

Computing and Optimizing Over All Fixed-Points of Discrete Systems on Large Networks

James R. Riehl^{1*}, Maxwell I. Zimmerman², Matthew F. Singh¹,
Gregory R. Bowman², ShiNung Ching^{1,3,4}

¹Department of Electrical and Systems Engineering, Washington University in St. Louis, 1 Brookings Dr, St. Louis, MO 63130,

²Department of Biochemistry and Molecular Biophysics, Washington University in St. Louis, 660 S Euclid Ave, St. Louis, MO 63110

³Department of Biomedical Engineering, Washington University in St. Louis, 1 Brookings Dr, St. Louis, MO 63130

⁴Division of Biology and Biomedical Sciences, Washington University School of Medicine, 660 S Euclid Ave, St. Louis, MO 63110

*To whom correspondence should be addressed; E-mail: jrriehl@wustl.edu.

Abstract

Equilibria, or fixed points, play an important role in dynamical systems across various domains, yet finding them can be computationally challenging. Here, we show how to efficiently compute all equilibrium points of discrete-valued, discrete-time systems on sparse networks. Using graph partitioning, we recursively decompose the original problem into a set of smaller, simpler problems that are easy to compute, and whose solutions combine to yield the full equilibrium set. This makes it possible to find the fixed points of systems on arbitrarily large networks meeting certain criteria. This approach can also be used without computing the full equilibrium set, which may grow very large in some cases. For example, one can use this method to check the existence and total number of equilibria, or to find equilibria that are optimal with respect to a given cost function. We demonstrate the potential capabilities of this approach with examples in two scientific domains: computing the number of fixed points in brain networks and finding the minimal energy conformations of lattice-based protein folding models.

Introduction

A fundamental element in the study of physics and dynamical systems is the notion of fixed points or equilibria. In the absence of exogenous input or disturbance, a system will not deviate from a fixed point, and the dynamics in the neighborhood of that point determines its stability, i.e. whether it is an attractor, repeller, or other type of critical point. Indeed, the state of a system (for our purposes, defined on a network) evolves in time on a state space that is landmarked by equilibria. Knowledge of these equilibria can therefore help to predict long term evolution of system trajectories. It is often particularly useful to know the equilibria that are optimal with respect to some cost function, for example, to find ground states or minimum energy configurations of materials, protein sequences, or brain networks, or to determine the most cost efficient or beneficial states in social or economic networks.

We focus here on dynamical systems where the dependency structure of the components of the state is described by a graph, and each component can take one of a small number of values. These include ferromagnetic models, simplified neural networks, models of decision or opinion dynamics on social and economic networks, foodwebs and ecological networks, and protein folding models. The notion of an equilibrium or fixed point has wide-ranging interpretations across these applications.

Finding critical points of generalized Ising models is a long-standing problem in theoretical physics [1] that has remained active over the years [2, 3]. Expanding the scale and types of structures for which we can completely characterize the critical points has value in the study of magnetism, phase transitions, fluid dynamics, and other topics [4].

In biochemistry, finding the minimum energy protein sequence is a sought after task for which extraordinary computational resources and schemes have been devoted [5, 6]. Moreover, the ability to find all critical points can contribute to a greater understanding of the protein folding process, another target of large-scale computational projects [7]. Comparing the structure of proteins to the underlying fixed-point landscape can also give clues about the evolution of a protein sequence and provide greater insight into the impact of protein topology on evolutionary pathways.

In neural network models, fixed points represent invariant patterns of activation that can be associated with memories, cognitive task goals, or default states of awareness [8]. This is perhaps most prevalent in the study of memory encoding, where attractor networks have long been a popular model for the formation and recall of memoranda and other types of associations [9, 10]. More recently, notions of free energy minimization have been used as a theoretical schema within which to understand brain function at multiple spatial scales [11]. Testing such hypotheses has been difficult, due to the taxing nature of estimating energy landscapes from data. Perhaps as a result, the state of the art in brain network analyses, especially at whole-brain scales, has been largely limited to analyses of inter-region correlation and subsequent (static) graph theoretical analyses [12]. Nonetheless, key efforts in computational modeling and analysis are now underway to bridge architectural characterizations of brain networks to more faithful understanding of temporal dynamics over such networks [13, 14]. Faster characterization of fixed points can enable novel assessments of how brain network dynamics – via attractor and energy landscapes, which ultimately mediate the input to

output lability of these networks [15] – are altered during cognition.

Elucidating fixed points is also a key issue in ecological network and food web models. Here, an equilibrium is a state in which species populations remain constant due to balance in synergistic and antagonistic interactions [16] similar to the dynamics of the classical Lotka Volterra model.

A tool for computing the fixed points in these systems would thus be of great consequence across several scientific domains. Unfortunately, even for discrete valued systems, finding all fixed points is computationally infeasible [17, 18]. Special cases of finding all fixed points of system (1) such as deciding on the existence of a Nash equilibrium for certain systems [18] and calculating the minimum-energy state of ferromagnetic spin models [17] are proven to be *NP-hard*. This generally means that the time required by any existing algorithm to compute a solution is exponential in the size of the problem (e.g., number of nodes in the network). For example, the naive approach of checking for fixed points among all possible configurations in a network in which nodes can take one of two states would require a time proportional to 2^n , and it is often not trivial to improve much on this brute force technique. When approximation is sufficient, methods such as simulated annealing [19, 20] and genetic algorithms [21, 22] can be effective in cleverly exploring the enormous state space, but these algorithms can generally only provide probabilistic guarantees on finding fixed points or global extrema. Otherwise, branch-and-bound algorithms [23] can reduce computation time by safely eliminating large regions of the total search space without compromising optimality. However, these methods often require significant effort in tailoring to a specific problem and their objective is typically to find a single global extremum. A recent approach to finding the global energy minimum on Ising lattices involves pairing combinatorial with convex optimization to converge upper and lower bounds on the solution, and was demonstrated on three-dimensional lattices containing up to 50 nodes, each of which can take two states [3]. In comparison, the method proposed here found the global minima and all fixed points on 3-dimensional lattices of 64 nodes, each of which can take three states (see **Case Study 2**), and is demonstrated in on other types of networks having hundreds and thousands of nodes.

Significant effort has also been put toward finding attractors in *Boolean networks*, in which nodes take one of two values (0 or 1) and local dynamics are based on logic rules. This work is most closely associated with genetic regulatory network models, where nodes represent genes that are either expressed (1) or not (0),

and edges indicate causal links between genes. Indeed, for the special case when each node can take exactly two values, the system used in our framework is equivalent to a Boolean network. Several approaches have been proposed to find attractors in Boolean networks, including simulation [24], aggregation [25], binary decision diagrams [26], and satisfiability algorithms [27]. In particular, satisfiability rules paired with a partition of the Boolean network are used in [28] to find attractors in sequential and parallel implementations. Since Boolean networks represent just one class of applicable systems, we leave an explicit performance comparison of these approaches to future work dedicated to this problem domain. Rather, we emphasize here three conceptual advancements over previous aggregation or partition-based techniques for finding attractors [25, 28]: (i) One level of partitioning may not be sufficient to make computation on large networks tractable – the extension to arbitrary levels of recursive partitioning yields what is to our knowledge the first provably tractable algorithm for computing fixed points on a particular class of networks, namely, those with local connectivity structure captured by the measure we call partition separability. As will be seen in our results, partition separability describes the situation where a network can be separated into groups of nodes (neurons, proteins, species, etc) that have a high degree of interconnectivity and interaction within group, and fewer, less dense interaction between groups. (ii) Computing the energy function at each stage of the decomposition allows for useful characterizations of the attractor landscape without explicitly computing the entire set of attractors, which may become prohibitively large. (iii) Extending to more general discrete-valued networks makes it possible to analyze networks in which node states can take more than two values.

To highlight the practical versatility of this method, we provide two case studies. In the first, we analyze the dynamics of functional brain networks extracted from the Human Connectome Project database. We show that such networks exhibit exponential growth in the number of fixed points as a function of interconnected brain regions, a prediction made in theoretical studies but here verified empirically for the first time. In the second case study, we address the decades-old challenge of finding low-energy configurations of proteins, toward facilitating a deeper understanding of protein evolution and design. We consider a prototypical protein lattice model and are able to efficiently find the global energy minima as well as perform an exact characterization of the ruggedness of the fixed-point landscape. Each of these case studies features a stochastic real world system for which a deterministic system is constructed and used as an analytical tool.

Specifically, we define the deterministic system such that its fixed points contain the set of energy extrema of the corresponding stochastic system.

Discrete systems and their fixed points

We consider discrete-time and discrete-valued dynamical systems in the general form

$$\mathbf{x}(t+1) = f(\mathbf{x}(t)), \quad (1)$$

with state vector $\mathbf{x}(t) := [x_1(t), \dots, x_n(t)]^\top$. Each component of the state $x_i(t)$ can take one of a finite set \mathcal{S}_i of possible values. Define a directed graph $\mathbb{G} = (\mathcal{V}_0, \mathcal{E}_0)$ in which each node in \mathcal{V}_0 corresponds to a component of the state, and the edges \mathcal{E}_0 define dependencies in the system dynamics. Namely, an edge $(i, j) \in \mathcal{E}_0$ indicates that $x_i(t+1)$ may depend on $x_j(t)$. Conversely, the lack of an edge means there is no dependency. The dynamics are therefore separable into the components $x_i(t+1) = f_i(\mathbf{x}(t))$, in which f_i only depends on the states of node i and its neighbors in the network.

A fixed point or *equilibrium* \mathbf{x}^* of the system is a state in which no node will change between time steps, i.e. $f(\mathbf{x}^*) = \mathbf{x}^*$. Let $\Omega(\mathcal{V}_0)$ denote the set of all such equilibria. Since $f(\mathbf{x})$ is in general nonlinear, computing $\Omega(\mathcal{V}_0)$ is a challenging task, and the main contribution of this work is a method for exploiting the network topology and local dynamics to compute a sparse representation of $\Omega(\mathcal{V}_0)$ in an efficient way.

Results

Our main result is an algorithm that computes and optimizes over the global equilibrium set of system (1). In this section, we describe how to do this using a combination of local dynamical analysis and recursive graph partitioning, discuss the computational complexity and related literature, and demonstrate the approach on practical applications in biochemistry and computational neuroscience.

Finding and optimizing over all fixed points

The approach begins by computing the set of *local equilibrium* (LEQ) states for each node in the network. A local equilibrium for a given node can be thought of as any system state such that the values of the node's neighbors will not cause the node to change its value. Formally, we define the LEQ set for a given node i as

the set of all system states \mathbf{x} such that $x_i(t)$ will not change at time $t + 1$ according to the local update rule $f_i(\mathbf{x}(t))$:

$$\mathcal{L}_i := \{\mathbf{x} \in \mathcal{S} : x_i = f_i(\mathbf{x})\}.$$

For sparsely connected networks in which the numbers of discrete values $|\mathcal{S}_i|$ that each node can take is not too large, these sets are fast to compute and can be efficiently stored using the sparse representations described in **Methods**. An important observation is that the global equilibrium set is equal to the intersection of all LEQ sets. In other words, when all nodes are at equilibrium with their respective neighbors, the entire network is at equilibrium. Unfortunately, computing all of these intersections may be too difficult for large networks. However, by using graph partitioning while retaining necessary information for each partitioned region, we can decompose the problem in a way that allows for exact construction of the full set of equilibrium points (see **Methods: Recursive partition-based intersections**). Specifically, by intersecting the LEQ sets within each partitioned subnetwork, we form a set of *regional equilibrium* (REQ) states. This effectively reduces the problem to finding all compatible REQ states on the partitioned subnetworks, which can be achieved by intersecting neighboring LEQ sets if the subnetworks are small enough, or otherwise by performing an additional partition and proceeding recursively. The end result is a top-level REQ set along with a hierarchical data structure \mathcal{T} that efficiently encodes all information necessary to generate the full equilibrium set $\Omega(\mathcal{V}_0)$.

What renders this algorithm tractable for k -separable networks is both the recursive decomposition and the fact that we need only a compact representation of the sets (including only the nodes on the partition boundaries) in order to construct higher level REQ sets and ultimately the global equilibrium set. The computational complexity thus depends primarily on the node degrees (number of connections to each node) and the number of edges that cross partition boundaries. Another key element in the approach is the construction of maps between the various stages of partitioning, which allow for quick access to the hierarchy of global, regional, and local equilibrium sets. Finally, if an energy or cost function is provided, the algorithm computes costs at each stage, which can be used to efficiently extract the optimal equilibrium states with respect to the given energy function. See **Methods** for a detailed presentation of this approach.

Computational complexity

The key characteristic in determining how much advantage will be gained by using the proposed approach on this NP-hard problem is how easily the network can be decomposed into smaller parts. We define the *partition separability* of a network as the maximum number of edges that cross partition boundaries over all partitions used in the decomposition. For network topologies such that the partition separability is independent of network size, the computation time of the proposed algorithm is exponential only in local network properties, namely, the number of local connections or *degree* of the nodes. Fig. 1 shows the time to compute all fixed points for three different network types of various size, governed by two-state best-response dynamics (See **SI: Best-response and linear threshold dynamics**). The network diagrams each depict one of the fixed points of their respective systems, which correspond to Nash equilibria. The top panel of Fig. 1 shows the substantial reduction in computational complexity when using the recursive partitioning approach compared to exhaustive search or intersecting all LEQ sets to compute all fixed points. We note that the local intersection method is already a steep improvement over exhaustive search. While ring networks are particularly well-suited to this approach since they are 2-separable, perhaps even better suited are tree networks, which are 1-separable. Computation times for optimizing over the entire equilibrium space on tree networks are shown in the middle panel of Fig. 1.

Figure 1: Time required to compute energy-extremal fixed points for three different network types: ring (top), tree (middle), and geometric random network (bottom). Node states are governed by linear threshold dynamics and energy is evaluated using a generalized Ising model (see **Supplementary Information (SI): Best-response and linear threshold dynamics**). Computation times for exhaustive search and intersecting all LEQ sets on ring networks are shown in the top panel. Black and white node colorings in the network diagrams respectively depict the states -1 and +1 in energy-extremal fixed points.

Indeed, we can prove that the algorithm is tractable for k -separable networks. Let \hat{d} denote the maximum degree and \hat{s} the maximum number of discrete values $|\mathcal{S}_i|$ for any node $i \in \mathcal{V}_0$ in the network, while q denotes the number of partition groups at each level, and r is the maximum number of nodes in a group at the lowest level of partitioning. The following theorem provides the parameterized computational complexity of the proposed algorithm, namely, that it is linear in the size of the network and exponential only in the node degree and the partition separability of the network (proof in **Appendix**).

Theorem 1. *The complexity of computing the top-level REQ set $\Omega^p(\mathcal{V}_0)$ and the results tree \mathcal{T} for a k -separable network is at most $O(\alpha n)$, where $\alpha := kq\hat{s}^{4kq} + r\hat{s}^{rd}$.*

For classes of networks that are not k -separable, although no longer linear, the complexity is reduced to the degree and separability properties, allowing solutions on much larger networks than is possible using more naive methods. For example, the algorithm can find all fixed points of these dynamics on geometric random networks, such as in the bottom panel of Fig. 1, provided they are sufficiently sparse and separable. However, some other standard graph topologies are less well-suited to this approach. For example, since Erdos-Renyi random graphs have no local connectivity bias, the partition separability will always grow with the size of the network, and while scale free networks might have some local connectivity, the large degrees of hub nodes can be prohibitive even for the calculation of local equilibrium states. Determining the partition separability of a given network is itself a computationally complex problem, but there exist fast graph partitioning algorithms, e.g. [29, 30], which can be used to approximate this quantity. See **SI: Computational complexity analysis** for a comparison of the separability of several different network types, as well as a detailed analysis and proof of the computational complexity.

Case Study 1: Finding the critical points in brain networks

A long-standing model of neural networks endows each neuron with a binary state representing whether it is active (firing) or inactive (not firing) [31]. Edges in such a model represent synaptic connections, which can be either excitatory, i.e. the firing of one neuron causes another to be more likely to fire, or inhibitory, i.e. the firing of one neuron causes another to be less likely to fire. Activation of a neuron is determined by a threshold on its presynaptic (incoming) activity.

In theoretical neuroscience, this class of model at the neuronal scale was initially used as a simplified descriptor of associative memories [8], where each fixed point in the system represents a complex but stationary pattern of neural activity corresponding to a particular memory, and, in the case of attractors, the region of attraction around the fixed point determines how easily the memory is recalled. Later work has examined stochastic variants of this model, closely related to the Ising model, to assess the role of correlations within neural populations and their effect on neural coding and information processing [32]. More recently,

such models have also been used as the basis of theoretical studies to show how properties of neuronal networks might arise as a consequence of associative learning, wherein associations are encoded as equilibria of the underlying networks [10]. The notion of energy via the Ising Hamiltonian is central to these analyses.

However, despite the relative simplicity of this model, many contemporary approaches rely on simulation studies and post-hoc statistical analyses [33]. Because the number of parameters in the model increases quickly with network size, such simulation-based approaches are inherently limited. Effort is thus being directed towards methods that can simplify the analyses of such models, especially as it pertains the model dynamics [34].

The overwhelming majority of functional network characterizations rely on static, graphical descriptions based on calculation of the pairwise correlation of the activity over a set of brain regions. Using dynamical systems models at such scales offers the potential for greater explanatory power relative to purely statistical descriptions [35, 36]. Such models span different spatiotemporal levels of description, from detailed biophysical models at neuronal scales, to mesoscopic, mean-field approximations at the level of brain regions. Often, these models incur a tradeoff between model complexity (including physiological interpretability/abstraction) and analytical tractability. The Ising framework provides a more abstract description of brain network activation, and has primarily been used as a means of analysis, especially in regards to establishing correlation between brain regions. However, the model itself does carry dynamical interpretability and recent efforts have highlighted its potential to generate hypotheses regarding the dynamical underpinnings of observed statistical outputs. For example, [37] uses the Ising framework to suggest how the architecture of a whole-brain scale network might impact its ‘repertoire’ of achievable states, as encoded through the network’s equilibria. Such model-based analysis and hypothesis generation is in the spirit of what we highlight in our Case Study. Indeed, an emerging area of study in functional neuroimaging pertains to nonstationary fluctuations in network covariance [38]. At a dynamical level, such fluctuations are fundamentally mediated by the attractor landscape of the system. However, given the high dimensionality, finding and characterizing all attractors is generally intractable. Here, as a first step to an attractor landscape characterization, we quantify the fixed points of 783 functional brain networks from the Human Connectome Project by ascribing linear threshold dynamics to the functional connectivity networks. The

average node degree resulting from this processing pipeline is $0.05n^2$ (see **Methods: Processing of data from Human Connectome Project**).

We observe in Fig. 2 that the number of fixed points increases exponentially with network size, with an exponent that varies for different resting state networks (RSNs). RSNs are subnetworks of the brain attributed to particular task specializations, including those closely associated with sensory inputs (Auditory, Visual) and others involved in more abstract cognitive tasks [39]. For example, there is substantial evidence of an RSN called Default Mode that seem to be more active during rest and less active during the performance of externally cued tasks [40]. In this study, the largest exponent corresponds to the frontoparietal network, suggestive of a large encoding capacity relative to networks such as the ventral attention network, which possesses fewer stationary states. The frontoparietal network in particular is thought to be a key mediator of cognitive control [41, 42], though the extent to which its dynamics facilitate such is as yet opaque. Though our case study is an exploratory proof of concept, these results are intriguing as they suggest that this network may have a favorable dynamical landscape for encoding a diversity of cognitive states.

Figure 2: The number of critical points of discrete dynamical systems constructed from resting-state functional brain networks grows exponentially with network size, with an exponent that varies over the different resting state networks. The data comes from 783 participants in the Human Connectome Project (see **Methods**).

Case Study 2: Analysis of a coarse-grained lattice model for protein folding

The biological function of a protein is highly dependent on its structure, i.e. how the linear sequence of amino acids folds onto itself in 3D space. While this can occur in many different ways, it is well-established that the most probable structures are those that minimize free energy, which is a function of hydrophobic and other intramolecular forces. Although a real protein sequence is composed of 20 amino acids, a reduced-order model can provide valuable insight for protein analysis and design. In particular, we use a model with 3 types of amino acids or residues: those that have positive charge, those that have negative charge, and those in which hydrophobic forces dominate. We consider a 3D lattice structure in which the nodes take values $x_i \in \mathcal{S}_i := \{-1, 0, +1\}$, subject to the following energy function, which is a modification of the

Hydrophobic Polar model, first introduced in [43]:

$$\Phi(\mathbf{x}) := \sum_{i=1}^n \phi_i(x_i) = \sum_{i=1}^n \sum_{j \in \mathcal{N}_i} x_i x_j + h_i(x_i),$$

where $\mathcal{N}_i := \{j \in \mathcal{V}_0 : (i, j) \in \mathcal{E}_0\}$ for each $i \in \mathcal{V}_0$ and

$$h_i(x_i) := \begin{cases} 1 & \text{if } x_i = 0 \text{ and } |\mathcal{N}_i| < d_{\max}, \\ b & \text{if } |x_i| = 1 \text{ and } |\mathcal{N}_i| = d_{\max}, \\ 0 & \text{otherwise,} \end{cases}$$

where d_{\max} denotes the maximum node degree (six for the 3D lattice) and b is a model parameter. We define the dynamics of the system such that at each time k , a random node updates to the state that minimizes the local energy $\phi_i(x_i)$. The point of this is not to directly model any physical dynamics, which are inherently stochastic, but rather to construct a deterministic system in which the set of fixed points is equivalent to the set of local minima of the energy function. We used the proposed algorithm to find all fixed points of

Figure 3: Optimal protein configurations on a $4 \times 4 \times 4$ lattice. Blue and red indicate positive and negative charges, respectively, and green corresponds to hydrophobic residues.

Figure 4: The 32 global energy minima in the irregular model, represented in the 32 rows of the color matrix, can be decomposed into 8 components, the nodes of which always take the same residue type, indicated by the color of the corresponding column entry. For example, the configuration shown on the left appears in row 12 of the color matrix, indicated by the arrow.

this model for two different cases: a cubic lattice of dimension $4 \times 4 \times 4$, and a more irregular structure also having 64 nodes, with $b = 10$. Since each node can take three values, there are a total of 3^{64} possible configurations in both models. For the lattice model, we find that 587,636,902 are fixed points (a fraction 1.7×10^{-22} of the total), and of these, only two are global minima. For the irregular model, there are 60,594,240 fixed points (a fraction of 1.8×10^{-23}), 32 of which are global minima. Both optimal folding sequences for the regular lattice are depicted in Fig. 3, and the 32 global minima for the irregular model are shown in terms of eight independent components in Fig. 4. We can also plot the energy distributions of all fixed points for both models, as shown in Fig. 5. Perhaps counterintuitively, we observe a broader and more rugged fixed-point landscape in the cubic lattice than in the irregular lattice.

Note that cubic lattice structures are not k -separable for some constant value of k as the dimensions grow; rather, they are s^2 -separable, where s denotes the dimension of one side. As a result, regular lattice

Figure 5: The distribution of energy minima in the cubic lattice (top) is both broader and more rugged than that of the irregular lattice (bottom).

structures are not ideally suited to this approach. The advantage of this method becomes more striking for structures with some irregularities, in which the network topology is more suitable for partitioning. Indeed, the energy distribution of the irregular lattice was computed in approximately $\frac{1}{30}$ the time needed for the cubic lattice.

Despite its simple appearance, the lattice model captures an essential challenge in protein folding and design: tackling sequence space. Even on this simplified model, the complexity of the problem is sufficient that an exhaustive search through sequence space is not possible. The standard practice for accessing sequence space relies on Monte Carlo searches, which fail to capture a vast majority of the fixed points for a given topology [44]. In the case of protein design, vast computational resources are spent sampling sequence space with the hopes of finding a sufficiently low local energy minimum. The proposed approach has the potential to significantly enhance this effort by providing a tractable means to map entire fixed-point landscapes, to within limits on protein size (see also **Discussion**).

Discussion

The methods presented in this paper provide significant advances in the ability to computationally characterize the dynamics of complex networked systems with discrete state spaces. In the context of whole-brain modeling (Case Study 2), the discrete-state model may be informative on second-to-second time-scales, upon which populations can be thought of as being active or not. Here, however, the coarseness of the state space precludes analysis over finer time-scales, when brain regions undergo more nuanced patterns of activation. Findings such as those we presented in Case Study 2 should thus be interpreted at the appropriate spatial and temporal scale, and not as a characterization of dynamics over continuous state spaces. Because our proposed approach is not restricted to binary models, somewhat richer analyses could be performed by augmenting the Ising-framework with additional (discrete) states.

Thus, from a practical standpoint, the proposed method will perform best for systems that can be modeled with a few discrete states per unit and with relatively sparse connection structure. When either the

number of discrete states or node degrees become large, as in the case of scale-free networks for example, then even the first step of computing the local equilibrium sets becomes intractable. For example, analysis of a full 20-amino-acid protein-folding model or a micro-scale model of brain activity, where neurons have thousands of synaptic connections, would currently lie outside the capabilities of the proposed algorithm as implemented on standard personal computers.

Methods

Sparse set notations and operations

It is convenient to represent a subset of the state space $\mathcal{S} := \prod_{i=1}^n \mathcal{S}_i$ as the set of global states in which some of the nodes take specified values while others are allowed to take any value. This can be expressed as a sparse vector or a set of (*index, value*) pairs. For example, define $\mathbf{a} = \{\mathbf{x} \in \{-1, 1\}^5 : x_3 = 1, x_5 = -1\}$. We could alternatively denote this as $\mathbf{a} = \{(*, *, 1, *, -1) \in \mathcal{S}\}$ in which $*$ indicates that an entry may take any value. Although illustrative, this notation may become cumbersome for higher dimensional systems, so we introduce the following more compact sparse representation:

$$\mathbf{a} = \{(i_k, v_k)_{k=1}^w\}_{\mathcal{S}} := \{\mathbf{x} \in \mathcal{S} : x_{i_k} = v_k \forall k \in \{1, \dots, w\}\},$$

where w is the number of nodes whose state is defined. According to this notation, the example above would be expressed as $\mathbf{a} = \{(3, 1), (5, -1)\}_{\mathcal{S}}$, where $\mathcal{S} = \{-1, 1\}^5$. We also use the notation $\mathbf{a}[i]$ to access the defined element $i \in \mathcal{V}_{\mathbf{a}}$, where $\mathcal{V}_{\mathbf{a}}$ denotes the set of all nodes assigned in \mathbf{a} . In the example, $\mathbf{a}[3] = 1$ and $\mathbf{a}[5] = -1$. We will achieve our objective primarily through the use of intersections and unions of these sparse representations of large subsets of the state space \mathcal{S} . The following example demonstrates these operations on two sparsely represented subsets $\mathbf{a}, \mathbf{b} \subset \{-1, 1\}^5$.

$$\mathbf{a} := \{(3, 1), (5, -1)\}_{\mathcal{S}}$$

$$\mathbf{b} := \{(2, 1), (3, 1), (4, 1)\}_{\mathcal{S}}$$

$$\mathbf{a} \cap \mathbf{b} = \{(2, 1), (3, 1), (4, 1), (5, -1)\}_{\mathcal{S}}$$

$$\mathbf{a} \cup \mathbf{b} = \{(3, 1)\}_{\mathcal{S}}$$

We see here that computing the union of two sets is achieved by performing an intersection of the sparse representations and vice versa, which allows for efficient computation of these operations.

We sometimes need to represent ordered sets of these sparsely represented subsets. For example, define $\mathcal{A} = (\mathbf{a}_1, \dots, \mathbf{a}_\alpha)$ and $\mathcal{B} = (\mathbf{b}_1, \dots, \mathbf{b}_\beta)$ for some arbitrary integers α and β . Commonly used operations on these sets include the set of all pairwise intersections and the set of all pairwise unions, which we respectively denote by

$$\mathcal{A} \cap \mathcal{B} := \{\mathbf{a}_i \cap \mathbf{b}_j, \forall i \in \{1, \dots, \alpha\}, \forall j \in \{1, \dots, \beta\}\},$$

$$\mathcal{A} \cup \mathcal{B} := \{\mathbf{a}_i \cup \mathbf{b}_j, \forall i \in \{1, \dots, \alpha\}, \forall j \in \{1, \dots, \beta\}\}.$$

To denote the set of all m -way intersections and unions, we write

$$\bigcap_{i=1}^m \mathcal{A}_i := \mathcal{A}_1 \cap \mathcal{A}_2 \cap \dots \cap \mathcal{A}_m, \quad \bigcup_{i=1}^n \mathcal{A}_i := \mathcal{A}_1 \cup \mathcal{A}_2 \cup \dots \cup \mathcal{A}_n.$$

It is straightforward to confirm that the commutative, associative, and distributive properties hold for these operations. It will also be convenient to relate whether or not elements of an ordered set \mathcal{A} are subsets of elements of \mathcal{B} . For this purpose, we define the matrix of size $\alpha \times \beta$:

$$\Delta(\mathcal{A}, \mathcal{B}) := [\delta_{ij}(\mathcal{A}, \mathcal{B})]_{\alpha \times \beta}, \quad (2)$$

in which each entry is given by

$$\delta_{ij}(\mathcal{A}, \mathcal{B}) := \begin{cases} 1, & \text{if } \mathbf{a}_i \supseteq \mathbf{b}_j \\ 0, & \text{otherwise} \end{cases}. \quad (3)$$

In addition, let $\delta_{i*}(\mathcal{A}, \mathcal{B}) := \{j : \delta_{ij}(\mathcal{A}, \mathcal{B}) = 1\}$ denote the index set of elements of \mathcal{B} that are contained in \mathbf{a}_i , and let $\delta_{*j}(\mathcal{A}, \mathcal{B}) := \{i : \delta_{ij}(\mathcal{A}, \mathcal{B}) = 1\}$ denote the index set of elements of \mathcal{A} that contain \mathbf{b}_j .

Recursive partition-based intersections

The next step after computing the LEQ sets (as described in **Results: Finding and optimizing over all fixed points**) is to partition the network into several smaller subnetworks and compute the regional equilibria on those subnetworks. The corresponding REQ sets can then be intersected using only the states of their boundary nodes to obtain a sparse representation of the global equilibrium set. At the end, we have all the information needed to quickly access any equilibrium without explicitly computing the whole set.

This method relies heavily on graph partitioning, which is itself a complex problem and an active research area. Although any valid graph partitioning algorithm can be used here, the choice may have a significant effect on performance. For example, an algorithm that tries to balance subgroup sizes while minimizing the number of edges that cross partition boundaries will generally lead to faster computation times than one that does not. It is also computationally advantageous for our purposes if the nodes on the partition boundaries have small LEQ sets. Optimally balanced graph partitioning is itself a computationally difficult problem. Fortunately, we do not require an optimal partition; rather, a reasonably good approximation will suffice, and recent research has shown the existence of algorithms that are tractable with respect to a parameterization of the requirements[45]. Indeed there are several approaches to choose from that are fast to compute relative to finding the fixed points, and we encourage researchers to use one of the several high-performance open-source graph partitioning software packages [46, 47, 48]. In our case studies and simulations, we used a spectral partitioning algorithm based on a k-means clustering of the eigenvectors of the weighted adjacency matrix, with weights proportional to the size of the LEQ sets [29].

Before proceeding, we introduce some definitions and notations related to partitioning the network. Let $\bar{\mathcal{V}} := \{\mathcal{V}_1, \dots, \mathcal{V}_q\}$ denote a *partition* of the node set $\mathcal{V} \subseteq \mathcal{V}_0$ into q subsets, such that $\mathcal{V} = \bigcup_{k=1}^q \mathcal{V}_k$ and $\mathcal{V}_j \cap \mathcal{V}_k = \emptyset$ for all $j, k \in \{1, \dots, q\}, j \neq k$. Further, we define the set of external boundary nodes ($\mathcal{V}_k^{\text{out}}$) as the set of nodes that are not in \mathcal{V}_k but are adjacent (either to or from) a node in \mathcal{V}_k . In contrast, we define the internal boundary nodes ($\mathcal{V}_k^{\text{in}}$) as the set of nodes that are in \mathcal{V}_k and are adjacent to a node that is not in \mathcal{V}_k . We also define the union of these two sets as the set of combined boundary nodes ($\mathcal{V}_k^{\text{bound}}$). Formally, these sets are defined as follows:

$$\begin{aligned} \mathcal{V}_k^{\text{out}} &:= \bigcup_{i \in \mathcal{V}_k} \mathcal{N}_i \setminus \mathcal{V}_k, & \mathcal{V}_k^{\text{in}} &:= \bigcup_{i \in \mathcal{V}_k^{\text{out}}} \mathcal{N}_i \cap \mathcal{V}_k, \\ \mathcal{V}_k^{\text{bound}} &:= \mathcal{V}_k^{\text{in}} \cup \mathcal{V}_k^{\text{out}}, \end{aligned}$$

where $\mathcal{N}_i := \{j \in \mathcal{V} : (i, j) \in \mathcal{E} \text{ or } (j, i) \in \mathcal{E}\}$ denotes the set of all in- and out-neighbors of node i . We denote the set of interior nodes, having no out-of-group neighbors by

$$\check{\mathcal{V}}_k := \mathcal{V}_k \setminus \mathcal{V}_k^{\text{in}}.$$

Lastly, we respectively denote the union of all boundary nodes in a partition by

$$\mathcal{V}^{\text{part}} := \bigcup_{k=1}^q \mathcal{V}_k^{\text{bound}}. \quad (4)$$

By intersecting LEQ sets on a node set $\mathcal{V} \subseteq \mathcal{V}_0$, we form a set of *regional equilibrium* (REQ) states.

$$\Omega(\mathcal{V}) := \bigcap_{i=1}^{|\mathcal{V}|} \mathcal{L}_i. \quad (5)$$

Let r be a free parameter that defines the maximum size of a node set that we will not partition further, but rather compute the full equilibrium set using (5). The set $\mathcal{Z}(\mathcal{V})$ represents the full state space of the nodes in \mathcal{V} and the empty set everywhere else:

$$\mathcal{Z}(\mathcal{V}) := \{\mathbf{x} \in \mathcal{S} : x_i = \emptyset, \forall i \notin \mathcal{V}\}. \quad (6)$$

We are now ready to introduce the core of the proposed method, which is the following recursive definition of a compact representation of the REQ sets $\Omega(\mathcal{V})$:

Equilibrium points on partition boundary nodes

$$\Omega^p(\mathcal{V}) := \begin{cases} \Omega(\mathcal{V}) & \text{if } |\mathcal{V}| \leq r \\ \bigcap_{k=1}^q \Omega^b(\mathcal{V}_k) & \text{if } |\mathcal{V}| > r \end{cases}, \quad (7)$$

$$\Omega^b(\mathcal{V}_k) := \Omega^p(\mathcal{V}_k) \cup \mathcal{Z}(\check{\mathcal{V}}_k). \quad (8)$$

The set $\Omega^p(\mathcal{V})$ represents the set of all possible equilibrium values for nodes that lie on the boundary of the partition of $\bar{\mathcal{V}}$. At the base level where $|\mathcal{V}| \leq r$, the problem is small enough to solve by intersecting all contained LEQ sets as in (5), and $\Omega^p(\mathcal{V})$ is equal to the full regional equilibrium set. Otherwise, we obtain $\Omega^p(\mathcal{V})$ by intersecting the sets $\Omega^b(\mathcal{V})$ for each partitioned subgroup. $\Omega^b(\mathcal{V}_k)$ is the set of all possible regional equilibrium values at the nodes that lie on the boundary of \mathcal{V}_k with some other partitioned subnetwork \mathcal{V}_j , $j \neq k$. This set is obtained by taking the union of the set $\Omega^p(\mathcal{V}_k)$ (defined recursively) with $\mathcal{Z}(\check{\mathcal{V}}_k)$. Taking the union of $\mathcal{Z}(\mathcal{V})$ with some sparsely represented set has the effect of decreasing the amount of information needed to represent the resulting set, since it will no longer contain any information about the nodes in \mathcal{V} . Notice that a graph partition $\bar{\mathcal{V}}$ is needed to compute the lower expression in (7).

It is useful to store the results of this computation in the form of a tree graph $\mathcal{T} := (\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}})$, in which each node $V_{\mathcal{V}} \in \mathcal{V}_{\mathcal{T}}$ contains the partition $\bar{\mathcal{V}}$, the set $\Omega^p(\mathcal{V})$ of possible equilibrium states on the partition boundary nodes, and the sets $\Omega^b(\mathcal{V}_k)$ and corresponding Δ matrices for each partitioned subnetwork:

$$V_{\mathcal{V}} := (\bar{\mathcal{V}}, \Omega^p(\mathcal{V}), \{\Omega^b(\mathcal{V}_k), \Delta(\Omega^p(\mathcal{V}), \Omega^b(\mathcal{V}_k)), \\ \Delta(\Omega^b(\mathcal{V}_k), \Omega^p(\mathcal{V}_k))\}_{k=1}^q)$$

for each partition $\bar{\mathcal{V}}$ performed in the course of the recursion. The edges $\mathcal{E}_{\mathcal{T}}$ in the tree simply connect each tree node $V_{\mathcal{V}}$ to the nodes corresponding to the partition of \mathcal{V} , i.e. $(V_{\mathcal{V}}, V_{\mathcal{V}_k}) \in \mathcal{E}_{\mathcal{T}}$ for each $k \in \{1, \dots, q\}$. As a result, we obtain a compact representation of the global equilibrium set. Specifically, $\Omega^p(\mathcal{V}_0)$ is equivalent to the union of the global equilibrium set with all state configurations of nodes interior to the top level partition, as formalized in the following theorem (proof in Appendix).

Theorem 2. *For any partition $\bar{\mathcal{V}} = \{\mathcal{V}_1, \dots, \mathcal{V}_k\}$ of an arbitrary node set \mathcal{V} such that $|\mathcal{V}| > r$,*

$$\Omega^p(\mathcal{V}) = \Omega(\mathcal{V}) \uplus \mathcal{Z}(\mathcal{V}^{part}).$$

When paired with the information contained in \mathcal{T} , there is sufficient information in $\Omega^p(\mathcal{V})$ to efficiently compute the number of equilibria, the full equilibrium set, or only those equilibria that are optimal with respect to some cost function.

Figure 6: Example showing a network in which nodes can take one of two states marked by black or white (gray indicates nodes that can take any value), partitioned into two subnetworks \mathcal{V}_1 (blue) and \mathcal{V}_2 (red). We show two of the eight LEQ sets \mathcal{L}_1 and \mathcal{L}_8 . Below them are the sets $\Omega^p(\mathcal{V}_1)$ and $\Omega^p(\mathcal{V}_2)$, generated from the intersection of their constituent LEQ sets. Further below are the compactly represented sets $\Omega^b(\mathcal{V}_1)$ and $\Omega^b(\mathcal{V}_2)$ on the two partition boundary nodes. Finally, we show the top-level set $\Omega^p(\mathcal{V})$, and the global equilibrium set $\Omega(\mathcal{V}_0)$ constructed from the information in \mathcal{T} .

To provide further intuition into the approach, we illustrate a small example in Figure 6. The network shown in the center is partitioned into groups indicated by the colored node outlines. Best-response dynamics with an anticoordinating payoff matrix of $\begin{pmatrix} 0 & 3 \\ 1 & 0 \end{pmatrix}$ was used for all edges in the network (see **SI: Best response and linear threshold dynamics**). Two of the eight LEQ sets (\mathcal{L}_1 and \mathcal{L}_8) are shown towards the top, with the elements of the set in a stacked arrangement. The two PREQ sets ($\Omega^p(\mathcal{V}_1)$ and $\Omega^p(\mathcal{V}_2)$) corresponding to the intersections of the LEQ sets within their respective partitioned subgroups are shown at the far left and

right. Compact representations of these sets on their boundary nodes are shown in the BREQ sets ($\Omega^b(\mathcal{V}_1)$ and $\Omega^b(\mathcal{V}_2)$) just underneath. Finally, the top level PREQ set $\Omega^p(\mathcal{V})$ resulting from the intersection of the two BREQ sets and the expanded global equilibrium set $\Omega(\mathcal{V}_0)$ appear at the bottom.

Existence and number of equilibria

The number of equilibria can be computed using the same recursion structure as above and can occur either simultaneous to searching the equilibrium space or in post-processing. To do this, at the lowest level of partitioning (the top case of (10)), we simply count in (9) how many regional equilibrium states are included in each element of the PREQ set. These quantities are then passed to the higher level in the partition tree along with the requisite bookkeeping steps, ultimately resulting in the total number of equilibria at the top level, as specified in the following theorem (proof in the **Appendix**).

Theorem 3. *Given the tree \mathcal{T} resulting from the computation (7)-(6) on system (1), the total number of equilibria $|\Omega(\mathcal{V}_0)|$ is given by*

$$n_{eq}(\Omega^p(\mathcal{V}_0)) := \sum_{j=1}^{|\Omega^p(\mathcal{V}_0)|} \check{\lambda}_j(\mathcal{V}_0), \quad (9)$$

where

$$\begin{aligned} \check{\lambda}_j(\mathcal{V}) &:= \begin{cases} 1, & \text{if } |\mathcal{V}| \leq r \\ \prod_{k=1}^q \lambda_j(\mathcal{V}_k), & \text{otherwise} \end{cases}, \\ \lambda_j(\mathcal{V}_k) &:= \begin{cases} 0, & \text{if } \delta_{*j}(\Omega^p(\mathcal{V}), \Omega^b(\mathcal{V}_k)) = \emptyset \\ \gamma_i(\mathcal{V}_k), \quad i \in \delta_{*j}(\Omega^p(\mathcal{V}), \Omega^b(\mathcal{V}_k)), & \text{otherwise} \end{cases}, \\ \gamma_i(\mathcal{V}_k) &:= \sum_{j=1}^{|\Omega^p(\mathcal{V}_k)|} \check{\lambda}_j(\mathcal{V}_k) \delta_{ij}(\Omega^b(\mathcal{V}_k), \Omega^p(\mathcal{V}_k)). \end{aligned} \quad (10)$$

Optimization on the global equilibrium set

Within the set of all fixed points, it is often useful to find the equilibria that optimize a given cost or energy function. An advantage of this approach is that the optimal equilibria can be quickly extracted from the results of the computation without generating the full equilibrium set, which can be very large in some cases.

Optimization using this approach works for the class of global cost functions arising from the linear superposition of possibly nonlinear local cost functions. Let $\Phi(\mathbf{x}) : \mathcal{S} \rightarrow \mathbb{R}$ denote a cost function on the state space \mathcal{S} , which is the summation of local node and edge costs:

$$\Phi(\mathbf{x}) := \sum_{i=1}^n h_i(x_i) + \sum_{i=1}^n \sum_{j \in \mathcal{N}_i} \phi_{ij}(x_i, x_j). \quad (11)$$

Suppose that we seek the set of equilibria that minimize this function:

$$\Omega^*(\mathcal{V}_0) := \{\mathbf{x} \in \Omega(\mathcal{V}_0) : \Phi(\mathbf{x}) \leq \Phi(\mathbf{y}), \forall \mathbf{y} \in \Omega(\mathcal{V}_0)\}. \quad (12)$$

It is also useful to define the cost of a particular local or regional equilibrium set, which we denote by $\omega(\mathcal{V}) \in \Omega(\mathcal{V})$, where $\Omega(\mathcal{V}) := \{\omega_1(\mathcal{V}), \dots, \omega_\alpha(\mathcal{V})\}$ and $\alpha := |\Omega(\mathcal{V})|$. Recalling that $\mathbf{x}[i]$ denotes the fixed state of node i in the sparse representation set $\mathbf{x} \subseteq \mathcal{S}$, we note that $\omega(\mathcal{V})[i]$ has the same interpretation for the equilibrium set $\omega(\mathcal{V})$. We can now express the cost of $\omega(\mathcal{V})$ as follows:

$$\Phi(\omega(\mathcal{V})) := \sum_{i \in \mathcal{V}} h_i(\omega(\mathcal{V})[i]) + \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{N}_i \cap \mathcal{V}} \phi_{ij}(\omega(\mathcal{V})[i], \omega(\mathcal{V})[j]). \quad (13)$$

At the lowest level of the partition tree, the cost of each element $\omega(\mathcal{V})$ is computed according to (13). Notice that while the costs associated with all nodes of \mathcal{V} are included, only the costs of edges connecting pairs of nodes in \mathcal{V} are included, while costs related to edges containing nodes not in \mathcal{V} are excluded. These inter-region costs are and defined as follows:

$$\tilde{\Phi}(\omega(\mathcal{V})) := \sum_{k=1}^q \sum_{i \in \mathcal{V}_k^{\text{in}}} \sum_{j \in \mathcal{N}_i \cap \mathcal{V}_k^{\text{out}}} \phi_{ij}(\omega(\mathcal{V})[i], \omega(\mathcal{V})[j]). \quad (14)$$

Although the cost of a base-level equilibrium set ω is unique and well-defined, the cost of a PREQ element $\omega^p(\mathcal{V})$ or BREQ element $\omega^b(\mathcal{V})$ may take multiple values, since they encapsulate the set of fixed points for all configurations of the interior nodes. Hence, we define the costs of these elements as ordered sets:

$$\begin{aligned} \phi(\mathcal{V}) &:= (\{\phi : \exists i : \Phi(\omega_i^p(\mathcal{V})) = \phi\}, \rho), \\ &= (\{\phi : \exists i : \Phi(\omega_i^b(\mathcal{V})) = \phi\}, \rho). \end{aligned} \quad (15)$$

Note that the set of feasible cost values in the sets $\Omega^p(\mathcal{V})$ and $\Omega^b(\mathcal{V})$ are identical, since they include the same underlying REQ set $\Omega(\mathcal{V})$. Next, we define a matrix that maps the elements of a PREQ set to the

number of ways each unique cost value is attainable in the respective node groups, in which we have labeled the rows and columns for ease of interpretation:

$$W^p(\mathcal{V}) := \begin{matrix} & \omega_1^p(\mathcal{V}) & \cdots & \omega_\mu^p(\mathcal{V}) \\ \begin{matrix} \phi_1(\mathcal{V}) \\ \vdots \\ \phi_\nu(\mathcal{V}) \end{matrix} & \begin{bmatrix} w_{11}^p & \cdots & w_{1\mu}^p \\ \vdots & \cdots & \vdots \\ w_{\nu 1}^p & \cdots & w_{\nu\mu}^p \end{bmatrix} & , & w_{ij}^p := \left| \{ \mathbf{x} \in \omega_j^p(\mathcal{V}) : \Phi(\mathbf{x}) = \phi_i(\mathcal{V}) \} \right|, \end{matrix} \quad (16)$$

where $\mu := |\Omega^p(\mathcal{V})|$, $\nu := |\phi^p(\mathcal{V})|$, and the matrix $W^b(\mathcal{V})$ is defined similarly.

Clearly, computing the quantities (15)-(16) explicitly would require generating the entire set $\Omega(\mathcal{V})$. Fortunately, there is a more efficient way to compute these using the partition tree structure \mathcal{T} and the cost values attained by lower level REQ sets. To help with this task, it will be convenient to define two sets of matrices. First, the matrix $\hat{W}(\mathcal{V}_k)$ maps the unique costs achievable by elements of the lower level BREQ sets $\Omega^b(\mathcal{V}_k)$ to the current level PREQ elements $\omega^p(\mathcal{V})$ with which the costs are compatible:

$$\hat{W}(\mathcal{V}_k) := \begin{matrix} & \omega_1^p(\mathcal{V}) & \cdots & \omega_\mu^p(\mathcal{V}) \\ \begin{matrix} \phi_1(\mathcal{V}_k) \\ \vdots \\ \phi_{\nu_k}(\mathcal{V}_k) \end{matrix} & \begin{bmatrix} \hat{w}_{11} & \cdots & \hat{w}_{1\mu} \\ \vdots & \cdots & \vdots \\ \hat{w}_{\nu_k 1} & \cdots & \hat{w}_{\nu_k \mu} \end{bmatrix} & , & \hat{w}_{ij} := \left| \{ \mathbf{x} \in \Omega^b(\mathcal{V}_k) \cap \{ \omega_j^p(\mathcal{V}_k) \} : \Phi(\mathbf{x}) = \phi_i(\mathcal{V}_k) \} \right|, \end{matrix}$$

where $\nu_k := |\phi(\mathcal{V}_k)|$. Second, the matrix $\hat{Z}(\mathcal{V}_k)$ contains the costs associated with each entry in $\hat{W}(\mathcal{V}_k)$, or ∞ if the entry corresponds to an infeasible pairing:

$$\hat{Z}(\mathcal{V}_k) := \begin{matrix} & \omega_1^p(\mathcal{V}) & \cdots & \omega_\mu^p(\mathcal{V}) \\ \begin{matrix} \phi_1(\mathcal{V}_k) \\ \vdots \\ \phi_{\nu_k}(\mathcal{V}_k) \end{matrix} & \begin{bmatrix} \hat{z}_{11} & \cdots & \hat{z}_{1\mu} \\ \vdots & \cdots & \vdots \\ \hat{z}_{\nu_k 1} & \cdots & \hat{z}_{\nu_k \mu} \end{bmatrix} & , & \hat{z}_{ij} := \begin{cases} \phi_i(\mathcal{V}_k) & \text{if } \hat{w}_{ij} > 0 \\ \infty & \text{if } \hat{w}_{ij} = 0 \end{cases}.$$

Next, we construct expanded matrices $\mathring{W}(\mathcal{V}_k)$ and $\mathring{Z}(\mathcal{V}_k)$ for each subgroup k . Let $\mathring{\phi}(\mathcal{V})$ denote the sums of all configurations of the unique subgroup costs $\phi(\mathcal{V}_k)$, of which the total number is $|\mathring{\phi}(\mathcal{V})| = \prod_{k=1}^q |\phi(\mathcal{V}_k)|$.

Then the expanded matrix $\hat{W}(\mathcal{V}_1)$ is given by:

$$\hat{W}(\mathcal{V}_1) := \begin{array}{l} \phi_1(\mathcal{V}_1) + \cdots + \phi_1(\mathcal{V}_{q-1}) + \phi_1(\mathcal{V}_q) = \dot{\phi}_1(\mathcal{V}) \\ \vdots \\ \phi_1(\mathcal{V}_1) + \cdots + \phi_1(\mathcal{V}_{q-1}) + \phi_{\nu_q}(\mathcal{V}_q) = \dot{\phi}_{\nu_q}(\mathcal{V}) \\ \phi_1(\mathcal{V}_1) + \cdots + \phi_2(\mathcal{V}_{q-1}) + \phi_1(\mathcal{V}_q) = \dot{\phi}_{\nu_q+1}(\mathcal{V}) \\ \vdots \\ \phi_{\nu_1}(\mathcal{V}_1) + \cdots + \phi_{\nu_{q-1}}(\mathcal{V}_{q-1}) + \phi_{\nu_q}(\mathcal{V}_q) = \dot{\phi}_{\nu}(\mathcal{V}) \end{array} \begin{bmatrix} \omega_1^p(\mathcal{V}) & \cdots & \omega_\mu^p(\mathcal{V}) \\ \hat{w}_{11} & \cdots & \hat{w}_{1\mu} \\ \vdots & \cdots & \vdots \\ \hat{w}_{11} & \cdots & \hat{w}_{1\mu} \\ \hat{w}_{11} & \cdots & \hat{w}_{1\mu} \\ \vdots & \cdots & \vdots \\ \hat{w}_{\nu_1 1} & \cdots & \hat{w}_{\nu_1 \mu} \end{bmatrix},$$

where \hat{w}_{ij} is equal to the ij^{th} entry of $\hat{W}(\mathcal{V}_1)$. The successive matrices $\hat{W}(\mathcal{V}_k)$ are defined similarly, with \hat{w}_{ij} appearing in the same row where $\phi_i(\mathcal{V}_k)$ appears in the summations yielding the entries of $\dot{\phi}(\mathcal{V})$.

Likewise, the expanded matrix $\hat{Z}(\mathcal{V}_1)$ is given by:

$$\hat{Z}(\mathcal{V}_1) := \begin{array}{l} \phi_1(\mathcal{V}_1) + \cdots + \phi_1(\mathcal{V}_{q-1}) + \phi_1(\mathcal{V}_q) = \dot{\phi}_1(\mathcal{V}) \\ \vdots \\ \phi_1(\mathcal{V}_1) + \cdots + \phi_1(\mathcal{V}_{q-1}) + \phi_{\nu_q}(\mathcal{V}_q) = \dot{\phi}_{\nu_q}(\mathcal{V}) \\ \phi_1(\mathcal{V}_1) + \cdots + \phi_2(\mathcal{V}_{q-1}) + \phi_1(\mathcal{V}_q) = \dot{\phi}_{\nu_q+1}(\mathcal{V}) \\ \vdots \\ \phi_{\nu_1}(\mathcal{V}_1) + \cdots + \phi_{\nu_{q-1}}(\mathcal{V}_{q-1}) + \phi_{\nu_q}(\mathcal{V}_q) = \dot{\phi}_{\nu}(\mathcal{V}) \end{array} \begin{bmatrix} \omega_1^p(\mathcal{V}) & \cdots & \omega_\mu^p(\mathcal{V}) \\ \hat{z}_{11} & \cdots & \hat{z}_{1\mu} \\ \vdots & \cdots & \vdots \\ \hat{z}_{11} & \cdots & \hat{z}_{1\mu} \\ \hat{z}_{11} & \cdots & \hat{z}_{1\mu} \\ \vdots & \cdots & \vdots \\ \hat{z}_{\nu_1 1} & \cdots & \hat{z}_{\nu_1 \mu} \end{bmatrix},$$

where \hat{z}_{ij} is equal to the ij^{th} entry of $\hat{Z}(\mathcal{V}_1)$. Next, let $\tilde{\phi}(\mathcal{V}) := [\tilde{\Phi}(\omega_1^p(\mathcal{V})), \dots, \tilde{\Phi}(\omega_{|\Omega^p(\mathcal{V})|}^p(\mathcal{V}))]$ denote a row vector of inter-group costs. The combined expanded cost matrix $\hat{Z}(\mathcal{V})$ and the number of instantiations of each cost $\dot{W}(\mathcal{V})$ can be expressed as follows:

$$\begin{aligned} \dot{W}(\mathcal{V}) &:= \dot{W}(\mathcal{V}_1) \odot \dot{W}(\mathcal{V}_2) \odot \cdots \odot \dot{W}(\mathcal{V}_q), \\ \hat{Z}(\mathcal{V}) &:= \sum_{k=1}^q \hat{Z}(\mathcal{V}_k) + \tilde{\phi}(\mathcal{V}) \otimes \mathbf{1}_{|\hat{Z}_1(\mathcal{V}_1)|}, \end{aligned}$$

where \odot denotes element-wise multiplication and \otimes denotes the Kronecker product. Finally, the compact matrix $W^p(\mathcal{V})$ is constructed by extracting the unique cost values in $\hat{Z}(\mathcal{V})$ and counting the total number of ways each cost is attained in $\dot{W}(\mathcal{V})$. We again refer the reader to the Cost algorithm in **SI** for details on this procedure. The optimal equilibrium set $\Omega^*(\mathcal{V}_0)$ can then be obtained by following a similar procedure, as described in detail in the Expand algorithm (see **SI**). Note that by using the null cost function $\Phi(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{S}$, the set $\Omega^*(\mathcal{V}_0)$ is equivalent to the global equilibrium set $\Omega(\mathcal{V}_0)$ of (1).

Competing interests

The authors declare no competing interests.

Code availability

Code used to implement the methods described in this paper are available at the following link: <https://github.com/jrriehl/FindEQ>.

Acknowledgments

We are grateful to all the funders of this work, which was supported by grant 15RT0189 from the US Air Force Office of Scientific Research, grants 1537015, 1509342, 1653589 from the US National Science Foundation, and by the National Science Foundation CAREER Award MCB-1552471. S.C. and G.R.B. hold Career Awards at the Scientific Interface from the Burroughs-Wellcome Fund. M.I.Z. holds a Monsanto Graduate Fellowship and a Center for Biological Systems Engineering Fellowship. M.F.S. was partially supported by NSF-DGE-1143954.

Author contributions

J.R.R. and S.C. conceived and developed the technical theory, simulations, performed the case studies, and prepared the manuscript. M.I.Z. and G.R.B. conceived and analyzed the protein folding case study. M.F.S. prepared the data from the Human Connectome Project. All authors contributed to writing and editing the manuscript.

References

- [1] Makoto Kaburagi and Junjiro Kanamori. A method of determining the ground state of the extended-range classical lattice gas model. *Progress of Theoretical Physics*, 54(1):30–44, 1975.
- [2] Alexander K Hartmann. Ground states of two-dimensional ising spin glasses: fast algorithms, recent developments and a ferromagnet-spin glass mixture. *Journal of Statistical Physics*, 144(3):519, 2011.

- [3] Wenxuan Huang, Daniil A Kitchaev, Stephen T Dacek, Ziqin Rong, Alexander Urban, Shan Cao, Chuan Luo, and Gerbrand Ceder. Finding and proving the exact ground state of a generalized ising model by convex optimization and max-sat. *Physical Review B*, 94(13):134424, 2016.
- [4] P Chandra, P Coleman, and AI Larkin. Ising transition in frustrated heisenberg models. *Physical review letters*, 64(1):88, 1990.
- [5] Rhiju Das, Bin Qian, Srivatsan Raman, Robert Vernon, James Thompson, Philip Bradley, Sagar Khare, Michael D Tyka, Divya Bhat, Dylan Chivian, et al. Structure prediction for casp7 targets using extensive all-atom refinement with rosetta@ home. *Proteins: Structure, Function, and Bioinformatics*, 69(S8):118–128, 2007.
- [6] Firas Khatib, Frank DiMaio, Seth Cooper, Maciej Kazmierczyk, Mirosław Gilski, Szymon Krzywda, Helena Zabranska, Iva Pichova, James Thompson, Zoran Popović, et al. Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature Structural and Molecular Biology*, 18(10):1175, 2011.
- [7] Adam L Beberg, Daniel L Ensign, Guha Jayachandran, Siraj Khaliq, and Vijay S Pande. Folding@ home: Lessons from eight years of volunteer distributed computing. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–8. IEEE, 2009.
- [8] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [9] Daniel J Amit. *Modeling brain function: The world of attractor neural networks*. Cambridge university press, 1992.
- [10] Danke Zhang, Chi Zhang, and Armen Stepanyants. Robust associative learning is sufficient to explain structural and dynamical properties of local cortical circuits. *bioRxiv*, page 320432, 2018.
- [11] Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127, 2010.

- [12] Gustavo Deco, Viktor K Jirsa, and Anthony R McIntosh. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12(1):43, 2011.
- [13] Andrea Avena-Koenigsberger, Bratislav Mistic, and Olaf Sporns. Communication dynamics in complex brain networks. *Nature Reviews Neuroscience*, 19(1):17, 2018.
- [14] Michael Breakspear. Dynamic models of large-scale brain activity. *Nature neuroscience*, 20(3):340, 2017.
- [15] Nancy J Kopell, Howard J Gritton, Miles A Whittington, and Mark A Kramer. Beyond the connectome: the dynamome. *Neuron*, 83(6):1319–1328, 2014.
- [16] Peter C de Ruiter, Anje-Margriet Neutel, and John C Moore. Modelling food webs and nutrient cycling in agro-ecosystems. *Trends in Ecology & Evolution*, 9(10):378–383, 1994.
- [17] Francisco Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241, 1982.
- [18] Jeremy Kun, Brian Powers, and Lev Reyzin. Anti-coordination games and stable graph colorings. In *International Symposium on Algorithmic Game Theory*, pages 122–133. Springer, 2013.
- [19] Sergio Boixo, Troels F Rønnow, Sergei V Isakov, Zhihui Wang, David Wecker, Daniel A Lidar, John M Martinis, and Matthias Troyer. Evidence for quantum annealing with more than one hundred qubits. *Nature Physics*, 10(3):218, 2014.
- [20] Sergei V Isakov, Ilia N Zintchenko, Troels F Rønnow, and Matthias Troyer. Optimised simulated annealing for ising spin glasses. *Computer Physics Communications*, 192:265–271, 2015.
- [21] AZ Maksymowicz, JE Galletly, MS Magdon, and IL Maksymowicz. Genetic algorithm approach for ising model. *Journal of magnetism and magnetic materials*, 133(1-3):40–41, 1994.
- [22] Volker Blum, Gus LW Hart, Michael J Walorski, and Alex Zunger. Using genetic algorithms to map first-principles results to model hamiltonians: Application to the generalized ising model for alloys. *Physical Review B*, 72(16):165113, 2005.

- [23] Limor Drori and David Peleg. Faster exact solutions for some np-hard problems. *Theoretical Computer Science*, 287(2):473–499, 2002.
- [24] Alessandro Di Cara, Abhishek Garg, Giovanni De Micheli, Ioannis Xenarios, and Luis Mendoza. Dynamic simulation of regulatory networks using squad. *BMC bioinformatics*, 8(1):462, 2007.
- [25] Yin Zhao, Jongrae Kim, and Maurizio Filippone. Aggregation algorithm towards large-scale boolean network analysis. *IEEE Transactions on Automatic Control*, 58(8):1976–1985, 2013.
- [26] Desheng Zheng, Guowu Yang, Xiaoyu Li, Zhicai Wang, Feng Liu, and Lei He. An efficient algorithm for computing attractors of synchronous and asynchronous boolean networks. *PloS one*, 8(4):e60593, 2013.
- [27] Elena Dubrova and Maxim Teslenko. A sat-based algorithm for finding attractors in synchronous boolean networks. *IEEE/ACM transactions on computational biology and bioinformatics*, 8(5):1393–1399, 2011.
- [28] Wensheng Guo, Guowu Yang, Wei Wu, Lei He, and Mingyu Sun. A parallel attractor finding algorithm based on boolean satisfiability for genetic regulatory networks. *PloS one*, 9(4):e94258, 2014.
- [29] Joao P Hespanha. An efficient matlab algorithm for graph partitioning. *University of California*, pages 1–8, 2004.
- [30] Maurizio Filippone, Francesco Camastra, Francesco Masulli, and Stefano Rovetta. A survey of kernel and spectral methods for clustering. *Pattern recognition*, 41(1):176–190, 2008.
- [31] Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.
- [32] Elad Schneidman, Michael J Berry II, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007, 2006.

- [33] David B Kastner, Stephen A Baccus, and Tatyana O Sharpee. Critical and maximally informative encoding between neural populations in the retina. *Proceedings of the National Academy of Sciences*, page 201418092, 2015.
- [34] Rubén Herzog, Maria-Jose Escobar, Rodrigo Cofre, Adrian G Palacios, and Bruno Cessac. Dimensionality reduction on spatio-temporal maximum entropy models of spiking networks. *bioRxiv*, page 278606, 2018.
- [35] Michael Schirner, Anthony Randal McIntosh, Viktor Jirsa, Gustavo Deco, and Petra Ritter. Inferring multi-scale neural mechanisms with brain network modelling. *Elife*, 7:e28927, 2018.
- [36] Joana Cabral, Morten L Kringelbach, and Gustavo Deco. Exploring the network dynamics underlying brain activity during rest. *Progress in neurobiology*, 114:102–131, 2014.
- [37] Gustavo Deco, Mario Senden, and Viktor Jirsa. How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Frontiers in computational neuroscience*, 6:68, 2012.
- [38] Martin A Lindquist, Yuting Xu, Mary Beth Nebel, and Brain S Caffo. Evaluating dynamic bivariate correlations in resting-state fmri: a comparison study and a new approach. *NeuroImage*, 101:531–546, 2014.
- [39] Christian F Beckmann, Marilena DeLuca, Joseph T Devlin, and Stephen M Smith. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):1001–1013, 2005.
- [40] Michael D Greicius, Ben Krasnow, Allan L Reiss, and Vinod Menon. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, 100(1):253–258, 2003.
- [41] Michael W Cole, Tal Yarkoni, Grega Repovš, Alan Anticevic, and Todd S Braver. Global connectivity of prefrontal cortex predicts cognitive control and intelligence. *Journal of Neuroscience*, 32(26):8988–8999, 2012.

- [42] Todd S Braver, Jeremy R Reynolds, and David I Donaldson. Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron*, 39(4):713–726, 2003.
- [43] Ken A Dill. Theory for the folding and stability of globular proteins. *Biochemistry*, 24(6):1501–1509, 1985.
- [44] Po-Ssu Huang, Scott E Boyken, and David Baker. The coming of age of de novo protein design. *Nature*, 537(7620):320, 2016.
- [45] Daniel Lokshtanov and Dániel Marx. Clustering with local restrictions. *Information and Computation*, 222:278–292, 2013.
- [46] George Karypis and Vipin Kumar. Multilevel algorithms for multi-constraint graph partitioning. In *SC'98: Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, pages 28–28. IEEE, 1998.
- [47] Chris Walshaw and Mark Cross. Jostle: parallel multilevel graph-partitioning software—an overview.
- [48] Charalampos Tsourakakis, Christos Gkantsidis, Bozidar Radunovic, and Milan Vojnovic. Fennel: Streaming graph partitioning for massive scale graphs. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 333–342, 2014.
- [49] Daniel S Marcus, Michael P Harms, Abraham Z Snyder, Mark Jenkinson, J Anthony Wilson, Matthew F Glasser, Deanna M Barch, Kevin A Archie, Gregory C Burgess, Mohana Ramaratnam, et al. Human connectome project informatics: quality control, database services, and data visualization. *Neuroimage*, 80:202–219, 2013.
- [50] Jonathan D Power. A simple but useful way to assess fmri scan qualities. *Neuroimage*, 154:150–158, 2017.
- [51] Matthew F Glasser, Stamatios N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, et al. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.

- [52] Jonathan D Power, Bradley L Schlaggar, and Steven E Petersen. Recent progress and outstanding issues in motion correction in resting state fmri. *Neuroimage*, 105:536–551, 2015.
- [53] Ludovica Griffanti, Gholamreza Salimi-Khorshidi, Christian F Beckmann, Edward J Auerbach, Gwenaëlle Douaud, Claire E Sexton, Enikő Zsoldos, Klaus P Ebmeier, Nicola Filippini, Clare E Mackay, et al. Ica-based artefact removal and accelerated fmri acquisition for improved resting state network imaging. *Neuroimage*, 95:232–247, 2014.
- [54] Evan M Gordon, Timothy O Laumann, Babatunde Adeyemo, Jeremy F Huckins, William M Kelley, and Steven E Petersen. Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cerebral Cortex (New York, NY)*, 26(1):288, 2016.
- [55] Jonathan D Power, Alexander L Cohen, Steven M Nelson, Gagan S Wig, Kelly Anne Barnes, Jessica A Church, Alecia C Vogel, Timothy O Laumann, Fran M Miezin, Bradley L Schlaggar, et al. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011.
- [56] Enzo Tagliazucchi, Dante R Chialvo, Michael Siniatchkin, Enrico Amico, Jean-Francois Brichant, Vincent Bonhomme, Quentin Noirhomme, Helmut Laufs, and Steven Laureys. Large-scale signatures of unconsciousness are consistent with a departure from critical dynamics. *Journal of The Royal Society Interface*, 13(114):20151027, 2016.

Appendices

Processing of data from Human Connectome Project

The data used consisted of resting state fMRI scans collected from 783 subjects during the Human Connectome Project (HCP). These subjects were part of the s900 release which contained a total of 818 subjects. However, 35 of these subjects which had been marked as having instabilities in the head coil were removed [49, 50]. Each subject underwent a total of four 15-minute scans divided among two separate scanning sessions. The two scans per session corresponded to a left-right and a right-left acquisition. Data was acquired at 3T with a temporal resolution (TR) of 720ms. Detailed information on HCP scan protocols are contained in [51].

The publicly available s900 data release includes the following processing procedure. Following the standard minimal preprocessing pipeline [51], the motion correction strategies recommended by Power and colleagues were applied [52], with the exception of Frame Censoring (removing time points), which the authors acknowledge has proven controversial. The remaining motion correction procedure consists of the FSL ICA-FIX correction [53] and regressing out the twelve HCP motion parameters and global signal (cerebrospinal fluid, white matter, and grey matter). To remove signal drift, we applied a .009 Hz high-pass filter and detrended to remove respiratory artifact. Data were parcellated into the Gordon atlas [54] and functional connectivities were computed using the Pearson correlation between parcels for the combined data across scans. More details on the task specializations associated with the various resting-state networks can be found in [55].

To suppress weak correlations and to facilitate comparison between subjects that are independent of overall connectivity variations, we thresholded correlation coefficients to preserve 5% of all pairwise connections, following similar methods as in [56]. Fixed points were computed for best-response dynamics in which the payoff matrices between each pair of regions i and j is given by $J_{ij}I_2$, where J_{ij} denotes the Pearson correlation coefficient between respective regions and I_2 is the 2×2 identity matrix. These dynamics are guaranteed to converge to a local minimum of the Ising Hamiltonian:

$$\Phi(\mathbf{x}) := \sum_{i=1}^n \phi_i(\mathbf{x}), \quad \text{where} \quad \phi_i(\mathbf{x}) := - \sum_{j=1}^n J_{ij} x_i x_j.$$

Analysis of a coarse-grained lattice model for protein folding

For this case study, the recursive partitioning algorithm was augmented with two stages of trimming, one to eliminate infeasible entries in the LEQ sets, and a second to eliminate known infeasible entries in the intermediate REQ sets. These steps are possible due to the regularity of the lattice structure and the uniformity of the cost function across nodes, and are generally good practice when applicable because of the resulting computational savings.

Proof of Theorem 2

Proof. Since $|\mathcal{V}| > r$, according to (7) and (8),

$$\Omega^p(\mathcal{V}) = \bigcap_{k=1}^q (\Omega^p(\mathcal{V}_k) \uplus \mathcal{Z}(\check{\mathcal{V}}_k)). \quad (17)$$

Pulling out the first term in this expression yields

$$\Omega^p(\mathcal{V}) = (\Omega^p(\mathcal{V}_1) \uplus \mathcal{Z}(\check{\mathcal{V}}_1)) \cap \bigcap_{k=2}^q (\Omega^p(\mathcal{V}_k) \uplus \mathcal{Z}(\check{\mathcal{V}}_k))$$

which can be expanded to

$$\begin{aligned} \Omega^p(\mathcal{V}) &= \left(\Omega^p(\mathcal{V}_1) \cap \bigcap_{k=2}^q (\Omega^p(\mathcal{V}_k) \uplus \mathcal{Z}(\check{\mathcal{V}}_k)) \right) \\ &\quad \uplus \left(\mathcal{Z}(\check{\mathcal{V}}_1) \cap \bigcap_{k=2}^q (\Omega^p(\mathcal{V}_k) \uplus \mathcal{Z}(\check{\mathcal{V}}_k)) \right). \end{aligned}$$

Due to (6) and the fact that the partition is disjoint, we know that $\mathcal{Z}(\check{\mathcal{V}}_1) \subseteq \Omega^p(\mathcal{V}_k)$ and that $\mathcal{Z}(\check{\mathcal{V}}_1) \cap \mathcal{Z}(\check{\mathcal{V}}_k) = \emptyset$ for all $k \in \{1, \dots, q\}$. It follows that

$$\Omega^p(\mathcal{V}) = \left(\Omega^p(\mathcal{V}_1) \cap \bigcap_{k=2}^q (\Omega^p(\mathcal{V}_k) \uplus \mathcal{Z}(\check{\mathcal{V}}_k)) \right) \uplus \mathcal{Z}(\check{\mathcal{V}}_1).$$

Noticing that the middle term is in the same form as (17), we can repeatedly apply the previous two steps to obtain

$$\Omega^p(\mathcal{V}) = \bigcap_{k=1}^q \Omega^p(\mathcal{V}_k) \uplus \bigcup_{k=1}^q \mathcal{Z}(\check{\mathcal{V}}_k). \quad (18)$$

Next, suppose that $|\mathcal{V}_k| > r$ for each k and let $\bar{\mathcal{V}}_k$ denote a partition of each \mathcal{V}_k . We can then perform the exact same expansion of each $\Omega^p(\mathcal{V}_k)$ from (17) to (18), resulting in

$$\Omega^p(\mathcal{V}) = \bigcap_{k=1}^q \left(\bigcap_{j=1}^q \Omega^p(\mathcal{V}_{k,j}) \uplus \bigcup_{j=1}^q \mathcal{Z}(\check{\mathcal{V}}_{k,j}) \right) \uplus \bigcup_{k=1}^q \mathcal{Z}(\check{\mathcal{V}}_k).$$

Similarly as above, since $\mathcal{Z}(\check{\mathcal{V}}_{i,j}) \subseteq \Omega^p(\mathcal{V}_{k,j})$ and $\mathcal{Z}(\check{\mathcal{V}}_{i,j}) \cap \mathcal{Z}(\check{\mathcal{V}}_{k,j}) = \emptyset$ for all $i, j, k \in \{1, \dots, q\}$, $i \neq k$, we have

$$\Omega^p(\mathcal{V}) = \bigcap_{k=1}^q \left(\bigcap_{j=1}^q \Omega^p(\mathcal{V}_{k,j}) \right) \uplus \bigcup_{j=1}^q \mathcal{Z}(\check{\mathcal{V}}_{k,j}) \uplus \bigcup_{k=1}^q \mathcal{Z}(\check{\mathcal{V}}_k).$$

Since $\mathcal{Z}(\check{\mathcal{V}}_{k,j}) \subseteq \mathcal{Z}(\check{\mathcal{V}}_k)$ for each $j \in \{1, \dots, q\}$, we can simplify as follows:

$$\Omega^p(\mathcal{V}) = \bigcap_{k=1}^q \left(\bigcap_{j=1}^q \Omega^p(\mathcal{V}_{k,j}) \right) \cup \bigcup_{k=1}^q \mathcal{Z}(\check{\mathcal{V}}_k).$$

By repeatedly performing this expansion according to (8), the term on the left will eventually become the union of base-level regional equilibrium sets $\Omega(\mathcal{V}_{k,j,\dots})$, which is equal to $\Omega(\mathcal{V})$. Furthermore, since $\bigcup_{k=1}^q \mathcal{Z}(\check{\mathcal{V}}_k) = \mathcal{Z}(\bigcup_{k=1}^q \check{\mathcal{V}}_k) = \mathcal{Z}(\mathcal{V}^{\text{part}})$, due to (4), we obtain

$$\Omega^p(\mathcal{V}) = \bigcap_{k=1}^q \Omega(\mathcal{V}_k) \cup \mathcal{Z}(\mathcal{V}^{\text{part}}).$$

Finally, since

$$\bigcap_{k=1}^q \Omega(\mathcal{V}_k) = \bigcap_{k=1}^q \bigcap_{i=1}^{|\mathcal{V}_k|} \mathcal{L}_i = \Omega(\mathcal{V}),$$

we have the desired result $\Omega^p(\mathcal{V}) = \Omega(\mathcal{V}) \cup \mathcal{Z}(\mathcal{V}^{\text{part}})$ and the proof is completed. \square

Proof of Theorem 3

Proof. Suppose $|\mathcal{V}| \leq r$. Then

$$n_{\text{eq}}(\Omega^p(\mathcal{V})) = \sum_{j=1}^{|\Omega^p(\mathcal{V})|} 1 = |\Omega^p(\mathcal{V})| = |\Omega(\mathcal{V})|, \quad (19)$$

where we used the fact that $|\Omega^p(\mathcal{V})| = |\Omega(\mathcal{V})|$ from (7). Next, suppose $|\mathcal{V}| > r$. Then we have

$$n_{\text{eq}}(\Omega^p(\mathcal{V})) = \sum_{j=1}^{|\Omega^p(\mathcal{V})|} \prod_{k=1}^q \lambda_j(\mathcal{V}_k), \quad (20)$$

where $|\Omega^p(\mathcal{V})|$ is the number of unique configurations of partition boundary nodes in \mathcal{V} that appear in the global equilibrium set, as proved in Theorem 2. Given one of these configurations, the regional equilibria in each group are independent from each other. Therefore, it suffices to show that $\lambda_j(\mathcal{V}_k)$ is equal to the number of REQ states associated with the set \mathcal{V}_k that are contained in $\omega_j^p(\mathcal{V})$, i.e. $\lambda_j(\mathcal{V}_k) = n_{\text{eq}}(\omega_j^p(\mathcal{V}) \cap \Omega^p(\mathcal{V}_k))$.

We can check for compatibility between each entry of $\Omega^p(\mathcal{V})$ and the constituent BREQ states $\Omega^b(\mathcal{V}_k)$ using $\delta_{*j}(\Omega^p(\mathcal{V}), \Omega^b(\mathcal{V}_k))$. If there exists j such that this set is empty for some k , then $\lambda_j(\mathcal{V}_k) = 0$ and there are no equilibria. Otherwise, $\delta_{*j}(\Omega^p(\mathcal{V}), \Omega^b(\mathcal{V}_k))$ is a singleton since $\Omega^b(\mathcal{V}_k)$ is defined on a subset of the nodes defined in $\Omega^p(\mathcal{V})$. One can then check that the quantity $\gamma_i(\mathcal{V}_k)$ is precisely equivalent

to $n_{\text{eq}} \left(\omega_j^p(\mathcal{V}) \cap \Omega^p(\mathcal{V}_k) \right)$, which is then equal to $|\Omega \left(\omega_j^p(\mathcal{V}) \cap \Omega^p(\mathcal{V}_k) \right)|$ if $|\mathcal{V}_k| \leq r$, and otherwise the number of compatible lower-level equilibria, and the proof proceeds by induction. □

Proof of Theorem 1

Proof. The overall computation involves three types of operations: local equilibrium checks, sparse set intersections, and computing the Δ mappings.

Computing the local equilibria: Finding the local equilibrium states at each node involves evaluating the local update rule for a maximum of \hat{s}^d total state configurations. Multiplying by the number of nodes n yields a computational complexity of $O(n\hat{s}^d)$ for this step.

Sparse set intersections: The sparse set intersections occur at two levels. At the base level, the intersection of a pair of local equilibrium sets involves checking a maximum of $(\hat{s}^d)^2$ overlapping state configurations. The result is then checked against a third local equilibrium set, and the process repeats until the maximum group size r is reached, yielding a total for each base-level partition group of

$$\hat{s}^{2d} + 2\hat{s}^{3d} + \dots + (r-1)\hat{s}^{rd} \leq r^2\hat{s}^{rd}.$$

Assuming maximum group sizes, this computation is performed approximately $\frac{n}{r}$ times, yielding an upper bound of $O(rn\hat{s}^{rd})$ for the local intersections.

Intersections at all higher levels depend on the number of boundary nodes between neighboring partition groups, of which there are a maximum of $2k$. However, the states of all boundary nodes must be represented for each partition group. Since there up to q partition groups, all of which may be adjacent to each other, the maximum dimension in each group is $2kq$. Hence, each intersection involves checking up to $(\hat{s}^{2kq})^2$ overlapping state configurations, each of which may contain up to $2kq$ nodes, resulting in a complexity of $O(2kq\hat{s}^{4kq})$ for each partition that is performed. Since the total number of partitions is bounded above by n , we have a total complexity of $O(kqn\hat{s}^{4kq})$ for the higher level sparse set intersections. The computational complexity of all sparse set intersections is then no greater than $O(rn\hat{s}^{rd} + kn\hat{s}^{4kq})$.

Computing the Δ matrices: The Δ matrices defined in (2) in the main article map regional equilibrium states between neighboring levels in the partition hierarchy. Since the maximum number of configurations

on each level is $O(\hat{s}^{2kq})$, the number of entries in each Δ matrix is at most $O(\hat{s}^{4kq})$. Since each entry δ_{ij} requires at most $2kq$ computations (see (3) in the article), the computational complexity for constructing each Δ matrix is bounded by $O(kq\hat{s}^{4kq})$. Finally, since both the number of nodes and edges in the tree \mathcal{T} are less than n , the number of Δ matrices to compute is at most $2n$, resulting in a complexity of $O(kqn\hat{s}^{4kq})$.

The overall computational complexity of computing \mathcal{T} is therefore bounded by $O(kqn\hat{s}^{4kq} + rn\hat{s}^{r\hat{d}})$. \square