1 Introduction

1.1 Two Approaches: The Signal-flow Approach and the Network of Local LTI Systems Approach

The developments in hardware technology, wireless communication networks and computational technology contribute to the advancement of large-scale networks of interconnected systems. The large-scale nature of such networks has started to arouse curiosity within the systems & control community. A few cases may already be distinguished. The exchange of information between the systems interacting in a network can be known, partially known or be unknown. Examples of networks with known interconnections include networks of pipelines and electrical grids. Unknown and abstract connections occur in spatial-temporal systems that are governed by partial differential equations (PDEs). After discretization of the equations, the dynamics can be viewed as an array of local dynamical systems that exchange information via physical interactions between the local models that are each associated to local input and output signals. Examples include arrays of sensors for measuring the optical disturbance induced by a spatially and temporally varying windfield that hovers over a ground-based telescope, and in flow control to stabilize boundary layers and reduce the drag. A deformable mirror with non-negligible time response that may be used in optics also belongs to this class of spatial-temporal dynamical systems. In these cases, the sensor delivers at a high frequency rate a large number of measurements that will subsequently be processed for data-driven control.

The topic of this book is the system identification of interconnected systems operating in a network. System identification attempts to find a mathematical representation for the dynamics by which measured signals are related. Such a mathematical relationship is termed a model. Finding such a model requires a systematic approach towards designing experiments, specifying model classes, using mathematical methods to find estimates for the parameters in these models and validating these estimated models. These topics are well discussed in, for example, Ljung (1999); Verhaegen and Verdult (2007). This book focuses on the definition of model classes for networks of dynamical systems and presents some system identification methodologies to estimate such models from input-output measurements.

The main bottlenecks for the extension of classical identification methods to large-scale networks are the huge volume of data and the large number of parameters necessary to describe these models, even when considering compact local model dynamics.

A key challenge that will be addressed in this book is how to efficiently deal with this tremendous degree of complexity while obtaining models that are as accurate as possible in representing the system behaviour given, for example, by the input-output response.

Two major system identification approaches for identifying network connected systems are under full development. These developments occur very much along the lines of the two model classes that are in use for lumped parameter systems: transfer function models and state-space models. Many of the system identification approaches for networks have been inspired by particular applications.

The first main approach for identifying dynamic networks is based on the input-output transfer functions describing the network. This approach characterizes the network by employing dynamic signal flow diagrams and will be referred to as the *signal flow approach*. It starts from a global model representation that, in principle, can be formulated without knowledge of the way in which the local systems are connected to each other. If size allows, the global dynamics may be identified and a special parametrization in the global coefficient matrices may be looked for afterwards to find out or estimate a network topology. However, we usually strive for lean parametrizations of the coefficient matrices in these models in order to deal with the data complexity and overcome the computational bottleneck. Specifying such models and their consequences on the identification methods to find these models for data is a challenge to be pursued in this book. These parametrizations may come from physical insights such as localizability (each local system in the network has a very limited influence on its neighbours), spatial invariance (all local systems are identical) or separability of the system dynamics along the physical dimensions. Such models, which rely on transfer functions, are labelled in this book as *global* models. Typical applications are from biological or computer networks (see, for example, Gevers et al. [2019]; Gonçalves and Warnick [2008]).

The second main approach for identifying networks of dynamic systems is based on large state-space models with structured system matrices and are developed, among others, in the scope of designing data-sparse controllers for network connected systems. For examples of the latter we refer to D'Andrea and Dullerud (2003); Massioni and Verhaegen (2009); Rice and Verhaegen (2009). Typical areas of application of such controller design methods are formation flying, cross-directional control in paper processing, automated highway systems, micro-cantilever array control for massive parallel data storage and lumped parameter approximations of PDEs. We refer to the work in D'Andrea and Dullerud (2003) for more details on these applications. Local models may be related to a physical entity such as cars on an automated highway, or can be virtual as is often the case when modelling spatial-temporal dynamical systems. Key for these models is the interconnection pattern of the network, i.e. how the local dynamical systems are connected to each other. Priors from physics such as spatial invariance and separability may also be formulated here although at a local scale, that is, on the system matrices of the local models. These may eventually translate into parametrizations of the global matrices. This second approach is referred to as the *network of local LTI systems*-approach, and the mathematical methods derived in this context can be viewed as *identification for distributed control* methods that aim to be efficient for distributed control design. The efficiency envisaged is that the model structure of the identified model can be exploited to arrive at low-cost controller design, both in terms of calculating and implementing these controllers.

The examples given so far are illustrative and follow existing literature that uses either transfer functions or state-space models for subsequent control. Both approaches may nonetheless be applied succesfully to an application mentioned for the other approach.

An appropriate model structure is of prime importance for system identification as has been discussed in Ljung (1999); Verhaegen and Verdult (2007). If the model structure is not included within the set of candidates that accurately represents the true system, the estimates derived will be biased irrespective of the size of the dataset or the convergence properties of the dedicated algorithm. This is the case for small and medium size models as demonstrated by the bias-variance analysis given, for example, in Ljung (1999). For large-scale models, this will be even more prominent because the number of data points available for identification (and validation) is small compared to the number of parameters to be estimated. One often seeks (data-)sparse model parametrizations and the enforcing of structure in the system matrices of the model to be estimated. In general, this may help to reduce the variance of the estimates compared to the unstructured case for the same number of temporal samples. However, the other side of the coin is that the model quality may degrade: an incorrect parametrization of the matrices is likely to increase the bias, e.g. when the physical insights are too crude. Instead of conducting a bias-variance trade-off analysis like that in Ljung (1999), we develop identification methodologies to estimate model structures where structural parameters, such as the order of local models, are derived from the given data. These structural parameters offer the freedom to the user to get the best model within that model class. This will be the case for the subspace-like identification methods to be developed in this book.

In this chapter, we will start by providing motivating examples additional to those already mentioned that demonstrate how a network of dynamical systems arises through the use of multi-dimensional sensor grids. We will then briefly outline the remaining chapters of this book.

1.2 Examples of Networks of Dynamical Systems

Dynamical systems with multi-dimensional sensor/actuator arrays consisting of a large number of nodes appear, or will appear, in the near future in various engineering applications. These applications range from optics to flow control, but also in many other applications in economics, systems biology, social network analysis, etc. In this book, the examples considered are mainly restricted to engineering problems.

978-1-316-51570-9 — Data-Driven Identification of Networks of Dynamic Systems
Michel Verhaegen , Chengpu Yu , Baptiste Sinquin
Excerpt
More Information

1.2.1 Data-Driven Predictive Control for Large-Scale and Extreme Adaptive Optics

In the current generation of large-scale adaptive optics (AO) and the forthcoming generation of extreme AO (XAO) being designed for the inspection of exoplanets, use will be made of (extreme) large multi-dimensional sensor/actuator arrays with thousands (ten-thousands) of nodes, each processing local signals at the kHz rate. The imaging of exoplanets using XAO is challenging due to the high contrast ratio and the small angular separation relative to the star; see Guyon (2018) for more details. In these AO systems, the turbulence induced in the atmosphere above the telescope induces a stochastic disturbance that greatly deteriorates the image quality and, without (real-time) compensation, reduces these extremely expensive telescopes to something comparable to Newton's telescope developed in the seventeenth century. The AO application is further discussed in Chapter 10. Here, we restrict ourselves to highlighting some identification challenges. Figure 1.1 illustrates the turbulence fields flowing over a telescope aperture. Such turbulence fields introduce optical aberrations that reduce the quality of the image recorded by the telescope. This would limit the science that can be done using such images such as investigating the atmosphere of exoplanets. Only the *wavefront* of the optical field aberrations is of interest as the intensity can often be assumed to be uniform. In order to improve the image resolution, use can be made of AO in order to reduce the effects of the optical aberrations. An important element in an AO system is the sensor that measures the optical aberrations.

The sensor measuring these wavefront aberrations in a ground telescope is a two-dimensional array, as depicted in Figure 1.2. The sensor, for example a Shack–Hartmann sensor, measures the combined effect of all the layers above the telescope, with each layer having a different wind speed and direction. A direct measurement of the optical wavefront aberration is not possible and a reconstruction always needs to be done. The reconstruction step may benefit from knowledge of model of spatial-temporal dynamics of the induced wavefront aberrations in order to predict these disturbances in time. Deriving such a model from the large number of measurements on an (extreme) large telescope is a challenging problem. The methods to be developed in this book serve as possible solutions for these types of problems. Based on the measured wavefront aberrations, a deformable mirror can then compensate the wavefront of the optical beam. Such a deformable mirror consists of a grid of actuators that operate under the mirror surface to deform that surface into a shape that compensates the overall wavefront distortation. Linking the sensor to this actuator is done by the controller. Modelling this link is also an identification problem that can be considered with the methods presented in this book.

Many instruments, such as a coronograph or a spectrograph, could benefit from scalable large-scale AO control algorithms. Examples of such instruments are HARMONI (high angular resolution monolithic optical and near-infrared integral field spectrograph) for the ESO's extreme large telescope (ELT) (Neichel et al., 2016), NFIRAOS (narrow field infrared adaptive optics system) for the Thirty Meter Telescope (Ellerbroek, 2011), and GPI (Gemini planet imager) (Poyneer et al., 2016).

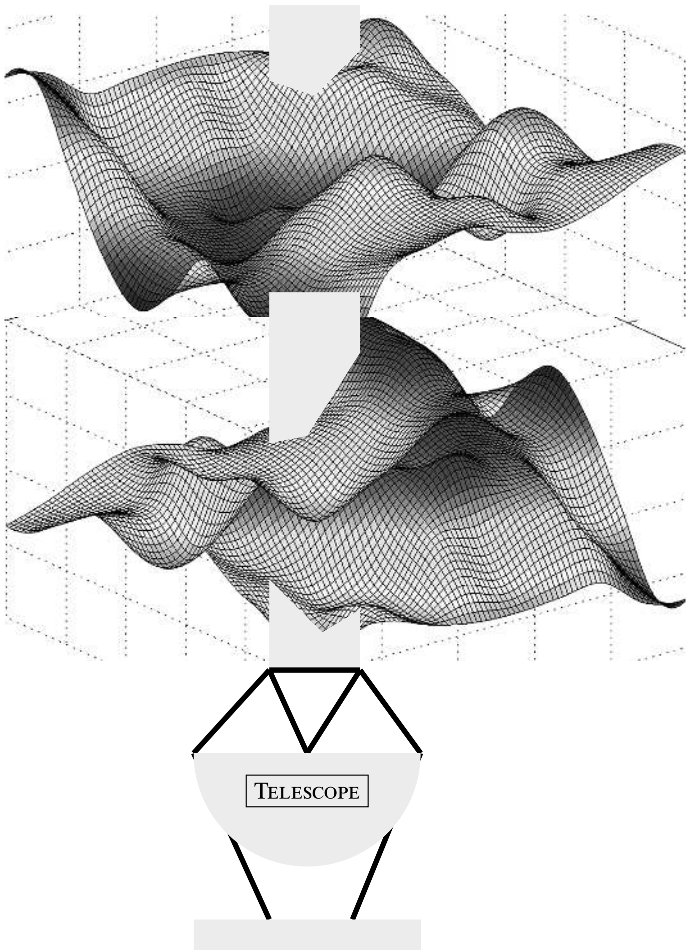

Figure 1.1 A schematic representation of a telescope with two layers of the turbulence above the telescope. Each layer in this plot schematically represents the surface of the optical wavefront. In this figure we do not plot the isoplanatic angle between the star and the object of interest. For more details on the nomenclature refer to Chapter 10.

AO is not limited to astronomy. It has been widely used in other application areas, such as microscopy (Pozzi et al., 2020) and optical coherence tomography for in vivo high resolution imaging of the human retina (Pircher and Zawadzki, 2017). Although the control loop usually operates at a few hertz in these applications, there has not been much development regarding the data-driven minimum variance control of such systems.

1.2.2 System Identification for Wind Farm Control

Another example of the application of system identification of large-scale systems is the area of offshore wind farms (Gebraad, 2014). To illustrate, we refer to Figure 1.3

Cambridge University Press & Assessment
978-1-316-51570-9 — Data-Driven Identification of Networks of Dynamic Systems
Michel Verhaegen , Chengpu Yu , Baptiste Sinquin
Excerpt
More Information

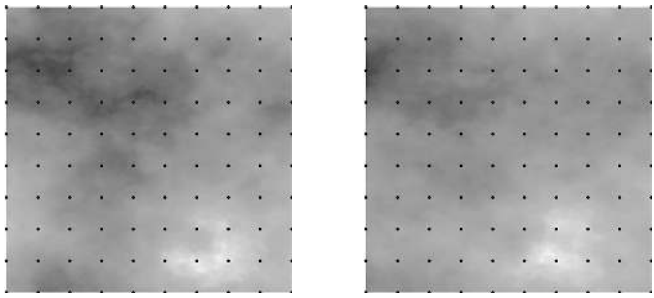

Figure 1.2 Example of a turbulence field consisting of the sum of two frozen layers propagating at different speeds: (left) $t = t_0$; (right) $t = t_1 > t_0$. The sensor array is of size 10×10 and the sensor nodes are represented by black dots. The correlations are not only spatial but also temporal. Reprinted/adapted from Sinquin (2019) with permission from TU Delft

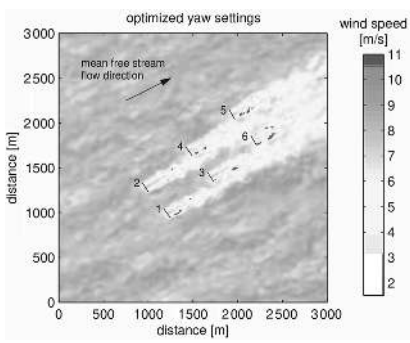

Figure 1.3 A 3×2 wind plant rotated 10° with respect to wind direction. Hub-height wind field at 800 s simulated time as calculated by the software SOWFA. The black lines indicate the rotor positions and yaw orientation of each turbine. Reprinted/adapted from Gebraad (2014) with permission from TU Delft

where six closely spaced wind turbines operate in a wind corridor. The control objective is to maximize the energy production of these turbines. This requires knowledge of the spatial-temporal interaction between them. Mathematical models are of use in the design of model-based controllers for maximizing the energy yield. Even though the sampling frequency for such applications is in the range of one Hertz (and much smaller than the AO applications considered in the previous subsection), the challenge for a mathematical model is that it should allow for an efficient controller design as well as its real-time operation. One approach to designing such a model is based on discretizing the Navier–Stokes equation (Boersma et al., 2018), and using this discretized model in ensemble Kalman filtering techniques (Doekemeijer et al., 2018). A second approach is to identify a dynamic model based on measurements of the wind velocity in time and space (e.g. on a 3D regular grid). Such models, can then be used to predict the future evolution of the wind field (such as its spatially distributed wind velocity and direction) in a predictive controller. The development of these

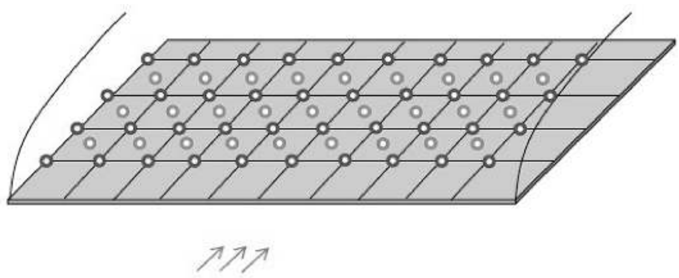

Figure 1.4 A 2D schematic of a multi-dimensional sensor and actuator array for active boundary control. The actuators are indicated by light circles, while the sensors by dark circles. The undisturbed flow is indicated by the arrows. The actuators and sensors are arranged in a regular grid. Reprinted/adapted from Sinquin (2019) with permission from TU Delft

models may be possible using the approaches outlined in this book, though the actual realization still needs to be done.

1.2.3 Active Boundary Layer Control

The goal of active boundary layer control on, for example, the wing of an airplane is to reduce the drag induced by the wind flow around the wing and fuselage by delaying the flow separation and preventing (as far as possible) the laminar flow (low drag) from transitioning into turbulent flow (high drag). This goal is achieved by using a multi-dimensional array of pressure sensors and a multi-dimensional array of actuators, such as blowing or suction actuators, to actively control the boundary flow around the surface. A possible configuration of these arrays is shown in Figure 1.4.

A trade-off in active boundary control is to determine a model for the temporal and spatial dynamics of the wind speed and vorticity that on one hand is compact, so that efficient control design and implementation becomes possible, and on the other that it is accurate enough to enable high performance control. In a number of occasions, such as the flow in a pipe, the dynamics of the flow close to the boundary of the pipe can be accurately described around a steady-state wind speed by the linearized Navier–Stokes equations. This yields a linear PDE for small variations around a steady state. The system identification methods developed in this book pursue the trade-off between model complexity and model accuracy using real-life measured data with multi-dimensional arrays of sensors and actuators.

For preliminary results on the development of compact models for boundary layer control via system identification we refer to Kim and Bewley (2007) and Inigo (2015).

1.2.4 Varied Applications Ranging from Weather Prediction to Sociology

Even without actuators to control, some datasets can be recast as a multi-dimensional networks, for example in weather prediction (Tsiligkaridis and Hero, 2013). There are many examples of geographic information systems that collect data via an array of sensors distributed geographically. The notion of dimension may also be more abstract, for example, in sociology and for studying relational networks (Hoff, 2015) as it does not feature a sensor grid. This area motivated the introduction of self-replicating patterns (Leskovec et al., 2010): the network is made up of different clusters that each replicate the same interconnection pattern, the clusters are themselves made of clusters of subsystems that again interact in a similar pattern. This may also occur in biology.

Adaptive optics systems, wind farms and flow control all rely on large arrays of sensors to measure wind field-induced aberrations and are potential applications of the system identification methods that are developed in this book. These examples may be recast into the signal flow approach or into the network of local LTI systems approach. In adaptive optics, identifying the spatial-temporal dynamics of the wave-front aberrations may be recast into a Signal Flow approach whereas identifying those of a deformable mirror would be seen as a set of interconnected subsystems with very limited neighbourhood allowing the use of the local approach. In such cases, there is already a strong a priori knowledge of the topology of the network that can be used to enhance the computational efficiency of the system identification. Both the signal flow approach and the network of local LTI systems may be used to represent discretized PDEs. This applies to the deterministic behaviour as well as the possible stochastic disturbances that are modelled by PDEs, such in modelling turbulence via the Navier–Stokes equations. The models derived may be used in the design of large-scale observers or Kalman filters.

In Section 1.3, we review the heat equation in the context of distributed systems to illustrate notions such as topology, subsystems, and how these arise from a discretized PDE.

1.3 The Spatio-Temporal Impulse Response

For particular type of waves, including the heat conduction or optical wave propagation in an empty medium, their behaviour is governed by a linear PDE featuring both spatial and temporal derivatives. At each discretization node of a PDE, a subsystem is assigned with its own input and output. In other words, the discretization and the interconnection pattern of the subsystems are highly related. Each subsystem shares information with its neighbours in time and space that can be pictured by its spatial-temporal impulse response.

A heat conduction example is used as an illustration. Here, we consider a thin metal plate with homogeneous material density and a known temperature distribution $T_0(\xi)$. In this notation, ξ represents the spatial variable and $t = t_0$ is the initial time point. Let the spatial boundary be denoted by Ω. The temporal propagation of the temperature over time $t > t_0$ is governed by thermal conduction principles and subjected to boundary conditions. Appropriate boundary conditions are the homogeneous Dirichlet boundary conditions. This results into the following PDE:

$$
\begin{cases} \frac{\partial}{\partial t} T(t,\xi) = -c\nabla^2 T(t,\xi) + u(t,\xi), \\ T(t_0,\xi) = T_0(\xi), \\ T(t,\Omega) = 0. \end{cases} \tag{1.1}
$$

where c is a positive constant, ∇^2 is the Laplacian operator and $u(t,\xi)$ is a spatially (shown by its dependency on ξ) temporally (shown by its dependency on t) distributed input quantity. The discretisation of the PDE in Equation (1.1) can be obtained by finite-differencing in both time and space. This gives rise to a lumped parameter system model. When considering a uniform two-dimensional spatial grid of size $N \times N$, the discrete spatial-temporal model of Equation (1.1) becomes,

$$
\begin{cases} T_{i_1,i_2}(k+1) = (1+4\alpha)T_{i_1,i_2}(k) - \alpha \sum_{(\bar{i}_1,\bar{i}_2)\in\mathcal{N}_{(i_1,i_2)}} T_{\bar{i}_1,\bar{i}_2}(k) + \Delta t u_{i_1,i_2}(k), \\ T_{i_1,i_2}(0) = T_0(i_1,i_2), \\ T_\Omega(k) = 0. \end{cases} \tag{1.2}
$$

where $\Delta t, \Delta\xi$ are the temporal and spatial discretisation steps, respectively, α is equal to $c\frac{\Delta t}{\Delta \xi^2}$ and $\mathcal{N}_{(i_1,i_2)}$ includes the four closest neighbours of node (i_1,i_2). This is a dynamic model (either in state-space form at or in difference equation format) that shows how local variables (temperatures) are modified by communication with its neighbours and the environment (input). Such dynamic models are local dynamic models, and the global model is a network between these local models. Their interaction can, for example, be displayed by considering the spatial-temporal impulse response of the dynamic systems. This is shown for two sets of model parameter values in Figure 1.5. Here, we restricted the spatial dimension to one in order to simplify the illustration.

In Figure 1.5, this spatial-temporal impulse response is displayed for a rod that is discretized spatially in 50 equidistant points (nodes). A non-zero initial condition is only provided at node 25 (while keeping the other initial conditions at zero).

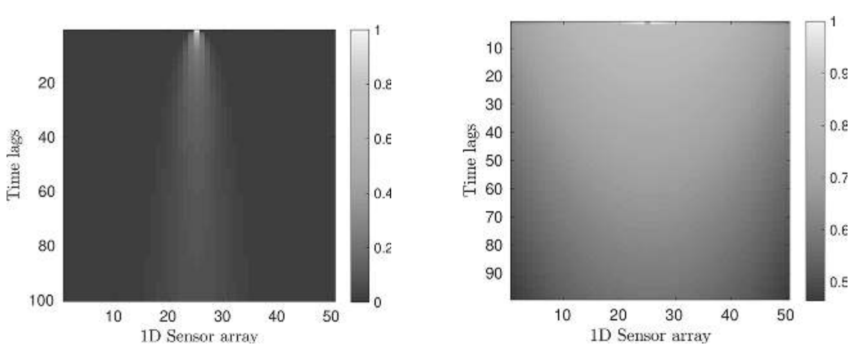

Figure 1.5 Spatial-temporal map for a one-dimensional string of subsystems $xi(k+1) = ax_i(k) + \alpha(x_{i-1}(k) + x_{i+1}(k))$, with (Left) $(a,\alpha) = (0.7, 0.15)$, and (Right) for a case with no spatial delay when propagating the state that corresponds to sequentially semi-separable systems that we will review in Chapter 3. Reprinted/adapted from Sinquin (2019) with permission from TU Delft

978-1-316-51570-9 — Data-Driven Identification of Networks of Dynamic Systems
Michel Verhaegen , Chengpu Yu , Baptiste Sinquin
Excerpt
More Information

From Figure 1.5, the spatial and temporal dynamics become clear. This figure illustrates how fast information travels from one node to the other. The further away in both time and space, the less it matters to the local dynamics. The larger the value in the map (indicated by the lighter colour), the more the state of the neighbour (possibly in the past) contributes to the temperature of the system at node 25. The further away in time, the more neighbouring nodes contribute to the value of the state of the system at node 25.

To make the spatial-temporal connectivity analysis of this example more systematic, use can be made of the so-called funnel causality for a spatially-invariant system (that is when all subsystems in the grid are identical). This notion was introduced in Bamieh and Voulgaris (2005) and is a function defined for every possible distance between two nodes that is equal to the first time at which a node is affected by a change of another node located at a given distance. This distribution in time is fixed for the PDE in Equation (1.2) and resembles a cone. As the model in Equation (1.2) belongs to the class of sequentially semi-separable (SSS) systems (Rice and Verhaegen, 2009), Figure 1.5 shows that the wider the displayed cone (like in the figure in the right), the further away the local states are instantaneously shared with all the neighbours in that cone. This notion of funnel causality is closely related to the parametrization of the system matrices, some of which are detailed in Chapter 3, for the signal flow approach, and Chapter 4 for the network of local LTI systems approach.

As a conclusion to this introduction, the computational limitations of system identification algorithms in handling large input-output datasets (e.g. from large sensor arrays) have spurred the analysis of alternatives for deriving algorithms with linear computational complexity with respect to the number of sensor measurements for data-driven control. It is then relevant to introduce further parametrization on the matrices that stems from physical insights, for example whether the system is deterministic or stochastic, which is strongly linked to whether the network topology is known or not, or whether it has invariances, such that the best trade-off between scalability and model accuracy is found.

1.4 Organization of the Book

The core part of this book consists of nine chapters in addition to this introduction and the conclusion. There are three main parts. The first part (Chapters 2–5) discusses the models. The second part (Chapters 6–9) discusses identification methods. The third part (Chapter 10) illustrates the potential of some of the methods using the challenging case study of large-scale adaptive optics.

The chapters dedicated to the models start with a refresher, in Chapter 2, of different generic models that have been used for the identification of (classical) lumped parameter systems. These include input-output transfer function models and state-space models. As a preparation for network systems, a variant of input-output models indicated in the literature as dynamic or structure function (DNF or DSF) models is also discussed. Though these models have been introduced for the purpose of modelling