learning like sparse coding (Barak et al., 2015) or sparse PCA (Ma and Wigderson, 2015). The main idea is to relax the problem of checking nonnegativity of a polynomial $p$ – which is NP-hard for polynomials with a degree of at least 4 (Murty and Kabadi, 1987) – by checking whether $p$ can be expressed as a sum of squared polynomials $f_i$, i.e. $p = \sum_i f_i^2$, which obviously implies nonnegativity. Conversely, not every nonnegative polynomial is SOS. Note that $p$ is required to be of even degree, that is $deg(p) = 2d$. Furthermore, $p$ is SOS if and only if there exists a psd. matrix $Q$ such that $p$ can be written in the form $p = m^T Q\, m$ with $m$ being the vector of all monomials in $p$ having a degree up to $d$ (Choi et al., 1995). Therefore, when $p$ is defined on $\nu$ variables, the size of $Q$ is $\binom{\nu+d}{d} \times \binom{\nu+d}{d}$. It has been shown that the search for $Q$ can be performed efficiently by semidefinite programming (SDP) (Parrilo, 2000). The primal form of an SDP is

$$\min_w \; b^T w, \tag{12}$$

$$\text{s.t.} \;\; Q = Q_0 + \sum_i w_i Q_i, \tag{13}$$

$$Q \succeq 0, \tag{14}$$

where we are given a cost vector $b$ and symmetric matrices $Q_i$ and where $\succeq$ denotes positive semidefiniteness. Loosely

speaking, (13) describes the affine equality constraints induced by requiring equality of the coefficients of $p$ and $m^T Q\, m$. Note that SOS in our context does not involve a cost function, but reduces the above SDP to the feasibility conditions only. For a thorough treatment of SDP, we refer the reader to (Vandenberghe and Boyd, 1996).

When constructing polynomial LFs, the inequality conditions (6) and (7) can be replaced by the SOS conditions

$$V(\boldsymbol{z}) - \epsilon \cdot q(\boldsymbol{z}) \text{ is SOS}, \tag{15}$$

$$-\Delta V(\boldsymbol{z}) - \epsilon \cdot q(\boldsymbol{z}) \text{ is SOS}. \tag{16}$$

Note that we need to subtract a nonnegative polynomial $q(\boldsymbol{z})$ that vanishes in the FP and is multiplied by a small constant $\epsilon > 0$ to guarantee the required strict positivity of $V$ and $-\Delta V$. The above SOS relaxations facilitate the task of proving nonnegativity considerably. Moreover, if the coefficients of $V$ are unknown, they are incorporated as parametric variables in the corresponding SDP (Papachristodoulou and Prajna, 2005). For a candidate function $V$ with a predefined algebraic structure, one can therefore automatically search for a psd. matrix $Q$ such that $V = m^T Q\, m$ and $V$ is a LF for a FP. Consequently, LFs do not have to be constructed analytically but are obtained from solving an SDP.

When dealing with multiple FPs, global LFs do not exist. Therefore, it is required to verify polynomial nonnegativity only on certain subsets of the domain. The following result (Parrilo, 2000) is a useful consequence of the Positivstellensatz from real algebraic geometry (Bochnak et al., 1998) and provides an elegant remedy for this problem:

**Lemma 2.** *Let $R_i = \{\boldsymbol{z} \in \mathbb{R}^m \,|\, r_i(\boldsymbol{z}) \leq 0\}$ and $T_i = \{\boldsymbol{z} \in \mathbb{R}^m \,|\, t_i(\boldsymbol{z}) = 0\}$ be a finite family of semialgebraic subsets[9] of $\mathbb{R}^m$ and let $S = \{\boldsymbol{z} \in \mathbb{R}^m \,|\, s(\boldsymbol{z}) \leq 0\}$ be another semialgebraic subset of $\mathbb{R}^m$, where $r_i$, $t_i$ and $s$ are arbitrary polynomials. Let $U$ be the union of all $R_i$ and $T_i$. Then $U \subseteq S$, if there exist SOS polynomials $g_i$ and (not necessarily SOS) polynomials $h_i$ such that*

$$-s(\boldsymbol{z}) + \sum_i g_i(\boldsymbol{z}) \cdot r_i(\boldsymbol{z}) + \sum_i h_i(\boldsymbol{z}) \cdot t_i(\boldsymbol{z}) \text{ is SOS}. \tag{17}$$

Lemma 2 reveals that semialgebraic set containment problems can be cast in terms of SOS. Hence, whenever one aims to check whether a certain semialgebraic set is contained in another, one can instead try to solve a corresponding SOS problem – which, in turn, is reformulated as an SDP.

**Example.** *Assume that we aim to check the positivity of $-\Delta V$ on a sphere with radius $r$ defined by the semialgebraic set $B = \{\boldsymbol{z} \in \mathbb{R}^m \mid \sum z_i^2 \leq r^2\}$, that is $B \subseteq \{\boldsymbol{z} \in \mathbb{R}^m \mid -\Delta V(\boldsymbol{z}) > 0\}$. With help of Lemma 2, we can try to prove this by searching for a SOS polynomial $g$ such that*

$$-(\Delta V + \epsilon \cdot q) + g \cdot (\sum z_i^2 - r^2) \text{ is SOS}. \tag{18}$$

[9] A (real) semialgebraic set is a finite union of subsets of $\mathbb{R}^m$ that are defined by polynomial equalities or inequalities.

### 4.2 SOS FOR BELIEF PROPAGATION

Equipped with the relevant theoretical framework, we can formulate an algorithmic estimation of BP's regions of attraction. Still, one important technical subtlety remains:

An application of SOS does only make sense, if we analyze polynomial systems. Hence, we have to deal with the rational normalization terms $\alpha_{ij}^{(n)}$ in BP's message update equation (9). To approach this, we substitute the right-hand side in (9) by auxiliary variables. After a few rearrangements, we obtain the following constrained update equations:

$$\mu_{ij}^{(n+1)} = \beta_{ij}^{(n)}, \text{ s.t.} \tag{19}$$

$$\frac{\beta_{ij}^{(n)}}{\alpha_{ij}^{(n)}} - e^{J_{ij}+\theta_i} \prod_{k \in N(i)\setminus j} \mu_{ki}^{(n)} - e^{-J_{ij}-\theta_i} \prod_{k \in N(i)\setminus j} (1 - \mu_{ki}^{(n)}) = 0. \tag{20}$$

If we insert the definition (10) of $\alpha_{ij}^{(n)}$ into (20), we obtain polynomial update equations subject to polynomial equality constraints. Note that the exponential terms are constants, as all $\theta_i$ and $J_{ij}$ are assumed to be fixed. Now let

$$b_{ij}(\boldsymbol{\mu}, \boldsymbol{\beta}) = \frac{\beta_{ij}^{(n)}}{\alpha_{ij}^{(n)}} - e^{J_{ij}+\theta_i} \prod_{k \in N(i)\setminus j} \mu_{ki}^{(n)} - e^{-J_{ij}-\theta_i} \prod_{k \in N(i)\setminus j} (1 - \mu_{ki}^{(n)}), \tag{21}$$

where the auxiliary variables $\beta_{ij}^{(n)}$ are collected in a vector $\boldsymbol{\beta}$, and define for all edges $(i,j) \in \mathbf{E}$ the semialgebraic sets $R_{ij} = \{\boldsymbol{\mu}, \boldsymbol{\beta} \in \mathbb{R}^{2|\mathbf{E}|} \,|\, b_{ij}(\boldsymbol{\mu}, \boldsymbol{\beta}) = 0\}$ to be the sets of all message (plus auxiliary) variables satisfying (20). Let further $q(\boldsymbol{\mu})$ be any nonnegative polynomial whose level sets $\{\boldsymbol{\mu} \in \mathcal{M} \,|\, q(\boldsymbol{\mu}) = r^2\}$ determine the geometric shape of an LF domain and that vanishes in the FP, e.g., the spherical polynomial $q(\boldsymbol{\mu}) = \sum_{(i,j) \in \mathbf{E}} (\mu_{ij} - \mu_{ij}^\circ)^2$. Then, in view of our preceding discussion, the following theorem is an immediate application of Lemma 2 to BP:

**Theorem 3.** *Let $\boldsymbol{\mu}^\circ$ be a FP of BP,*

- *$V(\boldsymbol{\mu})$ be a SOS polynomial with $V(\boldsymbol{\mu}^\circ) = 0$,*
- *$g(\boldsymbol{\mu}, \boldsymbol{\beta})$ be a SOS polynomial,*
- *$p_{ij}(\boldsymbol{\mu}, \boldsymbol{\beta})$ be (not necessarily SOS) polynomials,*
- *$b_{ij}(\boldsymbol{\mu}, \boldsymbol{\beta})$ be given by* (21)*,*
- *$q(\boldsymbol{\mu})$ be a nonnegative polynomial with $q(\boldsymbol{\mu}^\circ) = 0$,*
- *$r > 0$ be a constant (i.e., the radius of $q$) and*
- *$\epsilon > 0$ be an arbitrary small constant*

*such that*

$$-(\Delta V + \epsilon \cdot q) + \sum_{(i,j) \in \mathbf{E}} p_{ij} \cdot b_{ij} + g \cdot (q - r^2) \text{ is SOS}. \tag{22}$$

*Then $V$ is a LF for BP on the set $B_q(r) = \{\boldsymbol{\mu} \in \mathcal{M} \,|\, q(\boldsymbol{\mu}) \leq r^2\}$ and all sets of messages $\boldsymbol{\mu}^{(0)}$ initialized in a sublevel set $\Omega_{V,c} \subseteq B_q(r)$ converge to $\boldsymbol{\mu}^\circ$.*

To compute an ROA estimate, we can, e.g., take $q$ to be of spherical shape and state the SOS program (22). All polynomial coefficients are parametric variables in the corresponding SDP. We maximize the radius $r$ of the sphere $B_q(r)$ (e.g., by bisection) such that the SDP (13), (14) admits a feasible solution $Q$. With a monomial vector chosen in advance, it follows that $V = m^T Q\, m$ is a LF whose largest sublevel set $\Omega_c \subseteq B_q(r)$ provides a valid ROA estimate for a fixed point.

### 4.3 EXPANDING THE ROA ESTIMATES

To expand the estimated ROA, we can utilize an alternating update scheme between the LF and the auxiliary polynomials (Jarvis-Wloszek et al., 2003)). Concretely, one aims for solving the following optimization problem:

$$\max r \quad \text{s.t.} \tag{23}$$

$$-(\Delta V + \epsilon \cdot q) + \sum_{(i,j)\in\mathbf{E}} s_{ij} \cdot b_{ij} + g \cdot (V - c) \text{ is SOS}, \tag{24}$$

$$-(V - c) + h \cdot (q - r^2) \text{ is SOS}, \tag{25}$$

where the maximization is performed over the polynomial variables $V, g$ (both SOS), $s_{ij}$, and $h$. More precisely, we maximize the radius of a sphere that is inscribed in a sublevel set of the current Lyapunov function, that is $B_q(r) \subseteq \Omega_c$. This is guaranteed by condition (25). At the same time, we must ensure that the current sublevelset is always included in the set where $\Delta V < 0$. This is guaranteed by condition (24). The increasing radius of the inner sphere enforces the sublevel set and therefore the estimated ROA to grow over time. Note however, that we can not solve both SOS programs at the same time. This is due to bilinearities with respect to the SDP variables in $g \cdot V$ and $h \cdot r^2$. Specifically, we need to iteratively keep some of the polynomials fixed, while updating the others. The algorithmic procedure can be sketched as follows:

1. First, initialize $V$ (e.g., with the LF obtained by Theorem 1 in Sec. 3.3) and take $q$ to be a nonnegative shape polynomial (e.g., a sphere) that vanishes in the FP.
2. With $V$ fixed, maximize $c$ (e.g., by bisection) and compute $g$ such that (24) is satisfied. In this step, we find the largest valid sublevelset $\Omega_c$ of $V$.
3. With $V$ fixed, maximize $r$ (e.g., by bisection) and compute $h$ such that (25) is satisfied. This adapts the chosen shape to $\Omega_c$ (see Fig. 3a).
4. Check if $r$ is converged.
   - 4.1 If yes, stop the procedure and return the sublevelset $\Omega_c$ obtained in step 2 as ROA estimate.
   - 4.2 Otherwise, keep $g$, $c$ from step 2 and $h$, $r$ from step 3 fixed. Update $V$ by simultanously solving (24) and (25); this guarantees the relation $B_q(r) \subseteq \Omega_c \subseteq \{\boldsymbol{\mu} \,|\, \Delta V(\boldsymbol{\mu}) < 0\}$ (see Fig. 3b). Then go back to step 2.

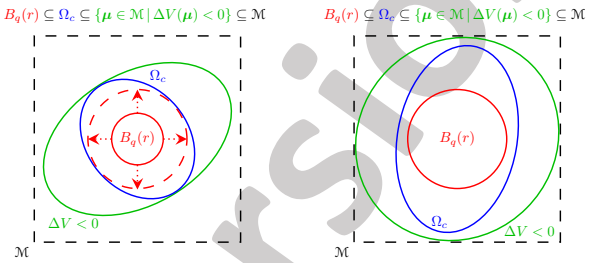


(a) Enlarge $B_q(r)$. (b) Update $V$.

Figure 3: An iterative ROA optimization procedure: inscribing a sphere $B_q(r)$ inside $\Omega_c$ enforces $\Omega_c$ to grow over time.

We apply the described procedure in Sec. 5.

### 4.4 PROPERTIES AND LIMITATIONS OF SOS

The ROA estimates are affected by two particular aspects of BP: first, the graph structure determines the computational complexity; and second, analytical properties of BP limit the maximum size of the estimated ROAs.

To understand the complexity of SOS, consider a polynomial of degree $2d$ in $\nu$ independent variables. When testing for non-negativity, the corresponding monomial vector $m$ determines the size of the SDP. In worst case, the matrix $Q$ (see Sec. 4.1) has a dimension of $\binom{\nu+d}{d} \times \binom{\nu+d}{d}$. Thus, if one quantity is fixed the SDP grows polynomially in the number of variables $\nu$ or the degree $d$. If both quantities increase, the SDP experiences exponential growth.

SOS for BP consists of $4|\mathbf{E}|$ independent variables (see Sec. 4.2). The corresponding SDP is of size at least $\binom{4|\mathbf{E}|+\lceil \frac{d_{\max}}{2} \rceil}{\lceil \frac{d_{\max}}{2} \rceil} \times \binom{4|\mathbf{E}|+\lceil \frac{d_{\max}}{2} \rceil}{\lceil \frac{d_{\max}}{2} \rceil}$, where $d_{\max}$ is the maximum degree of the graph. This renders solving the SDP problematic – particularly for dense graphs. Thus we must decrease the complexity, e.g., by the Newton polytope or block diagonalization. This helps to reduce the monomial vector and exploits symmetries in the SDP (Löfberg, 2009).

Moreover, note that the size of the estimated ROAs is inherently limited by the analytical properties of BP. This stems from the scenario that two incoming messages might completely disagree about the state of a RV, e.g., if $\mu_{ki}(1) = 1$ whereas $\mu_{li}(-1) = 1$. Formally, the normalization term $\alpha_{ij}$ in (9) goes to infinity in that case, i.e., BP has a pole for the corresponding message values. Unfortunately, our construction of ROA estimates suffers from poles in the message

space. In particular this limits the maximum possible radius as the spherical estimates must not include any poles. Considering more flexible ROA shapes would be one possible remedy; this, however, also increases the computational complexity by enlarging the degree of the LFs.

# 5 SIMULATIONS

We now provide a proof of concept and compute the ROAs for a range of models in Sec 5.1- 5.2. Moreover, we assess if the estimated ROAs reflect the true ones well in Sec. 5.3.

We consider homogenous 3-regular models on $4$ nodes. Note, however, that our approach generalizes well to models with arbitrary pairwise potentials. We mainly consider homogenous models for their theoretical properties that make the interpretation and visualization of the results particularly pleasant. We explain this as follows:

On the one hand, one can compute the critical values $J^*, \theta^*$ that induce a phase transition analytically (Georgii, 1988). On the other hand, all messages are identical at a fixed point, i.e., $\mu^\circ_{ij} = \mu^\circ$. Thus, all FPs lie on the main diagonal of the message space $\mathcal{M}$ (the 12-dimensional unit cube). For visualization purposes, we project the FPs and estimated ROAs on the one-dimensional subspace corresponding to the main diagonal. We rescale all projections by $\frac{1}{\sqrt{12}}$ to the $[0, 1]$. This makes the results independent of the model-size.

All FPs are visualized by solid (stable) or dashed (unstable) lines; the ROA estimates are illustrated as shaded regions. For computational reasons, we estimate elliptical ROAs (see Sec. 4.4). Note that these correspond to the sublevel sets of quadratic LFs (see Fig 2). We further optimize the estimated ROAs according to the algorithm introduced in Sec. 4.3.[10] We do, however, not visualize the estimated elliptical ROAs since projections of non-spherical objects may not reflect the object's geometric properties well and lead to wrong conclusions. Therefore, we illustrate the largest spheres $B_q(r)$ contained in the estimated ROAs instead.

## 5.1 ZERO LOCAL POTENTIALS ($\theta = 0$)

First, we consider models without local potentials, i.e. $\theta = 0$. The lack of any prior knowledge renders this a worst-case scenario in terms of BP's convergence properties (Knoll and Pernkopf, 2017). We vary the couplings $J$ in the interval $[-1, 2]$ and estimate the ROAs for all FPs (see Fig. 4a).

For $J \geq 0$, the unique stable FP $\boldsymbol{\mu}^\circ_0$ exists up to a critical value of $J^*$. The ROA is bounded by the poles of BP that are nearest to this FP (see Sec. 4.4). The nearest poles are those that keep all values from $\boldsymbol{\mu}^\circ_0$ except for two messages. The distance between $\boldsymbol{\mu}^\circ_0$ and these poles is given by $\sqrt{0.5}$, yielding an upper bound on each spherical ROA estimate. Similar bounds hold for all estimated ROAs in this section.

As the couplings $J$ increase, the estimated ROA of $\boldsymbol{\mu}^\circ_0$ shrinks slowly but steadily, until it rapidly decreases as $J$ approaches the phase transition boundary. This conforms with the observation that the number of BP iterations peaks at the phase transition (Mooij and Kappen, 2005). We conjecture that this behavior of BP leads to less robust ROA estimates in the vicinity of a phase transition.

By further increasing $J$, the FP $\boldsymbol{\mu}^\circ_0$ becomes unstable and two stable FPs $\boldsymbol{\mu}^\circ_1$ and $\boldsymbol{\mu}^\circ_2$ arise, each preferring another state of the marginals. The ROAs for both FPs grow rapidly after the phase transition. Moreover, both FPs are naturally separated by the – henceforth unstable – FP $\boldsymbol{\mu}^\circ_0$.

For $J \leq 0$, the unique stable FP $\boldsymbol{\mu}^\circ_0$ exists up to $-J^*$. The estimated ROAs are symmetric to those obtained for $J \leq 0$. It is well known that there do not arise any further stable fixed points and that BP does not converge for $J < J^*$ (Mooij and Kappen, 2005; Knoll and Pernkopf, 2017).

## 5.2 NON-ZERO LOCAL POTENTIALS ($\theta \neq 0$)

Second, we consider the more general case of non-zero local potentials. Due to the symmetry of the model, we can assume $\theta_i = \theta > 0$. In general, the existence of non-zero local potentials improves the convergence behavior of BP (Knoll and Pernkopf, 2017). We consider three scenarios with $J \in \{0.6, 0.7, 0.8\}$, vary $\theta$ in the interval $[0, 0.2]$ and estimate the ROAs for all FPs (see Fig. 4b - 4d)

For small values of $\theta < \theta^*$, two stable FPs $\boldsymbol{\mu}^\circ_1$ (green) and $\boldsymbol{\mu}^\circ_2$ (blue) exist and are accompanied by one unstable FP $\boldsymbol{\mu}^\circ_0$ (red). We refer to $\boldsymbol{\mu}^\circ_2$ as the *state-preserving* FP (as its marginals are in accordance with the local potentials).

As $\theta$ increases, the estimated ROA of $\boldsymbol{\mu}^\circ_2$ grows constantly. In contrast, the ROA of $\boldsymbol{\mu}^\circ_1$ shrinks constantly until the FP vanishes at $\theta^*$. Beyond the phase transition only a unique FP (the state-preserving) remains, with its ROA being bounded by the poles. Also note, how for a given local strength, the estimated ROAs increase with the couplings $J$ for all FPs.

## 5.3 ROBUSTNESS OF BP FIXED POINTS

Even though the volumes of the estimated ROAs are inherently limited in size, we show that their relative size reflects BP's convergence behavior. Therefore, we consider models with two stable FPs and compare the ratio of the largest elliptical ROA volumes to the ratio of convergent trajectories for the respective FP (we use $10^5$ random initializations).

We keep $J \in \{0.6, 0.7, 0.8\}$ fixed and vary $\theta$ in the interval $[0, \theta^*]$. The results are shown in Fig. 5. For $J \in \{0.7, 0.8\}$, the proportion of estimated ROA volumes provides a rough

[10] We used the SOS modul of the software YALMIP (Löfberg, 2009) for formulating the SOS programs and the interior-point optimizer MOSEK (MOSEKApS, 2021) for solving the SDPs.

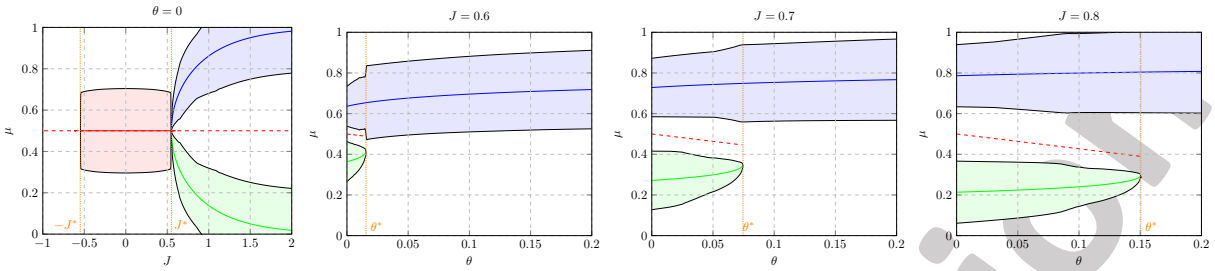


(a) (b) (c) (d)

Figure 4: Spherical ROA estimates (shaded regions) projected onto the main diagonal of $\mathcal{M}$. The phase transitions are illustrated by $\pm J^*$ and $\theta^*$ respectively. (a) $\theta = 0$ and $J \in [-1, 2]$: the ROA of $\boldsymbol{\mu}_0^\circ$ (red) shrinks until it becomes unstable; for $J > 0$, two FPs (blue and green) arise with increasing ROAs. (b),(c),(d) $\theta \in [0, 0.2]$ and $J \in \{0.6, 0.7, 0.8\}$: the ROA of the state-preserving FP (blue) grows with $J$ whereas the ROA of the opposing FP (green) shrinks until it vanishes.

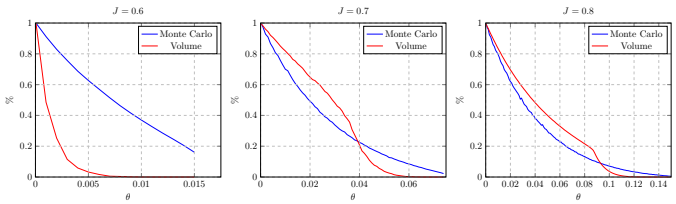


(a) (b) (c)

Figure 5: Comparison between ratio of volumes of elliptical ROA estimates (red) and ratio of convergent Monte-Carlo runs (blue) for FPs $\boldsymbol{\mu}_1^\circ$ and $\boldsymbol{\mu}_2^\circ$.

approximation to the proportion of the actual ROAs (obtained by the Monte Carlo runs). For $J = 0.6$, the ROA of the non state-preserving FP $\boldsymbol{\mu}_1^\circ$ is clearly underestimated; hence, an approximation is not reliable in that case. We explain this behavior with the lack of robustness of the ROA estimates in the vicinity of a phase transition (see Sec. 5.1).

The above observations suggest that the size of the estimated ROA provides a measure for the robustness of a FP. That is, the larger a ROA estimate, the more likely it remains stable under slight variation of the parameters.

## 6 CONCLUSION

In this paper, we have explained how to estimate the regions of attraction (ROAs) for belief propagation (BP), i.e., initial message values for which BP provably converges. Therefore, we have reformulated BP and introduced Lyapunov functions to compute lower bounds on the ROAs. Moreover, we provided an algorithm that utilizes the sum-of-squares method and iteratively enlarges the estimated ROAs.

We have estimated the ROAs for various models and observed: how the potentials impact the ROAs; that unstable fixed points confine the ROAs, if all messages are initialized identically; and that the estimated ROAs reflect the exact ones well despite being conservative.

Estimating the ROAs provides further insights into the behavior of BP and improves our understanding of the role of the initialization. We anticipate that our results inspire problem-tailored initialization strategies for improving BP.

**Acknowledgements**

This work was supported by the Graz University of Technology LEAD project "Dependable Internet of Things in Adverse Environments". We further acknowledge support from the Wireless Lab, Huawei Technologies Sweden. Finally, the financial support by the Christian Doppler Research Association, the Austrian Federal Ministry for Digital and Economic Affairs, and the National Foundation for Research, Technology and Development is gratefully acknowledged.

## References


Vitaly Aksenov, Dan Alistarh, and Janne H. Korhonen. Relaxed scheduling for scalable belief propagation. In *Proceedings of NeurIPS*, 2020.

Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of STOC*, 2015.

Jacek Bochnak, Michel Coste, and Marie-Francoise Roy. *Real Algebraic Geometry*. Springer, 1998.

Mi D. Choi, Kee Y. Lam, and Bruce Reznick. Sums of squares of real polynomials. *Proceedings of Symposia in Pure Mathematics*, 58(2):103–126, 1995.

Gregory F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial intelligence*, 42(2-3):393–405, 1990.

Gal Elidan, Ian McGraw, and Daphne Koller. Residual belief propagation: Informed scheduling for asynchronous message passing. In *Proceedings of UAI*, 2006.

Hans-Otto Georgii. *Gibbs Measures and Phase transitions*. De Gruyter, 1988.

Amir Globerson and Tommi Jaakkola. Convergent propagation algorithms via oriented trees. In *Proceedings of UAI*, 2007.

Tamir Hazan and Amnon Shashua. Convergent message-passing algorithms for inference over general graphs with convex free energies. In *Proceedings of UAI*, 2008.

Tom Heskes. On the uniqueness of loopy belief propagation fixed points. *Neural Computation*, 16(11):2379–2413, 2004.

Tom Heskes et al. Stable fixed points of loopy belief propagation are minima of the Bethe free energy. In *Proceedings of NIPS*, 2003.

Alexander T. Ihler, John W. Fisher, and Alan S. Willsky. Loopy belief propagation: convergence and effects of message errors. *Journal of Machine learning research*, 6(1):905–936, 2005.

Zachary Jarvis-Wloszek, Ryan Feeley, Weehong Tan, Kunpeng Sun, and Andrew Packard. Some controls applications of sum of squares programming. In *IEEE International Conference on Decision and Control*, 2003.

Hassan K. Khalil. *Nonlinear Systems*. Pearson, 2002.

Genshiro Kitagawa. An algorithm for solving the matrix equation $X = FXF^T + S$. *International Journal of Control*, 25(5): 745–753, 1977.

Christian Knoll and Franz Pernkopf. On loopy belief propagation–local stability analysis for non-vanishing fields. In *Proceedings of UAI*, 2017.

Christian Knoll and Franz Pernkopf. Belief propagation: accurate marginals or accurate partition function – where is the difference? In *Proceedings of UAI*, 2019.

Christian Knoll, Michael Rath, Sebastian Tschiatschek, and Franz Pernkopf. Message Scheduling Methods for Belief Propagation. In *Machine Learning and Knowledge Discovery in Databases*, pages 295–310. Springer, 2015.

Christian Knoll, Dhagash Mehta, Tianran Chen, and Franz Pernkopf. Fixed points of belief propagation–an analysis via polynomial homotopy continuation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 40(9):2124–2136, 2017.

Christian Knoll, Florian Kulmer, and Franz Pernkopf. Self-guided belief propagation – a homotopy continuation method. In *arXiv:1812.01339*, 2018.

Frederic Koehler. Fast convergence of belief propagation to global optima: Beyond correlation decay. In *Proceedings of NeurIPS*, 2019.

Johan Löfberg. Pre- and post-processing sum-of-squares programs in practice. *IEEE Transactions on Automatic Control*, 54(5): 1007–1011, 2009.

Tengyu Ma and Avi Wigderson. Sum-of-squares lower bounds for sparse PCA. In *Proceedings of NIPS*, 2015.

Victorin Martin, Jean-Marc Lasgouttes, and Cyril Furtlehner. The role of normalization in the belief propagation algorithm. *preprint arXiv:1101.4170*, 2011.

Talya Meltzer, Amir Globerson, and Yair Weiss. Convergent message passing algorithms: a unifying view. In *Proceedings of UAI*, 2009.

Ofer Meshi, Ariel Jaimovich, Amir Globerson, and Nir Friedman. Convexifying the Bethe free energy. In *Proceedings of UAI*, 2009.

Marc Mezard and Andrea Montanari. *Information, Physics, and Computation*. Oxford Univ. Press, 2009.

Joris M Mooij and Hilbert J Kappen. On the properties of the Bethe approximation and loopy belief propagation on binary networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2005 (11):P11012, 2005.

Joris M. Mooij and Hilbert J. Kappen. Sufficient conditions for convergence of the sum–product algorithm. *IEEE Trans. on Information Theory*, 53(12):4422–4437, 2007.

MOSEKApS. *The MOSEK optimization toolbox for MATLAB manual. Version 9.2.*, 2021. URL `http://docs.mosek.com/9.2/toolbox/index.html`.

Kevin P. Murphy, Yair Weiss, and Michael I. Jordan. Loopy belief propagation for approximate inference: an empirical study. In *Proceedings of UAI*, 1999.

Katta G. Murty and Santosh N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.

Antonis Papachristodoulou and Stephen Prajna. A tutorial on sum of squares techniques for systems analysis. In *American Control Conference*, 2005.

Pablo Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.

Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, 1988.

Richard Seeber, Markus Reichhartinger, and Martin Horn. A Lyapunov function for an extended super-twisting algorithm. *IEEE Trans. on Automatic Control*, 63(10):3426–3433, 2018.

Jinwoo Shin. Complexity of Bethe approximation. In *Proceedings of AISTATS*, 2012.

Charles Sutton and Andrew McCallum. Improved dynamic schedules for belief propagation. In *Proceedings of UAI*, 2007.

Weehong Tan and Andrew Packard. Searching for control Lyapunov functions using sums of squares programming. In *Allerton Conference on Communication, Control and Computing*, 2004.

Gerald Teschl. *Ordinary Differential Equations and Dynamical Systems*. 2004.

Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.

Martin J. Wainwright, Tommi S. Jaakkola, and Alan S. Willsky. Tree-based reparameterization framework for analysis of sum-product and related algorithms. *IEEE Trans. on Information Theory*, 49(5):1120–1146, 2003.

Yusuke Watanabe and Kenji Fukumizu. Graph zeta function in the Bethe free energy and loopy belief propagation. In *Proceedings of NIPS*, 2009.

Yair Weiss. Correctness of local probability propagation in graphical models with loops. *Neural Computation*, 12(1):1–41, 2000.

Adrian Weller and Tony Jebara. Approximating the Bethe partition function. In *Proceedings of UAI*, 2014.

Adrian Weller, Kui Tang, Tony Jebara, and David Sontag. Understanding the Bethe approximation: when and how can it go wrong? In *Proceedings of UAI*, 2014.

Max Welling and Yee W. Teh. Approximate inference in Boltzmann machines. *Artificial Intelligence*, 143(1):19–50, 2003.

Jonathan S. Yedidia, William T. Freeman, and Yair Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. on Information Theory*, 51(7):2282–2312, 2005.

Alan L. Yuille and Anand Rangarajan. The concave-convex procedure. *Neural Computation*, 15(4):915–936, 2003.