

How Wires Shape Volumes

On the relation between network topology and nonlocal power grid stability

DISSERTATION

zur Erlangung des akademischen Grades

doctor rerum naturalium

(Dr. rer. nat.)

im Fach Physik

eingereicht an der

Mathematisch-Naturwissenschaftlichen Fakultät I

Humboldt-Universität zu Berlin

von

Dipl.-Phys. Peter Johannes Menck

Präsident der Humboldt-Universität zu Berlin:

Prof. Dr. Jan-Hendrik Olbertz

Dekan der Mathematisch-Naturwissenschaftlichen Fakultät I:

Prof. Stefan Hecht PhD

Gutachter:

1. Prof. Dr. Dr. h.c. mult. Jürgen Kurths

2. Prof. Dr. Marc Timme

3. Dr. Michael Zaks

eingereicht am: 08. Oktober 2013

Tag der mündlichen Prüfung: 04. März 2014

Für Annbritt

Abstract

In electrical energy systems, hundreds of power plants are coupled via thousands of kilometres of transmission lines in order to reliably serve millions of customers. The sheer complexity of interactions between all these components is a never-ending source both of trouble for system operators who strive to keep their power grid stable, and of challenges for complex systems scientists seeking fundamental understanding.

This thesis addresses the question how the stability of a power grid depends on the network topology. Two pieces of methodological equipment are therefore required: first, a concept of stability; and, second, a framework for the description of network structures. Stability is assessed here by means of basin stability, a nonlocal and nonlinear inspection scheme that allows to quantify how stable a power grid is against large perturbations such as short circuits or major load switchings. The other part of the toolbox consists of the instruments of complex network theory that have been widely used to characterize networked systems.

To the best of the author's knowledge, basin stability has not yet been applied in the context of complex systems science. Therefore, a rigorous definition is provided that clarifies its relation to other stability concepts. Basin stability quantifies the stability of an attractor in the middle ground between instability and global stability by estimating the volume of its basin of attraction in state space. It has several favourable properties: First, the numerical costs for estimating basin stability are not too high and, crucially, independent of the system's dimension. Second, the estimation procedure appears to work properly even when fractal basin boundaries are present. Third, basin stability can be used in a component-wise fashion particularly suitable for the study of networked systems.

Basin stability is used to assess the relation between network topology and power grid stability once in an abstract and once in a more applied way. The abstract approach is motivated by the fact that power grids rely on the synchronous operation of their components, and focusses on the problem how general dynamical networks should be structured to support a stable synchronous state. There exists a huge body of literature on this subject. Yet the topologies of real-world synchronizing networks such as power grids or neural networks have eluded previous theoretical approaches that were based on local stability concepts. Here, a complementary nonlocal-stability perspective is offered. It turns out that networks with maximum basin stability are located at the opposite end of a certain network-theoretical spectrum than networks optimal with respect to local stability. This suggests that, during the evolution of synchronizing networks, the optimization for local stability and the simultaneous optimization for nonlocal stability have acted as opposing forces. The real-world topologies actually found may be explained as a compromise between these forces.

The more applied approach draws on a power grid model from the engineering literature. A component-wise version of basin stability is used to assess how a grid's degree of stability against large single-node perturbations is influenced by certain patterns in the wiring topology. Simulations in an ensemble of power grids yield various statistics that all support one main finding: The widespread and cheapest of all connection schemes, so-called dead ends and dead trees, strongly diminish stability. A case study of the Northern European power system demonstrates that the inverse is also true: 'Healing' dead ends by addition of transmission lines substantially enhances stability. This indicates a basic design principle for future power grids: Add just a few more lines to avoid dead ends.

Zusammenfassung

Hunderte Kraftwerke sind in Stromnetzen durch Tausende Kilometer von Leitungen miteinander verknüpft, um die Versorgung von Millionen von Kunden sicherzustellen. Die Wechselwirkungen zwischen all diesen Komponenten sind hochkomplex und damit eine unerschöpfliche Quelle sowohl von Beschwerden für Ingenieure, die sich um die Stabilität ihres Stromnetzes bemühen, als auch von Herausforderungen für Wissenschaftler, die komplexe Systeme von Grund auf verstehen wollen.

Diese Arbeit ist der Frage gewidmet, wie die Stabilität eines Stromnetzes mit seiner Topologie zusammenhängt. Der methodische Ansatz muss also aus zwei Teilen bestehen: Erstens aus einem geeigneten Stabilitätskonzept; und zweitens aus einem Begriffssatz zur Beschreibung von Netzwerktopologien. Stabilität wird hier mit Hilfe von Bassin-Stabilität vermessen, einer nichtlokalen und nichtlinearen Methode, mit der sich quantifizieren lässt, wie stabil ein Stromnetz ist gegen große Störungen wie etwa Kurzschlüsse oder Lastsprünge. Den anderen Teil des methodischen Ansatzes machen die Instrumente der Theorie komplexer Netzwerke aus, die schon in vielen Forschungsfeldern eingesetzt wurden, um netzartige Systeme zu beschreiben.

Nach Kenntnis des Autors ist Bassin-Stabilität im Kontext komplexer Systeme noch nie verwendet worden. Darum wird hier zunächst eine grundlegende Definition zur Verfügung gestellt, um Bassin-Stabilität in die Reihe anderer Stabilitätskonzepte einzuordnen. Bassin-Stabilität quantifiziert die Stabilität eines Attraktors zwischen Instabilität und globaler Stabilität durch Vermessung des Volumens seines Anziehungsbereichs im Zustandsraum. Das Konzept hat einige vorteilhafte Eigenschaften: Erstens sind die numerischen Kosten zur Schätzung von Bassin-Stabilität nicht allzu hoch und, was besonders wichtig ist, unabhängig von der Dimension des Zustandsraums. Zweitens funktioniert die numerische Schätzung anscheinend sogar dann, wenn fraktale Bassinränder vorliegen. Drittens kann Bassin-Stabilität auch komponentenweise angewendet werden, was besonders nützlich für die Untersuchung von netzartigen Systemen ist.

Bassin-Stabilität wird hier benutzt, um den Zusammenhang zwischen der Stabilität eines Stromnetzes und seiner Netztopologie zu untersuchen, und zwar einmal abstrakt und einmal konkret. Der abstrakte Ansatz ist dadurch motiviert, dass die Funktion eines Stromnetzes auf dem synchronen Zusammenwirken seiner Komponenten beruht. Er betrifft die Frage, wie generelle dynamische Netzwerke strukturiert sein sollten, um stabile Synchronisation zu unterstützen. Dieses Thema hat in der Literatur schon viel Aufmerksamkeit gefunden. Dennoch konnte man die Topologien von synchron funktionierenden Systemen aus der Wirklichkeit, etwa von neuronalen Netzen oder Stromnetzen, durch bisherige theoretische Ansätze nicht erklären. Diese Ansätze basierten allerdings auf einem lokalen Stabilitätsbegriff. Hier wird eine nichtlokale Perspektive eingenommen. Es stellt sich heraus, dass Netzwerke mit maximaler Bassin-Stabilität am entgegengesetzten Ende eines gewissen netzwerktheoretischen Spektrums angesiedelt sind als Netzwerke, die optimal sind in Hinsicht auf lokale Stabilität. Dies deutet darauf hin, dass in der Entwicklung von synchron funktionierenden Systemen die Optimierung auf lokale Stabilität hin und die gleichzeitige Optimierung auf nichtlokale Stabilität hin als entgegengesetzte Kräfte gewirkt haben. Die tatsächlich gefundenen Topologien wären dann als Kompromiss zwischen diesen beiden Optimierungsprozessen erklärbar.

Der konkrete Ansatz geht aus von einem Stromnetzmodell aus der Ingenieursliteratur. Bassin-Stabilität wird komponentenweise eingesetzt, um zu untersuchen, wie der Grad der Stabilität eines Stromnetzes gegen große Einzelknoten-Störungen beeinflusst wird von gewissen Mustern in der Netztopologie. Simulationen in einem Ensemble von Stromnetzen bringen diverse Statistiken hervor, die alle eine wesentliche Beobachtung stützen: Ausgerechnet das kostengünstigste und weitverbreitetste Anschlussschema – bestehend aus Stichleitungen – vermindert die Netzstabilität beträchtlich. In einer Fallstudie des nordeuropäischen Stromnetzes zeigt sich, dass die Umkehrung auch zutrifft: Beseitigung von Stichleitungen durch Zubau neuer Leitungen erhöht die Netzstabilität deutlich. Daraus lässt sich ein fundamentales Designprinzip für zukünftige Stromnetze ableiten: Man baue immer ein paar Verbindungen mehr, um Stichleitungen zu vermeiden.

Contents

Contents	x
List of Figures	xiv
List of Tables	xv
List of frequently used symbols	xviii
1 Introduction	1
2 Stability concepts	5
2.1 Dynamical systems terminology	8
2.1.1 States and state space	8
2.1.2 Contractiveness and minimal invariant sets	9
2.1.3 Illustration	9
2.2 Local stability and strong attractors	12
2.2.1 Asymptotic stability	12
2.2.2 Lyapunov exponents	12
2.2.3 Illustration	15
2.3 Nonlocal stability and weak attractors	16
2.3.1 Lyapunov functions	16
2.3.2 Definition of basin stability	16
2.3.3 Illustration	18
2.3.4 Basin stability and the choice of ρ	20
2.4 Basin stability estimation	22
2.4.1 Numerical procedure	22
2.4.2 High-dimensional systems	22
2.4.3 Systems with fractal basin boundaries	23
2.5 Inferring stability from data	27
2.6 Conclusions and outlook	29
3 Synchronization on complex networks	31
3.1 A motivating example	34
3.2 Network terminology	36
3.2.1 Nodes, edges and the adjacency matrix	36
3.2.2 Shortest paths, connectedness and betweenness	36
3.2.3 Degree and clustering	38
3.2.4 The Laplacian matrix	39

3.3	Local stability and the puzzle it poses	40
3.3.1	Master stability function formalism	40
3.3.2	Revisiting the motivating example	43
3.3.3	Synchronizability	43
3.3.4	The puzzle	44
3.4	The nonlocal stability perspective	48
3.4.1	Basin stability in Rössler networks	48
3.4.2	The puzzle – solved!?	50
3.5	Conclusions and outlook	54
3.6	Supporting material	55
3.6.1	Another choice of the reference subset \mathcal{Q}	55
3.6.2	2-dimensional Watts-Strogatz networks	56
3.6.3	Small-World networks with different link length distributions	59
3.6.4	Non-convexity of the basin in Rössler-networks	61
3.6.5	Edge-list of the example network	62
3.6.6	Numerical methods	62
4	Structure and stability of power grids	63
4.1	The infinite-grid model	66
4.1.1	Derivation	66
4.1.2	Perturbation analysis	68
4.2	The finite-grid model	71
4.2.1	Zhukov’s aggregation	71
4.2.2	Effects of large perturbations	73
4.2.3	Definition of single-node basin stability	78
4.3	Power grid ensembles	79
4.3.1	Basin stability statistics	81
4.3.2	Effects of dead trees on grid dynamics	85
4.4	Case study: power grid of Northern Europe	90
4.5	Conclusions and outlook	93
4.6	Supporting material	94
4.6.1	Model parameters	94
4.6.2	Parameter sensitivity	95
4.6.3	Numerical methods	95
5	Conclusions and outlook	97
	Acknowledgements	101
	Appendix	103
	Glossary	107
	Bibliography	112

List of publications

The author has submitted several of the findings and analyses presented in this thesis for publication in scientific journals. They have appeared in one conference contribution

Menck, P.J. & Kurths, J. Topological identification of weak points in power grids. *Proceedings of Nonlinear Dynamics of Electronic Systems 2012*, 1 – 4, 2012.

cited below as [Menck and Kurths, 2012], and two peer-reviewed papers

Menck, P.J., Heitzig, J., Marwan, N. & Kurths, J. How basin stability complements the linear-stability paradigm. *Nature Physics* **9**, 89–92, 2013.

Menck, P.J., Heitzig, J., Kurths, J. & Schellnhuber, H.J. Sustainable power grids need smart wiring. *Under review at Nature Communications*.

cited below as [Menck et al., 2013a] and [Menck et al., 2013b].

List of Figures

2.1	Marble on a marble track.	6
2.2	A damped driven pendulum.	10
2.3	Minimal invariant sets of the damped driven pendulum.	11
2.4	Different types of attractor in 3-dimensional state spaces.	14
2.5	Basins of attraction of the damped driven pendulum.	18
2.6	Basin stability of the damped driven pendulum.	19
2.7	Local Stability of the damped driven pendulum.	20
2.8	Basin stability of the damped pendulum with Gaussian ρ	21
2.9	Error of basin stability estimation.	23
2.10	Damped pendulum with fractal basin boundaries.	25
2.11	Basin stability in the pendulum with fractal basin boundaries.	26
2.12	State diagram of a bistable stylized forest-savanna model.	28
3.1	Synchronization of a network of chaotic Rössler oscillators.	35
3.2	Example network to illustrate network definitions.	37
3.3	Master stability function of x -coupled Rössler oscillators.	42
3.4	Watts-Strogatz graph generation model.	45
3.5	Synchronizability and basin stability in Watts-Strogatz networks of chaotic oscillators.	47
3.6	Explanatory characteristics.	49
3.7	Basin stability in Rössler networks.	50
3.8	Topological comparison of ensemble results with real-world networks.	53
3.9	Basin stability with respect to another reference subset \mathcal{Q} in Watts-Strogatz networks of chaotic oscillators.	55
3.10	Initial grid for 2-dimensional Watts-Strogatz network generation.	56
3.11	Synchronizability and basin stability in 2-dimensional Watts-Strogatz networks of chaotic oscillators.	57
3.12	Topological comparison of results for the 2-dimensional ensemble with real-world networks.	58
3.13	Synchronizability and basin stability in networks of chaotic oscillators with different link length distribution.	60
3.14	Two-dimensional details of the Rössler attractor's basin	61
4.1	Generator infinite-grid model.	66
4.2	State space of the infinite-grid model.	68
4.3	Basin stability of the generator in the infinite-grid model.	70
4.4	Zhukov's aggregation.	72
4.5	Example 8-node power grid.	74
4.6	Perturbation space analysis.	76

List of Figures

4.6	Perturbation space analysis (continued).	77
4.7	Example grid from the ensemble.	80
4.8	Basin stability in the power grid ensemble.	81
4.9	Relation between basin stability and topological properties of a node.	83
4.10	Effects of dead ends and dead trees.	84
4.11	Statistics on non-synchronous states in the power grid ensemble.	88
4.11	Statistics on non-synchronous states in the power grid ensemble (continued).	89
4.12	Northern European power grid.	91
4.13	‘Healed’ Northern European power grid.	92
4.14	Supplement: Northern European power grid with increased damping.	96

List of Tables

3.1	Topological properties of synchronizing real-world networks.	51
3.2	More topological properties of synchronizing real-world networks. . . .	52

List of frequently used symbols

$\langle \cdot \rangle$	Ensemble average
$\ \cdot \ $	Norm
\mathbf{A}	Adjacency matrix
α	Damping constant
b	Shortest-path betweenness
C	Clustering coefficient
d	Degree
E	Number of edges
e	Error of basin stability estimation
θ	Phase
K	Coupling constant
L	Average shortest-path length
\mathbf{L}	Laplacian matrix
λ_2	Second smallest eigenvalue of Laplacian matrix
λ_N	Largest eigenvalue of Laplacian matrix
N	Number of nodes
P	Net active power input
p	Watts-Strogatz model parameter
\mathcal{Q}	Reference subset for basin stability estimation
R	Synchronizability
$\text{Re}\{\cdot\}$	Real part of a complex number
ρ	Probability density of perturbations in basin stability estimation
S_B	Basin stability
T	Number of initial states for basin stability estimation
t	Time
ω	Frequency

1 Introduction

From Gallipoli, an Italian seaside town, the German city of Papenburg is about 2,000 kilometres away. Therefore, the Gallipolini may have found it a little far-fetched when they first heard that the power blackout which had darkened their town on 4 November 2006, had actually had its root cause in Papenburg.

It was meant to be a routine operation: Engineers of a German power company decided to deenergize a transmission line over the river Ems near Papenburg to allow the safe passage of a large cruise ship. At a total length of about 60 kilometres, the line constitutes only a tiny segment of the more than 220,000 kilometres long pan-European power grid, and its shut-down was expected to cause no trouble. However, it triggered a series of events in which numerous further transmission lines all over Europe overloaded and automatically switched off. The grid was on the brink of collapse. More than 10 million households including in the Netherlands, Spain – and Gallipoli – had to be disconnected for up to two hours to avoid a complete blackout [UCTE, 2007].

“All this happened because a single transmission line went missing,” commented Jürgen Kurths, an expert on complex systems and the supervisor of this thesis. “Isn’t it impressive that such a seemingly small change in the network topology can make the whole grid unstable?” He thus suggested the overarching question of this thesis:

How does the network topology influence the stability of a power grid?

To tackle it, I am going to need two pieces of methodological equipment: first, a framework for dealing with network structures; and, second, a suitable concept of power grid stability. As for the first, I will employ complex network theory [Watts and Strogatz, 1998; Strogatz, 2001; Albert and Barabási, 2002; Boccaletti et al., 2006; Newman, 2003, 2010], a versatile toolkit that has found fruitful application in various fields of research including the social sciences (see, e.g., Amaral et al. [2000]), infrastructure engineering (see, e.g., Crucitti et al. [2004b]), and biology (see, e.g., Sporns and Zwi [2004]). As for the second, I will assess power grid stability by means of basin stability, a nonlocal and nonlinear concept that allows to quantify how stable a grid is.

Facets of power grid stability

Power grid stability refers to the capacity of a power system not to suffer a blackout. We will presently see that grid stability is a multi-faceted problem. In contrast, blackouts all follow the same pattern [Ewart, 1978]: At the beginning, an unplanned event causes a single grid component such as a transmission line, transformer or generator to break down. This increases the stress on other components and provokes some of them to fail, too, thus further weakening the grid. A vicious cycle is set

in motion, and in an avalanche of component failures and deteriorating operating conditions, ever more parts of the grid become disconnected until, in the worst case, the system collapses entirely.

The operators of a power grid work hard to avoid such cascading failures, trying to ensure that their system is always $(n - 1)$ -stable, that is, able to cope with the loss of a single component. Apart from that, they strive to prevent component failure in the first place. As there are various different causes due to which a component might fail, power grids need to be stable in many ways [Machowski et al., 2008].

Take, for instance, a transmission line that is overloaded. Overheating will cause the line's wires to expand and to sag lower between the poles. This reduces the distance between wires and nearby objects (such as trees) and increases the risk of a flashover that could only be cleared by switching the line off. Thus, to ensure a grid's long-term stability, overloading has to be avoided.

Yet component failures can also have short-term causes. For example, when an industrial motor is turned on, it initially draws a large current that induces a temporary drop in the voltage. If this sudden undervoltage cannot be cleared fast enough, it may grow into a severe disturbance and induce protection schemes to shut parts of the grid down. Voltage stability is achieved through the provision of reactive power, a quantity that has to be supplied by nearby equipment instead of by transfers over the grid. Hence the overall network structure, whose influence on power system performance I aim to study in this thesis, is not highly relevant for this very local aspect of grid stability.

I will therefore place the focus on another, truly nonlocal challenge for stable grid operation which is related to the balance of active power, that is, to the nonlocal requirement that, across the whole grid and at any time, the total amount of power supplied by all generators must exactly equal the total amount of power demanded by all consumers. If this delicate balance is satisfied, power flows on the grid are constant and all components operate synchronously at the grid's rated frequency (50 or 60 Hz). However, perturbations such as load switchings, short circuits or weather-induced fluctuations in renewable generation constantly threaten to push the grid away from this synchronous balance. If a component indeed becomes desynchronized, it has to be disconnected, thereby weakening the grid and possibly triggering a cascading blackout. Consequently, power grids must be designed as stable as possible against desynchronization-prone perturbations.

A nonlocal stability perspective

When assessing grid stability, one needs to distinguish between small perturbations and large perturbations. Both kinds do take place in power grids [Machowski et al., 2008]: the former, for instance, when a consumer switches on their coffee machine [Dobson, 2013], and the latter, for example, when a short circuit forces a transmission line out of service. In recent times, the relation between network topology and small-perturbation stability has been investigated a lot (e.g., in Witthaut and Timme [2012]; Rohden et al. [2012]; Dörfler et al. [2013]).

Small perturbations are highly *local* affairs, and the corresponding analysis tools, such as linearization [Argyris et al., 2010] or convergence rate measurements [Scheffer

et al., 2009], are focussed on the surroundings of a system’s current state. They can be used to assess whether or not the system will return to this state after having been pushed *a tiny bit away* by some external influence.

By contrast, large perturbations – on which I will focus here – may drive the system *far away* from its current state, and any analysis has to take into account *nonlocal* information. Indeed, the system will only return to the current state after having been hit by a large perturbation, if it is not pushed outside of that state’s so-called basin of attraction, a truly nonlocal entity.

I will use basin stability, a nonlinear measure based on the basin’s volume, to quantify how stable a state is against large perturbations. Although volumes of basins have been studied before (see, e.g., Wiley et al. [2006]), basin stability has – to the best of my knowledge – not yet been applied in the context of complex systems science.

Contents and main findings

The contents of this thesis are arranged as follows. In Chapter 2, I will set basin stability into the context of other, particularly local stability concepts, providing a rigorous definition of it as well as a numerical estimation procedure for it. Basin stability gauges by a number $S_B \in [0, 1]$ the stability of an attractor between instability ($S_B = 0$) and global stability ($S_B = 1$). It will turn out that basin stability has several advantageous properties. First, the computational costs of its estimation are independent of the system dimension and, in general, not too high. Second, basin stability appears reliable even when studying systems whose basin boundaries are fractal. These two facts render basin stability a widely applicable tool.

I will then apply basin stability to investigate the relation between power grid stability and network structure once in an abstract and once in a more applied way. The abstract approach, presented in Chapter 3, is inspired by the fact that the reliable function of a power grid crucially depends on the synchronous operation of its components. Because of this, the chapter is dedicated to the problem how the topologies of general dynamical networks should be optimally shaped to support synchrony. This subject has received considerable scientific attention. Yet the topologies of synchronizing networks found in the real world, such as power grids or neural networks, could not be explained by previous, local-stability-based approaches. Offering a complementary nonlocal-stability perspective, I will demonstrate that networks with maximum basin stability reside at the opposite end of a certain network-theoretical spectrum than networks optimized for local stability. This suggests that the optimization for local stability and the simultaneous optimization for nonlocal stability may have acted as antagonistic evolutionary forces. My conjecture is that the real-world topologies actually found, so called small-world topologies, may be explained as a compromise between the two.

In the more applied approach, spelled out in Chapter 4, I will investigate a power grid model from the engineering literature and use a component-wise version of basin stability to evaluate the effects of certain topological patterns on a grid’s degree of stability against large single-node perturbations. A comprehensive statistical analysis

1 Introduction

of simulations in an ensemble of randomly generated power grids produces one main finding: The widespread and cheapest of all connection schemes, so-called dead ends and dead trees, severely reduce stability. In a case study of the Northern European power system that confirms this result, I will demonstrate that the inverse also holds: When dead ends are ‘healed’ through addition of some transmission lines, grid stability increases substantially. This hints at a basic design principle for tomorrow’s power grids: Do not go for the cheapest option: dead ends, but build a few extra lines to avoid them.

Finally, Chapter 5 contains a summary of the main results of this thesis and a discussion of possible starting points for further research.

Please note that the theoretical terms and concepts used in this text will be defined *en passant*. However, a Glossary is offered after the main part of the text.

2 Stability concepts

Stability is a crucial property of various natural and man-made systems. When you, for instance, slightly perturb a biological cell by altering the expression level of one of its genes, the cell's internal processes usually quickly return to the phenotype they were in before the perturbation [Huang and Ingber, 2007]. Likewise, when you turn on the lights in your household and thereby push the electricity grid a tiny bit away from balanced operation, this deviation dampens out in a fraction of a second, and the grid becomes balanced again [Dobson, 2013]. These are examples of dynamical regimes to which the respective system returns after small perturbations. Such regimes are called *stable*.

A system can possess more than one stable dynamical regime. In the human brain, for instance, healthy operation co-exists with pathological regimes that correspond to ailments such as epilepsy [Babloyantz and Destexhe, 1986; Lytton, 2008]. Similarly, the thermohaline circulation, a system of ocean currents important for the Earth's climate, is believed to possess an alternative stable regime in which it would cease to transport the huge amounts of heat it conveys today [Rahmstorf, 1995]. This illustrates that a system's different regimes typically reveal qualitatively very distinct modes of behaviour. Transitions between two regimes are therefore abrupt, sometimes catastrophic events whose mechanisms have received considerable scientific attention.

There are two common types of transition. The first type occurs when the environmental conditions change in a way that renders the currently occupied regime unstable. This is called a *bifurcation*, and inevitably forces the system into another regime. The second type, called *critical transition*, can only take place in *multistable* systems. In these systems, multiple stable regimes exist at the *same* environmental conditions, and a shift between them can be triggered by a large external shock [Scheffer, 2009]. The two examples cited above, the human brain and the thermohaline circulation, are actually multistable systems. Other examples include arrays of coupled lasers [Erzgräber et al., 2005] and the Amazon rainforest [Da Silveira Lobo Sternberg, 2001; Oyama and Nobre, 2003; Hirota et al., 2011]. Power grids, the real-world systems from which this thesis mainly draws its motivation, are also multistable [Machowski et al., 2008; Chiang, 2011].

Bifurcation-related transitions can be studied by means of local stability methods, in which a dynamical regime is assessed basically by inspecting the dominant curvature of the potential energy function in the regime's surroundings (as expressed by Lyapunov exponents). The curvature's absolute value measures the speed of convergence or divergence after a small perturbation, and its sign qualifies the regime as stable or unstable (Fig. 2.1). When a regime is losing stability, the convergence speed gradually decreases until the curvature finally changes sign. This *critical slowing-down* can be detected in observational data and was suggested as an early-warning signal of an approaching bifurcation-related transition [Scheffer et al., 2009].




Figure 2.1: **Marble on a marble track.** The track is immersed in a highly viscous fluid (honey) to make the system’s state space one-dimensional. Dashed arrows indicate where the marble would roll from each position. A, B and C label fixed points, at which the marble would remain forever. Fixed point B is a stable dynamical regime. The green bar indicates B’s basin of attraction \mathcal{B} . If the marble is perturbed from B to a state within the basin, it will return to B. Such perturbations are ‘permissible’. Perturbations to states outside the basin are ‘impermissible’. The dashed parabola shows the local curvature around B, fitting the true marble track only badly in most of the basin. Figure published in [Menck et al., 2013a].

However, as we will see below, critical slowing-down does not in general warn of the other type of abrupt regime shift: critical transitions. Indeed, these are nonlocal events to which the toolbox of local stability does not apply. One needs to employ stability concepts that are based on global objects [Chiang, 2011]: *basins of attraction*. Indeed, a system will only return to its former dynamical regime after a given large external shock, if it is not pushed out of the basin of attraction of that regime (Fig. 2.1). As basins can be extremely complicated entities [Nusse and Yorke, 1996], we will here focus on a single fundamental property: the basin’s volume. Based on it, we shall define a concept called *basin stability* which can be used to quantify *how stable* a regime is against transition-prone large external shocks.

The following pages provide the conceptual background for and the rigorous definition of basin stability. We will also study how basin stability can be estimated in numerical simulations. It turns out that the numerical costs associated with this estimation are relatively low and independent of the system’s dimensionality. This means that basin stability can be applied even to high-dimensional systems. Furthermore,

we will see evidence that basin stability may be applicable even when fractal basin boundaries are present.

Contents

The chapter is structured as follows. Section 2.1 introduces the basic terminology of dynamical systems theory that will be used all over this text. Then, in Section 2.2, we will discuss the concept of local stability and its relation to Lyapunov exponents. Section 2.3 is about nonlocal stability concepts and includes the definition of basin stability. In Section 2.4, we will consider a numerical procedure to estimate basin stability and assess its applicability to high-dimensional systems and systems with fractal basin boundaries. In Section 2.5, we will deal with ways to infer stability from observational data. Section 2.6 contains conclusions and outlook.

Some of the findings presented here have been published in [Menck et al., 2013a,b].

2.1 Dynamical systems terminology

Throughout all chapters of this thesis, we will investigate time-evolving systems with the concepts of dynamical systems theory. So let us first introduce some of these concepts following the very lucid expositions by Strogatz [2001] and Argyris et al. [2010].

2.1.1 States and state space

Dynamical systems theory starts by assuming that *all aspects* of a system can be captured by a specifiable set of numbers: the system's *state*

$$\mathbf{x} = (x_1, x_2, \dots, x_m). \quad (2.1)$$

In general, a system may be able to assume a vast number of different states, and its state may vary in time. So we denote by $\mathbf{x}(t) \in \mathcal{S}$ the system's state at time t , where \mathcal{S} is the set of all possible states which is commonly referred to as the system's *state space*. The dimensionality of \mathcal{S} is m .

Our definition of a state implies that, in the absence of any external or stochastic influences, $\mathbf{x}(t_0)$ at any time t_0 incorporates the *complete information* about all future states. We can hence write the system's time evolution as

$$\mathbf{x}(t) = \Phi_t(\mathbf{x}(t_0)), \quad (2.2)$$

and define the *trajectory* pertaining to the initial state $\mathbf{x}(t_0)$ as the set

$$\{\Phi_t(\mathbf{x}(t_0)) \mid t \geq t_0\}. \quad (2.3)$$

A trajectory *converges* to a set $\mathcal{V} \subset \mathcal{S}$, if for every open neighbourhood \mathcal{U} of \mathcal{V} there is a $t_{\mathcal{U}}$ so that $\Phi_t(\mathbf{x}(t_0)) \in \mathcal{U}$ for all $t > t_{\mathcal{U}}$. For brevity, we denote convergence by $\Phi_t(\mathbf{x}(t_0)) \rightarrow \mathcal{V}$.

Eq. (2.2) can equivalently¹ be expressed in a differential way as

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}) \quad (2.4)$$

which says that the change of the system's state \mathbf{x} in time, given by its time derivative $\dot{\mathbf{x}} := d\mathbf{x}/dt$, is determined entirely by a function \mathbf{F} that only depends on \mathbf{x} .² A system that obeys an equation of the form (2.4) is called *deterministic*.³ We shall discuss many such systems in what follows.

¹Provided Φ_t is differentiable w.r.t. t . In that case, write $\dot{\mathbf{x}} = \lim_{\Delta t \rightarrow 0} [\mathbf{x}(t + \Delta t) - \mathbf{x}(t)]/\Delta t = \lim_{\Delta t \rightarrow 0} [\Phi_{t+\Delta t}(\mathbf{x}(t_0)) - \Phi_t(\mathbf{x}(t_0))]/\Delta t = \lim_{\Delta t \rightarrow 0} [\Phi_{\Delta t}(\mathbf{x}(t)) - \Phi_0(\mathbf{x}(t))]/\Delta t =: \mathbf{F}(\mathbf{x}(t))$.

²A system for which \mathbf{F} does not explicitly depend on any independent variable such as the time variable t is called *autonomous*. However, an m -dimensional *non-autonomous system* for which \mathbf{F} is t -dependent can be transformed into an autonomous one by adding an $m + 1$ -th variable $x_{m+1} = t$ and setting $F_{m+1} = 1$. So we do not lose generality by limiting ourselves to autonomous systems here.

³We only consider systems in which \mathbf{F} gives unique trajectories.

2.1.2 Contractiveness and minimal invariant sets

One has to distinguish different kinds of deterministic system [Argyris et al., 2010]. In a *conservative* system, the total amount of energy is conserved, which implies $\operatorname{div} \mathbf{F} = 0$. However, we will solely study *dissipative* systems, in which friction-like processes give rise to a continuous loss of energy to the system's surroundings. Dissipative systems are characterized by $\operatorname{div} \mathbf{F} < 0$, which has the important consequence that volumes of sets of states contract in time. More specifically, if $\Delta\mathcal{V}$ denotes a set of states that occupies the state space volume (i.e., Lebesgue measure) $\Delta V := \operatorname{Vol}(\Delta\mathcal{V}) > 0$ and is small enough to be characterized by the same value of $\operatorname{div} \mathbf{F}$ throughout, then

$$\frac{d\Delta V}{dt} = \operatorname{div} \mathbf{F} \cdot \Delta V < 0 \quad (2.5)$$

implies contraction. Consequently, bigger volumes made up of many small volumes also contract. Hence, from every individual initial state, a dissipative dynamical system will converge to a zero-volume *minimal invariant set* \mathcal{I} , that is, to a non-empty set that has the invariance property

$$\Phi_t(\mathcal{I}) := \{\Phi_t(\mathbf{x}) \mid \mathbf{x} \in \mathcal{I}\} = \mathcal{I} \text{ for all } t \quad (2.6)$$

and whose proper non-empty subsets are themselves not invariant. Indeed, the zero-volume minimal invariant sets, objects of dimensionality lower than m , play an important role for understanding the long-term behaviour of dissipative deterministic systems.

2.1.3 Illustration

Consider a classical damped pendulum, such as the one depicted in Fig. 2.2, that is driven by a constant angular acceleration X . Assume its dynamics obey

$$\dot{\phi} = \omega \quad (2.7)$$

$$\dot{\omega} = -\alpha\omega + X - K \sin \phi, \quad (2.8)$$

where ϕ denotes the pendulum's angle and ω its angular frequency. Furthermore, $\alpha > 0$ is the dissipation coefficient and $K = g/\ell$, with the gravitational acceleration g and the pendulum's length ℓ . The system has the state space $\mathcal{S} = [0, 2\pi] \times \mathbb{R}$, and $\operatorname{div} \mathbf{F} = -\alpha < 0$ implies that it is dissipative.

So we now turn to the minimal invariant sets of the pendulum, choosing $\alpha = 0.1$, $X = 0.5$, and $K = 1$. It turns out that there are several such sets. First, the pendulum possesses two *fixed points*, that is: states with $\dot{\phi} = 0$ and $\dot{\omega} = 0$. We denote them by $\mathbf{x}_i^* = (\phi_i, \omega_i)$, $i = 1, 2$, with coordinates

$$\phi_{1,2} = \arcsin\left(\frac{X}{K}\right) \quad (2.9)$$

$$\omega_{1,2} = 0 \quad (2.10)$$

where we take ϕ_1 to be the solution of \arcsin inside $[0, \pi/2]$ and $\phi_2 = \pi - \phi_1$. These




Figure 2.2: **A damped driven pendulum.** The angle of the pendulum is measured by ϕ , and its length is ℓ . X indicates a constant angular acceleration that acts on the pendulum.

two solutions correspond to the angles at which the gravitational angular acceleration exactly balances the external angular acceleration X . Besides them, there is another invariant set: a periodic limit cycle, which corresponds to an incessant rotation of the pendulum at the approximate angular frequency

$$\omega_{\text{cycle}}(t) \approx \Omega + \frac{K}{\Omega} \cos(\Omega t + \phi_0) \quad (2.11)$$

with $\Omega = X/\alpha$. These three minimal invariant sets are depicted in Fig. 2.3.

We will need Eq. (2.11) in a later chapter, so let us quickly derive it here. Note that, at $K = 0$, the pendulum has a limit cycle solution with constant frequency $\omega(t) = X/\alpha =: \Omega$ and phase $\phi(t) = \Omega t + \phi_0$. In numerical simulations one observes that, for a range of K -values larger than zero, a similar solution exists which still approximately obeys $\phi_{\text{cycle}}(t) \approx \Omega t + \phi_0$. To specify it, we write $\omega_{\text{cycle}} = \Omega + f(t)$ where $f(t)$ is to be determined. Inserting this into Eq. (2.8) yields

$$\dot{f} = -\alpha f - K \sin(\Omega t + \phi_0). \quad (2.12)$$

The specific solution to this Eq. is

$$f(t) = -\frac{\alpha K}{\Omega^2 + \alpha^2} \left(\sin(\Omega t + \phi_0) - \frac{\Omega}{\alpha} \cos(\Omega t + \phi_0) \right). \quad (2.13)$$

For $|\Omega|/\alpha = |X|/\alpha^2 \gg 1$, this is approximated by

$$f(t) \approx \frac{K}{\Omega} \cos(\Omega t + \phi_0) \quad (2.14)$$

which yields Eq. (2.11). Integrating ω_{cycle} gives

$$\phi_{\text{cycle}}(t) \approx \Omega t + \frac{K}{\Omega^2} \sin(\Omega t + \phi_0) + \phi_0, \quad (2.15)$$

which means that $\phi(t) \approx \Omega t + \phi_0$ will indeed be observed if $\Omega^2 \gg K$, as is the case for our choice of parameters $X = 0.5, K = 1, \alpha = 0.1$.




Figure 2.3: **Minimal invariant sets of the damped driven pendulum.** **a**, shown are the angles $\phi_{1,2}$ of the two fixed points $\mathbf{x}_i^* = (\phi_i, \omega_i)$, $i = 1, 2$, at which the gravitational angular acceleration exactly balances the external angular acceleration X . **b**, time series of the limit cycle's angular frequency ω .

2.2 Local stability and strong attractors

As in the example of the pendulum we just discussed, a dynamical system of the form (2.4) is often intended as a model for a real-world system. The minimal invariant sets pertaining to such a model are candidates for capturing possible long-term dynamical regimes. As real-world systems usually settle in regimes that are stable against perturbations, we next define stability for invariant sets.

2.2.1 Asymptotic stability

We define a *perturbation* $\Delta \mathbf{x}$ as a push of the system from one state to another, that is

$$\mathbf{x} \mapsto \mathbf{x} + \Delta \mathbf{x}. \quad (2.16)$$

Stability against perturbations has several key aspects that are addressed by the following definitions [Argyris et al., 2010]. A compact invariant set \mathcal{I} is called *Lyapunov stable*, if for every open neighbourhood \mathcal{U} of \mathcal{I} one can find an open neighbourhood $\mathcal{V} \subset \mathcal{U}$ so that $\Phi_t(\mathcal{V}) \subset \mathcal{U}$ for all t . Lyapunov stability implies that, if a perturbation pushes the system from $\mathbf{x} \in \mathcal{I}$ to a state $\mathbf{x} + \Delta \mathbf{x}$ that is still close enough to \mathcal{I} , the ensuing trajectory will also remain close to \mathcal{I} forever. Building upon this, a compact invariant set \mathcal{I} is called *asymptotically stable* or *locally stable*, if it is Lyapunov stable and has an open neighbourhood \mathcal{U} so that for any open neighbourhood \mathcal{V} of \mathcal{I} there is a $t_{\mathcal{V}}$ so that $\Phi_t(\mathcal{U}) \subset \mathcal{V}$ for $t > t_{\mathcal{V}}$. In this case, the system not only stays close to \mathcal{I} after a small perturbation but its trajectory also finally converges back to it. We refer to an asymptotically stable minimal invariant set as a (*strong*) *attractor*. If an attractor loses stability through a change in the system parameters or if, conversely, a formerly unstable minimal invariant set becomes stable, this is called a *bifurcation*.

2.2.2 Lyapunov exponents

The convergence and divergence of trajectories nearby and on minimal invariant sets can instructively be characterized in terms of Lyapunov exponents. Following the exposition in Argyris et al. [2010], let us consider a dynamical system that evolves along a reference trajectory on a compact minimal invariant set \mathcal{I} , with its state at time t denoted by $\mathbf{x}_r(t)$ and starting at $\mathbf{x}_r(t_0)$. We will compare to this trajectory the evolution along another trajectory $\mathbf{x}(t)$ that starts at the slightly perturbed initial state $\mathbf{x}(t_0) = \mathbf{x}_r(t_0) + \Delta \mathbf{x}(t_0)$. Writing $\mathbf{x}(t) = \mathbf{x}_r(t) + \Delta \mathbf{x}(t)$, this gives

$$\dot{\mathbf{x}}(t) = \dot{\mathbf{x}}_r(t) + \Delta \dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}_r(t) + \Delta \mathbf{x}(t)) \quad (2.17)$$

$$= \mathbf{F}(\mathbf{x}_r(t)) + D\mathbf{F}(\mathbf{x}_r(t))\Delta \mathbf{x}(t) + \dots \quad (2.18)$$

where $D\mathbf{F}$ denotes the Jacobian matrix of \mathbf{F} . If $\Delta \mathbf{x}(t)$ is infinitesimally small, we can neglect higher-order terms and obtain a time-dependent linear system of differential

equations that determines $\Delta \mathbf{x}$, namely

$$\Delta \dot{\mathbf{x}} = D\mathbf{F}(\mathbf{x}_r(t))\Delta \mathbf{x}(t). \quad (2.19)$$

The Lyapunov exponent of $\Delta \mathbf{x}(t)$ w.r.t. the reference trajectory $\mathbf{x}_r(t)$ measures the average rate at which $\mathbf{x}_r(t) + \Delta \mathbf{x}(t)$ diverges from or converges to \mathbf{x}_r , and is defined as

$$\sigma_{\mathbf{x}_r}(\Delta \mathbf{x}) = \limsup_{t \rightarrow \infty} \frac{1}{t} \log \frac{|\Delta \mathbf{x}(t)|}{|\Delta \mathbf{x}(t_0)|} = \limsup_{t \rightarrow \infty} \frac{1}{t} \log |\Delta \mathbf{x}(t)|. \quad (2.20)$$

Provided the system's time evolution is ergodic, $\sigma_{\mathbf{x}_r}$ will be *the same* for almost all reference trajectories $\mathbf{x}_r(t)$ on \mathcal{I} . So we henceforth drop the subscript \mathbf{x}_r when stating σ . For an m -dimensional dynamical system, σ can assume m different values at maximum for different $\Delta \mathbf{x}$ [Argyris et al., 2010].

Lyapunov exponents are an important tool for investigating the local stability of invariant sets. Specifically, a *hyperbolic* fixed point, that is a state \mathbf{x}^* with $\mathbf{F}(\mathbf{x}^*) = 0$ for which the eigenvalues of $D\mathbf{F}(\mathbf{x}^*)$ have non-zero real parts, is an attractor if and only if all its Lyapunov exponents are negative, as in the neighbourhood of such a fixed point the linearized dynamics are equivalent to the full non-linear dynamics (theorem of Hartman-Grobman [Argyris et al., 2010]). A general compact minimal invariant set is asymptotically stable if all perturbations *transversal to it* are damped out. Hence the negativity of all Lyapunov exponents corresponding to transversal perturbations is a necessary condition for being an attractor. Nevertheless, attractors can possess non-negative Lyapunov exponents corresponding to *non-transversal* perturbations.

For example, Fig. 2.4 displays the four different types of attractor that can appear in 3-dimensional dynamical systems [Argyris et al., 2010]. The first type (Fig. 2.4a) is a stable fixed point, that is, a single state which is attractive in all directions and therefore has all three Lyapunov exponents negative.

The second type (Fig. 2.4b) is a stable limit cycle. On such an attractor, the dynamics return to exactly the same state after a certain period. It possesses one zero Lyapunov exponent, corresponding to small perturbations along its trajectory. The other two exponents correspond to transversal perturbations and are negative.

The third type (Fig. 2.4c) is a stable torus, that is, an attractor on which the dynamics are characterized by two periods that are incommensurate. In contrast to a limit cycle, a trajectory on a torus never returns to exactly the same state. A torus has two zero Lyapunov exponents and a negative one. The first zero exponent corresponds to small perturbations along a trajectory, and the second to small perturbations transversal to it but still to states inside the attractor.

Finally, the fourth type (Fig. 2.4d) is a chaotic attractor. Such an attractor characteristically displays a sensitive dependence on the initial state, which means that trajectories starting from two neighbouring states inside the attractor in general diverge at an exponential rate. This gives rise to a positive Lyapunov exponent [Pikovsky et al., 2003]. The second Lyapunov exponent of a chaotic attractor is zero, corresponding to small perturbations along the trajectory. The third exponent is negative, indicating asymptotic stability.




Figure 2.4: **Different types of attractor in 3-dimensional state spaces.** **a**, a stable fixed point. **b**, a limit cycle. **c**, a torus. **d**, a chaotic attractor (Lorenz). The symbols in the braces indicate the types of Lyapunov exponents the attractor has, with +, 0, or – representing a positive, zero or negative value. The figure is modelled on a similar figure in Argyris et al. [2010]. Panels **c**,**d** contain elements of free graphics files⁴.

⁴Origin: Wikimedia Commons. Author: Cepheus. Filenames: `Simple_Torus.svg`, `Lorenz_attractor2.svg`.

2.2.3 Illustration

In Section 2.1.3 we analyzed the dynamics of a damped pendulum driven by a constant angular acceleration X . With the pendulum's angle ϕ and frequency ω obeying

$$\dot{\phi} = \omega \quad (2.21)$$

$$\dot{\omega} = -\alpha\omega + X - K \sin \phi, \quad (2.22)$$

we found that, at the parameter settings $\alpha = 0.1$, $X = 0.5$, $K = 1$, there are at least three minimal invariant sets in state space: the two fixed points and the limit cycle shown in Fig. 2.3a,b. Let us now assess their stability properties.

First, note that

$$D\mathbf{F} = \begin{pmatrix} 0 & 1 \\ -K \cos \phi & -\alpha \end{pmatrix}. \quad (2.23)$$

For each of the two fixed points $\mathbf{x}_i^* = (\phi_i, \omega_i)$, $i = 1, 2$, with coordinates $\phi_{1,2} = \arcsin(X/K)$ and $\omega_{1,2} = 0$, where $\phi_1 \in [0, \pi/2]$ and $\phi_2 = \pi - \phi_1$, the Jacobian $D\mathbf{F}(\mathbf{x}_i^*)$ is a constant matrix. Therefore, for fixed point \mathbf{x}_i^* , we can write $\Delta \mathbf{x} = \beta_+ \mathbf{p}_+^i + \beta_- \mathbf{p}_-^i$, where \mathbf{p}_\pm^i are eigenvectors of $D\mathbf{F}(\mathbf{x}_i^*)$ to the eigenvalues

$$\xi_\pm^i = -\frac{\alpha}{2} \pm \frac{\sqrt{\alpha^2 - 4K \cos \phi_i}}{2}, \quad (2.24)$$

and obtain

$$\dot{\beta}_\pm = \xi_\pm^i \beta_\pm. \quad (2.25)$$

This implies that the two Lyapunov exponents σ_\pm^i of fixed point \mathbf{x}_i^* are just given by the real parts of ξ_\pm^i ,

$$\sigma_\pm^i = \text{Re } \xi_\pm^i. \quad (2.26)$$

For $0 < X < K$, both $\sigma_-^1 < 0$ and $\sigma_+^1 < 0$, because $\cos \phi_1 > 0$. By the theorem of Hartmann-Grobmann, this means that \mathbf{x}_1^* is attracting in all directions and hence a (strong) attractor. For \mathbf{x}_2^* , we also obtain $\sigma_-^2 < 0$, but $\sigma_+^2 > 0$, as $\cos \phi_2 < 0$. Therefore the second fixed point \mathbf{x}_2^* is an unstable saddle and no attractor. This is, of course, in agreement with how one would have judged from mere inspection of Fig. 2.3a: In contrast to \mathbf{x}_1^* , \mathbf{x}_2^* looks clearly unstable.

Now we turn to the limit cycle that we also found to be an invariant set at the parameter settings $\alpha = 0.1$, $X = 0.5$, and $K = 1$. The first of its two Lyapunov exponents is zero, corresponding to perturbations along the trajectory. Its second Lyapunov exponent is negative, and it turns out that the limit cycle is asymptotically stable. Indeed, it can be shown that, for a range of parameter settings, the stable fixed point and the limit cycle are competing attractors in the pendulum's state space [Strogatz, 2001]. Which of them the pendulum finally ends up in only depends on the initial state.

2.3 Nonlocal stability and weak attractors

Lyapunov exponents are a powerful tool for investigating the local stability of a minimal invariant set against *small* perturbations, that is, against perturbations for which the linearization in Eq. (2.19) is a valid approximation. However, Lyapunov exponents in general do not allow to assess the stability against *non-small* perturbations. Let us now contemplate this nonlocal stability problem.

2.3.1 Lyapunov functions

One way to assess the nonlocal stability of an attractor, at least for a fixed point, is by means of Lyapunov functions [Argyris et al., 2010]. Consider a fixed point \mathbf{x}^* . If there exists a function V whose partial derivatives are continuous in a neighbourhood \mathcal{U} of \mathbf{x}^* and that fulfils $V(\mathbf{x}^*) = 0$ and $V(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathcal{U} \setminus \{\mathbf{x}^*\}$, then the stability of \mathbf{x}^* against *all perturbations* to states inside \mathcal{U} is related to the evolution of V along trajectories. More specifically, if one defines

$$\dot{V}(\mathbf{x}(t)) = \text{grad } V \cdot \dot{\mathbf{x}} = \text{grad } V \cdot \mathbf{F}(\mathbf{x}), \quad (2.27)$$

and if $\dot{V}(\mathbf{x}) < 0$ for all $\mathbf{x} \in \mathcal{U} \setminus \{\mathbf{x}^*\}$, then the system converges to \mathbf{x}^* from every perturbed state $\mathbf{x} \in \mathcal{U}$.

Energy functions are promising candidates for Lyapunov functions. The problem is that there is no systematic way to construct such functions for general (dissipative) dynamical systems. To find them, “divine inspiration is usually required” in the words of Strogatz [2001]. We will here follow a different approach to characterize nonlocal stability.

2.3.2 Definition of basin stability

In a dissipative system, a trajectory started from any initial state will eventually either converge to a compact minimal invariant set or diverge. Based on the notion of compactification [tom Dieck, 2000], let us understand divergence as meaning convergence to ∞ , and treat this ‘point at infinity’ as a compact minimal invariant set as well.

The *basin of attraction* \mathcal{B} of a compact minimal invariant set \mathcal{I} is defined as the set of all states that converge to \mathcal{I} [Argyris et al., 2010],

$$\mathcal{B} = \{\mathbf{x} \in \mathcal{S} \mid \Phi_t(\mathbf{x}) \rightarrow \mathcal{I}\}. \quad (2.28)$$

A minimal invariant set whose basin of attraction has a positive volume (Lebesgue measure) is called a *weak attractor*. Note that every strong attractor is a weak attractor, as it has an open neighbourhood from which all states converge to the attractor and whose volume is, by definition of openness, positive. But a weak attractor is not necessarily a strong attractor. In the following, we will mostly deal with weak attractors and for simplicity just refer to them as attractors. When a statement only applies to strong attractors, we will explicitly mention that.

Complete knowledge of an attractor’s basin would allow ultimate assessability of its nonlocal stability [Chiang, 2011]: One could classify perturbations into the permissible (those pushing the system to a state within the basin) and the impermissible. However, basins are intricate entities [Nusse and Yorke, 1996] and especially hard to explore in high dimensions. Throughout this thesis, we will therefore focus on a single but fundamental property: the basin’s volume. Wiley et al. [2006] interpret the volume of an attractor’s basin as a measure of the likelihood of *arrival at* this attractor, i.e., as a measure of the attractors’s *relevance*. Almost equivalently, let us here understand the volume of the basin as an expression of the likelihood of *return to* the attractor after any random – possibly non-small – perturbation. This yields a second interpretation: The basin’s volume *quantifies how stable* an attractor is. We refer to the quantification of stability based on the basin’s volume as *basin stability* [Menck et al., 2013a].

In an alternative approach, C.S. Holling suggested to quantify nonlocal stability (or *resilience*, in his nomenclature) in terms of the width of the basin of attraction in a particular direction [Holling, 1973; Scheffer, 2009]. This concept has inspired a host of studies on complex socioecological systems. Yet the measure of stability it implies “is not so easily quantified, even in models” [van Nes and Scheffer, 2007]. One reason is that, in systems with many state variables, it is impossible to identify the single most relevant direction along which the width of the basin should be gauged. Basin stability offers two important improvements: (i) It follows a volume-based probabilistic approach that is compatible with the natural uncertainty about the strength and direction of perturbations. (ii) It provides a measure of stability that is clearly defined and easily quantified even in high-dimensional systems (see 2.4.2).

In order to formally define basin stability, note that the basins of all attractors fill out the volume of state space completely⁵. If a system has only one attractor to whose basin almost all states belong, this attractor is called *globally stable* [Argyris et al., 2010]. In contrast, if a system is *multistable*, i.e., possesses multiple attractors, each basin occupies only a certain portion of state space. So we define the basin stability of a compact minimal invariant set \mathcal{I} with basin \mathcal{B} as

$$S_{\mathcal{B}} = \int_{\mathcal{S}} \chi_{\mathcal{B}}(\mathbf{x}) \rho(\mathbf{x}) \, d\mathbf{x}, \quad (2.29)$$

where

$$\chi_{\mathcal{B}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in \mathcal{B} \\ 0 & \text{otherwise} \end{cases} \quad (2.30)$$

is the indicator function of \mathcal{B} and ρ is a density with $\int_{\mathcal{S}} \rho(\mathbf{x}) \, d\mathbf{x} = 1$ and for which $\rho(\mathbf{x}) \, d\mathbf{x}$ reflects how likely it is that a perturbed state close to \mathbf{x} occurs. The value of basin stability is bound to lie between zero and one,

$$S_{\mathcal{B}} \in [0, 1]. \quad (2.31)$$

An invariant set that is no attractor always has $S_{\mathcal{B}} = 0$. On the other extreme, a globally stable attractor always has $S_{\mathcal{B}} = 1$. For any other attractor, the value of basin

⁵For the treatment of divergence, see the remark at the beginning of this subsection.

2 Stability concepts

stability depends on the choice of ρ and measures the likelihood that the system will return to the attractor after a perturbation randomly drawn according to probability density ρ .

2.3.3 Illustration

We observed in Sections 2.1.3 and 2.2.3 that the damped driven pendulum, whose angle ϕ and frequency ω obey

$$\dot{\phi} = \omega \quad (2.32)$$

$$\dot{\omega} = -\alpha\omega + X - K \sin \phi, \quad (2.33)$$

has two attractors: the stable fixed point \mathbf{x}_1^* and a limit cycle. These two are indeed the only attractors that exist for a range of parameter settings [Strogatz, 2001], and their basins of attraction fill the volume of state space to its entirety. Fig. 2.5 displays how these basins look like.




Figure 2.5: **Basins of attraction of the damped driven pendulum.** The filled (open) circle marks the stable (unstable) fixed point \mathbf{x}_1^* (resp. \mathbf{x}_2^*). The green area indicates the basin of attraction of \mathbf{x}_1^* . From points in the white area, the pendulum converges to the limit cycle. Parameters: $\alpha = 0.1$, $K = 1$, and $X = 0.5$.

Let us now take a closer look at the basin stability of the stable fixed point $\mathbf{x}_1^* = (\phi_1, 0)$. It is well known that, at $X = K$, a *fold bifurcation* occurs in which the two fixed points \mathbf{x}_1^* and \mathbf{x}_2^* collide and disappear. Let us track how the stable fixed point's basin stability changes towards this bifurcation. To do that, we for simplicity set ρ to be a uniform distribution on $\mathcal{Q} = [0, 2\pi] \times [-20, 20]$ and zero elsewhere, choose $\alpha = 0.1$ and $K = 1$, and tune up the external angular acceleration X from zero to K . The result, obtained via the numerical procedure specified in Section 2.4.1, is shown in Fig. 2.6. Basin stability $S_{\mathcal{B}}$ starts at 1, when \mathbf{x}_1^* is indeed the only attractor. At $X \approx 4\alpha/\pi \approx 0.127$ [Strogatz, 2001], the limit cycle becomes stable and its basin starts to claim an ever larger portion of state space. $S_{\mathcal{B}}$ declines until hitting zero at $X = 1$, after which the limit cycle is the only stable attractor.




Figure 2.6: **Basin stability of the damped driven pendulum.** Basin stability $S_{\mathcal{B}}$ of the stable equilibrium \mathbf{x}_1^* versus X . Parameters: $\alpha = 0.1$, $K = 1$.

It is interesting to compare the basin stability characteristic in Fig. 2.6 to the change of the maximum Lyapunov exponent σ_+^1 of \mathbf{x}_1^* as X increases towards K . In every fold bifurcation, the maximum Lyapunov exponent of the stable fixed point becomes zero at the bifurcation point, and approaches it from below before [Strogatz, 2001]. Its declining magnitude corresponds to a slowing down of the rate at which a system recovers after a small perturbation. In climate science, this is called *critical slowing down* and has been suggested as an early-warning signal to detect the nearby collapse of climatic tipping elements (see, e.g., Scheffer et al. [2009] and the discussion in Section 2.5). Is there critical slowing down in the pendulum's dynamics?

Fig. 2.7 shows that, as X goes up, σ_+^1 is *constant* and equal to $-\alpha/2$ for all values except very close to $X = 1$. Only at $X = 0.999997$, when the square root term contributing to the eigenvalues ξ_{\pm}^1 becomes real (cf. Eq. (2.24)), does σ_+^1 start moving towards 0, where it reaches at $X = 1$. Does this render critical slowing down detectable *experimentally*? The maximum perturbation in positive ϕ -direction that the stable

2 Stability concepts

fixed point \mathbf{x}_1 can withstand is $\phi_2 - \phi_1$ (cf. Fig. 2.5), which at $X = 0.999997$ amounts to approximately $5 \cdot 10^{-3} \text{ rad} = 0.28^\circ$. For a pendulum of length $\ell = 10 \text{ cm}$, this corresponds to a displacement $\Delta = 0.5 \text{ mm}$. Therefore, only if fluctuations induced by the environment are significantly less than Δ will critical slowing down be detectable in an experiment.

If critical slowing down is detectable, it occurs *very late* as X increases from 0 to 1. Hence critical slowing down appears not to be a reliable *early*-warning signal for the pendulum's fold bifurcation.




Figure 2.7: **Local Stability of the damped driven pendulum.** **a,b:** Maximum Lyapunov exponent σ_+^1 of the stable fixed point versus X . **b** zooms in close to $X = 1$. Parameters: $\alpha = 0.1$, $K = 1$.

2.3.4 Basin stability and the choice of ρ

By definition (cf. Eq. (2.29)), the exact value of an attractor's basin stability hinges on the choice of ρ , the probability density of random perturbations. So far, we have employed a density that is uniform on a certain subset \mathcal{Q} of state space, which gives an impression of how a basin contracts or expands in a certain, possibly important region \mathcal{Q} . But such a choice of ρ will certainly not count as *realistic*. However, defining what would make a choice of ρ realistic and then finding it is beyond the scope of this thesis.

Just a few thoughts on this issue: From inspecting the marble track in Fig. 2.1 one might suggest that a realistic ρ should incorporate the “height of the landscape”, which corresponds to the magnitude of that system's energy function. Indeed, perturbations that involve the crossing of high energy barriers should occur less often than perturbations across low barriers. Now for the toy example of this chapter, the dissipative pendulum, there is no simple energy function. But we know that the kinetic energy of a perturbed state increases with the square of ω . Hence a Gaussian

distribution

$$\rho(\mathbf{x}) = \rho(\phi, \omega) = \frac{1}{(2\pi)^{3/2}b} \exp\left(-\frac{\omega^2}{2b^2}\right) \quad (2.34)$$

with zero mean and standard deviation b could be considered a simple candidate for a realistic choice of ρ .

Fig. 2.8 shows how $S_{\mathcal{B}}$ changes as the distribution's width b is increased at $\alpha = 0.1$, $K = 1$, and $X = 0.5$. For very small b up to about 0.8, basin stability assumes an almost constant value around 0.55. In this regime, typical perturbations in the ω -direction are not yet strong enough to drive the pendulum out of the basin of the stable fixed point (cf. Fig. 2.5). When b passes 0.8, strong enough perturbations become ever more likely and $S_{\mathcal{B}}$ declines sharply. After about $b = 20$, basin stability again remains almost constant, now at a value close to 0.1. This indicates that the basin is apparently unbounded and occupies a considerable portion of the negative- ω half of state space (cf. Fig. 2.5).




Figure 2.8: **Basin stability of the damped pendulum with Gaussian ρ .** Shown is the dependence of $S_{\mathcal{B}}$ on the width of the Gaussian distribution ρ , cf. Eq. (2.34). Parameters: $\alpha = 0.1$, $K = 1$, and $X = 0.5$

2.4 Basin stability estimation

Here we discuss how basin stability can be estimated in numerical experiments. Furthermore, we assess its applicability to high-dimensional systems and to systems with fractal basin boundaries.

2.4.1 Numerical procedure

Consider a compact minimal invariant set \mathcal{I} whose basin of attraction is \mathcal{B} . To estimate the corresponding basin stability $S_{\mathcal{B}}$ (cf. Eq. (2.29)), we T times draw a random initial state $\mathbf{x}(t_0)$ from \mathcal{S} according to the probability distribution determined by the density ρ and integrate the system equations. Then we count the number M of initial states for which the trajectory has converged to \mathcal{I} . Finally, we estimate $S_{\mathcal{B}}$ as M/T .

This procedure is a Monte Carlo rejection method [Von Neumann, 1951], and was also used by Wiley et al. [2006] to study the volumes of basins. As it consists of repeated Bernoulli experiments, we infer that the estimate M/T carries a standard error of [Menck et al., 2013a]

$$e := \frac{\sqrt{S_{\mathcal{B}}(1 - S_{\mathcal{B}})}}{\sqrt{T}}. \quad (2.35)$$

Throughout this thesis, we use at least $T = 500$, for which $e < 0.023$ in absolute terms.

2.4.2 High-dimensional systems

Note that the estimation error e for basin stability does not depend on the dimension of the dynamical system under investigation, cf. Eq. (2.35). Hence basin stability can as easily be applied to high-dimensional systems as to low-dimensional systems.

To check whether this is indeed the case, let us study a 300 dimensional system of $N = 100$ Rössler oscillators [Rössler, 1976] that are coupled through a network of interactions. The dynamics of oscillator i , $i = 1, \dots, N$, obey

$$\dot{x}_i = -y_i - z_i + K \sum_{j=1}^N A_{ij}(x_j - x_i) \quad (2.36)$$

$$\dot{y}_i = x_i + ay_i \quad (2.37)$$

$$\dot{z}_i = b + z_i(x_i - c) \quad (2.38)$$

where x_i, y_i, z_i are the state variables of unit i and the matrix \mathbf{A} indicates whether there is a direct influence of unit i on unit j and vice versa (in which case $A_{ij} = 1$) or not ($A_{ij} = 0$). We assume a uniform interaction strength K and choose the parameters $a = b = 0.2$ and $c = 7.0$. We will analyze systems like this one in depth in the next chapter, and here prematurely consume some of the results on that chapter's example network (depicted in Fig. 3.1a).

There exists a minimal invariant set \mathcal{M}_s on which all oscillators synchronously follow the same trajectory on the chaotic Rössler attractor (Fig. 3.1b). We set $K = 0.185$,

where \mathcal{M}_s is known to be an attractor for this network (see Section 3.3.2). How does the standard error in the estimate of its basin stability $S_B \approx 0.396$ depend on T ?

We set ρ to be a uniform distribution on $\mathcal{Q} := q^N$ with $q = [-15, 15] \times [-15, 15] \times [-5, 35]$ and perform the estimation procedure 1,000 times for each of different values of T . From these 1,000 realizations, we estimate the dependence of e on T . As shown in Fig. 2.9, the resulting characteristic agrees very well with Eq. (2.35), thus supporting the claim that e is independent of the system's dimension.




Figure 2.9: **Error of the basin stability estimation.** Red crosses indicate the standard error of the numerical basin stability estimation for different values of T . The dashed line shows the theoretical curve $e(T)$ as given by Eq. (2.35).

2.4.3 Systems with fractal basin boundaries

In Fig. 2.5 we saw that the basins of different attractors can have intricately intertwined shapes. But the boundary between the two basins shown there is still an easily discernible *smooth* curve. For other systems, however, basin boundaries can be *fractal* [McDonald et al., 1985]. If two basins are indeed separated by a fractal boundary and one draws a line across it from one basin to the other, the line leaves and enters each basin infinitely many times [Nusse and Yorke, 1996].

The nature of the basin boundaries influences the predictability of a system's behaviour in the long run, as McDonald et al. [1985] argued. Imagine we randomly draw initial states from a box through which the boundary between the basins of two attractors runs. Suppose each initial state is specified up to a certain numerical error ϵ . Then for an initial state that is closer to the boundary than ϵ , it is uncertain to which of the two attractors the system will converge. Denote by $f(\epsilon)$ the fraction of initial states for which the outcome is uncertain. If the boundary is a smooth curve, then these states are all located inside a strip of width 2ϵ along this curve, and $f(\epsilon)$ is just proportional to ϵ . However, if the boundary is fractal, then $f(\epsilon) \propto \epsilon^\alpha$, where typically $\alpha < 1$. This means that an improvement in numerical precision yields only

a rather small decrease in the share of uncertain initial states. Hence even at high precision, the long-term behaviour of systems with fractal basin boundaries may be hard to predict [McDonald et al., 1985].

The trouble is that predicting the long-term behaviour for initial states is the essence of basin stability estimation. So we will now investigate how fractal basin boundaries impact the accuracy of the numerical estimation procedure. For that purpose, we again turn to the damped pendulum, but this time apply a time-dependent forcing, dictating that the dynamics of the pendulum's angle ϕ and frequency ω obey

$$\dot{\phi} = \omega \quad (2.39)$$

$$\dot{\omega} = -\alpha\omega + X \cos t - K \sin \phi \quad (2.40)$$

$$\dot{t} = 1 \quad (2.41)$$

For $\alpha = 0.1$, $K = 1$ and $X = 7/4$, this system has several attractors [Kennedy and Yorke, 1991]. The four dominant of them, all limit cycles with period 2π , are shown in Fig. 2.10a: The *black* and the *red* attractor correspond to *rotations* of the pendulum, and the *orange* and the *yellow* attractor are *librations*. Their respective basins of attraction at $t = 0$ are shown in Fig. 2.10b. Strikingly, certain regions in this figure appear *sprinkled* with dots belonging to the different basins. Indeed, the boundary between the basins is not easily discernible, and remains so when zooming in (Fig. 2.10c). This is a consequence of the *Wada property* the basins have.

Here is the definition: Three subsets of state space are said to have the *Wada property* if any point on the boundary of one subset is also on the boundary of the two others. For the pendulum, the black basin, the red basin and the union of the orange and the yellow basin have the Wada property [Kennedy and Yorke, 1991; Nusse and Yorke, 1996]. This means that a trajectory starting at an initial state closer than the numerical error ϵ to the boundary of these sets could in principle converge to *any of the four* attractors.

To check this in numerical simulations, we write $\epsilon = 10^{-p}$ with p denoting *precision*, and forcefully discard all information after the p -th significant digit in the `double` variables used in the simulation code. A `float` variable has $p \approx 7$, and a full `double` variable has $p \approx 16$. For different values of p we then integrate from $t = 0$ a fixed set of fifty initial states, drawn uniformly at random from $\mathcal{Q} := [-\pi, 0] \times [-1, 2]$. The outcome, shown in Fig. 2.11a, reveals that some initial states, particularly the ones indicated by arrows, indeed lead to different outcomes for different values of p .

Now we estimate how basin stability $S_{\mathcal{B}}$ depends on p , drawing $T = 1,000$ initial states⁶ randomly from \mathcal{Q} with uniform ρ and again integrating from $t = 0$. As depicted in Fig. 2.11b, there seems to be no systematic influence of p on basin stability estimation. Indeed, most of the individual values of $S_{\mathcal{B}}(p)$ are within one standard deviation of the most precise value $S_{\mathcal{B}}(16)$. This suggests that, in contrast to long-term prediction for individual initial states (cf. Fig. 2.11a), volume estimation is quite robust under variation of p . Consequently, basin stability may even be applicable to systems with fractal basin boundaries.

⁶For other choices of T such as 500 or 5,000 the results look basically the same.




Figure 2.10: **Damped pendulum with fractal basin boundaries.** **a**, attractors of the damped pendulum with time-dependent forcing, cf. Eqs. (2.39)–(2.41). **b**, state space of the pendulum at $t = 0$. Black/red/orange/yellow colouring indicates convergence to the black/red/orange/yellow attractor. Convergence to other attractors is indicated by white colouring. **c**, zoom into the dashed square in panel **b**.




Figure 2.11: **Basin stability in the pendulum with fractal basin boundaries.**

a, shown are the outcomes of numerical integrations for a fixed set of fifty initial states at different values of the numerical precision p . The squares in each column correspond to the same initial state, and their respective colours indicate which state the system converges to from there at given precision p . Black/red/orange/yellow colouring indicates convergence to the black/red/orange/yellow attractor. **b**, basin stability S_B of the four attractors at different levels of p . The basin stability of the black/red/orange/yellow attractor is shown by the height of the black/red/orange/yellow bar. The grey shadows indicate \pm one standard deviation around the respective value of $S_B(16)$ (computed from multiple estimations thereof).

2.5 Inferring stability from data

Transitions between different attractors can drastically change the appearance and functionality of a system. When that system is only one part in a large assembly of systems, as is often the case in ecology and climatology, this can have grave consequences for the whole. Therefore scientists have put considerable research effort into inferring warning-signals of approaching transitions from observational data.

A typical cause of a transition is a shift in environmental conditions that makes the attractor currently occupied by a system lose its stability. This means that one of its transversal Lyapunov exponents changes from negative to positive. Hence, on the assumption of continuity, the magnitude of this exponent strives towards zero before the transition, implying that the convergence rate of the system after small perturbations becomes ever slower.

In theory, this *critical slowing down* can be detected in observed time series in several ways: On the one hand, a decline of the small-perturbation convergence rate tends to make successive states more similar to one another, thus increasing autocorrelation [Held and Kleinen, 2004]. On the other hand, it makes deviations due to noise decay more slowly, giving them the chance to accumulate and thus increasing variance [Carpenter and Brock, 2006], though this signal may be distorted [Dakos et al., 2012]. Critical slowing down was suggested as an early-warning signal of approaching transitions [van Nes and Scheffer, 2007; Dakos et al., 2008; Scheffer et al., 2009].

Warning of transitions is particularly critical for tipping elements in the Earth's climate system [Lenton et al., 2008; Schellnhuber, 2009]. One such tipping element is the Amazon rainforest, which presumably possesses two stable states: the present fertile forest state and a barren savanna state [Da Silveira Lobo Sternberg, 2001; Oyama and Nobre, 2003; Hirota et al., 2011]. A transition would emit huge amounts of carbon dioxide captured in the rich vegetation. At a range of environmental conditions, both Amazonian states are stable due to a positive feedback: Deep-rooting trees take up water stored in the soil and transpire it to the atmosphere. Thereby forest cover in an area increases overall precipitation and improves its own growing conditions. Consequently, a rather arid area (i.e., an area with weak precipitation inflow) may still be supportive of forest growth if its forest cover exceeds a certain critical threshold; were forest cover pushed below this threshold, the area would lose all of its trees [Menck et al., 2013a].

This is summarized in a conceptual model (the Levins model [Levins, 1969] with an additional discontinuity, see also Appendix)

$$\frac{dC}{dt} = F(C) = \begin{cases} r(1 - C)C - xC & \text{if } C > C_{\text{crit}}, \\ -xC & \text{if } C < C_{\text{crit}}. \end{cases}$$

Here, C is the relative forest cover that grows with the saturating rate r if $C > C_{\text{crit}}$ and dies with rate x (assuming $r > x > 0$). C_{crit} is the critical forest cover threshold. This model has two equilibria, the forest state $C_F = 1 - \frac{x}{r}$ and the savanna state $C_S = 0$. The equilibrium C_F (resp. C_S) exists and is stable if $C_F > C_{\text{crit}}$ (resp. $C_{\text{crit}} > 0$). Assuming that C_{crit} increases linearly with aridity A , we obtain the state diagram in Fig. 2.12.




Figure 2.12: **State diagram of a bistable stylized forest-savanna model.**

Dashed arrows show where the system state C would move from a certain initial state at given aridity A . The green (light yellow) area is the basin of attraction of the forest (savanna) state C_F (C_S). Solid circles indicate the emergence or disappearance of an equilibrium. Figure published in [Menck et al., 2013a].

Global warming may drive up aridity, pushing it eventually beyond the bifurcation point A_{crit} where the forest state vanishes. Here, a bifurcation-related transition would be enforced. Furthermore, as aridity increases, the volume of C_F 's basin of attraction shrinks, indicating that the forest state becomes less stable against perturbations. Indeed, due to this reduced basin stability, a large perturbation such as strong deforestation might induce a critical transition to the savanna state long before aridity reaches A_{crit} . Crucially, *none* of this is reflected by linear stability methods, which only judge the forest state C_F by the coefficient $F'(C_F) = x - r$ that *remains constant* as aridity goes up.

This implies [Menck et al., 2013a]: (i) There is no critical slowing down. Indeed, linear early-warning signals can be absent in systems with strong non-linearities such as this model and fail to announce an approaching bifurcation-related transition. (ii) Because of their highly local perspective, linearization-based stability methods such as the small-perturbation convergence rate do not indicate how stable C_F is against non-small perturbations. Thus they do not signal the increasing risk of a critical transition. Note that both these effects came up in a similar way in our above discussion of the damped driven pendulum, where slowing down did emerge but set in *too late* to be an *early*-warning signal (see Section 2.3.3).

This suggests that the small-perturbation convergence rate, which may or may not warn of an approaching bifurcation-related transition, is not a reliable indicator of an attractor's basin stability. Indeed, inferring basin stability from observational data poses a daunting task, as time series inevitably have Lebesgue measure zero in higher dimensions, whereas the object one is seeking information on, the basin of an attractor, always has a positive Lebesgue measure.

2.6 Conclusions and outlook

We have introduced the terminology of dynamical systems theory and discussed several of the different stability concepts it provides, including local stability, global stability, and multistability. In this context, we developed a clear definition of basin stability, a nonlocal concept based on the volume of the basin of attraction that can be used to quantify by a number $S_B \in [0, 1]$ the stability of a (weak) attractor between instability ($S_B = 0$) and global stability ($S_B = 1$).

We established several important facts about basin stability. First, the numerical costs of its estimation are relatively low and independent of the state space dimension. Second, we found evidence that basin stability is applicable even when fractal basin boundaries are present.

However, we observed a drawback, too: There might be no simple way to infer basin stability from observational data. In particular, the small-perturbation convergence rate, a possible indicator for the degree of local stability, does not reliably correlate with basin stability.

Several questions derive from this chapter's investigations that could be addressed in future work:

(i) Basin stability crucially depends on the choice of the probability density ρ of random large perturbations. In this thesis, we mostly limit ourselves to simple choices, where ρ is uniform on a certain subset of state space and zero elsewhere. This provides insights into the contraction or expansion behaviour of a basin inside this subset. However, ρ could be chosen in many other ways. Particularly, in many applications one may seek to make a realistic choice of ρ . It needs to be specified what realism would mean in this context and how it could be achieved.

(ii) In the way we have defined basin stability here, it is applicable solely to deterministic dynamical systems. Stochasticity enters only via the randomized perturbations that are applied. However, one may want to study stochastic dynamical systems. To apply basin stability, one would have to incorporate an appropriate stochastic notion of basins. This needs to be explored further.

(iii) Basin stability quantifies how stable an attractor is against perturbations in state space at fixed environmental conditions, i.e., at fixed system parameters. However, there may be parameter uncertainty or parameter perturbations. It should be investigated how this would best be taken into account when estimating an attractor's basin stability.

(iv) Time series may contain more or less robust signals that indicate local stability. As warning of approaching transitions can be crucial in complex systems particularly from the realms of ecology or climatology, it would be worth trying to find ways how (a decline of) basin stability could be detected from observational data.

3 Synchronization on complex networks

Synchronization refers to the observation of processes taking place together (συν-) in time (χρόνος). Defined like this, synchronization is a rather mundane phenomenon. For example, a car cruising smoothly on a highway clearly has its four wheels spinning at the same speed, thus showing synchronous behaviour. But there are also scores of systems in which the emergence of synchronization, or *order in time* [Strogatz, 2012], is intriguingly far from self-explaining. These systems have attracted considerable scientific interest, particularly in the past decades [Pikovsky et al., 2003].

The first account of spontaneous synchronization has been credited to Christiaan Huygens, a seventeenth century dutch astronomer, physicist and mathematician, who was also patent-holder for the first pendulum clocks. Huygens, reportedly while lying sick in bed [Pikovsky et al., 2003], observed that two such clocks, hanging from the same supporting beam, always adjusted their movements in a synchronous way, regardless of how he had started them. More specifically, he found the two clocks to always converge to a state in which both ticked at exactly the same rate while moving in opposite directions, with one pendulum assuming its left-most position when the other assumed its right-most position. This phenomenon was later termed *synchronization in anti-phase*. Huygens explained that it is caused by a small, hardly perceptible motion of the supporting beam by which the two pendulums influence one another. This motion gradually becomes smaller until, in anti-phase synchronization, it completely ceases. Huygens' observation was recently reproduced in a study of two metronomes placed on a freely moving base [Pantaleone, 2002]. Interestingly, besides anti-phase synchronization, this experimental setup turned out to be also capable of producing *in-phase synchronization*.

Since Huygens' description, synchronization has been reported in many disciplines. A fascinating example from entomology is synchrony in populations of fireflies. Indeed in some species, groups of male fireflies adjust their signaling rhythms with millisecond precision, apparently in order to better attract the females' attention [Moiseff and Copeland, 2010]. In the technological sphere, synchronization is, for instance, central to the function of power grids, the real-world systems that provide the main motivation for this thesis. A power grid's components have to be operated at the same synchronous frequency (usually 50 or 60 Hz) to achieve steady power flows and to avoid destruction of connected equipment [Machowski et al., 2008].

Synchronization also plays an important role in the bodies of mammals, and of humans in particular. For example, a mammal's internal biological clocks synchronize themselves to the 24-hours period of the Earth's rotation. Loss of synchronization brings bodily functions into temporary disorder, as is a common experience of travelers suffering jet lag after a long-distance flight [Pikovsky et al., 2003]. Furthermore, the frequencies of locomotion and breathing synchronize when a mammal is running [Bramble and Carrier, 1983]. Finally, both neural communication [Fries, 2005] and

memory processes [Fell et al., 2001; Fell and Axmacher, 2011] rely vitally on the synchronous firing of neurons. This means that, although synchronization is also associated with pathological states like Parkinsonian tremor [Hammond et al., 2007], the functional ability to support synchrony is pivotal for mammalian brains.

All these examples of spontaneous synchrony share two crucial ingredients [Pikovsky et al., 2003]. First, the individual systems that become synchronized are all *self-sustained oscillators*, that is, they produce a self-repeating process of a certain internal frequency also when *isolated* from their surroundings. Plants, for example, usually have their internal biological clocks close to the daily 24-hours period, but can be gene-manipulated to self-sustain oscillations with periods of 20 hours to more than 30 hours [Dodd et al., 2005]. Importantly, the phase of a self-sustained oscillator is free, meaning that its position on the internal oscillatory cycle can be adjusted by external influences. These external influences on individual oscillators are brought about by the second crucial ingredient to synchronization: some form of *coupling*. In Huygens' observations, coupling was mediated by the supporting beam from which both pendulum clocks suspended.

Many studies focussed on how the properties of the individual oscillators, particularly their internal frequencies, support or hinder synchronization [Kuramoto, 1975; Strogatz and Mirollo, 1991; Strogatz, 2000]. Fascinatingly, synchronization was even observed in systems of chaotic oscillators, despite the notorious divergence of nearby initial conditions such systems are known for [Pecora et al., 1997; Boccaletti et al., 2002]. Indeed, special kinds of chaotic oscillator for whose attractors phases can be defined were found capable of reaching a state of wild, uncorrelated amplitude variations but synchronized frequencies and locked phases [Rosenblum et al., 1996]. This phenomenon is called *chaotic phase synchronization*. Furthermore, sets of identical oscillators were found capable of reaching a *completely synchronous* state in which the units totally align their dynamics and follow exactly the same trajectory [Pecora and Carroll, 1990].

In recent years, researchers have paid considerable attention to scrutinizing how the nature of the coupling influences a system's propensity for synchronization. Among the above examples, the entrainment of biological clocks by the external day-and-night cycle illustrates that coupling can be *unidirectional*. In Huygens' pendulum clocks, the coupling is *bidirectional*. Furthermore, in large populations of oscillators such as neurons in the brain or generators in a power grid, the coupling is best described as a highly structured, possibly time-dependent network of interactions between the individual units. Aided by the theory of complex networks [Watts and Strogatz, 1998; Strogatz, 2001; Albert and Barabási, 2002; Boccaletti et al., 2006; Newman, 2003, 2010], many studies have therefore elaborated on how the topology of this network affects a system's ability to synchronize [Arenas et al., 2008]. We will also pursue investigations on this topic here.

Of particular interest to what follows below is the work on master stability functions by Pecora and Carroll [1998] who were able to establish a rigorous relation between the *local stability* of a network's completely synchronous state and the network's topological properties. Based on this result, a host of studies investigated what kinds of networks are the most suitable for synchronization [Arenas et al., 2008]. Notably, Hong et al. [2004] found that synchronizability is optimal in networks with so-called

random topologies. This is in sharp contrast to the observation that many synchronizing real-world systems, such as neural networks or power grids, have *small-world* topologies that are comparatively *regular* [Watts and Strogatz, 1998], and has left networks research with a long standing puzzle.

It is important to note that the theoretical insights by Pecora and Carroll and Hong et al. are based on local stability and hence only address the synchronous state’s resilience against *small perturbations*. Seeking to complement the theory, we will here attempt to identify networks that are optimal w.r.t. *basin stability*, the nonlinear stability concept defined in Section 2.3.2 that measures the synchronous state’s stability against *large perturbations*. Interestingly, we will discover that basin stability improves as topological *regularity* increases. This provokes to conjecture a solution to the said puzzle: The evolution of real-world synchronizing networks has been shaped by *two counteracting processes*: First, by the optimization w.r.t. local stability that promoted topological randomness; and, second, by the simultaneous optimization w.r.t. basin stability that favoured topological regularity. This contest ended in a topological tradeoff: small-worldness.

Contents

We will start in Section 3.1 by studying an example network of chaotic oscillators that can be tuned to achieve complete synchronization and that we will revisit throughout this chapter. After this, we will introduce in Section 3.2 several terms and concepts from the theory of complex networks that are needed for what follows. In Section 3.3, we will review the important studies by Pecora and Carroll [1998] and Hong et al. [2004], and elaborate on the puzzle they pose. In Section 3.4, we will explore how basin stability complements the local-stability-based approach of these studies and discuss the conjectured solution to the puzzle. Section 3.5 provides conclusions and outlook. Finally, Section 3.6 contains some additional material that supports the main findings but, if included earlier, would have made their presentation rather lengthy.

Some of the findings presented here have been published in [Menck et al., 2013a].

3.1 A motivating example

As this chapter’s motivating example, consider the network of 100 x -coupled *Rössler oscillators* [Rössler, 1976] sketched in Fig. 3.1a. In the figure, each lightblue circle represents one of the oscillators, and a line drawn between two circles indicates that the corresponding oscillators exert a direct influence on one another. The dynamics of unit i , $i = 1, \dots, N$, are given by

$$\dot{x}_i = -y_i - z_i + K \sum_{j=1}^N A_{ij}(x_j - x_i) \quad (3.1)$$

$$\dot{y}_i = x_i + ay_i \quad (3.2)$$

$$\dot{z}_i = b + z_i(x_i - c) \quad (3.3)$$

where x_i, y_i, z_i are the state variables of unit i and the matrix \mathbf{A} captures if there is a direct influence between units i and j (in which case $A_{ij} = 1$) or not ($A_{ij} = 0$). The interaction strength, assumed uniform throughout the network, is given by the constant $K \geq 0$. Finally, we choose the parameters $a = b = 0.2$ and $c = 7.0$, which, in the case of no coupling ($K = 0$), produce the (stable) chaotic *Rössler attractor* \mathcal{R} depicted in Fig. 3.1b in each individual unit’s 3-dimensional state space.

Can K be tuned so that all units follow the same trajectory on the Rössler attractor *synchronously*? In principle, this *completely synchronous state* can be established artificially for any value of K : Just initiate all units to the same state on the attractor, $(x_i, y_i, z_i)(0) := (x, y, z) \in \mathcal{R}$ for all $i = 1, \dots, N$; then they will remain on the same trajectory on this attractor indefinitely, as the coupling terms in Eq. (3.1) have no chance of becoming non-zero. However, a minor perturbation might immediately destroy the synchronous state. So, for which values of K is it asymptotically stable?

To answer this question, let us start with some trial and error. We set $K = 0.15$ and initiate the network to a state on the Rössler attractor as just described. Then, after some time, we apply a random perturbation to all oscillators.¹ The outcome is shown in Fig. 3.1c: The perturbation dampens out rapidly, and the whole network converges back to the synchronous state, which hence appears stable. Now we set $K = 1.50$ and repeat the experiment. Intuition might suggest that a larger K -value can only benefit the synchronous state’s stability. But this is not so! As shown in Fig. 3.1d, the perturbation does not damp out, but instead the units’ state variables *diverge*.² The synchronous state is not stable at this value of K .

In a groundbreaking study, Pecora and Carroll [1998] uncovered that the values of K for which the synchronous state is stable are completely determined by the network’s wiring topology. We will examine their reasoning in Section 3.3.1.

¹For each unit, the perturbation is drawn independently and applied to the y -coordinate.

²For reasons of visualizability, the perturbation applied in Fig. 3.1d is rather large. However, the same outcome, divergence, is also observed at $K = 1.5$ when the synchronous state is affected by small perturbations.




Figure 3.1: **Synchronization of a network of chaotic Rössler oscillators.** **a**, coupling diagram of a network of 100 Rössler units (an edge list of the network can be found in Section 3.6.5). **b**, Rössler attractor for $a = b = 0.2$, $c = 7.0$. **c-d**, time series of 20 units' y -coordinates before and after a random perturbation is applied. **c**, coupling $K = 0.15$. **d**, coupling $K = 1.5$.

3.2 Network terminology

Before going into the details of how a dynamical network's ability to synchronize is influenced by the wiring topology, we first have to introduce some descriptive network language. For that purpose, let us employ the standard nomenclature of complex network theory [Watts and Strogatz, 1998; Strogatz, 2001; Albert and Barabási, 2002; Boccaletti et al., 2006; Newman, 2003, 2010], a set of concepts and tools that has been developed and put to fruitful use by a number of disciplines, including the social sciences (with applications such as friendship networks [Amaral et al., 2000] or scientific citation networks [Price, 1965]), economics (e.g., airline [Guimera et al., 2005] and cargo ship networks [Kaluza et al., 2010]), infrastructure engineering (e.g., the internet [Faloutsos et al., 1999], power grids [Crucitti et al., 2004b]), climate science (e.g., correlation networks [Donges et al., 2009]), and biology (e.g., neural networks [Sporns and Zwi, 2004], gene regulatory circuits [Huang and Ingber, 2007], food webs [Gross et al., 2009]).

3.2.1 Nodes, edges and the adjacency matrix

A network consists of *nodes* and *edges*. In a social-sciences context, for example, nodes often represent individual persons, and edges express the relationships between them. Similarly, in the internet, nodes are computers or routers, and edges represent physical connections. Edges can be *undirected*: When nodes A and B in a computer network are linked by a physical connection, packages can usually travel from A to B or from B to A. Conversely, edges can be *directed*: In a business network, individual A might be related to individual B by being a customer of them, but this does not imply that B is also a customer of A. Furthermore, edges can be *weighted*: If, in an airport network, the presence of an edge between nodes A and B implies that there is a direct connection from A to B, such an edge could carry additional information that indicates how many passengers could maximally embark on flights across this connection per day. Such information attributed to an edge is commonly called its *weight*. *Unweighted* edges do not carry information beyond their existence.

A network of N nodes can conveniently be represented in terms of the *adjacency matrix* \mathbf{A} , a matrix of size $N \times N$ that incorporates information on all edges in the network. To set it up, for every combination (i, j) , where $i, j = 1, \dots, N$, define $A_{ij} = w_{ij}$ if there is an edge from node i to node j with weight $w_{ij} > 0$ and $A_{ij} = 0$ otherwise. By this definition, the adjacency matrix of an *undirected network* (i.e., a network that contains only undirected edges) is *symmetric*. In contrast, the adjacency matrix of a *directed network* may be *non-symmetric*. In unweighted networks, every edge carries the same weight 1 and is represented by $A_{ij} = 1$.

3.2.2 Shortest paths, connectedness and betweenness

Paths have been studied extensively in the context of social networks. In an acquaintance network, for example, there is a path of length r between two individuals A and C if A is acquainted with another individual B_1 , B_1 is acquainted with B_2 , ..., B_{r-2}

is acquainted with B_{r-1} , and B_{r-1} , finally, is acquainted with the target C . Similarly, in a more general network, a path of length r between two nodes i and j is defined as a set of hops along r adjacent edges that lead from i to j , that is, an ordered set of pairs of nodes

$$(i = k_0, k_1), (k_1, k_2), \dots, (k_{r-2}, k_{r-1}), (k_{r-1}, k_r = j) \text{ with } A_{k_{s-1}, k_s} > 0 \text{ for } s = 1, \dots, r.$$

A path between nodes i and j is called a *shortest path* if, along it, node j can be reached from node i in a minimum number ℓ_{ij} of hops, where ℓ_{ij} is called the *shortest-path length* of nodes i and j . If there is no path from i to j in a given network, the convention is to set $\ell_{ij} = \infty$. The *average shortest-path length* of a network is defined as

$$L = \frac{1}{N(N-1)} \sum_{i,j, \ell_{ij} \neq \infty} \ell_{ij}. \quad (3.4)$$

In the network shown in Fig. 3.2, the dashed arrows show a path of length 5 from node 4 to node 1. The shortest path has only length 1.

Interestingly, shortest paths really are rather short in acquaintance networks. Indeed, in a famous experiment by Milgram and co-workers in the 1960s [Milgram, 1967; Travers and Milgram, 1969], it turned out that randomly chosen individuals from one US city were able to pass a message to a target individual in another US city along acquaintance paths of average length 6! This has been referred to as the *small-world effect* that was also found in a recent study of the social network Facebook [Ugander et al., 2011].




Figure 3.2: **Example network to illustrate network definitions.** The network consists of 15 nodes which can be grouped into two connected sets A and B (indicated by the dashed circles). Arrows indicate paths from node 4 to node 1.

A *connected set* of a network is a set C of nodes whose pairwise shortest-path lengths are all finite but whose shortest-path lengths to nodes outside the set are all infinite.³ The example network in Fig. 3.2 has two connected sets, A and B. A network that consists just of a single connected set, that is, a network in which there is a path from every node i to every other node j , is called a (*fully*) *connected* network. The mentioned study on Facebook found that, in 2011, more than 99.9% of the social network's $N = 721,000,000$ active users belonged to the largest connected set (and about 2,000, or 0,0003 %, to the second-largest) [Ugander et al., 2011]. Hence one might say that Facebook is almost fully connected.

Individual nodes may play different roles in a network. A characteristic that tries to quantify how central node i is to a network is *shortest-path betweenness* b_i , defined as the number of shortest paths between pairs of nodes in the network that run through node i . Clearly, this measure seems to be relevant for communication networks such as the internet in which, ideally, messages travel along shortest paths. Indeed, the efficiency of communication networks may be seriously diminished by attacks that specifically disable nodes with large b_i [Holme et al., 2002]. A formal definition of shortest-path betweenness b_i of node i reads

$$b_i = \sum_{j \neq i, k \neq i, k > j} \frac{M_{jk}^i}{M_{jk}}, \quad (3.5)$$

where M_{jk} is the number of shortest paths from node j to node k and M_{jk}^i is the number of shortest paths from node j to node k that go through node i . If there is no path between a pair (j, k) the associated ratio is, by convention, set to zero in the summation.

3.2.3 Degree and clustering

Another important nodal characteristic is the degree of node i , defined as the number of edges that go from node i to other nodes,

$$d_i = \sum_{j=1}^N A_{ij}. \quad (3.6)$$

In a directed network, this is called the *out-degree* of node i , and its *in-degree* is defined as $\sum_{j=1}^N A_{ji}$. In a weighted network, d_i as defined here is called *intensity*. Nodes with large degree, commonly referred to as *hubs*, have a particular functional relevance in many networks. For example, in air transportation networks hubs include the large airports at which domestic and international routes of multiple airlines intersect [Guimera et al., 2005].

³This definition is somewhat imprecise for directed networks. In such a network, we say that two nodes i and j belong to the same connected set if $d_{ij} \neq \infty$ or $d_{ji} \neq \infty$.

The *clustering coefficient* C_i of node i , also called *transitivity*, captures how many pairs among the nodes to which node i is connected are also connected, and can be interpreted as a measure of local structure. In a social network, it expresses how many friends of a person's friends are also friends. The clustering coefficient is defined as

$$C_i = \frac{1}{d_i(d_i - 1)} \sum_{\substack{j=1 \\ A_{ij} \neq 0}}^N \sum_{\substack{k=1 \\ A_{ik} \neq 0}}^N A_{jk}. \quad (3.7)$$

The average of the nodal clustering coefficients, $C = 1/N \sum_i C_i$, is called *the network's clustering coefficient*.

3.2.4 The Laplacian matrix

An alternative matrix representation of an N -node network is the $N \times N$ Laplacian matrix \mathbf{L} . Its elements are defined as

$$L_{ij} = d_i \delta_{ij} - A_{ij}, \quad (3.8)$$

where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise. We will see below that the eigenvalues $\lambda_1, \dots, \lambda_N$ of \mathbf{L} can be very important to dynamical processes taking place on a network. Therefore we need to acknowledge two well-known facts about these eigenvalues.

The first fact is that the number of zero eigenvalues of a Laplacian matrix is equal to the number κ of connected sets of the corresponding network (provided all weights of existing edges are positive). The second fact is that, for a connected undirected network, the Laplacian's eigenvalues λ_2 to λ_N are all *positive* real numbers. Hence one can sort the eigenvalues, and write

$$0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_N.$$

The first fact can be proven as follows [Newman, 2010]. Denote the connected sets of the network under investigation by C_1, \dots, C_κ , and multiply \mathbf{L} from both sides by an N -dimensional vector \mathbf{x} to obtain

$$P = \mathbf{x}^* \mathbf{L} \mathbf{x} = \sum_i \sum_j A_{ij} (|x_i|^2 - x_i^* x_j) = \frac{1}{2} \sum_i \sum_j A_{ij} |x_i - x_j|^2, \quad (3.9)$$

where $*$ denotes the complex conjugate. Now, for each connected set construct an eigenvector of \mathbf{L} with eigenvalue zero. This is not too hard: For connected set C_l , $l = 1, \dots, \kappa$, and for $i = 1, \dots, N$, set $x_i = 1$ if node $i \in C_l$ and $x_i = 0$ otherwise. As the connected sets represent a disjunct decomposition of $\{1, \dots, N\}$, the set of eigenvectors $\{\mathbf{x}_1, \dots, \mathbf{x}_\kappa\}$ thus created is linearly independent. Hence the number of zero eigenvalues is at least κ .

Second, assume that \mathbf{x} is an eigenvector of \mathbf{L} with eigenvalue zero. Then the product P in Eq. (3.9) vanishes. However, as the terms summed over are all non-negative,

3 Synchronization on complex networks

P can only be zero if $x_i = x_j$ for every pair i, j of elements for which $A_{ij} \neq 0$. But this means that, for any two nodes i and j that are in the same connected set C_l of the network, $x_i = x_j = \xi_l$, where ξ_l is a set-specific constant. Hence \mathbf{x} is of the form $\mathbf{x} = \sum_{l=1}^{\kappa} \xi_l \mathbf{x}_l$, where $\{\mathbf{x}_1, \dots, \mathbf{x}_{\kappa}\}$ are the eigenvectors constructed above. This means that the number of zero eigenvalues is exactly κ .

To prove the second fact [Newman, 2010], note that a connected network has exactly one connected set and hence exactly one zero eigenvalue. Denote this eigenvalue by λ_1 . For any eigenvector \mathbf{x} of \mathbf{L} with eigenvalue λ , it follows from Eq. (3.9) that

$$\mathbf{x}^t \mathbf{L} \mathbf{x} = \lambda \mathbf{x}^t \mathbf{x} = \lambda \|\mathbf{x}\|^2 = \lambda \geq 0,$$

where we assumed without loss of generality that $\|\mathbf{x}\| = 1$. As λ_1 is the only zero eigenvalue, it follows that all the other eigenvalues are > 0 . Hence we can sort them in ascending order and write $0 = \lambda_1 < \lambda_2 \leq \lambda_3 \leq \dots \leq \lambda_N$.

3.3 Local stability and the puzzle it poses

3.3.1 Master stability function formalism

Thus equipped with network terminology, we can turn to the groundbreaking study of Pecora and Carroll on synchronization in dynamical networks [Pecora and Carroll, 1998]. Consider a system of N identical oscillators which are coupled through a connected network. Its dynamics are governed by the equations

$$\dot{\mathbf{r}}_i = \mathbf{F}(\mathbf{r}_i) + K \sum_j A_{ij} [\mathbf{H}(\mathbf{r}_j) - \mathbf{H}(\mathbf{r}_i)] = \mathbf{F}(\mathbf{r}_i) - K \sum_j L_{ij} \mathbf{H}(\mathbf{r}_j), \quad (3.10)$$

where $\mathbf{r}_i = (r_i^1, \dots, r_i^m)^t$ is the m -dimensional state vector describing the processes at node i . \mathbf{A} is the $N \times N$ adjacency matrix and \mathbf{L} is the Laplacian matrix (see Section 3.2). Finally, K denotes the overall coupling constant and $\mathbf{H}(\mathbf{r})$ is the coupling function prescribing through which of their m coordinates the connected nodes interact. \mathbf{F} determines the evolution of each individual oscillator in the case of no coupling ($K = 0$). Because \mathbf{L} has zero row sum by definition (cf. Eq. (3.8)), there always exists a synchronous state $\mathcal{M}_s = \{\mathbf{r}_1 = \mathbf{r}_2 = \dots = \mathbf{r}_N = \mathbf{s}(t) | t \in \mathbb{R}\}$ in the Nm -dimensional state space in which all individual oscillators follow the same trajectory $\mathbf{s}(t)$ on an attractor of the m -dimensional uncoupled system $\dot{\mathbf{r}} = \mathbf{F}(\mathbf{r})$. Technically, \mathcal{M}_s is a compact invariant set of states rather than a state (cf. Section 2.1). However, for notational simplicity, we keep on calling it *the synchronous state*.

Is \mathcal{M}_s stable? Based on local-stability methods (cf. Section 2.2), Pecora and Carroll showed that the answer to this question breaks up into two parts: First, \mathbf{F}, \mathbf{H} define a master stability function $\text{MSF}_{\mathbf{F}, \mathbf{H}}$ that is independent of the network. Second, K and the network define a set of numbers at which $\text{MSF}_{\mathbf{F}, \mathbf{H}}$ has to be evaluated to find out whether \mathcal{M}_s is stable. Indeed, \mathcal{M}_s is stable if K and the eigenvalues $\lambda_1, \dots, \lambda_N$ of \mathbf{L} satisfy $\text{MSF}_{\mathbf{F}, \mathbf{H}}(K\lambda_i) < 0$ for all $i = 2, \dots, N$.

To understand this, we follow the expositions in Pecora and Carroll [1998] and Arenas et al. [2008] and, employing a small perturbation analysis around \mathcal{M}_s , write $\mathbf{r}_i = \mathbf{s}(t) + \delta\mathbf{r}_i$ in order to expand $\mathbf{F}(\mathbf{r}_i) \approx \mathbf{F}(\mathbf{s}) + D\mathbf{F}(\mathbf{s})\delta\mathbf{r}_i$ and $\mathbf{H}(\mathbf{r}_i) \approx \mathbf{H}(\mathbf{s}) + D\mathbf{H}(\mathbf{s})\delta\mathbf{r}_i$ in Taylor series, where $D\mathbf{F}$ and $D\mathbf{H}$ are the Jacobian matrices of \mathbf{F} and \mathbf{H} respectively. This gives

$$\delta\dot{\mathbf{r}}_i = D\mathbf{F}(\mathbf{s})\delta\mathbf{r}_i - K D\mathbf{H}(\mathbf{s}) \sum_j L_{ij} \delta\mathbf{r}_j. \quad (3.11)$$

These equations can be used to calculate the Lyapunov exponents of \mathcal{M}_s (cf. Section 2.2.2). If the exponents corresponding to perturbations *transversal* to \mathcal{M}_s are negative, then \mathcal{M}_s is stable.

Pecora and Carroll found that this linearization-based stability assessment can be significantly simplified by exploiting the block form of Eqs. (3.11), which becomes apparent in direct-product notation

$$\delta\dot{\mathbf{r}} = [\mathbf{1}_N \otimes D\mathbf{F}(\mathbf{s}) - K \mathbf{L} \otimes D\mathbf{H}(\mathbf{s})] \cdot \delta\mathbf{r}, \quad (3.12)$$

where $\mathbf{1}_N$ is the N -dimensional identity matrix and $\delta\mathbf{r}$ is shorthand for the Nm -dimensional state vector $(\delta\mathbf{r}_1, \dots, \delta\mathbf{r}_N)^t$. This equation can be diagonalized by projecting $\delta\mathbf{r}$ onto the eigenvectors of \mathbf{L} ,

$$\delta\mathbf{r} = \sum_{i=1}^N \mathbf{e}_i \otimes \delta\mathbf{r}_i = \sum_{i=1}^N \left(\sum_{j=1}^N \alpha_{ij} \mathbf{v}_j \right) \otimes \delta\mathbf{r}_i = \sum_{j=1}^N \mathbf{v}_j \otimes \left(\sum_{i=1}^N \alpha_{ij} \delta\mathbf{r}_i \right) =: \sum_{j=1}^N \mathbf{v}_j \otimes \delta\mathbf{w}_j,$$

where \mathbf{e}_i has elements $e_{ij} = \delta_{ij}$ and denotes the N -dimensional unit vector pointing into direction i , $i = 1, \dots, N$, and \mathbf{v}_i is the eigenvector corresponding to the i th eigenvalue λ_i of \mathbf{L} , $i = 1, \dots, N$. Then, in contrast to the single *coupled* equation (3.12), one is left with N *decoupled* equations

$$\delta\dot{\mathbf{w}}_i = [D\mathbf{F}(\mathbf{s}) - K \lambda_i D\mathbf{H}] \delta\mathbf{w}_i, \quad i = 1, \dots, N, \quad (3.13)$$

into which the network topology enters only via the eigenvalues of \mathbf{L} . In a connected network, \mathbf{L} always has a single vanishing eigenvalue $\lambda_1 = 0$ with eigenvector $\mathbf{v}_1 = (1, 1, \dots, 1)$ (see Section 3.2.4). As this mode corresponds to perturbations *in parallel* to the synchronous state \mathcal{M}_s , it is not relevant for the assessment of its local stability. In contrast, the eigenmodes belonging to $\lambda_2, \dots, \lambda_N$ are *transversal* to \mathcal{M}_s ; if this state is to be asymptotically stable, then all of these modes must be damped out, which means that the corresponding maximum Lyapunov exponents have to be negative (cf. Section 2.2.2).

Pecora and Carroll noted that these $N - 1$ Lyapunov exponents can be determined from a *single* function that is *independent of the network*. Observing that Eqs. (3.13) are all of the same form

$$\delta\dot{\mathbf{w}} = [D\mathbf{F}(\mathbf{s}) - \alpha D\mathbf{H}] \delta\mathbf{w} \quad (3.14)$$

which only depends on \mathbf{F} and \mathbf{H} , they defined the *master stability function* $\text{MSF}_{\mathbf{F}, \mathbf{H}}$ as the function that assigns to the number α the maximum Lyapunov exponent of the m -dimensional dynamical system in (3.14). This reduces the local-stability assessment of the synchronous state \mathcal{M}_s of the dynamical network (3.10) to just a few evaluations

of the network-independent function $\text{MSF}_{\mathbf{F},\mathbf{H}}$. More specifically, \mathcal{M}_s is stable if

$$\text{MSF}_{\mathbf{F},\mathbf{H}}(K\lambda_i) < 0 \quad \text{for } i = 2, \dots, N. \quad (3.15)$$

For general directed or asymmetrically weighted networks, \mathbf{L} may be non-symmetric so that the eigenvalues λ_i are complex numbers, and $\text{MSF}_{\mathbf{F},\mathbf{H}}$ is a function defined on the complex plane. This case was discussed by Pecora and Carroll [1998]. In the following, we will limit ourselves to the less general case that has received far more attention [Arenas et al., 2008] and consider only connected, undirected networks whose edges carry symmetric or homogeneous weights. For these, \mathbf{L} is symmetric, all its eigenvalues are non-zero real numbers that can be sorted in an ascending fashion (cf. Section 3.2.4),

$$0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_N,$$

and $\text{MSF}_{\mathbf{F},\mathbf{H}}$ is an easily plottable function on the real line.




Figure 3.3: **Master stability function of x -coupled Rössler oscillators.** Parameters of the Rössler oscillators were chosen to be $a = b = 0.2$ and $c = 7.0$.

3.3.2 Revisiting the motivating example

Recall the motivating example network of x -coupled Rössler oscillators from Section 3.1, in which we observed the synchronous state to be stable for $K = 0.15$ but to be unstable for $K = 1.50$. As Eqs. (3.1)–(3.3) are of the form (3.10), we can now employ the master stability function formalism to understand this outcome. To cast Eqs. (3.1)–(3.3) into the form of (3.10), we write

$$\mathbf{F}(\mathbf{r}) = \begin{pmatrix} -y - z \\ x + ay \\ b + z(x - c) \end{pmatrix} \quad \text{and} \quad \mathbf{H}(\mathbf{r}) = \begin{pmatrix} x \\ 0 \\ 0 \end{pmatrix},$$

where $\mathbf{r} = (x, y, z)^t$, $a = b = 0.2$ and $c = 7.0$. If the network is in the synchronous state \mathcal{M}_s , all nodes follow the same trajectory on the chaotic Rössler attractor (cf. Fig. 3.1b,c). To determine the values of K for which \mathcal{M}_s is stable, we note that

$$D\mathbf{F}(\mathbf{r}) = \begin{pmatrix} 0 & -1 & -1 \\ 1 & a & 0 \\ z & 0 & (x - c) \end{pmatrix} \quad \text{and} \quad D\mathbf{H}(\mathbf{r}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

and then estimate $\text{MSF}_{\mathbf{F},\mathbf{H}}(\alpha)$ by determining the maximum Lyapunov exponent of (3.14) as a function of α . The result is shown in Fig. 3.3. $\text{MSF}_{\mathbf{F},\mathbf{H}}(\alpha)$ starts at $\alpha = 0$ with a positive value, reflecting the fact that the Rössler attractor is chaotic and therefore has a positive maximum Lyapunov exponent (cf. Section 2.2.2). As α is increased, $\text{MSF}_{\mathbf{F},\mathbf{H}}$ crosses zero at $\alpha_1 \approx 0.1232$, remains negative until $\alpha_2 \approx 4.663$ and then stays above zero indefinitely [Pecora, 1998].

We can now assess the local stability of the synchronous state \mathcal{M}_s by computing the eigenvalues $\lambda_2, \dots, \lambda_N$ of the network's Laplacian \mathbf{L} and checking whether $\text{MSF}_{\mathbf{F},\mathbf{H}}(K\lambda_i) < 0$ for all of them (cf. Eq. (3.15)). As $\text{MSF}_{\mathbf{F},\mathbf{H}}$ is negative only in the interval (α_1, α_2) , it actually suffices to check whether

$$\alpha_1 < K\lambda_2 \leq \dots \leq K\lambda_N < \alpha_2.$$

The two extremal non-zero eigenvalues of the example network read $\lambda_2 = 1.236$ and $\lambda_N = 13.871$. This means: The synchronous state is indeed *unstable* for $K = 1.50$ (as observed in Fig. 3.1d), as then $K\lambda_N = 20.807 > \alpha_2$. In contrast, the synchronous state is *stable* for $K = 0.15$ (as seen in Fig. 3.1c), because then both $K\lambda_2 = 0.185$ and $K\lambda_N = 2.081$ are within the interval (α_1, α_2) .

3.3.3 Synchronizability

Master stability functions can display a host of different shapes [Huang et al., 2009]. Indeed, for some choices of \mathbf{F} and \mathbf{H} , the function $\text{MSF}_{\mathbf{F},\mathbf{H}}(\alpha)$ is positive for all α and no network will ever spontaneously synchronize. For other choices, $\text{MSF}_{\mathbf{F},\mathbf{H}}(\alpha)$ crosses the zero line several times. Particularly interesting are choices of \mathbf{F} and \mathbf{H} for which $\text{MSF}_{\mathbf{F},\mathbf{H}}(\alpha)$ falls below zero at α_1 and then remains negative. In this case, the stability of a network's synchronous state \mathcal{M}_s depends only on $K\lambda_2$: indeed, \mathcal{M}_s is

stable if $K\lambda_2 > \alpha_1$. As networks with a larger λ_2 are able to synchronize at smaller values of K , they are regarded as more synchronizable w.r.t. such so-called *unbounded* master stability functions. The eigenvalue λ_2 , which graph theory calls *algebraic connectivity* [Fiedler, 1973; de Abreu, 2007], is also related to the convergence speed towards synchronization [Almendral and Díaz-Guilera, 2007].

The other particularly interesting shape of master stability functions that has received a lot of attention and that we will concentrate on here is illustrated in Fig. 3.3. In this *bounded* case, $\text{MSF}_{\mathbf{F},\mathbf{H}}(\alpha)$ is negative only in an interval $\alpha \in (\alpha_1, \alpha_2)$. Therefore, the synchronous state \mathcal{M}_s is stable in every network for which

$$\alpha_1 < K\lambda_2 \leq \dots \leq K\lambda_N < \alpha_2. \quad (3.16)$$

The final term, $K\lambda_N < \alpha_2$, implies that – somewhat counterintuitively – an increase in the coupling strength K can be *detrimental* for stability: Whereas \mathcal{M}_s may be stable for intermediate values of K above α_1/λ_2 , it becomes unstable when the coupling is driven beyond α_2/λ_N . This is exactly what we witnessed in our example (cf. Section 3.1).

Condition (3.16) is equivalent to demanding that

$$R := \frac{\lambda_N}{\lambda_2} < \frac{\alpha_2}{\alpha_1}, \quad (3.17)$$

provided K is chosen from the *stability interval*

$$I_s = (\alpha_1/\lambda_2, \alpha_2/\lambda_N). \quad (3.18)$$

As inequality (3.17) can more easily be met by networks for which the eigenratio R is small, such networks are regarded as *more synchronizable* w.r.t. bounded master stability functions. Hence R is called the *synchronizability* of a network.

3.3.4 The puzzle

In the past decade, a host of studies have investigated the relation between network topology and synchronizability R . Examples include the papers Hong et al. [2004]; Barahona and Pecora [2002]; Nishikawa et al. [2003]; Zhao et al. [2006]; Zhou et al. [2006]; Zhou and Kurths [2006]. Notably, Barahona and Pecora [2002] found that synchronizability R improves markedly as random edges are added to an initially regular lattice. It was argued that this improvement in R could be caused by two parallel topological effects of random link addition: first, by the plain increase in the number of edges; and, second, by the increase in the number of short cuts. Differentiating between the two, Hong et al. [2004] performed a similar study in a slightly different setting. They employed the Watts-Strogatz graph generating model that induces short cuts as well, but does so at a fixed number of edges [Watts and Strogatz, 1998]. The results of Hong et al. are highly important to what follows. So let us dive into the details here.

The Watts-Strogatz algorithm creates a network by initially setting up a regular lattice of N nodes of which each one is connected to its k nearest neighbours. Such a lattice is shown in Fig. 3.4a. Then each edge is rewired independently with probability p by re-drawing one of its end points uniformly at random.⁴ For large p , this yields heavily rewired networks similar to completely random graphs (Fig. 3.4c). For small non-zero p , in contrast, the resulting networks still basically look like the initial regular lattice, differing only by a few rewiring-induced short cuts (Fig. 3.4b).




Figure 3.4: **Watts-Strogatz graph generation model.** Example networks with $N = 16$ and $k = 4$. **a**, initial 1-dimensional configuration: a regular ring lattice in which the degree $d_i = k$ for every node $i = 1, \dots, N$. This is also the outcome of the model for rewiring probability $p = 0$. **b, c**, typical outcomes for rewiring probability $p = 0.2$ and $p = 0.9$.

It turns out that a few such short cuts already have a strong effect on the *global* network structure [Watts and Strogatz, 1998]: In comparison to the initial lattice, they drastically lower the average shortest-path length L . Indeed, as the number of rewiring-induced short cuts increases, L quickly reaches the very low value it is known to have in random graphs. In contrast, the effect of a few short cuts on the *local* network structure is comparatively minute. This is illustrated by the network in Fig. 3.4b, whose clustering coefficient C is still about as large as that of the initial lattice and hence far larger than that of the random graph in Fig. 3.4c.

Watts and Strogatz noted that, very much like the networks produced from their model at small non-zero p , many important real-world networks also reveal a small average shortest-path length combined with a high clustering coefficient. This class of networks has since been called *small-world* networks. Explicitly, one says that a network has the small-world property if its average shortest-path length L is close to the value L_R in a random graph of the same size and if its average clustering coefficient C is much larger than the value C_R in such a random graph. That is [Watts and Strogatz, 1998]:

$$L \approx L_R \text{ and } C \gg C_R \text{ defines } \textit{small-worldness}. \quad (3.19)$$

⁴Self-loops and multi-edges are typically not allowed.

3 Synchronization on complex networks

From the Watts-Strogatz model’s perspective, small-world networks reside *in the middle ground* between regular lattices and random graphs.

As mentioned above, Hong et al. [2004] applied the Watts-Strogatz model in their search for optimally synchronizable networks. They found that, as the model is tuned from regular lattices (model parameter $p = 0$) to random graphs ($p = 1$), synchronizability R shows a strong, monotonical improvement (reproduced in Fig. 3.5a). Consequently, according to linearization-based synchronizability, real-world networks whose function relies on synchronization should ideally look like random graphs. But neural networks and power grids display small-world topologies that, from the Watts-Strogatz model’s perspective, are *far more regular* than random graphs [Watts and Strogatz, 1998; Sporns and Zwi, 2004; Wang et al., 2010]. Indeed, when building well functioning synchronizable networks, nature and civilization appear to shun the predicted randomness (for details, see Section 3.4.2). This discrepancy between theory and observation has left networks research with a long-standing puzzle.




Figure 3.5: **Synchronizability and basin stability in Watts-Strogatz networks of chaotic oscillators.** **a**, expected synchronizability $\langle R \rangle$ versus the Watts-Strogatz model's parameter p . The scale of the y -axis was reversed to indicate improvement upon increase in p . **b**, expected basin stability $\langle \bar{S}_B \rangle$ versus p . The grey shade indicates \pm one standard deviation. The dashed line shows an exponential fitted to the ensemble results for $p \geq 0.15$. Solid lines are guides to the eye. The plots shown were obtained for $N = 100$ oscillators of Rössler type, each having on average $k = 8$ neighbours. Choices of larger N and different k produce results that are qualitatively the same, see Fig. 3.7. Figure published in [Menck et al., 2013a].

3.4 The nonlocal stability perspective

3.4.1 Basin stability in Rössler networks

In an attempt to complement the theory, let us apply basin stability to ensembles of Watts-Strogatz networks at whose nodes we place Rössler oscillators. The dynamics at node i obey

$$\dot{x}_i = -y_i - z_i - K \sum_{j=1}^N L_{ij} x_j \quad (3.20)$$

$$\dot{y}_i = x_i + a y_i \quad (3.21)$$

$$\dot{z}_i = b + z_i(x_i - c) \quad (3.22)$$

with coupling constant K , Laplacian matrix \mathbf{L} , $a = b = 0.2$, and $c = 7.0$. From the master stability function formalism (Section 3.3.1), we know that every such network has a completely synchronous state \mathcal{M}_s in which all nodes follow the same trajectory on the Rössler attractor (Fig. 3.1b). A network's synchronous state is stable if its synchronizability $R < \alpha_2/\alpha_1 = 37.88$ and $K \in I_s = (\alpha_1/\lambda_2, \alpha_2/\lambda_N)$, where $\alpha_1 = 0.1232$ and $\alpha_2 = 4.663$. However, the level of R does not quantify *how stable* the synchronous state is against large perturbations. To address this yet unasked question, we will here consider the synchronous state's basin stability $S_{\mathcal{B}}$ [Menck et al., 2013a].

Let \mathcal{B} denote the basin of attraction of a network's synchronous state and define its basin stability as

$$S_{\mathcal{B} \cap \mathcal{Q}} := \text{Vol}(\mathcal{B} \cap \mathcal{Q}) / \text{Vol}(\mathcal{Q}) \in [0, 1] \quad (3.23)$$

where \mathcal{Q} is a subset of state space with finite volume. This corresponds to the definition of basin stability in Section 2.3.2 with ρ chosen as a uniform distribution on \mathcal{Q} and zero elsewhere; so we can use the numerical procedure described in Section 2.4.1 to estimate $S_{\mathcal{B} \cap \mathcal{Q}}$. We choose $T = 500$, so that the estimate carries a standard error $e < 0.023$, and employ

$$\mathcal{Q} := q^N \text{ with } q = [-15, 15] \times [-15, 15] \times [-5, 35]. \quad (3.24)$$

Note that the Rössler attractor (Fig. 3.1b) is included in q . It turns out that other choices of \mathcal{Q} – for instance $([-8, 8] \times [-8, 8] \times [-8, 8])^N$ – would not produce a qualitative difference in what we will find out below (cf. Section 3.6.1). Therefore we henceforth suppress the subscript \mathcal{Q} when stating $S_{\mathcal{B}}$.

For each of several different values of the Watts-Strogatz model parameter p , we now generate an ensemble of 100 networks. Then, for each ensemble network we perform an estimation of basin stability $S_{\mathcal{B}}$ for ten equally spaced values of the coupling constant K inside the network-specific stability interval I_s (cf. Eq. (3.18)) and compute their mean $\bar{S}_{\mathcal{B}} = \text{mean}_{K \in I_s} S_{\mathcal{B}}(K)$. For instance, our example network (see Sections 3.1 and 3.3.2), which is actually a member of our ensembles, has the stability interval $I_s = (0.010, 0.336)$. At ten equally spaced K -values in the interior of I_s one gets

K	0.121	0.143	0.164	0.185	0.207	0.229	0.250	0.272	0.293	0.315
$S_{\mathcal{B}}(K)$	0.203	0.266	0.323	0.388	0.452	0.538	0.600	0.662	0.684	0.747

These give the average basin stability $\bar{S}_{\mathcal{B}} \approx 0.49$. We compute $\bar{S}_{\mathcal{B}}$ like this for every ensemble network and finally average $\bar{S}_{\mathcal{B}}$ over the 100 networks belonging to each p -value to obtain expected basin stability $\langle \bar{S}_{\mathcal{B}} \rangle$ as a function of p .

The main finding is depicted in Fig. 3.5b (see also Fig. 3.7): In sharp contrast to expected synchronizability $\langle R \rangle$, expected basin stability $\langle \bar{S}_{\mathcal{B}} \rangle$ *declines* exponentially fast as networks become more random.⁵ Therefore, the synchronous state is much *more stable* in networks that are *more regular*.

Why is this? Fig. 3.6a displays that the expected basin stability $\langle S_{\mathcal{B}}(K) \rangle$ of a network at any fixed coupling K increases strongly as K grows and that it hardly depends on p . Hence a network's mean basin stability $\bar{S}_{\mathcal{B}}$ should be determined primarily by the location of its stability interval I_s . Indeed, for increasing p , the expected stability interval $\langle I_s \rangle = (\alpha_1/\langle \lambda_2 \rangle, \alpha_2/\langle \lambda_N \rangle)$ of a network simultaneously broadens and shifts to the left (see Fig. 3.6b). This qualitatively explains the rapidly decreasing behaviour of expected mean basin stability shown in Fig. 3.5b. It also implies that the mean basin stability $\bar{S}_{\mathcal{B}}$ of a Rössler network mainly depends on the absolute values of λ_2 and λ_N (which determine I_s), as opposed to their ratio which is central to synchronizability studies (cf. Eq. (3.17)).




Figure 3.6: **Explanatory characteristics.** **a**, expected basin stability $\langle S_{\mathcal{B}}(K) \rangle$ at coupling K , measured in the interior of the expected stability interval for $p = 0.2, 0.5, 0.8$. Solid lines are guides to the eye. **b**, expected stability interval $\langle I_s \rangle$ at rewiring probability p . The red solid line represents its left bound $\alpha_1/\langle \lambda_2 \rangle$, the green dashed line its right bound $\alpha_2/\langle \lambda_N \rangle$.

⁵ The results obtained are not qualitatively different for networks produced by a two-dimensional Watts-Strogatz model and another model that varies the link length distribution [Li et al., 2010]. Details on this are provided in Section 3.5.




Figure 3.7: **Basin stability in Rössler networks.** Expected basin stability $\langle \bar{S}_B \rangle$ versus p for Watts-Strogatz networks with different k and **a:** $N = 100$, **b:** $N = 200$. Solid lines are guides to the eye.

3.4.2 The puzzle – solved!?

As discussed above, the local-stability-based theory of Hong et al. [2004] suggests that networks with random topologies should be optimal for supporting synchronous operation. In contrast, synchronizing networks from the real world were reported to display *small-world* topologies that, from the Watts-Strogatz model’s perspective, are comparatively regular. We will see at the end of this subsection that our results on basin stability may provide a solution to this puzzle. Yet for concreteness, let us first take a close look at the network properties of several synchronizing real-world networks, including power grids and neural networks (see Table 3.1).

Recall from Eq. (3.19) that a network is said to have a small-world topology if $C \gg C_R$ and $L \approx L_R$, where C_R and L_R are the expected values of C and L in a random network with the same numbers N, E of nodes and edges [Watts and Strogatz, 1998]. Hence in order to assess the small-worldness of our networks, we have to estimate C_R and L_R . To do that, we apply the Watts-Strogatz algorithm with rewiring probability $p = 1$ to regular ring lattices of a given N and E and average over many realizations. The initial regular ring lattices are created in two steps. First, connect every node to its k nearest neighbours, where k is the largest even integer smaller than $\beta E/N$ with $\beta = 2$ ($\beta = 1$) for undirected (directed) networks. Second, add the remaining $E - kN/\beta$ edges randomly between nodes that are $k/2 + 1$ steps apart on the ring [Sporns and Zwi, 2004].

The values in Table 3.1 show that it is hard to judge just from the comparison of C to C_R and of L to L_R whether a network is small-world or not. Indeed, whereas for the Macaque Visual Cortex one has $C/C_R \approx 1.42$ and $L/L_R \approx 1.04$, the UK Power Grid gives $C/C_R \approx 11.76$ and $L/L_R \approx 1.61$. Does any of this count as small-worldness?

#	Network	N	E	C	L	C_R	L_R
1	Macaque Visual Cortex	30	311	0.53	1.73	0.36	1.66
2	Macaque Cortex	71	746	0.46	2.38	0.15	2.03
3	Cat Cortex	52	820	0.55	1.81	0.31	1.70
4	C. Elegans Neural Network	297	2345	0.18	3.99	0.027	2.98
5	Power Grid of Western US	4941	6594	0.08	18.99	0.00035	10.19
6	Power Grid of Central Europe	4335	5551	0.07	28.73	0.00020	10.73
7	Power Grid of Northern Europe	524	640	0.04	14.49	0.0030	8.38
8	Power Grid of the UK	393	484	0.04	12.54	0.0034	7.78

Table 3.1: **Topological properties of synchronizing real-world networks.** N specifies the number of nodes and E the number of edges in the network. Furthermore, L is the average shortest-path length and C the clustering coefficient. These two quantities have been widely used to characterize small-worldness [Watts and Strogatz, 1998]. X_R represent average values of X computed in random networks of the same N, E , where $X = L, C$. For networks 1 to 3, the values of L and C were taken from [Sporns and Zwi, 2004].

In the Watts-Strogatz model, C and L assume their lowest values (C_R, L_R) in random graphs ($p = 1$) and their largest values (C_L, L_L) in regular lattices ($p = 0$). So Sporns and Zwi [2004] suggested to also take into account C_L and L_L when assessing small-worldness. They defined

$$\zeta_C = \frac{C - C_R}{C_L - C_R} \text{ and } \zeta_L = \frac{L - L_R}{L_L - L_R} \quad (3.25)$$

and then referred to a network as small-world if ζ_C is high and ζ_L is low. Following their reasoning, we also obtain C_L and L_L as well as ζ_C and ζ_L for our networks. However, as can be seen in Table 3.2, the network size parameters N and E appear to influence how big ζ_C should be to count as high, and how small ζ_L should be to count low. To get rid of this network size effect, let us introduce [Menck et al., 2013a]

$$\xi_X = -\frac{\log(X/X_L)}{\log(X_L/X_R)} = 1 - \frac{\log(X/X_R)}{\log(X_L/X_R)} \quad (3.26)$$

with $X = C$ or L . ξ_X counts how many of orders of magnitude X is away from X_L in relation to the number of orders of magnitude between X_L and X_R .⁶ Lattices have $\xi_C = \xi_L = 0$, whereas random graphs have $\xi_C = \xi_L = 1$. From the above, small-world networks should display ξ_C not too far from 0 and ξ_L not too far from 1, and indeed the real-world networks studied here do so (see Table 3.2).

⁶Note that $\log C$ can be interpreted as a measure of dimensionality [Donner et al., 2011].

3 Synchronization on complex networks

We can now use ξ_C and ξ_L to topologically compare the real-world networks listed in Tables 3.1 and 3.2 to our ensemble findings on basin stability and synchronizability [Menck et al., 2013a]. This is done in Fig. 3.8. The figure summarizes what we have observed above: Whereas the local-stability-based synchronizability R improves as networks become more *random*, the synchronous state's basin stability, our nonlocal measure of how stable this state is, increases as networks become more *regular*. This adds a crucial piece to the puzzle about the topologies of synchronizing real-world networks and, I conjecture, makes its solution emerge: In synchronizing networks, the functional need for the synchronous state to be as stable as possible promotes topological regularity. Thus during network evolution, the optimization for synchronizability and the simultaneous optimization for basin stability have acted as two opposing forces. Their struggle resulted in a topological tradeoff: small-worldness.

#	Network	C_L	L_L	ζ_C	ζ_L	ξ_C	ξ_L
1	Macaque Visual Cortex	0.66	1.97	0.56	0.26	0.36	0.76
2	Macaque Cortex	0.66	3.85	0.60	0.19	0.24	0.75
3	Cat Cortex	0.69	2.14	0.63	0.25	0.28	0.73
4	C. Elegans Neural Network	0.64	19.25	0.25	0.06	0.40	0.84
5	Power Grid of Western US	0.39	925.7	0.20	0.01	0.22	0.86
6	Power Grid of Central Europe	0.35	846.8	0.20	0.02	0.21	0.77
7	Power Grid of Northern Europe	0.29	107.6	0.13	0.06	0.43	0.79
8	Power Grid of the UK	0.30	80.3	0.12	0.07	0.44	0.80

Table 3.2: **More topological properties of synchronizing real-world networks.**

X_L represents average values of X computed in lattices of the same N, E , where $X = L, C$. The measures ζ_L and ζ_C , defined in Eq. (3.25), were suggested in [Sporns and Zwi, 2004] to characterize small-worldness. ξ_L and ξ_C , defined in Eq. (3.26), are new measures less prone to the influence of the network size than ζ_L, ζ_C .




Figure 3.8: **Topological comparison of ensemble results with real-world networks.** A circle represents the results for Watts-Strogatz networks with $N = 100$, $k = 10$ and rewiring probability $p \in \{0.05, 0.1, 0.15, \dots, 1.0\}$ (p increasing from left to right). A circle's area is proportional to the expected basin stability $\langle \bar{S}_B \rangle$. A circle's colour indicates the expected synchronizability $\langle R \rangle$. Squares represent real-world networks reported to display a small-world topology (Tables 3.1 and 3.2). Plotting ξ_L against ξ_C allows to compare networks of different sizes w.r.t. average shortest-path L and clustering coefficient C , quantities that characterize small-worldness [Watts and Strogatz, 1998]. $(\xi_L, \xi_C) = (0, 0)$ labels a regular network whereas $(\xi_L, \xi_C) = (1, 1)$ labels a random network. Small-world networks reside in the top left quadrant. Figure published in [Menck et al., 2013a].

3.5 Conclusions and outlook

We have studied how general dynamical networks should be structured to optimally support the stable synchronous operation of their components. Previous theory suggests that synchrony could most easily be achieved in networks that have highly random topologies. However, real-world networks whose function relies on synchronization – such as neural networks or power grids – have long been known to possess small-world topologies that are comparatively regular. Addressing this puzzle, we noted that the previous theory is based on the inherently local linear approach to stability, and applied basin stability to add a nonlocal perspective.

Investigating ensembles of paradigmatic chaotic oscillators, we found that, in sharp contrast to linearization-based synchronizability, nonlocal basin stability improves substantially the more regular a network is. My conjecture is that this solves the puzzle: The qualitative optimization towards a locally stable synchronous state and the quantitative optimization towards a synchronous state that is as stable as possible nonlocally have acted as antagonistic evolutionary forces. The final outcome – small-worldness – emerged as a topological compromise.

Against the broadness of this conjecture, the setting of the investigations in this chapter may appear rather narrow: We have explored the stability properties of only one type of synchronization (complete synchronization) in simple networks (undirected, unweighted) of only one type of chaotic oscillator (Rössler). In the sense of Popper [1963], future work might test whether our findings are refuted or corroborated under different circumstances. Furthermore:

(i) In the networks we have studied, all nodes were identical oscillators and capable of producing a completely synchronous state. However, the nodes of real-world synchronizing networks are not, in general, identical. So a possible next step could be to investigate how topology influences the basin stability of the phase synchronized state in networks of non-identical chaotic oscillators.

(ii) Specifically, our networks were made up of chaotic Rössler oscillators that are well-known to be capable of producing rich collective dynamics. However, Rössler oscillators do certainly not constitute a realistic model of real-world synchronizing networks such as neural networks or power grids. It could be interesting to test our findings when actually employing realistic models.

(iii) The edges in our networks all carried the same, uniform weight. It is an open question whether basin stability can be systematically optimized by changing the edge weights of a fixed topology.

(iv) Similarly, we focussed on undirected topologies here. The effects of directed topological elements on basin stability are yet to be explored.

3.6 Supporting material

3.6.1 Another choice of the reference subset \mathcal{Q}

The findings reported in Fig. 3.5b on the relation between expected basin stability $\langle \bar{S}_{\mathcal{B}} \rangle$ and the Watts-Strogatz model parameter p are based on a specific choice of the reference subset \mathcal{Q} in the definition of basin stability, namely on $\mathcal{Q} = q^N$ with $q = [-15, 15] \times [-15, 15] \times [-5, 35]$ (cf. Section 3.4). Here, we perform the same ensemble study again, yet now with respect to

$$\tilde{\mathcal{Q}} = \tilde{q}^N \text{ with } \tilde{q} = [-8, 8] \times [-8, 8] \times [-8, 8]. \quad (3.27)$$

The outcome shown in Fig. 3.9. Apparently, there is no qualitative difference from the results shown in Fig. 3.5b: Expected basin stability becomes ever better the more regular a network is.




Figure 3.9: **Basin stability with respect to another reference subset \mathcal{Q} in Watts-Strogatz networks of chaotic oscillators.** Shown is expected basin stability $\langle \bar{S}_{\mathcal{B}} \rangle$ versus the Watts-Strogatz model's parameter p , with $N = 100$ oscillators of Rössler type, each having on average $k = 8$ neighbours. Here, a different reference subset \mathcal{Q} is used in the basin stability estimation than for the results reported in Fig. 3.5, namely $\tilde{\mathcal{Q}} = \tilde{q}^N$ with $\tilde{q} = [-8, 8] \times [-8, 8] \times [-8, 8]$ (cf. Section 3.4).

3.6.2 2-dimensional Watts-Strogatz networks

To obtain our main results on basin stability in Section 3.4, we used a 1-dimensional ring as the initial configuration of the Watts-Strogatz algorithm (cf. Fig. 3.4). Here, we will use the 2-dimensional lattice depicted in Fig. 3.10.




Figure 3.10: **Initial grid for 2-dimensional Watts-Strogatz network generation.** Each node is connected to its four nearest neighbours and the four diagonal ones among its next-nearest neighbours. So each node has 8 links.

We choose $N = 20 \cdot 20 = 400$ as the network size and, as before, place a Rössler oscillator at each node. Note that basin stability calculations on this lattice size approach the limits of what is computationally feasible today. For each network in the ensemble we estimate basin stability $S_B(K)$ for several $K \in I_s$ and compute their mean \bar{S}_B . Here, we use a slightly smaller reference subset, $\mathcal{Q} = q^N$ with $q = [-15, 15] \times [-15, 15] \times [-4, 35]$, than above as otherwise the values of S_B would be too small to be accurately measurable with $T = 500$ integrations (cf. Section 2.4.1). Note that the Rössler attractor (Fig. 3.1b) is still included in q . Finally, we average \bar{S}_B over the ensemble to obtain $\langle \bar{S}_B \rangle$.

The results are shown in Fig. 3.11. There is no qualitative difference to Fig. 3.5: Whereas synchronizability improves as networks become more random, expected basin stability is larger in networks that are more regular and falls off exponentially with increasing randomness.

Fig. 3.12 shows how the 2-dimensional results compare topologically to synchronizing networks from the real world (cf. Fig. 3.8). Again, there is no qualitative difference. Whereas the need for good synchronizability drives networks towards randomness, the need for large basin stability drives them towards regularity.




Figure 3.11: **Synchronizability and basin stability in 2-dimensional Watts-Strogatz networks of chaotic oscillators.** **a**, expected synchronizability $\langle R \rangle$ versus the Watts-Strogatz model's parameter p . The scale of the y -axis was reversed to indicate improvement upon increase in p . **b**, expected basin stability $\langle \bar{S}_B \rangle$ versus p . The grey shade indicates \pm one standard deviation. The dashed line shows an exponential fitted to the ensemble results for $p \geq 0.05$. Solid lines are guides to the eye. The plots shown were obtained for $N = 400$ oscillators of Rössler type, each having on average $k = 8$ neighbours. Figure published in [Menck et al., 2013a].




Figure 3.12: **Topological comparison of results for the 2-dimensional ensemble with real-world networks.** A circle represents the results for 2D Watts-Strogatz networks with $N = 400$, $k = 8$ and rewiring probability $p \in \{0.05, 0.1, 0.15, \dots, 1.0\}$ (p increasing from left to right). A circle's area is proportional to the expected basin stability $\langle \bar{S}_B \rangle$. A circle's colour indicates the expected synchronizability $\langle R \rangle$. Squares represent real-world networks reported to display a small-world topology (Tables 3.1 and 3.2). Plotting ξ_L against ξ_C allows to compare networks of different sizes w.r.t. average shortest-path L and clustering coefficient C , quantities that characterize small-worldness [Watts and Strogatz, 1998]. $(\xi_L, \xi_C) = (0, 0)$ labels a regular network whereas $(\xi_L, \xi_C) = (1, 1)$ labels a random network. Small-world networks reside in the top left quadrant. Figure published in [Menck et al., 2013a].

3.6.3 Small-World networks with different link length distributions

Here we perform another basin stability study on small-world networks with different link length distributions. The networks are created as follows: We start with a 2-dimensional lattice in which every node is connected to its 4 nearest neighbours. Then we add short cuts so that, in the end, the average number of links per node is 8. For short cut addition, we use the following procedure [Li et al., 2010]: A pair of nodes, (i, j) is chosen randomly. Then a new link connecting them is added to the network with probability

$$P(r) \propto r^{-\alpha},$$

where r is the Manhattan distance between i and j , i.e., the distance between them on the original 2-dimensional lattice with no short cuts.

In networks created like this, we expect the length of a typical short cut to decline when α increases. How do basin stability and synchronizability depend on α ?

To investigate this, we choose the network size to be $N = 20 \cdot 20 = 400$ and, as before, place a Rössler oscillator at each node. Then, for every network we compute $\bar{S}_{\mathcal{B}}$. Here, we use a slightly smaller reference subset, $\mathcal{Q} = q^N$ with $q = [-15, 15] \times [-15, 15] \times [-3, 35]$, than above as otherwise the values of $S_{\mathcal{B}}$ are too small to be accurately measurable with $T = 500$ integrations (cf. Section 2.4.1). Note that the Rössler attractor is still included in q . Finally, we average $\bar{S}_{\mathcal{B}}$ over the ensemble to obtain $\langle \bar{S}_{\mathcal{B}} \rangle$.

Fig. 3.13a shows that expected synchronizability $\langle R \rangle$ declines as α increases. This means that, from the linear-stability perspective, an optimal link length distribution should fall off slowly or not at all (corresponding to rather low α).

The perspective of basin stability, again, disagrees with this. In Fig. 3.13b, we see that $\langle \bar{S}_{\mathcal{B}} \rangle$ evolves in an opposite way to expected synchronizability in a large portion of the parameter interval: $\langle \bar{S}_{\mathcal{B}} \rangle$ grows as α increases. Finally, as α is tuned up further, basin stability reaches a maximum and then declines rapidly towards zero as synchronizability approaches the stability threshold.

These results appear not to contain considerably more information than Fig. 3.5 and Fig. 3.7: For large α as well as for small p , networks have an intense local structure that is reduced as α decreases and p increases. This may be the common reason behind the phenomena we observe.




Figure 3.13: **Synchronizability and basin stability in networks of chaotic oscillators with different link length distribution.** **a**, expected synchronizability $\langle R \rangle$ versus the link length distributions's parameter α . The scale of the y -axis was reversed. The grey shade indicates \pm one standard deviation. **b**, expected basin stability $\langle \bar{S}_B \rangle$ versus α . The grey shade indicates \pm one standard deviation. Solid lines are guides to the eye. The plots shown were obtained for $N = 400$ oscillators of Rössler type, each having on average $k = 8$ neighbours.

3.6.4 Non-convexity of the basin in Rössler-networks

One might have the idea to measure the volume of the basin \mathcal{B} of the completely synchronous state in a network of chaotic Rössler oscillators in an absolute rather than a relative sense. For this, high-dimensionality poses challenges. If \mathcal{B} were a bounded convex set, its volume could be estimated in $\mathcal{O}(n^4)$ time steps with today's best algorithm [Lovász and Vempala, 2006], where $n = 3N$ is the dimension of state space. We have $N \geq 100$, so this would be numerically very expensive.

In any case, it turns out that \mathcal{B} is not convex. This follows from the non-convexity of the chaotic attractor's basin \mathcal{B}_1 of a single Rössler oscillator, as can be seen in the two-dimensional details displayed in Fig. 3.14. As the set $\mathcal{B}_1^N := \{\mathbf{r}_1 = \dots = \mathbf{r}_N = \mathbf{r} \mid \mathbf{r} \in \mathcal{B}_1\}$ is a non-convex slice of the basin \mathcal{B} of any network's synchronous state, \mathcal{B} is non-convex, too. Hence the powerful tools for volume estimation of convex bodies are not applicable here.




Figure 3.14: **Two-dimensional details of the Rössler attractor's basin.** **a** shows the xy -detail. A point (x, y) refers to the initial state $(x, y, 0)$. **b** shows the xz -detail. A point (x, z) refers to the initial state $(x, 0, z)$. In both panels the white region indicates the Rössler attractor's basin of attraction. The green shape depicts a two-dimensional projection of the Rössler attractor.

3.6.5 Edge-list of the example network

For the sake of reproducibility, here we study in detail the estimation of mean basin stability \bar{S}_B for the example network depicted in Fig. 3.1a and discussed in Sections 3.1 and 3.3.2. Its edge list reads:

0-1	0-2	0-3	4-61	2-91	1-3	1-4	1-5	3-20	2-4	2-5	2-6	3-4	3-5
3-6	3-7	4-5	4-6	7-73	4-8	5-6	5-7	5-8	5-9	7-88	8-45	6-9	6-10
7-8	7-9	7-10	7-11	8-9	8-10	8-11	8-12	9-10	9-11	9-12	9-13	10-11	10-12
10-13	10-14	12-93	11-13	14-3	11-15	12-13	12-14	12-15	12-16	14-49	13-15	13-16	13-17
14-15	14-16	14-17	14-18	15-16	15-17	15-18	19-65	16-17	16-18	16-19	20-94	17-18	17-19
17-20	17-21	18-19	20-65	18-21	18-22	19-20	19-21	19-22	19-23	20-21	22-16	20-23	20-24
21-22	21-23	24-0	21-25	22-23	22-24	22-25	22-26	24-47	23-25	23-26	27-75	24-25	26-34
27-17	28-90	25-26	25-27	28-8	25-29	26-27	26-28	26-29	26-30	27-28	27-29	27-30	27-31
28-29	30-46	31-36	28-32	30-70	29-31	29-32	29-33	30-31	30-32	30-33	30-34	31-32	31-33
31-34	35-52	32-33	32-34	35-6	36-48	33-34	33-35	33-36	33-37	34-35	34-36	34-37	34-38
35-36	35-37	35-38	35-39	36-37	36-38	39-73	36-40	37-38	37-39	40-84	37-41	38-39	38-40
38-41	42-93	39-40	41-46	39-42	43-13	40-41	40-42	43-86	40-44	41-42	41-43	41-44	45-67
42-43	42-44	45-77	46-95	43-44	43-45	43-46	43-47	44-45	44-46	44-47	44-48	45-46	47-60
45-48	49-75	46-47	46-48	46-49	46-50	47-48	47-49	47-50	51-96	48-49	48-50	48-51	52-34
49-50	49-51	49-52	49-53	50-51	50-52	50-53	50-54	51-52	51-53	51-54	51-55	53-28	54-21
52-55	56-83	54-1	53-55	56-88	57-88	55-79	54-56	54-57	58-78	56-85	55-57	55-58	55-59
56-57	56-58	56-59	56-60	57-58	57-59	60-45	57-61	58-59	58-60	61-36	62-75	60-71	59-61
59-62	63-12	61-33	60-62	60-63	64-7	61-62	61-63	61-64	61-65	62-63	62-64	62-65	66-77
63-64	63-65	66-74	67-47	64-65	64-66	67-95	68-19	65-66	65-67	65-68	65-69	66-67	66-68
66-69	66-70	67-68	67-69	67-70	67-71	68-69	68-70	68-71	68-72	69-70	69-71	69-72	73-14
70-71	70-72	73-18	74-16	71-72	71-73	71-74	71-75	72-73	72-74	72-75	72-76	74-47	73-75
73-76	77-56	75-10	74-76	74-77	74-78	75-76	75-77	75-78	75-79	77-10	76-78	79-97	76-80
77-78	77-79	77-80	77-81	78-79	78-80	81-93	78-82	79-80	79-81	79-82	79-83	80-81	80-82
83-30	80-84	81-82	81-83	84-14	81-85	82-83	82-84	82-85	86-52	84-50	83-85	83-86	83-87
85-51	86-72	84-87	84-88	85-86	85-87	85-88	85-89	86-87	86-88	86-89	86-90	87-88	87-89
87-90	91-12	88-89	88-90	88-91	88-92	90-97	91-25	89-92	89-93	91-10	92-65	90-93	90-94
92-54	91-93	94-15	91-95	92-93	92-94	92-95	92-96	93-94	93-95	93-96	93-97	95-18	94-96
94-97	94-98	95-96	95-97	95-98	99-59	96-97	96-98	96-99	0-34	98-60	97-99	97-0	97-1
98-99	98-0	98-1	98-2	99-0	99-1	99-2	99-3						

This network consists of $N = 100$ nodes and $E = 400$ edges (and was generated using the Watts-Strogatz model with rewiring probability $p = 0.2$). The minimum and maximum non-zero eigenvalues of its Laplacian matrix are $\lambda_2 = 1.236$ and $\lambda_N = 13.871$, respectively. Hence its stability interval $I_s = (0.010, 0.336)$. The synchronous state is stable for $K \in I_s$. We measure its basin stability at 10 different equally spaced values of K in the interior of I_s , obtaining

K	0.121	0.143	0.164	0.185	0.207	0.229	0.250	0.272	0.293	0.315
$S_B(K)$	0.203	0.266	0.323	0.388	0.452	0.538	0.600	0.662	0.684	0.747

To estimate the mean basin stability \bar{S}_B of this network, we compute the average of these values, which is $\bar{S}_B \approx 0.49$.

3.6.6 Numerical methods

For all the numerical simulations of Rössler networks, we employed a fourth order Runge-Kutta algorithm, integrating for 60.000 steps with step size $\tau = 0.01$. If very long transients were detected, we increased the number of steps up to 960.000.

4 Structure and stability of power grids

Transmission grids for electrical energy, or just *power grids*, are of huge importance for the functioning of today's societies. For example, the German grid delivers about 600 billion kWh of energy per year [Statistisches Bundesamt, 2012], which amounts to an average power supply of roughly 69 million kW, or 850 W for each of today's 80,2 million inhabitants. For comparison, a human being engaged in long-time endurance sports can achieve something like 85 W. So the German power grid provides an empowerment factor of ten, implying that, if the same service was to be delivered by human muscle effort alone, every inhabitant would have to employ ten energy workers drudging 24 hours a day, or 40 workers labouring 42 hours a week. Under the latter conditions, a total of 3.2 billion workers would be needed to meet the energy demand of this single country.

This phantastic number underscores the fact that a power grid fulfils a highly important task. The flipside is that its failure can impose enormous costs. For instance, the largest blackout on Northern American records, occurring in August 2003, "affected an area with an estimated 50 million people", and, in some regions, "power was not restored for 4 days" [U.S.-Canada Power System Outage Task Force, 2004]. The total costs of this event were estimated between \$5 billion and \$10 billion. Worryingly, such blackouts occur rather often: The distribution of blackout sizes reveals no exponential shape, as one might have hoped, but falls off extremely slowly with a power-law behaviour [Carreras et al., 2004; Newman et al., 2011].

Against this backdrop, it is no surprise that the scientific community has put considerable effort into understanding the complex dynamics of power grids. Carreras et al. [2004] adopted a system theory perspective and described a power grid as evolving through consecutive periods of load growth which are interrupted by sudden blackout events. After every blackout, grid operators strengthen the grid, only for the next period of load growth to follow. This model, which resembles the growth-avalanche dynamics of a continuously loaded sand pile [Fairley, 2004], reproduces the actual blackout size distribution remarkably well [Newman et al., 2011].

Other researchers investigated the mechanics of large-scale blackouts from a complex network theory perspective. They elaborated on the central observation that such events usually include cascading breakdowns of grid components. Motter and Lai [2002] and Crucitti et al. [2004a,b] investigated how the failure of a *single node* in a big network can bring about a long series of further failures, thus wrecking grid performance to a minimum. Buldyrev et al. [2010] included in their considerations that failures in power grids are actually not independent from, but interrelated with events in other complex systems, such as communication networks. They found that such interdependencies in *networks of networks* make power grids even more vulnerable to single node failure.

In all these studies, the dynamics of grid components and the power flows between them are modelled either in a long-time-scale, steady-state fashion or not at all. Here, we will study grid failures at short time scales. From this perspective, a power grid is normally operated in the synchronous state, in which all frequencies are equal to the rated frequency (50 or 60 Hz) and in which steady power flows equate supply and demand everywhere [Machowski et al., 2008; Hill and Chen, 2006; Spring, 2003]. Blackouts have often ensued from a loss of synchrony [Ewart, 1978]. Indeed, when parts of a power grid desynchronize, destructive power oscillations entail. To avoid damage, affected components must then be switched off. Such switchings can in turn desynchronize other grid components, possibly provoking a cascade of further shut-downs and ending in a large-scale blackout.

Fortunately, grids can be – and usually are – designed so that the synchronous state is locally stable, implying that desynchronization cannot be caused by *small perturbations*, such as somebody turning on a household device [Dobson, 2013]. Yet many intriguing questions on the relation between grid topology and local stability are still unanswered, and this is a highly active field of research. Recently, Rohden et al. [2012] found that decentralization of power supply can improve the local stability of the synchronous state. Witthaut and Timme [2012] reported that, counterintuitively, addition of transmission lines can decrease stability. Dörfler et al. [2013] formulated rigorous conditions of a locally stable synchronous state in terms of the wiring topology. And, finally, Motter et al. [2013] encountered a new way to improve local stability by tuning parameters of individual nodes in the grid.

However, even if the synchronous state is stable against *small perturbations*, a power grid’s state space is also populated by numerous stable non-synchronous states, to which the grid might be pushed by short circuits, fluctuations in renewable generation or other *large perturbations* [Machowski et al., 2008; Filatrella et al., 2008; Rohden et al., 2012; Chiang et al., 1995; Chiang, 2011]. Indeed, large perturbations occur so often that a whole sub-branch of power grid engineering, called *transient stability analysis*, has been dedicated to them. The standard transient stability toolbox, based on time domain simulations and Lyapunov function considerations [Chiang et al., 1995; Chiang, 2011], assesses whether or not a power grid will return to synchrony after a given large perturbation. That is: It *qualifies if* the synchronous state is stable against this particular perturbation.

Here we will *quantify how* stable the synchronous state is against large perturbations in general. As the traditional linearization-based approach to stability is too local for this assessment, we will employ a single-node version of basin stability [Menck et al., 2013a], the nonlinear concept defined in Section 2.3. It will enable us to scrutinize how the synchronous state’s stability against large localized perturbations is influenced by patterns of the network topology.

When investigating power grids, one must be aware that these are incredibly complex machines, containing thousands of interacting components, influenced by volatile weather conditions, and subject to sudden operator action. Modelling every detail of a grid’s behaviour therefore is an impossible task [Machowski et al., 2008], and we will refrain from attempting that here. Instead, we will employ a rather basic (but well founded) model of nodal dynamics which helps to have a clear view on our main subject: the influence on grid stability of the network topology. We will use this model

for simulations in an ensemble of artificially generated power grids. The characteristics we will obtain from a statistical analysis of this ensemble all support one main conclusion: The widespread and cheapest of all connection schemes, so-called dead ends and dead trees, strongly diminish stability. When testing this outcome in a case study of the Northern European power grid, we will observe that the inverse also holds true: ‘Healing’ dead ends by addition of transmission lines considerably enhances stability. This suggests that, when expanding today’s grids or planning tomorrow’s grids, dead ends should be avoided.

Contents

The chapter is organized as follows. In Section 4.1, we will analyze the effects of a typical large perturbation that hits a power grid. For the analysis, we will use a basic but powerful model from power grid engineering, called the infinite-grid model. As this model contains just a single dynamical node and is hence of no use for questions concerning topology, in Section 4.2 we will discuss the generalization of this model to N -node networks. This section also contains the introduction of *single-node basin stability*, a component-wise nonlocal stability concept. As an illustration, we will investigate the single-node basin stability of a small example network. In Section 4.3, we will study an ensemble of artificially generated power grids. Here, it will turn out that dead ends, and dead trees in general, have a crucial influence on grid stability. We will investigate why that is. Section 4.4 contains a case study of the Northern European power grid, in which we will find that ‘healing’ of dead ends significantly enhances stability. Section 4.5 concludes the chapter. Some supplementary material is provided in Section 4.6.

Some of the findings and analyses presented here have been published in [Menck and Kurths, 2012; Menck et al., 2013b].

4.1 The infinite-grid model




Figure 4.1: **Generator infinite-grid model.** V denotes the voltage behind the generator's transformer and V_{grid} is the voltage of the 'infinite' power grid. $Z = R + iX$ is the transmission line's total impedance, composed of resistance R and reactance X . The figure contains elements of free graphics files¹.

A large perturbation that hits a power grid typically involves a local power imbalance. Imagine, for example, a thermal generator that delivers a constant amount of power via a single high-voltage transmission line. Suddenly, the line suffers a short circuit. Automatic control devices immediately interrupt the line to clear the fault, and hence the mechanical power injected via the generator's turbine has no electrical way out anymore. By energy conservation, this power surplus needs to go somewhere; it flows into the turbine's rotational energy, thus driving up its frequency. Hence when the transmission line automatically re-closes after some delay, the generator has moved away from its pre-fault working point. From this perturbed state, will it return to the desired synchronous state?

4.1.1 Derivation

To illuminate the described situation we will now introduce the classical model of power grid engineering [Machowski et al., 2008; Hill and Chen, 2006; Filatrella et al., 2008; Rohden et al., 2012; Witthaut and Timme, 2012]. The model is based on two fundamental physical principles: the law of energy conservation and the law of induction. Energy conservation dictates that the amount of power² flowing into the generator must always equal the amount of power flowing out. There are several contributions to this power balance. On the *in*-side, there are the injections of local power generation devices, including thermal injection via the turbine. On the *out*-side, there is local consumption as well as the power flow across the transmission line. Another crucial contribution stems from the fact that the generator's turbine constitutes a considerable energy storage. Indeed, the amount of energy stored in

¹Origin: Wikimedia Commons. Author: MBizon. Filename: Electricity_Grid_Schematic_English.svg.

²We limit ourselves to active power here. Reactive power is assumed to be balanced at all times, corresponding to perfect voltage control. This assumption is called the *decoupling assumption* in power grid engineering.

the turbine's rotating mass typically reaches up to 10 seconds times maximum power output for a thermal generator [Horowitz and Phadke, 2008]. A positive change in this storage enters into the power balance as a negative flow. Hence we obtain

$$\frac{d}{dt} \left(\frac{1}{2} M \Omega^2 \right) = L + L_{\text{control}} - L_{\text{trans}}, \quad (4.1)$$

where Ω is the turbine's angular frequency, M is its moment of inertia, L represents *net power injection*, equalling local power injection minus local power consumption, and L_{trans} denotes the amount of power flowing from the generator to the grid. The role of the control term L_{control} will become clear below.

What is L_{trans} ? Let us denote the generator's AC voltage vector by $\mathbf{V} = V e^{i\theta}$, the grid's voltage vector by $\mathbf{V}_{\text{grid}} = V_g e^{i\theta_g}$ and the line impedance by $\mathbf{Z} = R + iX$. Then the current $\mathbf{I} = (\mathbf{V} - \mathbf{V}_{\text{grid}})/\mathbf{Z}$ and

$$\begin{aligned} L_{\text{trans}} &= \text{Re} [\mathbf{V} \mathbf{I}^*] = \text{Re} \left[\mathbf{V} \frac{\mathbf{V}^* - \mathbf{V}_{\text{grid}}^*}{\mathbf{Z}^*} \right] \\ &= \text{Re} \left[\frac{R + iX}{R^2 + X^2} (V^2 - V V_g e^{i(\theta - \theta_g)}) \right] \\ &= \frac{1}{R^2 + X^2} \left[R (V^2 - V V_g \cos(\theta - \theta_g)) + X V V_g \sin(\theta - \theta_g) \right]. \end{aligned}$$

In the high-voltage transmission part of the grid we will focus on here, $R \ll X$, so that

$$L_{\text{trans}} = \frac{V V_g}{X} \sin(\theta - \theta_g). \quad (4.2)$$

So far, we have not established a link between the electrical quantity θ and the mechanical quantity Ω , which means that plugging (4.2) into (4.1) would, at this stage, not really promise any insights. Fortunately, there is such a link: According to the law of induction, the time derivative of θ , namely the *electrical angular frequency* $\dot{\theta}$, is exactly the same as the *mechanical angular frequency* Ω of the generator's rotating turbine.³ As the generator is always operated as close as possible to the grid's rated frequency ω_r , we write $\Omega = \omega_r + \omega$ and from now on measure θ in a frame of reference that co-rotates with ω_r . This allows us to obtain a closed set of model equations

$$\dot{\theta} = \omega \quad (4.3)$$

$$M \omega_r \dot{\omega} = L - D \omega - \frac{V V_g}{X} \sin(\theta - \theta_g) \quad (4.4)$$

where we have assumed $|\omega| \ll \omega_r$ and set $L_{\text{control}} = -D \omega$. This common form of the control term expresses the effect of damper windings whose job it is to keep the deviation ω from the rated frequency ω_r as small as possible. For a slightly simpler nomenclature, we divide Eq. (4.4) by $M \omega_r$, define *net power input* $P := L/(M \omega_r)$, the *damping constant* $\alpha := D/(M \omega_r)$, and the *transfer capacity* $K := V V_g/(X M \omega_r)$,

³This is only true for a two-pole generator. For a four-pole generator, $\dot{\theta} = 2\Omega$. However, this only changes what follows by a cumbersome multiplicative constant, but does not induce any qualitatively different behaviour. So we stick to the simpler two-pole case here.

and obtain

$$\dot{\theta} = \omega \quad (4.5)$$

$$\dot{\omega} = P - \alpha\omega - K \sin(\theta - \theta_g). \quad (4.6)$$

For functional reasons, the transfer capacity $K > 0$ of the transmission line must be larger than $|P|$. For economic reasons, it is usually not much larger than $|P|$ (cf. Section 4.6.1). In this most basic version of the classical model, the grid is assumed to be ‘infinite’, meaning unaffected by anything that happens at the generator. Hence $\theta_g \equiv 0$. We shall study ‘finite’ grids below.




Figure 4.2: **State space of the infinite-grid model.** Shown is the state space of the model, with $\alpha = 0.1$, $P = 1$ and $K = 8$ (the choice of model parameters is elaborated on in Section 4.6.1). The solid black circle marks the desired synchronous state $(\theta_s, 0)$, and the solid red line shows the non-synchronous limit cycle attractor. The basin of attraction of $(\theta_s, 0)$ is coloured green, and that of the limit cycle is coloured white. The grey dashed line indicates the fault-on (i.e., $K = 0$) trajectory and its end point $(\theta, \omega)_{t_1}$. Figure published in [Menck et al., 2013b].

4.1.2 Perturbation analysis

Now let us recap the short-circuit fault described above. Initially, the generator is in the synchronous state $(\theta_s, \omega_s = 0)$, with $\theta_s = \arcsin(P/K) \in (-\pi/2, \pi/2)$, in which $K \sin \theta$, the power flow across the transmission line, exactly balances net power input P . Then suddenly, at time t_0 , K is switched to 0. This makes ω increase until the line re-connects at time t_1 (trajectory in Fig. 4.2). The crucial question is: Will the generator converge to $(\theta_s, 0)$ from the perturbed state $(\theta, \omega)_{t_1}$? If t_1 is small, then

$(\theta, \omega)_{t_1}$ constitutes a small perturbation w.r.t. $(\theta_s, 0)$ and the standard linearization-based stability analysis applies. It yields that the generator will return to $(\theta_s, 0)$ if $K > |P|$, as then $(\theta_s, 0)$ is locally stable.

Indeed, as discussed in more detail in Section 2.2.3, a standard small-perturbation analysis of Eqs. (4.5)–(4.6) leads to

$$\delta\dot{\theta} = \delta\omega \quad (4.7)$$

$$\delta\dot{\omega} = -\alpha\delta\omega - K\cos(\theta - \theta_g)\delta\theta \quad (4.8)$$

which, evaluated at the synchronous state, has the maximum Lyapunov exponent

$$\sigma_+ = \text{Re} \left\{ -\frac{\alpha}{2} + \frac{\sqrt{\alpha^2 - 4K\cos\theta_s}}{2} \right\}. \quad (4.9)$$

As $\theta_s \in (-\pi/2, \pi/2)$ so that $\cos(\theta_s) > 0$, σ_+ is negative. Hence the synchronous state is locally stable (cf. Section 2.2).

However, the *clearing time* t_1 is typically not that small (it can reach several hundred milliseconds [Machowski et al., 2008]), so that $(\theta, \omega)_{t_1}$ is *not a small perturbation* w.r.t. $(\theta_s, 0)$. This renders the local stability analysis insufficient. Indeed, the generator will only return to the synchronous state if $(\theta, \omega)_{t_1}$ is inside the *basin of attraction* of that state (green area in Fig. 4.2). Otherwise, it will not regain synchronization with the grid but converge to a different solution of (4.5)–(4.6): a non-synchronous limit cycle characterized by

$$\omega_{\text{ns}}(t) \approx \frac{P}{\alpha} + \frac{\alpha K}{P} \cos\left(\frac{P}{\alpha}t\right) \quad (4.10)$$

(provided $|P|/\alpha^2 \gg 1$, $|P|^2/\alpha^2 \gg K$, as derived in Section 2.1.3).

Other, possibly serial, faults may push the generator from the synchronous state to perturbed states anywhere in state space. If, for instance, during the above event a faltering of locally connected renewable generation had driven P below zero, the generator's state would have deviated into the lower half of Fig. 4.2. We will therefore quantify how stable the synchronous state is against general large perturbations in terms of *basin stability* $S \in [0, 1]$, a measure of the basin's volume [Menck et al., 2013a]. More specifically, we will measure S as described in Sections 2.3.2 and 2.4.1, employing the uniform distribution $\rho = 1/|\mathcal{Q}|$ on $(\theta, \omega) \in \mathcal{Q} = [0, 2\pi] \times [-100, 100]$.

Intuition suggests that the synchronous state should become more stable when the transfer capacity K increases. This is indeed what we find [Menck et al., 2013b]: The expanding green area in Figs. 4.3a–b (compared to Fig. 4.2) and the characteristic in Fig. 4.3c show that basin stability S , starting from $S = 0$ for $K < |P|$, improves substantially as K goes up, until finally synchrony becomes the only stable state ($S = 1$).




Figure 4.3: **Basin stability of the generator in the infinite-grid model.** **a-b**, state space of the model, with $\alpha = 0.1$, $P = 1$ and **a**: $K = 50$, **b**: $K = 65$ (see Section 4.6.1). The solid black circle marks the desired synchronous state $(\theta_s, 0)$, and the solid red line shows the non-synchronous limit cycle attractor. The basin of attraction of $(\theta_s, 0)$ is coloured green, and that of the limit cycle is coloured white. **c**, basin stability S of the synchronous state versus the transfer capacity K . Figure published in [Menck et al., 2013b].

4.2 The finite-grid model

Surely, power grids are not infinite as just assumed. There will be some interplay between multiple nodes after one of them has been hit by a large perturbation. Whether the grid will return to its synchronous state depends on the affected node's properties, particularly its position in the grid topology. Hence we now turn to a finite version of the model that consists of N nodes [Machowski et al., 2008; Hill and Chen, 2006; Filatrella et al., 2008; Rohden et al., 2012; Witthaut and Timme, 2012]. Derived analogously to Eqs. (4.3)–(4.4), it reads

$$\dot{\theta}_i = \omega_i \quad (4.11)$$

$$M_i \omega_r \dot{\omega}_i = L_i - D_i \omega_i - \sum_{j=1}^N Y_{ij} V_i V_j \sin(\theta_i - \theta_j) \quad (4.12)$$

where θ_i , ω_i , and V_i are the *phase*, *angular frequency*, and *magnitude* of the voltage vector at generator i , measured in a frame of reference that co-rotates with the grid's rated frequency ω_r . Furthermore, M_i is the (cumulative) *moment of inertia* of the turbine(s) rotating at node i , L_i is the *net power input*, and D_i is the damping constant. Finally, \mathbf{Y} is the *admittance matrix* that represents high-voltage transmission lines, with $Y_{ij} = 1/X_{ij}$ if there is a line between nodes i and j and $Y_{ij} = 0$ otherwise.⁴

To obtain simplified model equations in parallel to (4.5)–(4.6), we divide Eq. (4.12) by $M_i \omega_r$ and define $\alpha_i := D_i/(M_i \omega_r)$, $P_i := L_i/(M_i \omega_r)$ and $K_{ij} := Y_{ij} V_i V_j/(M_i \omega_r)$. The result is formally equivalent to a second-order Kuramoto model [Kuramoto, 1975]

$$\dot{\theta}_i = \omega_i \quad (4.13)$$

$$\dot{\omega}_i = P_i - \alpha_i \omega_i - \sum_{j=1}^N K_{ij} \sin(\theta_i - \theta_j) \quad (4.14)$$

4.2.1 Zhukov's aggregation

As mentioned in the introduction, the nodes of a power grid can be modelled at any level of detail. Here, we are mostly interested in the influence the wiring topology exerts on the collective dynamics of a grid's nodes. Hence we keep the nodal dynamics simple and employ Eqs. (4.13)–(4.14) as a basic model of the high-voltage transmission line network, aggregating details at lower voltage levels (Fig. 4.4a). More specifically, we apply Zhukov's aggregation (see Machowski et al. [2008]), lump together all generators (with power supplies $s_l^i > 0$) and consumers (with power demands $d_m^i < 0$) connected by lower voltage lines to a high-voltage node i and replace them by a single *equivalent generator* with net power input $P_i = \sum_l s_l^i + \sum_m d_m^i$ (see Fig. 4.4b). Nodes with $P_i > 0$ are referred to as *net generators*. Nodes with $P_i < 0$ are *net consumers*.

⁴Therein, X_{ij} is the reactance of the transmission line between nodes i and j . In high-voltage grids, line resistance is small and usually neglected. The line model used is a valid approximation for short transmission lines [Machowski et al., 2008].




Figure 4.4: **Zhukov's aggregation.** **a**, schematic illustration of a power grid, consisting of high-voltage transmission grid and lower-voltage distribution grid plus generators and consumers. **b**, the grid of **a** after Zhukov's aggregation. All generators (with power supplies $s_l^i > 0$) and consumers (with power demands $d_m^i < 0$) connected to node i of the high-voltage grid are represented by a single equivalent generator with net power injection $P_i = \sum_l s_l^i + \sum_m d_m^i$. Nodes at which $P_i > 0$ (resp. $P_i < 0$) are called net generators (net consumers). The figure contains elements of free graphics files⁵.

4.2.2 Effects of large perturbations

How do finite grids respond to large perturbations hitting a single node? For a start, let us consider the simple 8-node network in Fig. 4.4b. Our main interest are the effects of the wiring topology on the grid's collective dynamics. So we simplify other details, setting $P_i = +P > 0$ for the 4 net generators, $P_i = -P$ for the 4 net consumers, $\alpha_i = \alpha$ for all 8 nodes, and $K_{ij} = K > 0$ if there is a line connecting nodes i and j and $K_{ij} = 0$ otherwise. Finally, we choose K large enough ($> 2P$) so that there is a synchronous state with $\omega_i = 0$ and $\dot{\omega}_i = 0$ for all nodes and constant phases θ_i^s . This state⁶ is shown in Fig. 4.5a, in which arrows indicate constant power flows that balance the specific P_i everywhere in the grid. The synchronous state is stable.

However, there are also several non-synchronous states in the grid's state space which are also stable. Some of them are shown in Figs. 4.5b-e. They can be found in numerical simulations of Eqs. (4.13)–(4.14) when the initial state vector $(\theta_i(0), \omega_i(0))_{i=1,\dots,N}$ is chosen as the synchronous state plus a random large perturbation $(\Delta\theta_i, \Delta\omega_i)$ that hits a particular node i , that is: when the initial state is chosen as

$$\begin{pmatrix} \theta_1(0) \\ \omega_1(0) \\ \vdots \\ \theta_i(0) \\ \omega_i(0) \\ \vdots \\ \theta_N(0) \\ \omega_N(0) \end{pmatrix} = \begin{pmatrix} \theta_1^s \\ 0 \\ \vdots \\ \theta_i^s \\ 0 \\ \vdots \\ \theta_N^s \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \Delta\theta_i \\ \Delta\omega_i \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \quad (4.15)$$

Before we discuss the depicted non-synchronous states, a technical remark is in order: It is always hard to judge when a simulation has reached an attractor. Here, it helps to note that, from Eq. (4.14), $\sum_i \dot{\omega}_i = -\alpha \sum_i \omega_i + \sum_i P_i$ implies the convergence of $\sum_i \omega_i$ to zero for balanced power grids (for which $\sum_i P_i = 0$). Hence when $|\sum_i \omega_i|$ is very small, one can conclude with some confidence that the transient phase is over.

Let us now consider the non-synchronous state shown in Fig. 4.5b. It is the result of a large perturbation that has hit node 1, the node that also emerges as the most strongly desynchronized. Its frequency ω_1 oscillates around a mean of about $P_1/\alpha = +P/\alpha$, which is surprisingly similar to the frequency oscillation in the non-synchronous state of a generator that is coupled to an infinite grid (cf. Eq. (4.10)). Analogously, in Fig. 4.5c, a large perturbation that hits node 4 gives rise to a pronounced desynchronization of node 4 itself, with ω_4 oscillating around $P_4/\alpha = -P/\alpha$. These observations seem to indicate that P_i/α is the *natural frequency* of oscillations of node i .

A different scenario is shown in Fig. 4.5d, where a large perturbation hits node 8. This time, in the end it is not node 8 that is heavily desynchronized, but node 5 is pushed to oscillate about its natural frequency P_5/α . Fig. 4.5e depicts the outcome of a perturbation that has hit node 3. In the resulting non-synchronous state, *two*

⁵See footnote 1.

⁶There may be more than one synchronous state for a given grid. For this grid, however, only one stable synchronous state is reached by the numerics.




Figure 4.5: **Example 8-node power grid.** **a**, synchronous state of the 8-node grid. The tip (thickness) of an arrow indicates the direction (strength) of constant power flow. **b-f**, time series of the ω_i in a non-synchronous state. The state was reached in a numerical simulation after node i had been hit by a large perturbation, where **b**: $i = 1$, **c**: $i = 4$, **d**: $i = 8$, **e**: $i = 3$, and **f**: $i = 8$ (again). Colour marks (white, yellow, orange, red, black) are used to classify the states (see text). Simulation parameters: $P = 1, \alpha = 0.1, K = 8$ (see Section 4.6.1).

nodes are heavily desynchronized: ω_3 oscillates around P_3/α and ω_5 oscillates around P_5/α . Finally, in Fig. 4.5f a large perturbation to node 8 triggers a non-synchronous state with *three* heavily desynchronized nodes: Node 8 itself, node 3, and node 5 jiggle around their respective natural frequencies.

The first two observed non-synchronous states, Fig. 4.5b-c, are qualitatively similar: A single node i is affected by a large perturbation, which induces this node's frequency ω_i to oscillate about the natural frequency P_i/α , whereas the frequencies of all other nodes remain comparatively close to zero. Let us call this a *direct response* and refer to it as the *yellow non-synchronous state* of node i . The other non-synchronous states involve *indirect responses*, in which not (only) the node initially hit by the perturbation suffers strong desynchronization, but also other nodes in the grid. Let us call Fig. 4.5d, in which predominantly node 5 is indirectly desynchronized, the *orange state*, Fig. 4.5e, in which predominantly nodes 3 and 5 are desynchronized, the *red state*, and Fig. 4.5f, in which predominantly nodes 3, 5 and 8 are desynchronized, the *black state*.

A careful survey of the perturbation space $(\Delta\theta_i, \Delta\omega_i) \in [0, 2\pi] \times [-30, 30]$ for all $i = 1, \dots, 8$ reveals that about 99.9 % of initial states of the form (4.15) result in the grid assuming either the (white) synchronous state or one of the (i -specific) yellow non-synchronous states, the orange state, the red state, or the black state. The individual perturbation space analyses for the eight nodes are shown in Figs. 4.6a-h. Nodes 1, 5, and 7 show very similar outcomes that look a bit like the phase portrait in Fig. 4.3a of a generator coupled to an infinite grid. Large perturbations to the lower half of the panel yield convergence to the synchronous state, whereas perturbations to the upper half lead to the respective yellow non-synchronous state with frequency oscillations about the natural frequency P/α . Nodes 4 and 6 show the inverted picture, as they have negative natural frequencies ($P_4 = P_6 = -P$). The outcome for node 2 looks similar, just that, here, from certain initial states in the upper half not the yellow state of node 2 but the orange state or the red state is reached.

Whereas, as we just observed, non-synchronous states involving an indirect response are an exceptional outcome for nodes 1, 2, 4, 5, 6, 7, such states are dominant in the perturbation spaces of nodes 3 and 8. Indeed, for these nodes the synchronous state is reached from significantly less than half of perturbation space, and the respective yellow state is not reached *at all*.

For node 3, the majority of initial states from the lower half trigger the orange state, meaning strong desynchronization not of the initially hit node 3, but of node 5, and the majority of initial states from the upper half trigger the red state, in which both nodes 3 and 5 are pushed to jiggle about their natural frequencies P_3/α and P_5/α .

For node 8, large perturbations to the upper half also typically induce the red state. The majority of large perturbations to the lower half lead to the black state, which involves strong desynchronization of three nodes: 3, 5 and 8.

These observations demonstrate that the response of a grid to a large perturbation hitting a single node strongly depends on this node's topological properties. We will now employ basin stability to define a measure of this response for each node.




Figure 4.6: **Perturbation space analysis.** Perturbation space $(\Delta\theta_i, \Delta\omega_i) \in [0, 2\pi] \times [-30, 30]$ for node i of the 8-node grid, where **a**: $i = 1$, **b**: $i = 2$, **c**: $i = 3$, **d**: $i = 4$. Colour indicates to which (class of) state (see Fig. 4.5) the grid converges when initiated according to (4.15). Simulation parameters: $P = 1, \alpha = 0.1, K = 8$ (see Section 4.6.1).




Figure 4.6: **Perturbation space analysis (continued).** Perturbation space $(\Delta\theta_i, \Delta\omega_i) \in [0, 2\pi] \times [-30, 30]$ for node i of the 8-node grid, where **e**: $i = 5$, **f**: $i = 6$, **g**: $i = 7$, **h**: $i = 8$. Colour indicates to which (class of) state (see Fig. 4.5) the grid converges when initiated according to (4.15). Simulation parameters: $P = 1, \alpha = 0.1, K = 8$ (see Section 4.6.1).

4.2.3 Definition of single-node basin stability

The white areas in Figs. 4.6a-h are actually snippets from the basin of attraction \mathcal{B} of the grid's synchronous state. More specifically, they correspond to the intersection of the set

$$\mathcal{B}_i := \{(\theta_i, \omega_i) | (\theta_j, \omega_j)_{j=1, \dots, N} \in \mathcal{B} \text{ with } \theta_j = \theta_j^s \text{ and } \omega_j = 0 \text{ for all } j \neq i\} \quad (4.16)$$

with a two-dimensional box \mathcal{Q} . As the synchronous state appears to be more stable against large perturbations that hit nodes with a larger white area, we define *the single-node basin stability* of node i as [Menck et al., 2013b]

$$S_i := \text{Vol}(\mathcal{B}_i \cap \mathcal{Q}) / \text{Vol}(\mathcal{Q}) \in [0, 1]. \quad (4.17)$$

This corresponds to the basin stability definition of Section 2.3.2 with ρ chosen as a uniform distribution on $\{(\theta_i, \omega_i) | \theta_j = \theta_j^s \text{ and } \omega_j = 0 \text{ for all } j \neq i\} \cap \mathcal{Q}$ and zero elsewhere. So we can use the numerical procedure described in Section 2.4.1 to estimate S_i , again choosing $T = 500$ so that the error of the estimate $e < 0.023$. The number S_i expresses the probability that the grid will return to its synchronous state after node i has been hit by a random large perturbation. We use

$$\mathcal{Q} := [0, 2\pi] \times [-100, 100] \quad (4.18)$$

here and in the following, as this choice allows to clearly distinguish the three important cases in which \mathcal{B}_i covers (i) significantly less than half of state space (Fig. 4.2, Figs. 4.6c,h); (ii) about half of state space (Fig. 4.3a, Figs. 4.6a,d-g); or (iii) all of state space (Fig. 4.3b).

In the 8-node example grid discussed above (Fig. 4.4b), the single-node basin stabilities read:

$$\begin{array}{cccc} S_1 = 0.54 & S_2 = 0.47 & S_3 = \mathbf{0.08} & S_4 = 0.56 \\ S_5 = 0.54 & S_6 = 0.53 & S_7 = 0.52 & S_8 = \mathbf{0.10} \end{array}$$

Apparently, the synchronous state is particularly vulnerable against perturbations hitting nodes 3 and 8, which is plausible in view of the perturbation space portraits in Figs. 4.6c,h. In the next sections, we will see how the topological properties of these nodes might explain this observation.

4.3 Power grid ensembles

In the 8-node example network, we just observed the well-known fact that power grids are multistable and possess numerous stable non-synchronous states [Machowski et al., 2008; Filatrella et al., 2008; Rohden et al., 2012; Witthaut and Timme, 2012; Chiang et al., 1995; Chiang, 2011]. A synchronous state, with frequencies $\omega_i = 0$ and $\dot{\omega}_i = 0$ at all nodes i and constant phases θ_i^s , can only exist if two conditions are met: First, total injected power must equal total consumed power, so that $\sum_i P_i = 0$. Second, the network matrix \mathbf{K} must be such that net power input P_i can be balanced through power transfers over the grid at each node i , which, for instance, necessitates $\sum_j K_{ij} > |P_i|$.

It can be shown that a synchronous state is locally stable (cf. Section 2.2) if the phase differences $|\theta_i^s - \theta_j^s| < \pi/2$ for all pairs of nodes (i, j) with $K_{ij} > 0$ [Machowski et al., 2008]. Indeed, for grid dynamics given by

$$\dot{\theta}_i = \omega_i \quad (4.19)$$

$$\dot{\omega}_i = P_i - \alpha_i \omega_i - \sum_{j=1}^N K_{ij} \sin(\theta_i - \theta_j) \quad (4.20)$$

this can quite easily be seen from the master stability function formalism (see Section 3.3.1) in the case of identical damping, i.e. $\alpha_i = \alpha$ at all nodes. A small-perturbation analysis yields

$$\begin{pmatrix} \delta \dot{\theta}_i \\ \delta \dot{\omega}_i \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 0 & -\alpha \end{pmatrix} \begin{pmatrix} \delta \theta_i \\ \delta \omega_i \end{pmatrix} - \sum_j K_{ij} \cos(\theta_i^s - \theta_j^s) \begin{pmatrix} 0 \\ \delta \theta_i - \delta \theta_j \end{pmatrix} \quad (4.21)$$

$$= \begin{pmatrix} 0 & 1 \\ 0 & -\alpha \end{pmatrix} \begin{pmatrix} \delta \theta_i \\ \delta \omega_i \end{pmatrix} - \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \sum_j L_{ij} \begin{pmatrix} \delta \theta_j \\ \delta \omega_j \end{pmatrix}. \quad (4.22)$$

With

$$L_{ij} := \delta_{ij} \sum_l K_{il} \cos(\theta_i^s - \theta_l^s) - K_{ij} \cos(\theta_i^s - \theta_j^s) \quad (4.23)$$

this is exactly of the form (3.11). If $|\theta_i^s - \theta_j^s| < \pi/2$ for all connected pairs (i, j) , one has $\cos(\theta_i^s - \theta_j^s) > 0$ so that the matrix \mathbf{L} is symmetric and positive semidefinite if \mathbf{K} is symmetric and positive semidefinite. Hence the synchronous state is stable if the master stability function, i.e. the maximum Lyapunov exponent (cf. Section 2.2), of

$$\delta \dot{\mathbf{w}} = \left[\begin{pmatrix} 0 & 1 \\ 0 & -\alpha \end{pmatrix} - \lambda \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \right] \delta \mathbf{w} = \begin{pmatrix} 0 & 1 \\ -\lambda & -\alpha \end{pmatrix} \delta \mathbf{w} \quad (4.24)$$

is negative for every non-zero eigenvalue λ of \mathbf{L} . However, the maximum Lyapunov exponent of this expression is given by (observe the similarity to Eq. (4.9))

$$\sigma_+ = \text{Re} \left\{ -\frac{\alpha}{2} + \frac{\sqrt{\alpha^2 - 4\lambda}}{2} \right\}. \quad (4.25)$$

As for a synchronous state with $|\theta_i^s - \theta_j^s| < \pi/2$ across all transmission lines any non-zero eigenvalue of \mathbf{L} is strictly positive (cf. Section 3.2.4), it follows that such a state is locally stable against small perturbations.

A power grid can, in principle, have multiple stable synchronous states. However, all the grids we will consider here have a single synchronous state that is *dominant* in the sense that, in numerical experiments, the vast majority of initial conditions in a neighbourhood of $\{(\theta_i, \omega_i) \mid \omega_i = 0 \text{ for all } i\}$ converge to it. For brevity, we will refer to the dominant synchronous state of a grid as *the synchronous state*.

Assuming a grid has a stable synchronous state, defined by frequencies $\omega_i = 0$ and constant phases θ_i^s , how stable is this state against large perturbations that locally afflict a single node? And how does this depend on the network topology? Before turning to a case study of the Northern European power grid in Section 4.4, we address these questions statistically by studying an ensemble of 1,000 randomly generated power grids with N nodes among which we place E transmission lines uniformly at random (no loops, no double edges). By choosing $N = 100$ and $E = 135$ we ensure that the resulting simple random networks (see Fig. 4.7) have the average degree $\langle d \rangle = 2.7$, a value typical of power grids (cf. Sun [2005] and Section 3.2). As above, we impose the dynamics determined by Eqs. (4.19)–(4.20). To have a clear view on the effects of the topology, we simplify other details, assuming all generator properties and all transmission capacities to be equal, that is, $\alpha_i = \alpha$ for all i , and $K_{ij} = K_{ji} = K$ if nodes i and j are connected and $K_{ij} = 0$ otherwise. Furthermore, we select a load scenario with only two types of nodes, randomly choosing $N/2$ net generators with $P_i = +P$ and $N/2$ net consumers with $P_i = -P$.




Figure 4.7: **Example grid from the ensemble.** The grid has $N = 100$ nodes and $E = 135$ transmission lines. Squares (circles) represent $N/2$ net consumers with $P_i = -P$ (net generators with $P_i = +P$). The colour scale indicates a node's basin stability S_i . Figure published in [Menck and Kurths, 2012].

4.3.1 Basin stability statistics

For each ensemble grid, we measure for each node i its *single-node basin stability* $S_i \in [0, 1]$, a number that reflects the volume of the synchronous state's basin in a 2-dimensional subspace pertaining to the node's two state variables (see Section 4.2.3). S_i can be interpreted as the likelihood that the grid will return to its synchronous state after node i has been hit by a random large perturbation. An example grid from the ensemble and its single-node basin stabilities are shown in Fig. 4.7.

As the ensemble contains 1,000 grids with $N = 100$ nodes each, we obtain 100,000 individual measurements of single-node basin stability S_i . Fig. 4.8 depicts the histogram of all these S_i . Evidently, the majority of nodes show a *fair* value of basin stability. So what is special about the nodes with *poor* stability ($S_i < 0.30$) or *high* stability ($S_i > 0.95$)?




Figure 4.8: **Basin stability in the power grid ensemble.** Shown is the histogram of single-node basin stability S_i for all nodes in all ensemble networks. Dotted lines delimit nodes with poor stability ($S_i < 0.30$) and high stability ($S_i > 0.95$). Simulation parameters: $N = 100$, $E = 135$, $\alpha = 0.1$, $P = 1$, and $K = 8$ (see Section 4.6.1). Figure published in [Menck et al., 2013b].

In the infinite-grid model, we observed that the synchronous state's stability improves substantially when K increases (Fig. 4.3c). Hence, in finite grids, we might expect the synchronous state to be very stable against perturbations that hit nodes with large degree (cf. Section 3.2), as such a node is effectively coupled to the grid through the transfer capacity dK . To check this, we compute the average basin stability $\langle S \rangle$ of all nodes in the whole ensemble that have degree d . The resulting characteristic is shown in Fig. 4.9a. It is rather flat and has a large standard deviation. Hence, against

the initial guess, basin stability $\langle S \rangle$ does not increase with d . However, degree does have a second-order importance: If the neighbours of a node i have a large average degree

$$d_{\text{av}} = \frac{1}{d_i} \sum_{\substack{j \\ K_{ij} > 0}} d_j,$$

then the node's expected basin stability $\langle S \rangle$ is large (provided $d_i \geq 2$, Figs. 4.9b-c). Indeed, nodes with degree $d = 3$ whose neighbours have $d_{\text{av}} = 17/3 = 5.6$ possess $\langle S \rangle = 0.93$. An example node with these properties is labelled **A** in Fig. 4.7. It has $S_{\mathbf{A}} = 1.0$.

A major clue for understanding these observations, and the influence of topology on stability in general, comes from the betweenness characteristic that shows how the average basin stability $\langle S \rangle$ of nodes in the ensemble depends on shortest-path betweenness b (cf. Section 3.2). Although the curve (Fig. 4.10a) is insignificant for most values of b , it reveals some pronounced downward peaks that, according to the explanatory sketch (Fig. 4.10b), show that the synchronous state is particularly *unstable* against perturbations hitting nodes that are located inside *dead ends* or, more generally, *dead trees*. As illustrated in Fig. 4.10c, a *dead tree* (also just called a *tree*) is defined as a connected set of nodes none of which is contained in a closed loop of edges. A *dead end* is a dead tree whose nodes' maximum degree is at most 2. Every dead tree contains at least one dead end. In Fig. 4.7, the nodes marked **B**₁ and **B**₂ are both located in a dead end. They have $b_{\mathbf{B}_1} = N - 2$ and $b_{\mathbf{B}_2} = 2N - 6$ and both possess poor values of basin stability.

What is so special about dead trees? A detailed analysis of the grid dynamics (see next section) reveals that a large single-node perturbation can typically only desynchronize the node initially hit (cf. the *yellow* states of the 8-node grid discussed in Section 4.2.2). But when a dead end is near, perturbations tend to *creep into it*, rattling and desynchronizing the nodes it contains.

Therefore dead trees, as they incorporate dead ends, should drastically lower the basin stability of nodes adjacent to them. We check this speculation as above by computing the average basin stability $\langle S \rangle$ of all nodes in the ensemble that are or are not adjacent to a dead tree. The resulting curve (Fig. 4.10d) shows that nodes adjacent to dead trees are indeed far less stable than non-adjacent nodes. Furthermore, non-adjacent nodes do display the increasing dependence of $\langle S \rangle$ on degree d we had hypothesized earlier.

It is now clear why we could not already observe this dependence in Fig. 4.9a: Although a large degree d topologically benefits the stability of a node, it also increases the likelihood that this node connects to a stability-adverse dead tree. Conversely, as dead trees in our ensemble grids often consist of a single node with degree one, a randomly picked node whose neighbours have a large average degree is unlikely to connect to a dead tree. This explains the increase in $\langle S \rangle$ in Figs. 4.9b-c [Menck et al., 2013b].




Figure 4.9: **Relation between basin stability and topological properties of a node.** **a**, ensemble average basin stability $\langle S \rangle$ of nodes with degree d . **b-c**, $\langle S \rangle$ of nodes of degree d whose neighbours have the average degree d_{av} (**b**: $d = 1, 2$, **c**: $d = 3, 4$). In all panels, red and green shades indicate ± 1 standard deviation. Simulation parameters: $N = 100$, $E = 135$, $\alpha = 0.1$, $P = 1$, and $K = 8$ (see Section 4.6.1). Figure published in [Menck et al., 2013b].




Figure 4.10: Effects of dead ends and dead trees. **a**, ensemble average basin stability $\langle S \rangle$ of nodes with shortest-path betweenness b . **b**, shown are typical examples of nodes (coloured red) that have certain distinct values of betweenness b . They are located inside dead ends or dead trees (see also panel **c**). The dashed boxes indicate the respective remainder of the grid that contains the printed-on number of nodes. **c**, shown is a snippet from the Northern European power grid (see Fig. 4.12). Squares (circles) depict net consumers with $P_i = -P$ (net generators with $P_i = +P$). The colour scale indicates how large a node's basin stability S_i is. The set of nodes $\{1, 2, 3, 4\}$ make up a dead tree that includes two dead ends, $\{1\}$ and $\{3, 4\}$. Nodes 2 and 3 have the distinct betweenness values $b_2 = 3N - 10$ and $b_3 = N - 2$. As expected from panels **a-b**, they possess a poor basin stability. **d**, ensemble average basin stability $\langle S \rangle$ of nodes of degree d that are adjacent or non-adjacent to dead trees. Nodes inside dead trees are not included in the statistics. In panels **a** and **d**, red and green shades indicate ± 1 standard deviation. Simulation parameters: $N = 100$, $E = 135$, $\alpha = 0.1$, $P = 1$, and $K = 8$ (see Section 4.6.1). Figure published in [Menck et al., 2013b].

4.3.2 Effects of dead trees on grid dynamics

In the finite-grid model

$$\begin{aligned}\dot{\theta}_i &= \omega_i \\ \dot{\omega}_i &= -\alpha\omega_i + P_i - \sum_{j=1}^N K_{ij} \sin(\theta_i - \theta_j)\end{aligned}$$

there exist a multitude of non-synchronous states for the power grids in our ensemble. As $(\sum_i \dot{\omega}_i) = -\alpha(\sum_i \omega_i)$, all of them satisfy

$$\sum_i \omega_i = 0. \quad (4.26)$$

However, a typical non-synchronous state reached after a large perturbation has hit a node with *fair* basin stability is very different from a typical non-synchronous state reached after a large perturbation has hit a node with *poor* basin stability (cf. histogram in Fig. 4.8). It turns out that this difference has to do with the presence of dead ends. Indeed, it is their proximity that makes poor-stability nodes so unstable.

To see this, we statistically analyze the non-synchronous states that appeared in our ensemble studies. Recall that, to estimate single-node basin stability for a certain node in a grid, we randomly drew $T = 500$ initial states and integrated the model equations (cf. Sections 2.4.1 and 4.2.3). From each of these initial states, the grid either converged to synchrony or to a non-synchronous state. The latter outcome emerged quite often, and indeed, performing the basin stability estimation procedure for all nodes in all ensemble grids (totalling 50,000,000 simulation runs) yielded a huge number of different non-synchronous states. Yet strikingly, all of them are instances of just *eight different representative scenarios* [Menck et al., 2013b].

Classification according to these scenarios relies on the important observation that, in each of the non-synchronous states, the nodes divide themselves into two distinct groups: The vast majority remain *almost synchronous*, each having its frequency ω_i oscillating close to zero. However, a few nodes become *strongly desynchronized*, with ω_i oscillating close to P_i/α (cf. Eq. (4.10)). As observed in the 8-node example in Section 4.2.2, P_i/α seems to be the *natural frequency* of strongly desynchronized nodes.

Now let us look at scenario A in Figure 4.11, which shows four nodes that are part of a larger grid. Two of the nodes are generators ($P_i = +P > 0$, marked by circles) and two are consumers ($P_i = -P$, marked by squares). The time series next to the nodes depict their frequencies in a non-synchronous state that the grid has been pushed into by a large perturbation hitting the dark grey node. Evidently, the dark grey node itself is strongly desynchronized, having ω_i oscillating around P_i/α . In contrast, the three other nodes – and actually all unshown nodes of the grid – remain almost synchronous (with ω_i wiggling around zero). This *direct response*, in which only the node primarily hit suffers strong desynchronization, is what we take scenario A to represent. This scenario is very important for fair-stability nodes. Indeed, when a random large perturbation hitting a fair-stability node pushes a grid into a non-synchronous state, this state is of type A in 86.8% of the cases. For poor-stability nodes, the corresponding percentage is only 0.1%.

Now let us turn to scenario B1. Again, the corresponding panel in Figure 4.11 shows a non-synchronous state that has been triggered by a large perturbation hitting the dark grey node. As in scenario A, all but one of the nodes remain almost synchronous. However, this time there is an *indirect response*: Not the primarily hit (dark grey) node, but another *secondary* node becomes strongly desynchronized. We take scenario B1 to represent states in which the strongly desynchronized secondary node is of the same type as the node primarily hit (both are consumers in the picture) and forms a dead end directly neighbouring that node. The percentage of desynchronizing large perturbations leading to a B1-type non-synchronous state is 13.8% for poor-stability nodes and 9.3% for fair-stability nodes.

Scenario B2 is very similar to scenario B1: Again, it is not the primarily hit (dark grey) node, but a secondary node forming a directly adjacent dead end that suffers strong desynchronization. However, this time the primary node is of the opposite type to the secondary node (consumer and generator, in the picture). Scenario B2 is very important for poor-stability nodes: 24.6% of desynchronizing large perturbations hitting such nodes lead to this kind of non-synchronous state. The respective percentage for fair-stability nodes is just 1.9%.

Scenario B3 depicts another indirect type of response, in which the perturbation entails strong desynchronization of a single secondary node that is not a direct neighbour of the (dark grey) primary node but terminates a multi-node dead end adjacent to it. As scenario B2, this scenario is much more important for poor-stability nodes (8.5%) than for fair-stability nodes (0.3%).

Whereas the indirect scenarios B1-B3 all have one strongly desynchronized node, the indirect scenarios C1-C3 involve *two strongly desynchronized nodes*. Scenario C1 represents non-synchronous states in which the primary node *and* a neighbouring dead-end node of *the same type* are strongly desynchronized. Scenario C2, the most important scenario for poor-stability nodes, represents non-synchronous states in which the primary node and a neighbouring dead-end node of *unlike type* are strongly desynchronized. The respective percentages for this scenario are 31.0% for poor-stability nodes and 0.5% for fair-stability nodes. Scenario C3 shows a typical non-synchronous state in which the primary node is not among the two strongly desynchronized nodes. These are located in an adjacent dead end.

Finally, scenario D represents all non-synchronous states in which *more than two nodes are strongly desynchronized*. It turns out that these nodes are always located in dead trees. The respective percentages are 9.4% for poor-stability nodes and 0.5% for fair-stability nodes.

To sum up, 86.8% of large perturbations that hit *fair*-stability nodes and lead to a non-synchronous state can only strongly desynchronize the node primarily affected (scenario A). In contrast, 99.9% of large perturbations that hit *poor*-stability nodes and lead to a non-synchronous state involve the strong desynchronization of nodes in a nearby dead end (scenarios B1-3, C1-3, D).

This analysis suggests that poor-stability nodes are so unstable because of the dead trees adjacent to them. Indeed, the statistics confirm that nodes adjacent to dead trees are much less stable than non-adjacent nodes (Fig. 4.10d).

Note that these findings are consistent with what we observed for the 8-node grid

studied in Sections 4.2.2, 4.2.3. There, we measured *fair* stability values for nodes 1, 2, 4, 5, 6, and 7. None of these six nodes is located inside of or adjacent to a dead tree. For all of them, the most relevant non-synchronous state turned out to be the *yellow* state (cf. Figs. 4.6a,b,d-g), which corresponds to scenario A in Fig. 4.11.

In contrast, we measured *poor* stability values for nodes 3 and 8, which are both either located inside of or adjacent to a *dead end*. For node 3, the most relevant non-synchronous states are the *orange* state, corresponding to scenario B3, and the *red* state, corresponding to scenario C1 (cf. Fig. 4.6c). For node 8, apart from the orange and the red state, the *black* state is important (cf. Fig. 4.6h), which corresponds to scenario D in Fig. 4.11. The yellow state, or scenario A, has no relevance for these two nodes.




Figure 4.11: **Statistics on non-synchronous states in the power grid ensemble.** Scenarios A, B1-B3 sketch different typical non-synchronous states from the ensemble simulations. In each scenario, the dark grey node suffers a large perturbation. In the resulting non-synchronous state, nodes marked with a +, - or 0 have their frequency ω wiggling around $+P/\alpha$, $-P/\alpha$, or 0, respectively, as indicated by the depicted time series. The red (grey) percentage under each scenario gives the share of desynchronizing large perturbations that hit nodes with poor (fair) stability and produce a non-synchronous state similar to the one sketched. Figure published in [Menck et al., 2013b].




Figure 4.11: **Statistics on non-synchronous states in the power grid ensemble (continued).** Scenarios C1-C3 and D sketch different typical non-synchronous states from the ensemble simulations. In each scenario, the dark grey node suffers a large perturbation. In the resulting non-synchronous state, nodes marked with a +, - or 0 have their frequency ω wiggling around $+P/\alpha$, $-P/\alpha$, or 0, respectively, as indicated by the depicted time series. The red (grey) percentage under each scenario gives the share of desynchronizing large perturbations that hit nodes with poor (fair) stability and produce a non-synchronous state similar to the one sketched. Figure published in [Menck et al., 2013b].

4.4 Case study: power grid of Northern Europe

In the preceding sections we observed that dead ends and, more generally, dead trees have devastating effects on the synchronous state's stability against perturbations hitting nearby nodes. These results emerged from simulations in networks that were randomly generated and hence had somewhat artificial topologies. Indeed, the only network property that we tuned to be realistic was the average degree. To see whether the results carry over to real-world topologies, we now perform a case study of the Northern European power grid [Menck et al., 2013b].

The transmission part of this grid is shown in Fig. 4.12. As before, we investigate it using the model equations

$$\begin{aligned}\dot{\theta}_i &= \omega_i \\ \dot{\omega}_i &= -\alpha_i \omega_i + P_i - \sum_{j=1}^N K_{ij} \sin(\theta_i - \theta_j)\end{aligned}$$

and, as our focus are the effects of the topology, simplify other details, setting $\alpha_i = \alpha$ for all i and $K_{ij} = K_{ji} = K$ if nodes i and j are connected and $K_{ij} = 0$ otherwise. Furthermore, we impose a random load scenario with only two types of nodes, randomly choosing $N/2$ net generators with $P_i = +P$ and $N/2$ net consumers with $P_i = -P$. The load scenario shown in Fig. 4.12 is a typical outcome of this procedure.

The colour-scale in Fig. 4.12 indicates how large a node's basin stability is. Apparently, dead trees have the same strongly weakening effect as in the artificial ensemble grids. Indeed, all poor-stability nodes are either located inside of or adjacent to a dead tree. To check whether 'healing' of these appendices benefits stability, we add transmission lines in a parsimonious way.

To be specific, for each dead tree we identify the node in the grid that has the minimum Euclidean distance to any node inside the tree (but is not itself part of or adjacent to the tree). Add a transmission line between this node and the tree node it is so close to. Repetition of this step until all dead trees have been 'healed' yields 27 extra lines (marked red in Fig. 4.13).

The insets in Fig. 4.13 demonstrate that line addition substantially enhances the synchronous state's stability. Indeed, all nodes that previously showed poor basin stability have improved to at least fair stability.




Figure 4.12: **Northern European power grid.** The grid has $N = 236$ nodes and $E = 320$ transmission lines. The load scenario was chosen randomly, with squares (circles) depicting $N/2$ net consumers with $P_i = -P$ (net generators with $P_i = +P$). The colour scale indicates how large a node's basin stability S_i is. The simulation parameters, $\alpha = 0.1$, $P = 1$, and $K = 8$, imply the simplifying assumptions that all generators in the grid are of the same making and that all transmission lines are of the same voltage and impedance. These assumptions enable us to focus on the effects of the wiring topology (see Section 4.6.1). Figure published in [Menck et al., 2013b].




Figure 4.13: **‘Healed’ Northern European power grid.** Shown is the same grid as in Fig. 4.12 with $N = 236$ nodes, to whose originally $E = 320$ transmission lines further 27 lines have been added to ‘heal’ dead trees (see text). New lines are coloured red. As before, squares (circles) depict $N/2$ net consumers with $P_i = -P$ (net generators with $P_i = +P$). The colour scale indicates how large a node’s basin stability S_i is (computed after addition of new transmission lines). Insets I-III show the basin stability values before line addition. Simulation parameters: $\alpha = 0.1$, $P = 1$, and $K = 8$ (see Section 4.6.1). Figure published in [Menck et al., 2013b].

4.5 Conclusions and outlook

We have investigated the stability of power grids by means of a component-wise version of basin stability, assigning to each node of a grid a number called *single-node basin stability* that measures how stable the grid's synchronous state is against large perturbations hitting that node. Of the many functional aspects that presumably matter for grid stability, we have focussed on the influence the network topology has. Our statistical analysis of an ensemble of artificially generated power grids produced a clear message: The widespread and cheapest of all connection schemes, namely dead ends and dead trees, do serious harm to grid stability. A case study of the Northern European power grid confirmed this result and demonstrated that the inverse statement is also true: 'Healing' dead ends by means of some extra transmission lines substantially enhances stability.

These findings imply that, in today's power grids, special control equipment should better be at hand in the proximity of dead ends. In the design of future power grids, dead ends should be avoided altogether.

The detrimental effect of dead ends provides an alternative explanation for an intricate recent observation. Rohden et al. [2012] studied how a power grid's stability against a particular large perturbation changes as the grid's supply structure becomes increasingly decentralized. They found that the more decentralized a grid is, the smaller is its stability. In the process of decentralization they applied, large power plants were replaced by ten small plants each. Crucially, these new small plants were connected to the grid *as dead ends*. Our findings imply that decentralization schemes avoiding dead ends could be less harmful to power grid stability. Future work should investigate that.

Several more research questions arise from our results. In particular:

(i) We assumed unweighted networks, in which the transfer capacities of all transmission lines are equal. Investigating the effects of link weights on grid stability could be a worthwhile project.

(ii) Similarly, we placed only two types of nodes in our power grids: net generators and net consumers. Extending our study to multiple kinds of nodes would be important.

(iii) We studied grid stability against large perturbations that afflict a single node. This could straightforwardly be extended to multi-node perturbations.

(iv) We only investigated uniform distributions of large perturbations. Realistic distributions of perturbations might be hard to obtain, but would certainly benefit the understanding of power grid stability.

(v) Basin stability only captures the volumes of basins. However, as illustrated in Fig. 4.6, there may also be other stability-related characteristics of basins such as their possible fractality that might deserve research effort.

Finally, the single-node version of basin stability used here may be a helpful tool for investigations into other multistable complex systems, including ecosystems [May, 1977; Gross et al., 2009] and gene regulatory networks [Huang and Ingber, 2007].

4.6 Supporting material

4.6.1 Model parameters

Recall the more explicit version of the model equations, which reads

$$\begin{aligned}\dot{\theta}_i &= \omega_i \\ M_i \omega_r \dot{\omega}_i &= L_i - D_i \omega_i - \sum_{j=1}^N Y_{ij} V_i V_j \sin(\theta_i - \theta_j)\end{aligned}$$

where θ_i , ω_i , and V_i are the *phase*, *angular frequency*, and *magnitude* of the voltage vector at generator i , measured in a frame of reference that co-rotates with the grid's rated frequency ω_r . Furthermore, M_i is the (cumulative) *moment of inertia* of the turbine(s) rotating at node i , L_i is the *net power input*, and D_i is the damping constant. Finally, \mathbf{Y} is the *admittance matrix* that represents high-voltage transmission lines, with $Y_{ij} = 1/X_{ij}$ if there is a line between nodes i and j and $Y_{ij} = 0$ otherwise. The model incorporates frequency dynamics based on the balance of active power, but neglects voltage dynamics on the assumption of perfect reactive power control. Negligence of the interaction between the two parts of the power balance is called the *decoupling assumption* in power grid engineering.

In both the ensemble study (cf. Section 4.3) and the case study of the Northern European power grid (cf. Section 4.4), we choose the model parameters as follows. As we seek to focus on the effects of the (unweighted) topology, we assume that: (i) All generators are of the same making, so that $M_i = M$ and $D_i = D$ for all i ; (ii) the voltage level is the same everywhere, so that $V_i = V$ for all i ; (iii) all transmission lines have the same reactance, so that $X_{ij} = X$ if there is a line between i and j . Furthermore, we impose a load scenario in which half of the nodes are randomly selected to be net generators with $L_i = L > 0$ and the other half are net consumers with $L_i = -L$. This choice satisfies the synchronization condition of total generated power being equal to total consumed power, $\sum_i L_i = 0$.

We specify a load scenario by setting $L = 200$ MW and choose as the transmission capacity $V^2/X = 1,600$ MW, which corresponds [Spring, 2003] to a 400 km long line at a voltage of 380 kV. Furthermore, we set $M = 40 \cdot 10^3$ kg m², which amounts to assuming the inertia of a 400 MW power plant at each node [Horowitz and Phadke, 2008]. Note that the average installed generation capacity per node in the Northern European high-voltage transmission grid is indeed 400 MW (amounting to a total of 96 GW) [ENTSO-E, 2011].

For the simplified model

$$\begin{aligned}\dot{\theta}_i &= \omega_i \\ \dot{\omega}_i &= -\alpha \omega_i + P_i - \sum_{j=1}^N K_{ij} \sin(\theta_i - \theta_j)\end{aligned}$$

together with $\omega_r = 2\pi \cdot 50$ Hz $\approx 314.59 \cdot 1/\text{s}$ these settings give $P = L/(M\omega_r) \approx 16 \cdot 1/\text{s}^2$ and $K = (V^2/X)/(M\omega_r) \approx 128 \cdot 1/\text{s}^2$. $1/\alpha$ is the decay time of elec-

tromechanical transients induced by large perturbations and is typically of the order [Machowski et al., 2008] 1 - 10 s. Here we choose $\alpha = 0.4 \cdot 1/\text{s}$ so that $1/\alpha = 2.5 \text{ s}$. Finally, we measure time in units of 0.25 s, so that in our simulations $P = 1$, $K = 8$ and $\alpha = 0.1$. Note that the Northern Grid configuration shown in Fig. 4.12 can cope with up to $P = 1.86$ at $K = 8$, so that our setting $P = 1$ can be considered a modest load scenario.

4.6.2 Parameter sensitivity

In general, one observes that single-node basin stability increases in the simulated power grids when the damping constant α or the transmission capacity K increases. However, our main finding is retained: Nodes adjacent to (or inside of) dead trees typically reveal stability values that are significantly below the average stability of non-adjacent nodes. This is illustrated in Fig. 4.14 which displays another basin stability study of the Northern European transmission grid that is based on the parameters $\alpha = 0.2$ (doubled compared to before), $P = 1.6$ and $K = 8$.

4.6.3 Numerical methods

For all the numerical simulations of power grids, we employed a fourth-order Runge-Kutta algorithm, integrating for 60.000 steps with step size $\tau = 0.01$. If very long transients were detected, we increased the number of steps up to 960.000.




Figure 4.14: **Supplement: Northern European power grid with increased damping.** The grid has $N = 236$ nodes and $E = 320$ transmission lines. The load scenario is, topologically, the same as in Fig. 4.12, with squares (circles) depicting $N/2$ net consumers with $P_i = -P$ (net generators with $P_i = +P$). The colour scale indicates how large a node's basin stability S_i is. Here, we use different simulation parameters than for Fig. 4.12: $\alpha = 0.2$ (doubled), $P = 1.6$, and $K = 8$. Whereas in general grid stability increases compared to Fig. 4.12 (note the many white nodes), nodes adjacent to (or inside of) tree-like appendices still typically have significantly lower S_i than non-adjacent nodes. Figure published in [Menck et al., 2013b].

5 Conclusions and outlook

All the investigations reported in this thesis were motivated by the question how the stability of a power grid depends on the network topology. The conceptual equipment I have employed to address this question mainly consisted of complex network theory and its tools to characterize the topology of networked systems, and of basin stability, a nonlocal stability concept based on the basin's volume that allows to quantify how stable an attractor is against non-small perturbations.

It may not be obvious that there should be a significant relationship between network properties, which derive from a finite set of numbers stored in the adjacency matrix, and basin stability, a function of a high-dimensional geometrical object with positive volume. Yet some interesting results have turned up. Let me briefly summarize the main contributions of this thesis.

Summary

Although volumes of basins of attraction have been investigated before, including in the literature on complex networks [Wiley et al., 2006], I am not aware of any study that uses the stability interpretation I have employed here. The observation underlying this interpretation is that the volume of a basin can be understood as an expression of the likelihood that a system returns to a given attractor after having been hit by a random – possibly non-small – perturbation. This observation may appear rather straightforward, and it would thus be surprising if the stability interpretation had really not been used so far. But in this case, and although it can certainly be refined in many ways, basin stability could be considered a new stability concept.

In any case, this thesis contains several conceptual advances regarding the applicability of basin stability. First of all, I have set it into the context of dynamical systems theory, thereby endowing it with a firm theoretical basis that allows to explore its advantages and disadvantages. Basin stability can be used to quantify by a number $S_B \in [0, 1]$ the stability of a (weak) attractor in the middle ground between instability ($S_B = 0$) and global stability ($S_B = 1$). An important result I obtained is that the numerical costs for estimating basin stability are relatively low and, importantly, do not depend on the system's dimension. In addition, I found evidence that the estimation procedure works properly even when fractal basin boundaries are present. Furthermore, there is also a component-wise version of basin stability that may be particularly helpful for investigations into networked systems. These advantages make basin stability a widely applicable tool.

In any particular application, an issue that will have to be addressed is that basin stability requires the choice of a probability density of perturbations in state space. Whereas any choice can help to deepen the understanding of certain aspects of a system's behaviour, the freedom of choice may give rise to unwelcome ambiguities.

A possible drawback of basin stability is that it may be very hard to infer any information on it from observational data. In particular, I observed that there is no clear relationship with the small-perturbation convergence rate, a local-stability indicator popular in complex systems science [Scheffer et al., 2009].

I have applied basin stability to scrutinize the influence of network topology on power grid stability once in an abstract and once in a more applied way. In the abstract approach, I have contemplated topological conditions for stable synchronization on general complex networks and observed that optimization for local stability and optimization for nonlocal stability can require completely different properties from a dynamical system. This may seem natural from a mathematical point of view. Yet one might have hoped that the often easy-to-use methods of local stability could serve as dependable proxies for the nonlocal variety, too – which they do not.

More specifically, I have approached the long-standing puzzle of networks science that synchronizing networks from the real world such as power grids and neural networks are so different from what previous theory predicts. To be optimally synchronizable, these networks should possess random topologies according to the local-stability-based findings of Hong et al. [2004]. But they are actually structured in a small-world fashion that, from the Watts-Strogatz model’s perspective [Watts and Strogatz, 1998], is comparatively regular. I found in a study of networks of chaotic Rössler oscillators that basin stability improves markedly the more regular a topology is. I therefore conjecture a solution to the mentioned puzzle: During the evolution of synchronizing networks, the optimization for local stability and the simultaneous optimization for nonlocal stability have acted as opposing forces. Their struggle ended in a topological compromise: small-worldness.

In the more applied approach, I have finally analyzed a power grid model from the engineering literature and investigated the influence that the network-related properties of single components may have on the stability of the whole grid. For this purpose, I defined a single-node version of basin stability and employed it to a large ensemble of artificially generated power grids. The statistical analysis of the ensemble unveiled several clear relationships between nodal network properties and grid stability that all support one main finding: The widespread and cheapest of all connection schemes, so-called dead ends and dead trees, strongly reduce stability. I found this result confirmed in a case study of the Northern European power system, and established that the inverse is also true: ‘Healing’ of dead ends by addition of transmission lines substantially increases stability. This suggests a fundamental design principle for future power grids: Add just a few more lines to avoid dead ends.

As a conclusion, I believe that I have managed to collect some parts of an answer to the overarching question of this thesis how the stability of a power grid is influenced by the network topology.

Research outlook

Several new questions derive from the results presented here that could serve as starting points for future research. Regarding the development of novel conceptual approaches, it is an open problem how reliable information on nonlocal stability can be gathered from observational data. As discussed above, the small-perturbation convergence rate – suggested as a warning-signal for the loss of local stability [Scheffer et al., 2009] – may fail to signal anything when a basin of attraction contracts. It could be an interesting undertaking to try to spot traces of changing basin stability in time series.

Another important issue mentioned above is that basin stability depends on the choice of a probability density of perturbations. Whereas any choice may lead to intriguing insights into a system’s dynamics, one might attempt to find methods for choosing the probability density in a realistic way – albeit it is highly unclear what ‘realistic’ would mean.

Additional questions arise from my investigations into synchronization on complex networks, in which I studied only one type of systems of chaotic oscillators coupled through undirected, unweighted links. Against the backdrop of this rather narrow focus, it may appear rather bold that I have conjectured a solution to the long-standing puzzle of networks science about the small-worldness of synchronizable real-world networks. In the best sense of Popper [1963], future research could try to refute my conjecture, for instance by comparing the topological requirements of local stability and nonlocal stability under other circumstances. Particularly, it could be investigated whether usage of a proper model of neural networks or power grids would yield results that agree or disagree with my findings.

Also regarding the concrete application of basin stability to power grids my work has produced several open problems. I used a model on which I imposed several simplifications in order to have a maximally clear view on the effects of the network topology on grid stability. Specifically, I allowed only two types of nodes, consumers and generators, and completely neglected any diversity inside these groups. Furthermore, I assumed all transmission lines to have identical properties, thus disregarding the different voltage levels, conductor types and line lengths characteristic of real-world grids. It could be a worthwhile project to study how the stability of a power grid is actually affected by such inhomogeneities.

Apart from that, it might be exciting to extend the single-node version of basin stability used here to a multi-node version. This extension should be straightforward.

Finally, basin stability now stands ready as an aide for model-based investigations into all kinds of multistable systems, such as the human brain [Babloyantz and Destexhe, 1986; Lytton, 2008], coupled lasers [Erzgräber et al., 2005], multi-species ecosystems [May, 1977; Gross et al., 2009], gene regulatory networks [Huang and Ingber, 2007; Cookson et al., 2009], and climatic tipping elements [Rahmstorf, 1995; Da Silveira Lobo Sternberg, 2001; Lenton et al., 2008].

Acknowledgements

I want to thank Jürgen Kurths, supervisor of this thesis and my *Doktorvater*. I am highly grateful for the motivating topic he proposed, for innumerable discussions and stimulating thoughts, for the organizational framework he provided, and for the welcoming atmosphere surrounding him that made me feel comfortable in his group from day one.

I also thank Jobst Heitzig for the many hours of fruitful exchange, his admirable openness to discussions on basically any subject, his crisp way of thinking, and the splendid time we had at Pappelallee.

Furthermore, I feel grateful towards many other people I had the chance to meet and discuss with at Potsdam-Institute for Climate Impact Research, in particular: Jonathan Donges, Reik Donner, Naoya Fujiwara, Carsten Grabow, Ji Peng, Norbert Marwan, Kira Rehfeld, Hans Joachim Schellnhuber, and Paul Schultz. And I thank the proofreaders of this text for their comments and patience: Jonathan Donges, Reik Donner, Carsten Grabow, Jobst Heitzig, Kira Rehfeld, and Paul Schultz.

I am also grateful to the Konrad-Adenauer-Stiftung and its employees for the material and non-material support before and during the course of this thesis. I will never forget the wonderful seminars I was allowed to participate in, particularly the one on ‘Learning from the Past: Remembrance and Reconciliation in Europe’ that took place in Sarajevo in 2012.

Finally, I want to thank Annbritt that she has endured the ups and downs of my mood, particularly prior to conference talks and expected replies from editors. She always encouraged me to engage in my project with passion.

Appendix

Computational resources

Many of the results presented in this thesis are based on numerical simulations. The programmes used for these simulations were developed by the author, written in the language C++, and were run on the cluster computer at Potsdam-Institute for Climate Impact Research.

Analysis tools employed included the **igraph** package [Csardi and Nepusz, 2006] and several scripts in the languages R, Perl or Bash developed by the author. The Figures were generated and edited using **gnuplot** and **Inkscape**. Data on real-world networks was kindly provided by the research group of Prof. Dr. Dr. h.c. mult. Jürgen Kurths.

Derivation of the conceptual Amazonian vegetation model

To motivate the simple growth equation that we use in Section 2.5, we here review the equilibrium model of Amazonian vegetation presented in Da Silveira Lobo Sternberg [2001]. Consider a region of the Amazon basin that is subdivided into cells having different dry season (d.s.) precipitation requirements for forest establishment. Assume that the frequency of cells whose d.s. precipitation requirement is p follows a normal distribution,

$$f(p) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(p-\mu)^2}{2\sigma^2}} \quad (1)$$

with mean μ and standard deviation σ . Then the relative forest cover C in the region is related to the region's average d.s. precipitation P through

$$C(P) = \int_0^P f(p) dp. \quad (2)$$

Forest cover increases overall precipitation because trees take up water stored in the soil and release it to the atmosphere via evapotranspiration. The total amount of d.s. precipitation in the region can therefore be expressed as

$$P(C) = P_{\text{in}} + C \cdot \Phi, \quad (3)$$

where P_{in} is the precipitation inflow from other regions and Φ is the contribution of one unit of forest cover to overall precipitation. An arid region is characterized by low P_{in} .

Now imagine that the region has a certain level of forest cover C_{before} before global climatic climate change alters P_{in} . How will C change? A pictorial answer is given in Fig. 1.




Figure 1: **Equilibrium model of Amazonian vegetation.** The solid green line represents equilibrium forest cover C at d.s. precipitation P . The dashed black line is the overall d.s. precipitation P as a function of C . Step-stair sequences are shown for two different initial values of forest cover. Reproduced from Da Silveira Lobo Sternberg [2001] .

The solid green line represents equilibrium forest cover C at d.s. precipitation P , cf. Eq. (2). The dashed black line is the overall d.s. precipitation P at forest cover C , cf. Eq. (3). If $C_{\text{before}} = C_1$, forest cover is above the critical threshold C_{crit} , so that P can support even more trees. Hence C grows, therefore P increases, upon which C grows further, and so on. Finally, due to this positive feedback loop, C is propelled towards 1.0. If, on the contrary, $C_{\text{before}} = C_2$, forest cover is below C_{crit} . Forest cover at this level cannot be supported by P , starts to die back, again triggers a positive feedback loop, and eventually vanishes completely.

We describe the growth dynamics using the Levins model [Levins, 1969], a widely accepted basic vegetation model [Tilman, 1994], to which we add a non-smooth switch in the growth term (corresponding to the limit $\sigma \ll \Phi$):

$$\frac{dC}{dt} = F(C) = \begin{cases} r(1 - C)C - xC & \text{if } C > C_{\text{crit}}, \\ -xC & \text{if } C < C_{\text{crit}}. \end{cases}$$

According to this model, forest cover C grows with the saturating rate r if $C > C_{\text{crit}}$ and dies with rate x (assuming $r > x > 0$). This model has two equilibria, the forest state $C_F = 1 - \frac{x}{r}$ and the savanna state $C_S = 0$. The equilibrium C_F (resp. C_S) exists and is stable if $C_F > C_{\text{crit}}$ (resp. $C_{\text{crit}} > 0$). So, as in the equilibrium model reviewed above, C converges to a non-zero constant value if $C > C_{\text{crit}}$ and vanishes completely otherwise. When increasing aridity drives up C_{crit} : i) C_F 's basin of attraction of shrinks, implying that C_F becomes less stable against perturbations such as strong deforestation (cf. Fig. 2.12). ii) C_F vanishes at $C_{\text{crit}} = C_F$.

Our model is conceptual. We do not intend it to serve as a description of reality. Rather, we use it as a reality-inspired aide to theory to illustrate the difference between the realm of small perturbations and the realm of non-small perturbations. This difference manifests itself in two phenomena:

i) As the linear stability coefficient of the forest state, $F'(C_F) = x - r$, is a constant independent of C_{crit} , critical slowing down does not take place. This is reflected by a zero recovery exponent [Kuehn, 2011]. Indeed, critical slowing down may be absent in systems in which strong non-linearities such as switches exist.

ii) More important to us, because of their local nature, linear stability and the small-perturbation convergence rate do not sense that C_F 's basin of attraction of shrinks as C_{crit} goes up. Dakos et al. [2008] write that “transitions caused by a sudden large disturbance without a preceding gradual loss of [stability against non-small perturbations]¹ will not be announced by slowing down.” On top of that, our model shows that slowing down may be absent *even if* such a gradual loss of stability is going on.

We conclude from this that linear stability and the small-perturbation convergence rate are *unreliable proxies* of how stable a state is against non-small perturbations. Instead of employing such local proxies, one should study global stability concepts like basin stability. The same conclusions can be drawn from the study of a basic damped driven pendulum (see Section 2.3.3).

¹The original term they use is “resilience”, referring to Holling’s concept discussed in the main text.

Glossary

Adjacency matrix. A \triangleright *network's* adjacency matrix is a matrix representation of this network. See Section 3.2.

Attractor. Two notions of attractor are distinguished in this text. A \triangleright *strong attractor* is a compact \triangleright *minimal invariant set* that is asymptotically stable. A \triangleright *weak attractor* is a compact minimal invariant set whose \triangleright *basin of attraction* has a positive volume (Lebesgue measure). Whenever the attribute weak or strong is omitted, the term attractor refers to a weak attractor.

Average shortest-path length. The average shortest-path length of a \triangleright *network* is the average length of the \triangleright *shortest paths* between all \triangleright *nodes* in the network. For a rigorous definition, see Section 3.2.

Basin of attraction. The basin of attraction of a compact \triangleright *minimal invariant set* is the set of \triangleright *states* in a \triangleright *dynamical system's* \triangleright *state space* from which the system \triangleright *converges* to this minimal invariant set.

Basin stability. Basin stability is a \triangleright *nonlocal stability* concept based on the volume (Lebesgue measure) of the \triangleright *basin of attraction* of a compact \triangleright *minimal invariant set*. It is rigorously defined in Section 2.3.

Betweenness. See \triangleright *shortest-path betweenness*.

Bifurcation. A bifurcation refers to a qualitative change in the \triangleright *local stability* of a compact \triangleright *minimal invariant set*.

Clustering coefficient. The clustering coefficient of a \triangleright *node* refers to the number of its \triangleright *neighbours* that are also neighbours. The clustering coefficient of a \triangleright *network* is the average of its nodal clustering coefficients. For a rigorous definition, see Section 3.2.

Complete synchronization. See \triangleright *synchronization*.

Component. See \triangleright *network component*.

Connected network. A \triangleright *network* is called connected if the set of all its \triangleright *nodes* is a \triangleright *connected set*.

Connected set. A set of \triangleright nodes of a \triangleright network is called connected if there is a \triangleright path along the \triangleright edges of the network between any two \triangleright nodes inside this set.

Convergence. A \triangleright trajectory of a \triangleright dynamical system converges to a compact \triangleright minimal invariant set in this system's \triangleright state space if, loosely speaking, the trajectory comes ever closer to this set as time proceeds. A formal definition is provided in Section 2.1.

Dead end. A dead end is a \triangleright dead tree whose \triangleright nodes' maximum \triangleright degree is at most 2.

Dead tree. A dead tree in a \triangleright network is a \triangleright connected set of \triangleright nodes none of which is contained in a closed \triangleright loop of \triangleright edges.

Degree. The degree of a \triangleright node is the number of nodes to which this node is linked by an \triangleright edge.

Design. Design refers to the same as \triangleright network topology.

Divergence. A \triangleright trajectory of a \triangleright dynamical system is said to diverge if it \triangleright converges to the point at infinity. See Section 2.3.

Dynamical regime. Dynamical regime is vague term meant to capture the long-term behaviour observed in a natural system. When a natural system is modelled as a \triangleright dynamical system, \triangleright attractors are the analogues of a dynamical regimes.

Dynamical system. A dynamical system in this text is defined as a set of deterministic differential equations. All dynamical systems investigated here are dissipative. See definitions in Section 2.1.

Edge. An edge between two \triangleright nodes of a \triangleright network captures a direct relation between these nodes. As additional information, an edge may carry a weight or a direction. For details and examples, see Section 3.2.

Laplacian matrix. A \triangleright network's Laplacian matrix is a matrix representation of this \triangleright network. See Section 3.2.

Local stability. See \triangleright stability.

Link. A link is the same as an \triangleright edge.

Loop. A loop in a \triangleright network is a \triangleright path that runs through a single \triangleright node more than once without hopping along the same \triangleright edge more than once.

Lyapunov exponents. Lyapunov exponents are a tool for assessing the \triangleright local stability of a compact \triangleright minimal invariant set.

Minimal invariant set. A minimal invariant set in a \triangleright *system's* \triangleright *state space* does not change as time progresses. A formal definition is provided in Section 2.1.

Multistability. A \triangleright *dynamical system* is said to be multistable if there are multiple \triangleright *attractors* in its \triangleright *state space*.

Neighbour. A \triangleright *node* in a \triangleright *network* is the neighbour of another node, if there exists an \triangleright *edge* between these two nodes.

Network. A network is a set of \triangleright *nodes* whose interrelations are captured by a set of \triangleright *edges*. For details and examples, see Section 3.2.

Network component. Network component is an umbrella term and refers to a \triangleright *node* or \triangleright *edge* of a \triangleright *network*.

Network design. Network design refers to the same as \triangleright *network topology*.

Network structure. Network structure refers to the same as \triangleright *network topology*.

Network topology. Network topology labels the observation of how \triangleright *edges* and \triangleright *nodes* of a \triangleright *network* are organized. It can be seen as incorporating any information about a network that can be obtained by application of an analysis technique. Aspects of a network's topology include its \triangleright *average shortest-path length*, its \triangleright *clustering coefficient*, and the \triangleright *degrees* of its nodes.

Node. Nodes are symbols for concrete or abstract objects whose interrelations are described by means of a \triangleright *network*. They can represent various different entities, including power generators, human beings, self-sustained oscillators or airports.

Nonlocal stability. See \triangleright *stability*.

Path. A path between two \triangleright *nodes* A and B on a network refers to an ordered set of \triangleright *edges* by which node B can be reached from node A. For details and examples, see Section 3.2.

Perturbation. A perturbation is understood as a sudden shift in a \triangleright *system's* \triangleright *state space* from one \triangleright *state* to another, so-called \triangleright *perturbed state*. The \triangleright *stability* of a compact \triangleright *minimal invariant set* against \triangleright *small perturbations* can be assessed using \triangleright *local stability* methods. In contrast, the stability against \triangleright *large perturbations* requires application of \triangleright *nonlocal stability* methods.

Perturbed state. A perturbed state is the result of a \triangleright *perturbation*.

Phase. The phase of (a component of) a \triangleright *dynamical system* refers to the position on an oscillatory cycle at a certain moment in time.

Power. Power labels the rate at which energy is produced or consumed. In AC power grids, power has an active and a reactive component. Whenever the attribute is omitted, power refers to active power in this text.

Regime transition. A regime transition refers to the change of a natural system from one \triangleright *dynamical regime* to another. When a natural system is modelled as a \triangleright *dynamical system*, a regime transition corresponds to the transition from one \triangleright *attractor* to another. This can be induced by a \triangleright *bifurcation*, in which an attractor loses its \triangleright *local stability*. In \triangleright *multistable systems*, another possible cause of transition is a \triangleright *large perturbation* that pushes the system from one attractor into the \triangleright *basin of attraction* of another attractor.

Self-sustained oscillator. A self-sustained oscillator is a \triangleright *dynamical system* whose dynamics evolve at a characteristic internal frequency when isolated from its surroundings.

Short cut. Short cut refers to a new \triangleright *edge* whose addition to a \triangleright *network* significantly shortens the length of one or more \triangleright *shortest paths*.

Shortest path. A shortest path between two \triangleright *nodes* A and B in a \triangleright *network* is a path along a minimum number of \triangleright *edges*. For details and examples, see Section 3.2.

Shortest-path betweenness. The shortest-path betweenness of a \triangleright *node* is the number of \triangleright *shortest paths* in a \triangleright *network* that run through this node. For a rigorous definition, see Section 3.2.

Small-worldness. A \triangleright *network* is said to have the small-world property if it has a comparatively small \triangleright *average shortest-path length* and a comparatively large \triangleright *clustering coefficient*. See Sections 3.3.4 and 3.4.2.

Stability. A fundamental distinction is made in this thesis between \triangleright *local stability* concepts and \triangleright *nonlocal stability* concepts. A compact \triangleright *minimal invariant set* of a \triangleright *dynamical system* is said to be asymptotically stable – or as a synonym: locally stable – if the system returns to this set after small \triangleright *perturbations*. Details and definitions can be found in Section 2.2. In contrast, nonlocal stability concepts address the question whether system returns to a compact minimal invariant set after large perturbations. See definitions in Section 2.3.

State. The state of a (deterministic) \triangleright *dynamical system* contains the complete information about this system at a given point in time.

State space. \triangleright *State space* refers to the set of possible \triangleright *states* a \triangleright *dynamical system* may assume.

Strong attractor. See \triangleright *attractor*.

Structure. Structure refers to the same as \triangleright *network topology*.

Synchronizability. Synchronizability is a measure based on \triangleright *local stability* that assesses the ability of a \triangleright *network* to support the \triangleright *synchronous* operation of its \triangleright *components*. See Section 3.3.3 for a rigorous definition.

Synchronization. Synchronization refers to the observation that two or more \triangleright *self-sustained oscillators* have their dynamics evolving at the same frequency. For identical chaotic oscillators, \triangleright *complete synchronization* means that these oscillators show exactly the same dynamics.

Synchronous state. This term refers to a compact \triangleright *minimal invariant set* of a multi-component \triangleright *dynamical system* in which the dynamics of the components show \triangleright *synchronization*.

Topology. Topology refers to the same as \triangleright *network topology*.

Trajectory. A trajectory is a set of \triangleright *states* in a \triangleright *dynamical system's state space* along which this system evolves in time.

Undirected network. An undirected network is a \triangleright *network* whose \triangleright *edges* carry no direction.

Unweighted network. An unweighted network is a \triangleright *network* whose \triangleright *edges* carry no weights.

Watts-Strogatz network. A Watts-Strogatz network is a \triangleright *network* that is randomly generated according to the Watts-Strogatz algorithm [Watts and Strogatz, 1998]. See Section 3.3.4.

Weak attractor. See \triangleright *attractor*.

Bibliography

- Albert, R. and Barabási, A.-L. (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47.
- Almendral, J. A. and Díaz-Guilera, A. (2007). Dynamical and spectral properties of complex networks. *New Journal of Physics*, 9(6):187.
- Amaral, L. A. N., Scala, A., Barthélemy, M., and Stanley, H. E. (2000). Classes of small-world networks. *Proceedings of the National Academy of Sciences*, 97(21):11149–11152.
- Arenas, A., Díaz-Guilera, A., Kurths, J., Moreno, Y., and Zhou, C. (2008). Synchronization in complex networks. *Physics Reports*, 469(3):93 – 153.
- Argyris, J. H., Faust, G., Haase, M., and Rudolf, F. (2010). *Die Erforschung des Chaos: Eine Einführung in die Theorie nichtlinearer Systeme*. Springer, Heidelberg.
- Babloyantz, A. and Destexhe, A. (1986). Low-dimensional chaos in an instance of epilepsy. *Proceedings of the National Academy of Sciences*, 83(10):3513–3517.
- Barahona, M. and Pecora, L. M. (2002). Synchronization in small-world systems. *Physical Review Letters*, 89:054101.
- Boccaletti, S., Kurths, J., Osipov, G., Valladares, D., and Zhou, C. (2002). The synchronization of chaotic systems. *Physics Reports*, 366(1-2):1 – 101.
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., and Hwang, D.-U. (2006). Complex networks: Structure and dynamics. *Physics Reports*, 424(4-5):175–308.
- Bramble, D. M. and Carrier, D. R. (1983). Running and breathing in mammals. *Science*, 219(4582):251–256.
- Buldyrev, S. V., Parshani, R., Paul, G., Stanley, H. E., and Havlin, S. (2010). Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291):1025–1028.
- Carpenter, S. and Brock, W. (2006). Rising variance: a leading indicator of ecological transition. *Ecology letters*, 9(3):311–318.
- Carreras, B. A., Lynch, V. E., Dobson, I., and Newman, D. E. (2004). Complex dynamics of blackouts in power transmission systems. *Chaos*, 14(3):643–652.
- Chiang, H.-D. (2011). *Direct Methods for Stability Analysis of Electric Power Systems*. John Wiley & Sons, Chichester.

- Chiang, H.-D., Chu, C.-C., and Cauley, G. (1995). Direct stability analysis of electric power systems using energy functions: theory, applications, and perspective. *Proceedings of the IEEE*, 83(11):1497–1529.
- Cookson, N. A., Tsimring, L. S., and Hasty, J. (2009). The pedestrian watchmaker: Genetic clocks from engineered oscillators. *FEBS letters*, 583(24):3931.
- Crucitti, P., Latora, V., and Marchiori, M. (2004a). Model for cascading failures in complex networks. *Physical Review E*, 69:045104.
- Crucitti, P., Latora, V., and Marchiori, M. (2004b). A topological analysis of the Italian electric power grid. *Physica A*, 338(1–2):92–97.
- Csardi, G. and Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal, Complex Systems*, 1695(5).
- Da Silveira Lobo Sternberg, L. (2001). Savanna-forest hysteresis in the tropics. *Global Ecology and Biogeography*, 10(4):369–378.
- Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V., and Held, H. (2008). Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38):14308–14312.
- Dakos, V., Van Nes, E. H., D’Odorico, P., and Scheffer, M. (2012). Robustness of variance and autocorrelation as indicators of critical slowing down. *Ecology*, 93(2):264–271.
- de Abreu, N. M. M. (2007). Old and new results on algebraic connectivity of graphs. *Linear Algebra and its Applications*, 423(1):53 – 73.
- Dobson, I. (2013). Complex networks: Synchrony and your morning coffee. *Nature Physics*, 9(3):133–134.
- Dodd, A. N., Salathia, N., Hall, A., Kévei, E., Tóth, R., Nagy, F., Hibberd, J. M., Millar, A. J., and Webb, A. A. (2005). Plant circadian clocks increase photosynthesis, growth, survival, and competitive advantage. *Science Signaling*, 309(5734):630.
- Donges, J. F., Zou, Y., Marwan, N., and Kurths, J. (2009). The backbone of the climate network. *Europhysics Letters*, 87(4):48007.
- Donner, R. V., Heitzig, J., Donges, J. F., Zou, Y., Marwan, N., and Kurths, J. (2011). The geometry of chaotic dynamics – a complex network perspective. *European Physical Journal B*, 84(4):653–672.
- Dörfler, F., Chertkov, M., and Bullo, F. (2013). Synchronization in complex oscillator networks and smart grids. *Proceedings of the National Academy of Sciences*, 110(6):2005–2010.
- ENTSO-E (2011). Statistical Yearbook 2011. https://www.entsoe.eu/fileadmin/user_upload/_library/publications/entsoe/Statistical_Yearbook/SYB_2011/121216_SYB_2011_final.pdf. [Online; accessed 30 September 2013].

- Erzgräber, H., Lenstra, D., Krauskopf, B., Wille, E., Peil, M., Fischer, I., and Elsässer, W. (2005). Mutually delay-coupled semiconductor lasers: Mode bifurcation scenarios. *Optics Communications*, 255(4–6):286 – 296.
- Ewart, D. (1978). Whys and wherefores of power system blackouts: An examination of the factors that increase the likelihood and the frequency of system failure. *IEEE Spectrum*, 15(4):36–41.
- Fairley, P. (2004). The unruly power grid. *IEEE Spectrum*, 41(8):22–27.
- Faloutsos, M., Faloutsos, P., and Faloutsos, C. (1999). On power-law relationships of the internet topology. In *ACM SIGCOMM Computer Communication Review*, volume 29, pages 251–262. ACM.
- Fell, J. and Axmacher, N. (2011). The role of phase synchronization in memory processes. *Nature Reviews Neuroscience*, 12(2):105–118.
- Fell, J., Klaver, P., Lehnertz, K., Grunwald, T., Schaller, C., Elger, C. E., and Fernandez, G. (2001). Human memory formation is accompanied by rhinal-hippocampal coupling and decoupling. *Nature Neuroscience*, 4(12):1259–1264.
- Fiedler, M. (1973). Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305.
- Filatrella, G., Nielsen, A. H., and Pedersen, N. F. (2008). Analysis of a power grid using a Kuramoto-like model. *European Physical Journal B*, 61:485–491.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474 – 480.
- Gross, T., Rudolf, L., Levin, S. A., and Dieckmann, U. (2009). Generalized models reveal stabilizing factors in food webs. *Science*, 325(5941):747–750.
- Guimera, R., Mossa, S., Turtshi, A., and Amaral, L. N. (2005). The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles. *Proceedings of the National Academy of Sciences*, 102(22):7794–7799.
- Hammond, C., Bergman, H., and Brown, P. (2007). Pathological synchronization in Parkinson’s disease: networks, models and treatments. *Trends in Neurosciences*, 30(7):357–364.
- Held, H. and Kleinen, T. (2004). Detection of climate system bifurcations by degenerate fingerprinting. *Geophysical Research Letters*, 31(23).
- Hill, D. and Chen, G. (2006). Power systems as dynamic networks. In *Proceedings of the 2006 IEEE International Symposium on Circuits and Systems*, pages 722–725.
- Hirota, M., Holmgren, M., Van Nes, E. H., and Scheffer, M. (2011). Global resilience of tropical forest and savanna to critical transitions. *Science*, 334(6053):232–235.
- Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23.

- Holme, P., Kim, B. J., Yoon, C. N., and Han, S. K. (2002). Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109.
- Hong, H., Kim, B. J., Choi, M. Y., and Park, H. (2004). Factors that predict better synchronizability on complex networks. *Physical Review E*, 69:067105.
- Horowitz, S. and Phadke, A. (2008). *Power System Relaying*. John Wiley & Sons, Chichester.
- Huang, L., Chen, Q., Lai, Y.-C., and Pecora, L. M. (2009). Generic behavior of master-stability functions in coupled nonlinear dynamical systems. *Physical Review E*, 80:036204.
- Huang, S. and Ingber, D. E. (2007). A non-genetic basis for cancer progression and metastasis: Self-organizing attractors in cell regulatory networks. *Breast Disease*, 26(1):27–54.
- Kaluza, P., Kölzsch, A., Gastner, M. T., and Blasius, B. (2010). The complex network of global cargo ship movements. *Journal of the Royal Society Interface*, 7(48):1093–1103.
- Kennedy, J. and Yorke, J. A. (1991). Basins of Wada. *Physica D*, 51(1):213–225.
- Kuehn, C. (2011). A mathematical framework for critical transitions: Bifurcations, fast-slow systems and stochastic dynamics. *Physica D*, 240(12):1020 – 1035.
- Kuramoto, Y. (1975). Self-entrainment of a population of coupled non-linear oscillators. In *International symposium on mathematical problems in theoretical physics*, pages 420–422. Springer.
- Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J. (2008). Tipping elements in the Earth’s climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793.
- Levins, R. (1969). Some demographic and genetic consequences of environmental heterogeneity for biological control. *Bulletin of the ESA*, 15(3):237–240.
- Li, G., Reis, S. D. S., Moreira, A. A., Havlin, S., Stanley, H. E., and Andrade, J. S. (2010). Towards design principles for optimal transport networks. *Physical Review Letters*, 104:018701.
- Lovász, L. and Vempala, S. (2006). Simulated annealing in convex bodies and an $\mathcal{O}^*(n^4)$ volume algorithm. *Journal of Computer and System Sciences*, 72(2):392 – 417.
- Lytton, W. W. (2008). Computer modelling of epilepsy. *Nature Reviews Neuroscience*, 9(8):626–637.
- Machowski, J., Bialek, J. W., and Bumby, J. R. (2008). *Power System Dynamics: Stability and Control*. John Wiley & Sons, Chichester.

- May, R. M. (1977). Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269(5628):471–477.
- McDonald, S. W., Grebogi, C., Ott, E., and Yorke, J. A. (1985). Fractal basin boundaries. *Physica D*, 17(2):125–153.
- Menck, P. J., Heitzig, J., Marwan, N., and Kurths, J. (2013a). How basin stability complements the linear-stability paradigm. *Nature Physics*, 9(2):89–92.
- Menck, P. J., Heitzig, J., Kurths, J., and Schellnhuber, H. J. (2013b). How dead ends undermine power grid stability. *Under review at Nature Communications*.
- Menck, P. J. and Kurths, J. (2012). Topological identification of weak points in power grids. In *Proceedings of Nonlinear Dynamics of Electronic Systems 2012*, pages 1–4. VDE.
- Milgram, S. (1967). The small world problem. *Psychology today*, 2(1):60–67.
- Moiseff, A. and Copeland, J. (2010). Firefly synchrony: a behavioral strategy to minimize visual clutter. *Science*, 329(5988):181–181.
- Motter, A. E. and Lai, Y.-C. (2002). Cascade-based attacks on complex networks. *Physical Review E*, 66:065102.
- Motter, A. E., Myers, S. A., Anghel, M., and Nishikawa, T. (2013). Spontaneous synchrony in power-grid networks. *Nature Physics*, 9(3):191–197.
- Newman, D. E., Carreras, B. A., Lynch, V. E., and Dobson, I. (2011). Exploring complex systems aspects of blackout risk and mitigation. *IEEE Transactions on Reliability*, 60(1):134–143.
- Newman, M. (2003). The structure and function of complex networks. *SIAM Review*, 45:167 – 256.
- Newman, M. (2010). *Networks: An Introduction*. Oxford University Press, Oxford.
- Nishikawa, T., Motter, A. E., Lai, Y.-C., and Hoppensteadt, F. C. (2003). Heterogeneity in oscillator networks: Are smaller worlds easier to synchronize? *Physical Review Letters*, 91:014101.
- Nusse, H. E. and Yorke, J. A. (1996). Basins of attraction. *Science*, 271(5254):1376–1380.
- Oyama, M. D. and Nobre, C. A. (2003). A new climate-vegetation equilibrium state for tropical South America. *Geophysical Research Letters*, 30(23):2199.
- Pantaleone, J. (2002). Synchronization of metronomes. *American Journal of Physics*, 70:992.
- Pecora, L. M. (1998). Synchronization conditions and desynchronizing patterns in coupled limit-cycle and chaotic systems. *Physical Review E*, 58:347–360.

Bibliography

- Pecora, L. M. and Carroll, T. L. (1990). Synchronization in chaotic systems. *Physical Review Letters*, 64(8):821–824.
- Pecora, L. M. and Carroll, T. L. (1998). Master stability functions for synchronized coupled systems. *Physical Review Letters*, 80:2109–2112.
- Pecora, L. M., Carroll, T. L., Johnson, G. A., Mar, D. J., and Heagy, J. F. (1997). Fundamentals of synchronization in chaotic systems, concepts, and applications. *Chaos*, 7(4):520–543.
- Pikovsky, A., Rosenblum, M., and Kurths, J. (2003). *Synchronization: a universal concept in nonlinear sciences*, volume 12. Cambridge University Press, Cambridge.
- Popper, K. R. (1963). *Conjectures and refutations*. Routledge & Kegan Paul, London.
- Price, D. (1965). Statistical studies of networks of scientific papers. In *Statistical Association Methods for Mechanized Documentation: Symposium Proceedings*, volume 269, page 187. US Government Printing Office.
- Rahmstorf, S. (1995). Bifurcations of the Atlantic thermohaline circulation in response to changes in the hydrological cycle. *Nature*, 378(6553):145–149.
- Rohden, M., Sorge, A., Timme, M., and Witthaut, D. (2012). Self-organized synchronization in decentralized power grids. *Physical Review Letters*, 109:064101.
- Rosenblum, M. G., Pikovsky, A. S., and Kurths, J. (1996). Phase synchronization of chaotic oscillators. *Physical Review Letters*, 76:1804–1807.
- Rössler, O. (1976). An equation for continuous chaos. *Physics Letters A*, 57(5):397 – 398.
- Scheffer, M. (2009). *Critical Transitions in Nature and Society*. Princeton University Press, Princeton.
- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., and Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260):53–59.
- Schellnhuber, H. J. (2009). Tipping elements in the Earth System. *Proceedings of the National Academy of Sciences*, 106(49):20561–20563.
- Sporns, O. and Zwi, J. (2004). The small world of the cerebral cortex. *Neuroinformatics*, 2:145–162. 10.1385/NI:2:2:145.
- Spring, E. (2003). *Elektrische Energienetze: Energieübertragung und -verteilung*. VDE Verlag, Berlin.
- Statistisches Bundesamt (2012). <https://www.destatis.de/DE/ZahlenFakten/Wirtschaftsbereiche/Energie/Erzeugung/Tabellen/ErneuerbareEnergie.html>. [Online; accessed 30 September 2013].

- Strogatz, S. (2001). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry and engineering*. Perseus Books Group.
- Strogatz, S. H. (2000). From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D*, 143(1):1–20.
- Strogatz, S. H. (2012). *Sync: How order emerges from chaos in the universe, nature, and daily life*. Hyperion, New York.
- Strogatz, S. H. and Mirollo, R. E. (1991). Stability of incoherence in a population of coupled oscillators. *Journal of Statistical Physics*, 63(3-4):613–635.
- Sun, K. (2005). Complex networks theory: A new method of research in power grid. In *Transmission and Distribution Conference and Exhibition: Asia and Pacific, 2005 IEEE/PES*, pages 1 –6.
- Tilman, D. (1994). Competition and biodiversity in spatially structured habitats. *Ecology*, 75:2–16.
- tom Dieck, T. (2000). *Topologie*. Walter de Gruyter, Berlin.
- Travers, J. and Milgram, S. (1969). An experimental study of the small world problem. *Sociometry*, pages 425–443.
- UCTE (2007). Final report system disturbance on 4 November 2006. https://www.entsoe.eu/fileadmin/user_upload/_library/publications/ce/otherreports/Final-Report-20070130.pdf. [Online; accessed 5 September 2013].
- Ugander, J., Karrer, B., Backstrom, L., and Marlow, C. (2011). The anatomy of the Facebook social graph. *arXiv preprint arXiv:1111.4503*.
- U.S.-Canada Power System Outage Task Force (2004). Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations. <http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf>. [Online; accessed 5 September 2013].
- van Nes, E. H. and Scheffer, M. (2007). Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *The American Naturalist*, 169(6):pp. 738–747.
- Von Neumann, J. (1951). Various techniques used in connection with random digits. *Applied Math Series*, 12(36-38):1.
- Wang, Z., Scaglione, A., and Thomas, R. (2010). Generating statistically correct random topologies for testing smart grid communication and control networks. *IEEE Transactions on Smart Grid*, 1(1):28 –39.
- Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442.

Bibliography

- Wiley, D. A., Strogatz, S. H., and Girvan, M. (2006). The size of the sync basin. *Chaos*, 16(1):015103.
- Witthaut, D. and Timme, M. (2012). Braess’s paradox in oscillator networks, desynchronization and power outage. *New Journal of Physics*, 14(8):083036.
- Zhao, M., Zhou, T., Wang, B.-H., Yan, G., Yang, H.-J., and Bai, W.-J. (2006). Relations between average distance, heterogeneity and network synchronizability. *Physica A*, 371(2):773–780.
- Zhou, C. and Kurths, J. (2006). Dynamical weights and enhanced synchronization in adaptive complex networks. *Physical Review Letters*, 96(16):164102.
- Zhou, C., Motter, A. E., and Kurths, J. (2006). Universality in the synchronization of weighted random networks. *Physical Review Letters*, 96:034101.