

Distributed Kalman and Particle Filtering

1

Ali H. Sayed^{*}, Petar M. Djurić^{a,**} and F. Hlawatsch[†]

^{}Ecole Polytechnique Federale de Lausanne, EPFL, CH-1015 Lausanne; ^{**}Stony Brook University, Department of Electrical and Computer Engineering, Stony Brook, NY 11794, USA; [†]TU Wien, Institute of Telecommunications, 1040 Vienna, Austria*

^aCorresponding: petar.djuric@stonybrook.edu

ABSTRACT

This chapter discusses distributed Kalman and particle filtering algorithms for state estimation in decentralized multi-agent networks. It is assumed that the spatially distributed agents acquire local measurements with information about a time-varying state described by some underlying state-space model. The agents seek to estimate the time-varying state in a decentralized manner. They are only allowed to interact locally by sharing data or estimates with their immediate neighbors. It is shown how the agents can construct local estimates of the state trajectory through a cooperative process of interactions. Both diffusion- and consensus-based strategies are presented.

Keywords: distributed sequential estimation, distributed Kalman filtering, distributed particle filtering, diffusion, distributed proposal adaptation, likelihood consensus, target tracking, wireless agent network.

1.1 DISTRIBUTED SEQUENTIAL STATE ESTIMATION

Distributed strategies can be applied to solving state-space filtering and smoothing problems [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18]. In these applications, agents interconnected by communication links seek to track the state of some underlying state-space model based on their own measurements and information ex-

2 CHAPTER 1 Chapter Title

changed with their neighbors.

Many natural and man-made systems can be modeled as dynamic systems described by state-space models [19, 20, 21]. The *state* of a dynamic system comprises certain system-related quantities of interest, such as the location and velocity of a moving vehicle or the concentration of a pollutant in a chemical plume. The state is time-varying and unknown. Each agent acquires local measurements and benefits from sharing information with neighboring agents. In this chapter, we consider distributed time-sequential methods for estimating the state from recurring measurements acquired by spatially distributed agents, where each agent is able to communicate only with a limited set of spatially close neighboring agents [1, 2, 3, 4, 5, 6, 7, 8, 10, 11, 12, 13, 14, 15, 16, 17, 18].

The objective of this chapter is to provide an overview of distributed state-space estimation for linear and nonlinear models, with emphasis on filtering and prediction problems while noting that similar techniques can be applied to smoothing problems [12]. In the first part of our treatment, we motivate and examine the diffusion Kalman filter for one-step prediction by following the general approach outlined in [12, 15, 18]. Similar diffusion strategies for fixed-lag and fixed-point smoothing appear in [12]. We focus on the diffusion Kalman form because it has been shown in [15, 18, 22, 23] that diffusion implementations have superior stability and tracking performance when optimizing aggregate costs in response to streaming data. In the second part of the chapter, we describe distributed particle filtering, where we focus on both consensus-based and diffusion-based approaches. The methods we present can be used in a wide range of applications including surveillance, localization and navigation, environmental and agricultural monitoring, target tracking, exploration, search and rescue, the Internet of Things, and logistics.

1.1.1 THE SETUP

The problem of interest is the estimation of a state $\mathbf{x}_n \in \mathbb{R}^N$, where $n \in \{1, 2, \dots\}$ is a discrete time index, from sequences of measurements $\{\mathbf{z}_{1,k}, \mathbf{z}_{2,k}, \dots, \mathbf{z}_{n,k}\}$, with $\mathbf{z}_{i,k} \in \mathbb{R}^p$, observed up to the current time n at K different agents $k \in \{1, 2, \dots, K\}$. More specifically, we are interested in estimating \mathbf{x}_n in a *time-sequential* and *distributed* manner. Time-sequential processing means that the estimation of \mathbf{x}_n at time n recursively reuses relevant results from the estimation of \mathbf{x}_{n-1} at time $n-1$ while incorporating the new measurements $\mathbf{z}_{n,k}$, $k = 1, 2, \dots, K$. Distributed processing means that the estimation of \mathbf{x}_n is done locally at the individual agents, which exchange relevant information in a spatially local manner. Such a distributed (decentralized) mode of operation has important advantages over a centralized mode [17, 16]. For example, it does not require a fusion center collecting all the measurements, which would constitute a “single point of failure.” It also does not need communication between distant points or the use of complex routing strategies. The distributed solution is further robust to network node and link failures and is able to adapt to changing network topologies. Moreover, the computational complexity and communication cost per agent scale well with the network size (number of agents).

The topology of local information exchange is defined by a decentralized network of agents in which each agent k is able to communicate with a certain set of “neighbor” agents $\mathfrak{N}_k \subseteq \{1, 2, \dots, K\}$. This set \mathfrak{N}_k also includes agent k . Usually, the neighbors of an agent are located in its proximity. Practical examples of decentralized agent networks include robotic networks, networks of unmanned terrestrial, aerial or underwater vehicles, and networks of cameras. We assume that the agent network is connected, i.e., each agent is connected to any other agent via one or multiple “communication hops.”

Since the measurements are dispersed among the agents rather than available at a single processing unit, an important aspect of any distributed estimation algorithm is the dissemination of agent-related (and possibly other) information through the network via local communication. Various dissemination schemes are available, such as consensus [24, 25], gossip [26, 27], and diffusion [12, 15, 18]. A performance benchmark for a distributed algorithm is the performance that would be achieved by a centralized algorithm that has direct access to all the measurements. The amount of communication that is required for good performance is an important property of a distributed algorithm, since communication increases the power consumption and processing delay of the agents.

A summary of the notation used in this chapter is presented in Table 1.1.

1.2 DISTRIBUTED KALMAN FILTERING

In the first part of our treatment, we motivate and examine the diffusion Kalman filter for one-step prediction by following the general approach outlined in [12, 15, 18]. We start with a brief description of the network topology and an introduction of the data model.

1.2.1 NETWORK TOPOLOGY

We consider a network consisting of K agents connected according to some graph topology. A left-stochastic $K \times K$ matrix $A = (a_{\ell k})$ is associated with the topology, where each entry $a_{\ell k}$ denotes a nonnegative scaling factor that will be used to scale some of the data transitioning from agent ℓ to agent k (likewise for $a_{k\ell}$). The left-stochasticity of A means that it satisfies

$$A^T \mathbf{1}_K = \mathbf{1}_K, \quad (1.1)$$

so that the entries on each column of A add up to one (the symbol $\mathbf{1}_K$ represents a $K \times 1$ vector of ones). We assume that the graph is strongly-connected. The connectedness condition means that, for any two agents ℓ and k , there always exist paths with nonzero scaling weights in A linking ℓ and k in either direction; the path from ℓ to k does not need to agree with the reverse path from k to ℓ . Strong-connectedness additionally means that there exists at least one agent in the graph with $a_{kk} > 0$. That is,

4 CHAPTER 1 Chapter Title

Table 1.1 Summary of notation.

Notation	Description
$n \in \{1, 2, \dots\}$	discrete time index
$\mathbf{x}_n \in \mathbb{R}^N$	state vector at time n
$k \in \{1, 2, \dots, K\}$	agent index
K	number of agents
$\mathbf{z}_{n,k} \in \mathbb{R}^p$	measurement vector of agent k at time n
$\mathbf{z}_n \triangleq (\mathbf{z}_{n,1}^T \mathbf{z}_{n,2}^T \cdots \mathbf{z}_{n,K}^T)^T$	vector of all agent measurements at time n
$\mathbf{z}_{1:n} \triangleq (\mathbf{z}_1^T \mathbf{z}_2^T \cdots \mathbf{z}_n^T)^T$	all-agents measurement sequence up to time n
$\mathfrak{N}_k \subseteq \{1, 2, \dots, K\}$	set of neighboring agents of agent k (including agent k)
$d_k = \mathfrak{N}_k $	cardinality of set \mathfrak{N}_k
\mathbf{A}	a left stochastic matrix
\mathbf{J}	adjacency matrix
$\mathbf{1}_K$	$K \times 1$ vector with elements equal to 1
δ_{nj}	Kronecker delta function
\propto	equal up to a constant factor
\mathbf{F}_n and \mathbf{G}_n	state matrices of the linear state-space model
$\mathbf{H}_{n,k}$	local data matrix of the linear state-space model
\mathbf{Q}_n	covariance matrix of the state noise process
$\mathbf{R}_{n,k}$	covariance matrix of the observation noise process
$\mathbf{\Pi}_0$	covariance matrix of initial state
\mathbf{I}_N	$N \times N$ identity matrix
$\mathbf{g}_n(\mathbf{x}_{n-1}, \mathbf{u}_n)$	state-evolution function at time n
\mathbf{u}_n	state noise at time n
$f(\mathbf{x}_n \mathbf{x}_{n-1})$	state-transition probability density function
$\mathbf{h}_{n,k}(\mathbf{x}_n, \mathbf{v}_{n,k})$	observation function of agent k at time n
$\mathbf{v}_{n,k}$	observation noise of agent k at time n
$f(\mathbf{z}_{n,k} \mathbf{x}_n)$	local likelihood function of agent k
$f(\mathbf{z}_n \mathbf{x}_n)$	global (all-agents) likelihood function
$f(\mathbf{x}_n \mathbf{z}_{1:n})$	posterior probability density function
$\mathbb{E}[\cdot]$	expectation
$\mathbf{x}_n^{(m)}$	m th particle of the state \mathbf{x}_n
$w_n^{(m)}$	weight of the particle $\mathbf{x}_n^{(m)}$
M	number of particles
$q(\mathbf{x}_n \mathbf{x}_{n-1}^{(m)})$	importance probability density function
$\ln(\cdot)$	natural logarithm
$\mathcal{N}(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian probability density function

at least one agent places some trust or confidence on its local data. These conditions are satisfied by most graphs of interest. It follows from the strong-connectedness

condition that the left-stochastic matrix A is primitive [15].

1.2.2 LINEAR STATE-SPACE MODELS

We consider a linear state-space model defined by

$$\mathbf{x}_{n+1} = \mathbf{F}_n \mathbf{x}_n + \mathbf{G}_n \mathbf{u}_n, \quad n = 0, 1, \dots, \quad (1.2)$$

$$\mathbf{z}_{n,k} = \mathbf{H}_{n,k} \mathbf{x}_n + \mathbf{v}_{n,k}, \quad n = 0, 1, \dots; k = 1, 2, \dots, K, \quad (1.3)$$

where \mathbf{F}_n and \mathbf{G}_n are state matrices, and $\mathbf{H}_{n,k}$ is a local data matrix. The noise vectors \mathbf{u}_n and $\mathbf{v}_{n,k}$ are assumed to be realizations of zero-mean white processes with covariance matrices denoted by \mathbf{Q}_n and $\mathbf{R}_{n,k}$, respectively,

$$\mathbb{E} \left[\begin{pmatrix} \mathbf{u}_n \\ \mathbf{v}_{n,k} \end{pmatrix} \begin{pmatrix} \mathbf{u}_j \\ \mathbf{v}_{j,k} \end{pmatrix}^T \right] \triangleq \begin{pmatrix} \mathbf{Q}_n & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{n,k} \end{pmatrix} \delta_{nj}, \quad (1.4)$$

where δ_{nj} denotes the Kronecker delta: it is equal to one when $n = j$ and zero otherwise. The measurement noise vectors $\mathbf{v}_{n,k}$ are also assumed to be independent of each other across space (i.e., for different k), that is,

$$\mathbb{E} [\mathbf{v}_{n,k} \mathbf{v}_{j,\ell}^T] = \mathbf{R}_{n,k} \delta_{k\ell} \delta_{nj}. \quad (1.5)$$

The initial state vector, \mathbf{x}_0 , is assumed to have zero mean and covariance matrix

$$\mathbb{E} [\mathbf{x}_0 \mathbf{x}_0^T] = \mathbf{\Pi}_0 > \mathbf{0}, \quad (1.6)$$

and it is uncorrelated with \mathbf{u}_n and $\mathbf{v}_{n,k}$, for all n and k . The notation $\mathbf{\Pi}_0 > \mathbf{0}$ signifies that $\mathbf{\Pi}_0$ is a positive definite matrix. We further assume that $\mathbf{R}_{n,k} > \mathbf{0}$. The parameter matrices $\{\mathbf{F}_n, \mathbf{G}_n, \mathbf{Q}_n, \mathbf{\Pi}_0, \mathbf{H}_{n,k}, \mathbf{R}_{n,k}\}$ are considered known by each agent k . Note that no Gaussian assumptions are made.

1.2.3 NON-COOPERATIVE FILTERING

Assume initially that each agent k in the network acts individually and uses solely its local data stream $\mathbf{z}_{n,k}$, $n = 1, 2, \dots$, to track the state vector \mathbf{x}_n . In this case, agent k can run the well-known Kalman filter in any of its various forms (covariance, measurement and time-update, or information form) to carry out this task [19, 28, 29]. For instance, let $\widehat{\mathbf{x}}_{n,k|j}^{\text{ind}}$ denote the linear least-mean-squares error (LLMSE) estimate of \mathbf{x}_n by agent k using all the local measurements up to time j , i.e., $\{\mathbf{z}_{1,k}, \mathbf{z}_{2,k}, \dots, \mathbf{z}_{j,k}\}$. The superscript “ind” is used to emphasize that these estimators are based on individual (non-cooperative) behavior. We denote the corresponding estimation error and error-covariance matrix by

$$\widehat{\mathbf{x}}_{n,k|j}^{\text{ind}} \triangleq \mathbf{x}_n - \widehat{\mathbf{x}}_{n,k|j}^{\text{ind}}, \quad (1.7)$$

$$\mathbf{P}_{n,k|j}^{\text{ind}} \triangleq \mathbb{E} [\widehat{\mathbf{x}}_{n,k|j}^{\text{ind}} (\widehat{\mathbf{x}}_{n,k|j}^{\text{ind}})^T]. \quad (1.8)$$

6 CHAPTER 1 Chapter Title

Then, it is well known [28, 29] that predicted (when $j = n - 1$) and/or filtered (when $j = n$) state estimates can be computed by agent k by one of several forms, listed below.

Covariance form of the non-cooperative Kalman filter

start with $\widehat{\mathbf{x}}_{0,k|1}^{\text{ind}} = \mathbf{0}$, $\mathbf{P}_{0,k|1}^{\text{ind}} = \mathbf{\Pi}_0$.

repeat at every agent k for $n \geq 0$:

$$\begin{aligned} \mathbf{e}_{n,k}^{\text{ind}} &= \mathbf{z}_{n,k} - \mathbf{H}_{n,k} \widehat{\mathbf{x}}_{n,k|n-1}^{\text{ind}} \\ \mathbf{R}_{e,n,k}^{\text{ind}} &= \mathbf{R}_{n,k} + \mathbf{H}_{n,k} \mathbf{P}_{n,k|n-1}^{\text{ind}} \mathbf{H}_{n,k}^{\text{T}} \\ \mathbf{K}_{p,n,k}^{\text{ind}} &= \mathbf{F}_n \mathbf{P}_{n,k|n-1}^{\text{ind}} \mathbf{H}_{n,k}^{\text{T}} (\mathbf{R}_{e,n,k}^{\text{ind}})^{-1} \\ \widehat{\mathbf{x}}_{n+1,k|n}^{\text{ind}} &= \mathbf{F}_n \widehat{\mathbf{x}}_{n,k|n-1}^{\text{ind}} + \mathbf{K}_{p,n,k}^{\text{ind}} \mathbf{e}_{n,k}^{\text{ind}} \\ \mathbf{P}_{n+1,k|n}^{\text{ind}} &= \mathbf{F}_n \mathbf{P}_{n,k|n-1}^{\text{ind}} \mathbf{F}_n^{\text{T}} - \mathbf{K}_{p,n,k}^{\text{ind}} \mathbf{R}_{e,n,k}^{\text{ind}} (\mathbf{K}_{p,n,k}^{\text{ind}})^{\text{T}} + \mathbf{G}_n \mathbf{Q}_n \mathbf{G}_n^{\text{T}} \end{aligned} \quad (1.9)$$

end

Measurement and time-update form of the non-cooperative Kalman filter

start with $\widehat{\mathbf{x}}_{0,k|1}^{\text{ind}} = \mathbf{0}$, $\mathbf{P}_{0,k|1}^{\text{ind}} = \mathbf{\Pi}_0$.

repeat at every agent k for $n \geq 0$:

(measurement-update)

$$\begin{aligned} \mathbf{e}_{n,k}^{\text{ind}} &= \mathbf{z}_{n,k} - \mathbf{H}_{n,k} \widehat{\mathbf{x}}_{n,k|n-1}^{\text{ind}} \\ \mathbf{R}_{e,n,k}^{\text{ind}} &= \mathbf{R}_{n,k} + \mathbf{H}_{n,k} \mathbf{P}_{n,k|n-1}^{\text{ind}} \mathbf{H}_{n,k}^{\text{T}} \\ \widehat{\mathbf{x}}_{n,k|n}^{\text{ind}} &= \widehat{\mathbf{x}}_{n,k|n-1}^{\text{ind}} + \mathbf{P}_{n,k|n-1}^{\text{ind}} \mathbf{H}_{n,k}^{\text{T}} (\mathbf{R}_{e,n,k}^{\text{ind}})^{-1} \mathbf{e}_{n,k}^{\text{ind}} \\ \mathbf{P}_{n,k|n}^{\text{ind}} &= \mathbf{P}_{n,k|n-1}^{\text{ind}} - \mathbf{P}_{n,k|n-1}^{\text{ind}} \mathbf{H}_{n,k}^{\text{T}} (\mathbf{R}_{e,n,k}^{\text{ind}})^{-1} \mathbf{H}_{n,k} \mathbf{P}_{n,k|n-1}^{\text{ind}} \end{aligned} \quad (1.10)$$

(time-update)

$$\begin{aligned} \widehat{\mathbf{x}}_{n+1,k|n}^{\text{ind}} &= \mathbf{F}_n \widehat{\mathbf{x}}_{n,k|n}^{\text{ind}} \\ \mathbf{P}_{n+1,k|n}^{\text{ind}} &= \mathbf{F}_n \mathbf{P}_{n,k|n}^{\text{ind}} \mathbf{F}_n^{\text{T}} + \mathbf{G}_n \mathbf{Q}_n \mathbf{G}_n^{\text{T}} \end{aligned}$$

end

Information form of the non-cooperative Kalman filter

start with $\widehat{\mathbf{x}}_{k,0|1}^{\text{ind}} = \mathbf{0}$, $(\mathbf{P}_{k,0|1}^{\text{ind}})^{-1} = \mathbf{\Pi}_0^{-1}$.

repeat at every agent k for $n \geq 0$:

$$\begin{aligned} (\mathbf{P}_{n,k|n}^{\text{ind}})^{-1} &= (\mathbf{P}_{n,k|n-1}^{\text{ind}})^{-1} + \mathbf{H}_{n,k}^{\text{T}} \mathbf{R}_{n,k}^{-1} \mathbf{H}_{n,k} \\ (\mathbf{P}_{n,k|n}^{\text{ind}})^{-1} \widehat{\mathbf{x}}_{n,k|n}^{\text{ind}} &= (\mathbf{P}_{n,k|n-1}^{\text{ind}})^{-1} \widehat{\mathbf{x}}_{n,k|n-1}^{\text{ind}} + \mathbf{H}_{n,k}^{\text{T}} \mathbf{R}_{n,k}^{-1} \mathbf{z}_{n,k} \\ \widehat{\mathbf{x}}_{n,k+1|n}^{\text{ind}} &= \mathbf{F}_n \widehat{\mathbf{x}}_{n,k|n}^{\text{ind}} \\ \mathbf{P}_{n+1,k|n}^{\text{ind}} &= \mathbf{F}_n \mathbf{P}_{n,k|n}^{\text{ind}} \mathbf{F}_n^{\text{T}} + \mathbf{G}_n \mathbf{Q}_n \mathbf{G}_n^{\text{T}} \end{aligned} \quad (1.11)$$

end

Clearly, these non-cooperative forms assume that the agents act individually.

1.2 Distributed Kalman Filtering 7

However, since all the agents are tracking the same state vector \mathbf{x}_n , it is expected that the mean-square-error performance can improve through cooperation by having neighboring agents share with each other information about their own estimators. We consider one initial form of cooperation before extending it to arrive at the diffusion Kalman filter later in Subsection 1.2.7.

1.2.4 INCREMENTAL COOPERATION

Consider an arbitrary agent k and let \mathfrak{N}_k denote its neighborhood, i.e., the collection of all agents ℓ that are connected to agent k by edges and that can share information with it. This set includes also the agent k . The set comprises all agents ℓ for which $a_{\ell k} > 0$; its cardinality is denoted by $d_k = |\mathfrak{N}_k|$.

In addition to its own measurement vector $\mathbf{z}_{n,k}$, at each time instant n , agent k has also access to the measurement vectors $\{\mathbf{z}_{n,\ell}\}$ from its neighbors, $\ell \in \mathfrak{N}_k \setminus \{k\}$. In this way, at every time instant n , agent k has access to $|\mathfrak{N}_k|$ measurement vectors (its own measurements and the ones collected from its neighbors) and denoted by $\{\mathbf{z}_{n,k_1}, \mathbf{z}_{n,k_2}, \dots, \mathbf{z}_{n,k_{d_k}}\}$. In this notation, we are referring to the neighbors of agent k (including agent k itself) by the indices $\{k_1, k_2, \dots, k_{d_k}\}$. Let $\hat{\mathbf{x}}_{n,k|j}^{\text{loc}}$ denote the LLMSE estimate of \mathbf{x}_n obtained by agent k (including agent k) by using these local (or neighborhood) measurements from time 0 up to time j . The superscript “loc” is used to emphasize that these estimators are based on local neighborhood measurements. We denote the corresponding estimation error and error-covariance matrix by

$$\tilde{\mathbf{x}}_{n,k|j}^{\text{loc}} \triangleq \mathbf{x}_n - \hat{\mathbf{x}}_{n,k|j}^{\text{loc}}, \quad (1.12)$$

$$\mathbf{P}_{n,k|j}^{\text{loc}} \triangleq \mathbb{E} \left[\tilde{\mathbf{x}}_{n,k|j}^{\text{loc}} (\tilde{\mathbf{x}}_{n,k|j}^{\text{loc}})^{\text{T}} \right] \quad (1.13)$$

Then, as shown in listing (1.14), the predicted (when $j = n - 1$) and filtered (when $j = n$) state estimates can be computed by agent k by running multiple measurement-updates, one for each neighbor of k [28, 29]. In this listing, we are denoting the state estimate at the end of the incremental loop through the neighborhood by $\boldsymbol{\psi}_{n,k}$, which is equal to the filtered estimate $\hat{\mathbf{x}}_{n,k|n}^{\text{loc}}$ that results from the d_k measurement updates. This equality is highlighted by the symbol (\star) placed next to one of the equations in listing (1.14). The corresponding error covariance matrix of $\hat{\mathbf{x}}_{n,k|n}^{\text{loc}}$ is denoted by $\mathbf{P}_{n,k|n}^{\text{loc}}$. Although the variables $\boldsymbol{\psi}_{n,k}$ and $\hat{\mathbf{x}}_{n,k|n}^{\text{loc}}$ are identical in the incremental implementation, we keep separate notations for both symbols because they will be distinct and will play different roles in the diffusion form derived further ahead in Subsection 1.2.7. In particular, the (\star) assignment will be replaced by (1.48).

8 CHAPTER 1 Chapter Title

Measurement and time-update incremental form

start with $\widehat{\mathbf{x}}_{0,k|1}^{\text{loc}} = \mathbf{0}$, $\mathbf{P}_{0,k|1}^{\text{loc}} = \mathbf{\Pi}_0$.

repeat at every agent k for $n \geq 0$:

(incremental measurement-updates)

set auxiliary variables:

$$\boldsymbol{\psi}_{n,k}^{(0)} \leftarrow \widehat{\mathbf{x}}_{n,k|n-1}^{\text{loc}} \quad (N \times 1 \text{ vector})$$

$$\mathbf{P}_{n,k}^{(0)} \leftarrow \mathbf{P}_{n,k|n-1}^{\text{loc}} \quad (N \times N \text{ matrix})$$

for each neighbor $m = 1, 2, \dots, d_k$ **do**:

$$\mathbf{e}_{n,k_m} = \mathbf{z}_{n,k_m} - \mathbf{H}_{n,k_m} \boldsymbol{\psi}_{n,k}^{(m-1)}$$

$$\mathbf{R}_{e,n,k_m} = \mathbf{R}_{n,k_m} + \mathbf{H}_{n,k_m} \mathbf{P}_{n,k}^{(m-1)} \mathbf{H}_{n,k_m}^T$$

$$\boldsymbol{\psi}_{n,k}^{(m)} = \boldsymbol{\psi}_{n,k}^{(m-1)} + \mathbf{P}_{n,k}^{(m-1)} \mathbf{H}_{n,k_m}^T \mathbf{R}_{e,n,k_m}^{-1} \mathbf{e}_{n,k_m}$$

$$\mathbf{P}_{n,k}^{(m)} = \mathbf{P}_{n,k}^{(m-1)} - \mathbf{P}_{n,k}^{(m-1)} \mathbf{H}_{n,k_m}^T \mathbf{R}_{e,n,k_m}^{-1} \mathbf{H}_{n,k_m} \mathbf{P}_{n,k}^{(m-1)}$$

(1.14)

end

$$\mathbf{P}_{n,k|n}^{\text{loc}} \leftarrow \mathbf{P}_{n,k}^{(d_k)}$$

$$\boldsymbol{\psi}_{n,k} \leftarrow \boldsymbol{\psi}_{n,k}^{(d_k)}$$

$$\widehat{\mathbf{x}}_{n,k|n}^{\text{loc}} \leftarrow \boldsymbol{\psi}_{n,k}^{(d_k)} \quad (\star)$$

(time-update)

$$\widehat{\mathbf{x}}_{n+1,k|n}^{\text{loc}} = \mathbf{F}_n \widehat{\mathbf{x}}_{n,k|n}^{\text{loc}}$$

$$\mathbf{P}_{n+1,k|n}^{\text{loc}} = \mathbf{F}_n \mathbf{P}_{n,k|n}^{\text{loc}} \mathbf{F}_n^T + \mathbf{G}_n \mathbf{Q}_n \mathbf{G}_n^T$$

end

It can be verified that the following relations hold for the variables in the incremental form (1.14) — see the proof in [12, Appendix B]:

$$\left(\mathbf{P}_{n,k|n}^{\text{loc}}\right)^{-1} = \left(\mathbf{P}_{n,k|n-1}^{\text{loc}}\right)^{-1} + \mathbf{S}_{n,k}, \quad (1.15)$$

with

$$\mathbf{S}_{n,k} \triangleq \sum_{\ell \in \mathfrak{N}_k} \mathbf{H}_{n,\ell}^T \mathbf{R}_{n,\ell}^{-1} \mathbf{H}_{n,\ell}, \quad (1.16)$$

and

$$\left(\mathbf{P}_{n,k|n}^{\text{loc}}\right)^{-1} \widehat{\mathbf{x}}_{n,k|n}^{\text{loc}} = \left(\mathbf{P}_{n,k|n-1}^{\text{loc}}\right)^{-1} \widehat{\mathbf{x}}_{n,k|n-1}^{\text{loc}} + \mathbf{q}_{n,k}, \quad (1.17)$$

with

$$\mathbf{q}_{n,k} \triangleq \sum_{\ell \in \mathfrak{N}_k} \mathbf{H}_{n,\ell}^T \mathbf{R}_{n,\ell}^{-1} \mathbf{z}_{n,\ell}. \quad (1.18)$$

Recursions (1.14) compute the optimal local estimate $\widehat{\mathbf{x}}_{n,k|n}^{\text{loc}}$ of agent k by incorporating solely the measurements $\{\mathbf{z}_{n,\ell}\}$ from the neighborhood \mathfrak{N}_k . This is a limited form of cooperation since the recursions (1.14) are not exploiting the fact that be-

sides the measurements $\{z_{n,\ell}\}$, the neighbors of agent k also have their own estimates of x_n . These estimates are given by $\hat{x}_{n,\ell|n}^{\text{loc}}$. It is expected that by additionally exploiting the neighborhood estimates, agent k should be able to achieve two objectives: (a) enhance its state estimation accuracy and (b) have its local estimate $\hat{x}_{n,\ell|n}^{\text{loc}}$ approach the desired global estimate $\hat{x}_{n|n}$ of the state vector x_n (which would be produced by an optimum centralized estimator that has access to all the measurements in the network). We shall use this observation to motivate and derive the diffusion Kalman filter. First, however, we take a brief digression and review a well-known data fusion construction [28, 29], which will be used to form *approximate* local estimates.

1.2.5 DATA FUSION

A problem that is encountered often in applications deals with the need to fuse together data collected from several sources in order to enhance the accuracy of the estimation process. This is similar to the scenario we are facing in the network context, except that the graph topology and the time evolution of the data add a new level of complexity in the form of cross-correlations.

Thus, assume for a moment that we have a collection of K agents that are distributed over some region in space without paying particular attention to whether they are connected by a communication link or not. All agents are interested in estimating the same parameter vector of size $N \times 1$, denoted generically by \mathbf{x} , which is assumed to be zero-mean and to have a positive definite covariance matrix $\mathbf{\Pi}_x = \mathbb{E}[\mathbf{x}\mathbf{x}^T]$. For example, the agents could be tracking a moving object with the goal of estimating its speed and direction of motion. Each agent k collects a measurement vector \mathbf{z}_k of size $p \times 1$ that is related to the desired \mathbf{x} via the linear measurement model (observe that in the discussion in this section we are not dealing with sequential processing of the data and, therefore, the time index is dropped),

$$\mathbf{z}_k = \mathbf{H}_k \mathbf{x} + \mathbf{v}_k, \quad (1.19)$$

where \mathbf{H}_k is the model matrix that maps \mathbf{x} to \mathbf{z}_k at agent k and \mathbf{v}_k is zero-mean measurement noise with a positive definite covariance matrix $\mathbf{R}_k = \mathbb{E}[\mathbf{v}_k \mathbf{v}_k^T]$.

Assume initially that the measurements from all agents can be collected centrally at some fusion center. For example, each agent k could transmit its measurement vector \mathbf{z}_k and its model parameters $\{\mathbf{H}_k, \mathbf{R}_k\}$ to the fusion center. The data collected at the fusion center thus satisfy the linear model:

$$\underbrace{\begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \\ \vdots \\ \mathbf{z}_K \end{pmatrix}}_{\triangleq \mathbf{z}} = \underbrace{\begin{pmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \vdots \\ \mathbf{H}_K \end{pmatrix}}_{\triangleq \mathbf{H}} \mathbf{x} + \underbrace{\begin{pmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \vdots \\ \mathbf{v}_K \end{pmatrix}}_{\triangleq \mathbf{v}}. \quad (1.20)$$

We assume that the noises $\{\mathbf{v}_k\}$ across all agents k are *uncorrelated* with each other

10 CHAPTER 1 Chapter Title

so that the covariance matrix of the aggregate noise vector \mathbf{v} is block diagonal and given by

$$\mathbf{R}_v = \text{blkdiag}\{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_K\}. \quad (1.21)$$

Now, it is well known that the LLMSE estimate of \mathbf{x} that is based on \mathbf{z} is given by [28, 29]

$$\widehat{\mathbf{x}} = \left(\mathbf{\Pi}_x^{-1} + \mathbf{H}^T \mathbf{R}_v^{-1} \mathbf{H} \right)^{-1} \mathbf{H}^T \mathbf{R}_v^{-1} \mathbf{z}, \quad (1.22)$$

with the resulting minimum mean-square-error (MMSE) matrix equal to

$$\mathbf{P} \triangleq \mathbb{E}[\widehat{\mathbf{x}} \widehat{\mathbf{x}}^T] = \left(\mathbf{\Pi}_x^{-1} + \mathbf{H}^T \mathbf{R}_v^{-1} \mathbf{H} \right)^{-1}, \quad (1.23)$$

where $\widehat{\mathbf{x}} = \mathbf{x} - \widehat{\mathbf{x}}$ denotes the estimation error. Note that \mathbf{P} is a matrix of size $N \times N$.

The solution (1.22) requires all agents to transmit their *raw* data $\{\mathbf{z}_k, \mathbf{H}_k, \mathbf{R}_k\}$ to the fusion center. A more efficient fusion method, one that reduces the communications overhead, can be derived by allowing the agents to perform some local processing and to share the results of this processing step rather than their raw data. Specifically, assume that each agent estimates \mathbf{x} using first its own data \mathbf{z}_k . We denote the resulting estimator by $\widehat{\mathbf{x}}_k$ and it is given by

$$\widehat{\mathbf{x}}_k = \left(\mathbf{\Pi}_x^{-1} + \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k \right)^{-1} \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{z}_k. \quad (1.24)$$

The corresponding MMSE is

$$\mathbf{P}_k = \left(\mathbf{\Pi}_x^{-1} + \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k \right)^{-1}. \quad (1.25)$$

Next, assume that the agents share the processed data $\{\widehat{\mathbf{x}}_k, \mathbf{P}_k\}$ with the fusion center instead of the raw data $\{\mathbf{z}_k, \mathbf{H}_k, \mathbf{R}_k\}$. We now show that the desired global quantities $\{\widehat{\mathbf{x}}, \mathbf{P}\}$ can be recovered directly from the locally generated data $\{\widehat{\mathbf{x}}_k, \mathbf{P}_k\}$.

To begin with, observe that we can rework expression (1.23) for the global MMSE matrix as follows:

$$\mathbf{P}^{-1} = \mathbf{\Pi}_x^{-1} + \sum_{k=1}^K \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k = \sum_{k=1}^K \mathbf{P}_k^{-1} - (K-1) \mathbf{\Pi}_x^{-1}. \quad (1.26)$$

This expression allows the fusion center to determine \mathbf{P}^{-1} directly from knowledge of the quantities $\{\mathbf{P}_k^{-1}, \mathbf{\Pi}_x^{-1}\}$. Note further that the global estimate expression (1.22) can be rewritten as

$$\mathbf{P}^{-1} \widehat{\mathbf{x}} = \mathbf{H}^T \mathbf{R}_v^{-1} \mathbf{z} = \sum_{k=1}^K \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{z}_k = \sum_{k=1}^K \mathbf{P}_k^{-1} \widehat{\mathbf{x}}_k. \quad (1.27)$$

Therefore, we arrive at the following alternative method to fuse the data from multi-

ple agents:

$$\mathbf{P}^{-1} = \sum_{k=1}^K \mathbf{P}_k^{-1} - (K-1)\mathbf{\Pi}_x^{-1}, \quad (1.28)$$

$$\mathbf{P}^{-1}\widehat{\mathbf{x}} = \sum_{k=1}^K \mathbf{P}_k^{-1}\widehat{\mathbf{x}}_k. \quad (1.29)$$

Observe that the individual estimates are scaled by the inverses of their MMSE matrices. It is useful to note that the derivation of these expressions requires the noises across the agents to be uncorrelated with each other, so that \mathbf{R}_v is block diagonal.

1.2.6 APPROXIMATE FUSION RELATIONS

Let us now return to the network context and use the fusion expressions (1.28) and (1.29) to motivate the diffusion Kalman filter. Let $\widetilde{\mathbf{x}}_{n|n}$ denote the filtered estimate of \mathbf{x}_n that is based on the observations $\mathbf{z}_{j,k}$ from across all agents $k = 1, 2, \dots, K$ at all times j up to time n . Let $\mathbf{P}_{n|n}$ denote the corresponding error covariance matrix, i.e.,

$$\mathbf{P}_{n|n} \triangleq \mathbb{E}[\widetilde{\mathbf{x}}_{n|n}\widetilde{\mathbf{x}}_{n|n}^T], \quad \text{where } \widetilde{\mathbf{x}}_{n|n} \triangleq \mathbf{x}_n - \widehat{\mathbf{x}}_{n|n}. \quad (1.30)$$

These are the global quantities that we are interested in evaluating in a distributed manner. We can employ the fusion expressions (1.28) and (1.29) to approximate these global variables in terms of the non-cooperative variables $\{\widehat{\mathbf{x}}_{n,k|n}^{\text{ind}}, \mathbf{P}_{n,k|n}^{\text{ind}}\}$ as follows:

$$\mathbf{P}_{n|n}^{-1} \approx \sum_{k=1}^K (\mathbf{P}_{n,k|n}^{\text{ind}})^{-1} - (K-1)\mathbf{\Pi}_n^{-1}, \quad (1.31)$$

$$\mathbf{P}_{n|n}^{-1}\widehat{\mathbf{x}}_{n|n} \approx \sum_{k=1}^K (\mathbf{P}_{n,k|n}^{\text{ind}})^{-1}\widehat{\mathbf{x}}_{n,k|n}^{\text{ind}}, \quad (1.32)$$

where $\mathbf{\Pi}_n \triangleq \mathbb{E}[\mathbf{x}_n\mathbf{x}_n^T]$ denotes the covariance matrix of \mathbf{x}_n , which satisfies the recursion

$$\mathbf{\Pi}_{n+1} = \mathbf{F}_n\mathbf{\Pi}_n\mathbf{F}_n^T + \mathbf{G}_n\mathbf{Q}_n\mathbf{G}_n^T. \quad (1.33)$$

Relations (1.31) and (1.32) are approximate expressions, while relations (1.28) and (1.29) are exact fusion formulae. To see why, recall that the global estimate $\widehat{\mathbf{x}}_{n|n}$ is based on the observations $\mathbf{z}_{j,k}$ from time $j = 0$ up to time $j = n$ across all agents k . If we now appeal to the state-space model described by (1.2) and (1.3), it is easy to recognize that these observations cannot generally be related directly to the unknown \mathbf{x}_n via a linear model of the form (1.19). Exact fusion formulae in this case are cumbersome and require the use of error cross-covariance matrices (see, e.g., [30, 31, 32]). The above approximations, however, are sufficient for our purposes and they are commonly used in the literature on data fusion methods, including the

12 CHAPTER 1 Chapter Title

following simpler form:

$$\mathbf{P}_{n|n}^{-1} \approx \sum_{k=1}^K (\mathbf{P}_{k,n|n}^{\text{ind}})^{-1}, \quad (1.34)$$

where the term involving $\mathbf{\Pi}_n^{-1}$ is ignored. Expression (1.32) amounts to estimating $\widehat{\mathbf{x}}_{n|n}$ by using a covariance-weighted combination of the local estimates, $\widehat{\mathbf{x}}_{n,k|n}^{\text{ind}}$, while expression (1.34) estimates the error covariance matrix by pooling together the individual error covariances.

In principle, the same fusion expressions (1.28) and (1.29) could also be used to approximate the same global variables $\{\widehat{\mathbf{x}}_{n|n}, \mathbf{P}_{n|n}\}$ in terms of the incremental (neighborhood) estimates $\{\boldsymbol{\psi}_{n,k}, \mathbf{P}_{n,k|n}^{\text{loc}}\}$ defined in Subsection 1.2.4, say as:

$$\mathbf{P}_{n|n}^{-1} \widehat{\mathbf{x}}_{n|n} \approx \sum_{k=1}^K (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \boldsymbol{\psi}_{n,k}, \quad (1.35)$$

$$\mathbf{P}_{n|n}^{-1} \approx \sum_{k=1}^K (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1}, \quad (1.36)$$

where we recall that $\boldsymbol{\psi}_{n,k} = \widehat{\mathbf{x}}_{n,k|n}^{\text{loc}}$. However, observe now that the construction of the estimates $\boldsymbol{\psi}_{n,k}$ across agents relies on shared observations. For example, the observations from both agent k and its neighbors influence the value of $\boldsymbol{\psi}_{n,k}$; likewise, some of these same observations influence the value of other estimates because they belong to the neighborhoods of other agents as well. This data redundancy is not present in the computation of the individual estimates $\widehat{\mathbf{x}}_{n,k|n}^{\text{ind}}$, where each agent k relies solely on its local sequential data $\mathbf{z}_{n,k}$.

The presence of redundant information in the intermediate estimates $\boldsymbol{\psi}_{n,k}$ is due to the graph topology, which defines the neighborhoods over the network. We can exploit this topology to replace the approximations (1.35) and (1.36) by more revealing relations that bring forth the graph structure. To do so, we let \mathbf{J} denote the $K \times K$ adjacency matrix of the network graph. This matrix consists of unit and zero entries representing the connectivity of the agents, namely,

$$J_{\ell k} = \begin{cases} 1, & \text{if } a_{\ell k} > 0 \\ 0, & \text{otherwise} \end{cases}. \quad (1.37)$$

Although unnecessary in the final statement of the algorithm, we shall assume for the sake of the derivation that there exists a vector \mathbf{s} such that $\mathbf{J}\mathbf{s} = \mathbf{1}_K$. For example, if \mathbf{J} happens to be invertible, then this vector \mathbf{s} exists, is unique, and is given by $\mathbf{s} = \mathbf{J}^{-1}\mathbf{1}_K$. In general, however, it is not always guaranteed that a vector \mathbf{s} satisfying $\mathbf{J}\mathbf{s} = \mathbf{1}_K$ exists. It is nevertheless possible to verify that the matrix \mathbf{J} can be made invertible by flipping some of its diagonal entries (from zero to one and from one to zero). By turning a diagonal entry of \mathbf{J} from zero to one, we are in effect allowing the agent at that location to rely on its local measurement. On the other hand, by turning a diagonal entry of \mathbf{J} from one to zero, we are disabling this feature. This

1.2 Distributed Kalman Filtering 13

action is tolerable because this measurement is not discarded by the network since it is still used by the neighbors of the agent to update their estimates (as long as at least one entry of the row of \mathbf{J} with a diagonal entry equal to zero has a value equal to one). We denote the entries of \mathbf{s} by γ_k , $k = 1, 2, \dots, K$, or $\mathbf{s} = \text{col}\{\gamma_k\}$ so that

$$\mathbf{J}\mathbf{s} = \mathbb{1}_K \iff \sum_{k=1}^K \gamma_k J_{\ell k} = 1, \text{ for every } \ell = 1, 2, \dots, K. \quad (1.38)$$

In a manner similar to (1.32) and (1.34), we first use the fusion expressions (1.28) and (1.29) to approximate the incremental estimators $\{\boldsymbol{\psi}_{n,k}, \mathbf{P}_{n,k|n}^{\text{loc}}\}$ by using the non-cooperative variables $\{\widehat{\mathbf{x}}_{n,k|n}^{\text{ind}}, \mathbf{P}_{n,k|n}^{\text{ind}}\}$ as follows:

$$(\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \boldsymbol{\psi}_{n,k} \approx \sum_{\ell=1}^K J_{\ell k} (\mathbf{P}_{n,\ell|n}^{\text{ind}})^{-1} \widehat{\mathbf{x}}_{n,\ell|n}^{\text{ind}}, \quad (1.39)$$

$$(\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \approx \sum_{\ell=1}^K J_{\ell k} (\mathbf{P}_{n,\ell|n}^{\text{ind}})^{-1}. \quad (1.40)$$

Recall that the individual estimates $\widehat{\mathbf{x}}_{n,k|n}^{\text{ind}}$ rely on independent streams of data and that $J_{\ell k} = 1$ if information can flow from agent ℓ to agent k , and $J_{\ell k} = 0$ otherwise.

Using the above relations we can now relate the global variables $\{\widehat{\mathbf{x}}_{n|n}, \mathbf{P}_{n|n}\}$ to the local variables $\{\boldsymbol{\psi}_{n,k}, \mathbf{P}_{n,k|n}^{\text{loc}}\}$ and replace (1.35) and (1.36). Indeed, using the entries γ_k , we note that

$$\begin{aligned} \sum_{k=1}^K \gamma_k (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \boldsymbol{\psi}_{n,k} &\stackrel{(1.39)}{\approx} \sum_{\ell=1}^K \underbrace{\left(\sum_{k=1}^K \gamma_k J_{\ell k} \right)}_{=1} (\mathbf{P}_{n,\ell|n}^{\text{ind}})^{-1} \widehat{\mathbf{x}}_{n,\ell|n}^{\text{ind}} \\ &= \sum_{\ell=1}^K (\mathbf{P}_{n,\ell|n}^{\text{ind}})^{-1} \widehat{\mathbf{x}}_{n,\ell|n}^{\text{ind}} \stackrel{(1.32)}{\approx} \mathbf{P}_{n|n}^{-1} \widehat{\mathbf{x}}_{n|n} \end{aligned} \quad (1.41)$$

and, similarly,

$$\sum_{k=1}^K \gamma_k (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \stackrel{(1.40)}{\approx} \sum_{\ell=1}^K (\mathbf{P}_{n,\ell|n}^{\text{ind}})^{-1} \stackrel{(1.34)}{\approx} \mathbf{P}_{n|n}^{-1}. \quad (1.42)$$

In summary, we arrive at the following approximate relations between the global variables $\{\widehat{\mathbf{x}}_{n|n}, \mathbf{P}_{n|n}\}$ and the neighborhood variables $\{\boldsymbol{\psi}_{n,k}, \mathbf{P}_{n,k|n}^{\text{loc}}\}$:

$$\mathbf{P}_{n|n}^{-1} \widehat{\mathbf{x}}_{n|n} \approx \sum_{k=1}^K \gamma_k (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \boldsymbol{\psi}_{n,k}, \quad (1.43)$$

$$\mathbf{P}_{n|n}^{-1} \approx \sum_{k=1}^K \gamma_k (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1}. \quad (1.44)$$

These relations replace (1.35) and (1.36), where the graph structure is represented

14 CHAPTER 1 Chapter Title

through the scalars γ_k .

Remark (Covariance-intersection method). Expressions (1.43) and (1.44) have a form similar to another commonly used fusion technique that relies on the use of the covariance-intersection (CI) method [33, 34] and which avoids the need for error cross-covariance matrices. In the CI method, the global variables $\{\widehat{\mathbf{x}}_{n|n}, \mathbf{P}_{n|n}\}$ are estimated according to (1.43) and (1.44) where the scalars γ_k are instead treated as *design* parameters. Specifically, they are chosen as nonnegative convex combination coefficients satisfying $\sum_{k=1}^K \gamma_k = 1, \gamma_k \geq 0$, and their values are selected to optimize the resulting $\mathbf{P}_{n|n}$, say, by minimizing the trace or determinant of $\mathbf{P}_{n|n}$, i.e.,

$$\min_{\{\gamma_k\}} \text{Tr}(\mathbf{P}_{n|n}) \quad \text{or} \quad \min_{\{\gamma_k\}} \det(\mathbf{P}_{n|n}). \quad (1.45)$$

1.2.7 DIFFUSION COOPERATION

Continuing with (1.43) and (1.44), one difficulty with these expressions is that computation of the global estimator $\widehat{\mathbf{x}}_{n|n}$ requires access to the local estimators $\{\psi_{n,\ell}\}$ from across the entire network. We can obtain a decentralized solution that relies solely on local interactions as follows. Substituting (1.44) into (1.43) gives

$$\widehat{\mathbf{x}}_{n|n} \approx \left[\sum_{k=1}^K \gamma_k (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \right]^{-1} \sum_{k=1}^K \gamma_k (\mathbf{P}_{n,k|n}^{\text{loc}})^{-1} \psi_{n,k}, \quad (1.46)$$

which has the form of a convex weighted average, namely,

$$\widehat{\mathbf{x}}_{n|n} \approx \sum_{k=1}^K \mathbf{\Gamma}_k \psi_{n,k} \quad (1.47)$$

with nonnegative-definite coefficient matrices $\mathbf{\Gamma}_k$ that add up to the identity matrix, i.e., $\sum_{k=1}^K \mathbf{\Gamma}_k = \mathbf{I}_N$. The result (1.47) suggests one useful approximation by which the local estimators $\{\psi_{n,\ell}\}$ can be fused within the neighborhood of every agent k to obtain a local version for $\widehat{\mathbf{x}}_{n|n}$, which we shall denote by $\widehat{\mathbf{x}}_{n,k|n}$ (we are removing the “loc” superscript). This local version is obtained by limiting the convex combination operation (1.47) to the neighborhood of each agent, and by replacing the matrix combination weights $\mathbf{\Gamma}_\ell$ by convex combination scalars $a_{\ell k}$. In this way, the computation (1.47) is replaced locally by

$$\widehat{\mathbf{x}}_{n,k|n} \leftarrow \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \psi_{n,\ell}. \quad (1.48)$$

Comparing with the incremental listing (1.14), we see that the above calculation amounts to replacing the assignment $\widehat{\mathbf{x}}_{n,k|n}^{\text{loc}} \leftarrow \psi_{n,k}$ (marked by \star) in (1.14) by a local convex combination step. In view of this substitution, the error covariances of the local estimators $\widehat{\mathbf{x}}_{n,k|n}$ computed by (1.48) are not given anymore by the matrices $\mathbf{P}_{n,k|n}$ in the listings (1.49) and (1.52) below (where again we removed the superscript “loc”). In summary, we arrive at the listings (1.49) and (1.52) for the diffusion

1.2 Distributed Kalman Filtering 15

Kalman filter in two equivalent forms: the measurement and time-update form and the information form.

More generally, and in view of (1.47), a more enhanced fusion of the local estimators $\{\psi_{n,\ell}\}$ is possible by employing convex combination matrices in (1.48) rather than scalars; for example, these combination matrices could be defined in terms of the inverses $\mathbf{P}_{n,\ell|n}^{-1}$ as suggested by (1.46). This construction, however, would entail added complexity and would require sharing of additional information regarding these inverses. The implementation (1.49) from [12] employs scalar combination coefficients $\{a_{\ell k}\}$ in order to reduce the complexity of the resulting algorithm. Reference [35] studies the alternative fusion of the estimators $\{\psi_{n,\ell}\}$ in the diffusion Kalman filter by exploiting information about the inverses $\mathbf{P}_{n,\ell|n}^{-1}$ and using covariance-intersection combinations over neighborhoods in a manner similar to (1.45).

Diffusion Kalman filter (measurement and time-update form)

start with $\widehat{\mathbf{x}}_{0,k|k-1} = \mathbf{0}$, $\mathbf{P}_{0,k|k-1} = \mathbf{\Pi}_0$.

repeat at every agent k for $n \geq 0$:

(measurement-updates)

set auxiliary variables:

$$\psi_{n,k}^{(0)} \leftarrow \widehat{\mathbf{x}}_{n,k|n-1} \quad (N \times 1 \text{ vector})$$

$$\mathbf{P}_{n,k}^{(0)} \leftarrow \mathbf{P}_{n,k|n-1} \quad (N \times N \text{ matrix})$$

for each neighbor $\ell = 1, 2, \dots, d_k$ do:

$$\mathbf{e}_{n,k\ell} = \mathbf{z}_{n,k\ell} - \mathbf{H}_{n,k\ell} \psi_{n,k}^{(\ell-1)}$$

$$\mathbf{R}_{e,n,k\ell} = \mathbf{R}_{n,k\ell} + \mathbf{H}_{n,k\ell} \mathbf{P}_{n,k}^{(\ell-1)} \mathbf{H}_{n,k\ell}^T$$

$$\psi_{n,k}^{(\ell)} = \psi_{n,k}^{(\ell-1)} + \mathbf{P}_{n,k}^{(\ell-1)} \mathbf{H}_{n,k\ell}^T \mathbf{R}_{e,n,k\ell}^{-1} \mathbf{e}_{n,k\ell}$$

$$\mathbf{P}_{n,k}^{(\ell)} = \mathbf{P}_{n,k}^{(\ell-1)} - \mathbf{P}_{n,k}^{(\ell-1)} \mathbf{H}_{n,k\ell}^T \mathbf{R}_{e,n,k\ell}^{-1} \mathbf{H}_{n,k\ell} \mathbf{P}_{n,k}^{(\ell-1)}$$

(1.49)

end

$$\mathbf{P}_{n,k|n} \leftarrow \mathbf{P}_{n,k}^{(d_k)}$$

$$\psi_{n,k} \leftarrow \psi_{n,k}^{(d_k)}$$

(combination step)

$$\widehat{\mathbf{x}}_{n,k|n} = \sum_{\ell \in \mathfrak{N}_k} a_{\ell k} \psi_{\ell,n}$$

(time-update)

$$\widehat{\mathbf{x}}_{n+1,k|n} = \mathbf{F}_n \widehat{\mathbf{x}}_{n,k|n}$$

$$\mathbf{P}_{n+1,k|n} = \mathbf{F}_n \mathbf{P}_{k,n|n} \mathbf{F}_n^T + \mathbf{G}_n \mathbf{Q}_n \mathbf{G}_n^T$$

end

We also note that the combination step in (1.52) is similar to the update used in the consensus Kalman filter derived in [5] with one main difference. The consensus Kalman filter employs a specific combination construction at the local level of the

16 CHAPTER 1 Chapter Title

form:

$$\widehat{\mathbf{x}}_{n,k|n} = (1 + (1 - d_k)\epsilon)\boldsymbol{\psi}_{n,k} + \sum_{\ell \in \mathfrak{N}_k \setminus \{k\}} \epsilon \boldsymbol{\psi}_{n,\ell}, \quad (1.50)$$

where we recall that $d_k = |\mathfrak{N}_k|$, and $\epsilon > 0$ is a small weight that is *equal* for all neighbors. In comparison, the diffusion implementation (1.52), which can be written as

$$\widehat{\mathbf{x}}_{n,k|n} = a_{kk}\boldsymbol{\psi}_{n,k} + \sum_{\ell \in \mathfrak{N}_k \setminus \{k\}} a_{\ell k}\boldsymbol{\psi}_{n,\ell}, \quad (1.51)$$

allows more freedom in assigning generally different weights $a_{\ell k}$ to different neighbors [10, 12].

Diffusion Kalman filter (information form)

start with $\widehat{\mathbf{x}}_{0,k|-1} = \mathbf{0}$, $(\mathbf{P}_{0,k|-1})^{-1} = \mathbf{\Pi}_0^{-1}$.

repeat at every agent k for $n \geq 0$:

$$\mathbf{S}_{n,k} = \sum_{\ell \in \mathfrak{N}_k} \mathbf{H}_{n,\ell}^T \mathbf{R}_{n,\ell}^{-1} \mathbf{H}_{n,\ell}$$

$$\mathbf{q}_{n,k} = \sum_{\ell \in \mathfrak{N}_k} \mathbf{H}_{n,\ell}^T \mathbf{R}_{n,\ell}^{-1} \mathbf{z}_{n,\ell}$$

$$(\mathbf{P}_{n,k|n})^{-1} = (\mathbf{P}_{n,k|n-1})^{-1} + \mathbf{S}_{n,k} \quad (1.52)$$

$$\boldsymbol{\psi}_{n,k} = \widehat{\mathbf{x}}_{n,k|n-1} + \mathbf{P}_{n,k|n}(\mathbf{q}_{n,k} - \mathbf{S}_{n,k}\widehat{\mathbf{x}}_{n,k|n-1})$$

$$\widehat{\mathbf{x}}_{n,k|n} = \sum_{\ell \in \mathfrak{N}_k} a_{\ell k}\boldsymbol{\psi}_{n,\ell}$$

$$\widehat{\mathbf{x}}_{n+1,k|n} = \mathbf{F}_n \widehat{\mathbf{x}}_{n,k|n}$$

$$\mathbf{P}_{n+1,k|n} = \mathbf{F}_n \mathbf{P}_{n,k|n} \mathbf{F}_n^T + \mathbf{G}_n \mathbf{Q}_n \mathbf{G}_n^T$$

end

1.2.8 PERFORMANCE ANALYSIS

In view of the approximations (1.44) and (1.48), it is important to examine the mean-square error performance of the resulting diffusion Kalman filter. In this section, we examine how close the local estimates $\{\widehat{\mathbf{x}}_{n,k|n}\}$ get to the state variable \mathbf{x}_n by evaluating the size of the mean-square error in steady-state after the filter has had sufficient time to learn.

We start by collecting all state estimation errors from across the network into an extended $K \times 1$ block vector (whose individual entries are of size $N \times 1$ each):

$$\widetilde{\mathbf{X}}_{n|n} \triangleq \begin{pmatrix} \widetilde{\mathbf{x}}_{1,n|n} \\ \widetilde{\mathbf{x}}_{2,n|n} \\ \vdots \\ \widetilde{\mathbf{x}}_{K,n|n} \end{pmatrix}, \quad \text{where } \widetilde{\mathbf{x}}_{n,k|n} \triangleq \mathbf{x}_n - \widehat{\mathbf{x}}_{n,k|n}, \quad (1.53)$$

and introduce supporting block-diagonal and Kronecker product matrices:

$$\mathbf{U}_{n-1} \triangleq \mathbf{1}_K \otimes \mathbf{u}_{n-1}, \quad (1.54)$$

$$\mathbf{V}_n \triangleq \text{blkcol}\{\mathbf{v}_{n,1|n}, \mathbf{v}_{n,2|n}, \dots, \mathbf{v}_{n,K|n}\}, \quad (1.55)$$

$$\mathcal{P}_{n|n} \triangleq \text{blkdiag}\{\mathbf{P}_{n,1|n}, \mathbf{P}_{n,2|n}, \dots, \mathbf{P}_{n,K|n}\}, \quad (1.56)$$

$$\mathcal{H}_n \triangleq \text{blkdiag}\{\mathbf{H}_{n,1}, \mathbf{H}_{n,2}, \dots, \mathbf{H}_{n,K}\}, \quad (1.57)$$

$$\mathcal{S}_n \triangleq \text{blkdiag}\{\mathbf{S}_{n,1}, \mathbf{S}_{n,2}, \dots, \mathbf{S}_{n,K}\}, \quad (1.58)$$

$$\mathcal{R}_n \triangleq \text{blkdiag}\{\mathbf{R}_{n,1}, \mathbf{R}_{n,2}, \dots, \mathbf{R}_{n,K}\}, \quad (1.59)$$

$$\mathcal{J} \triangleq \mathbf{J} \otimes \mathbf{I}_N, \quad (1.60)$$

$$\mathcal{A} \triangleq \mathbf{A} \otimes \mathbf{I}_N, \quad (1.61)$$

where \otimes denotes the Kronecker product operation. Subtracting \mathbf{x}_n from both sides of the expression for $\boldsymbol{\psi}_{n,k}$ in (1.52), and using $\mathbf{z}_{n,\ell} = \mathbf{H}_{n,\ell}\mathbf{x}_n + \mathbf{v}_{n,\ell}$, gives

$$\begin{aligned} \tilde{\boldsymbol{\psi}}_{n,k} &= \tilde{\mathbf{x}}_{n,k|n-1} - \mathbf{P}_{n,k|n}(\mathbf{q}_{n,k} - \mathbf{S}_{n,k}\tilde{\mathbf{x}}_{n,k|n-1}) \\ &= \tilde{\mathbf{x}}_{n,k|n-1} - \mathbf{P}_{n,k|n} \sum_{\ell \in \mathfrak{Y}_k} \mathbf{H}_{n,\ell}^T \mathbf{R}_{n,\ell}^{-1} (\mathbf{z}_{n,\ell} - \mathbf{H}_{n,\ell}\tilde{\mathbf{x}}_{n,k|n-1}) \\ &= (\mathbf{I}_N - \mathbf{P}_{n,k|n}\mathbf{S}_{n,k})\tilde{\mathbf{x}}_{n,k|n-1} - \mathbf{P}_{n,k|n} \sum_{\ell \in \mathfrak{Y}_k} \mathbf{H}_{n,\ell}^T \mathbf{R}_{n,\ell}^{-1} \mathbf{v}_{n,\ell}, \end{aligned} \quad (1.62)$$

where $\tilde{\boldsymbol{\psi}}_{n,k} \triangleq \mathbf{x}_n - \boldsymbol{\psi}_{n,k}$. Using the combination step from (1.52) gives

$$\tilde{\mathbf{x}}_{n,k|n} = \sum_{\ell \in \mathfrak{Y}_k} a_{\ell k} \left((\mathbf{I}_N - \mathbf{P}_{n,\ell|n}\mathbf{S}_{n,\ell})\tilde{\mathbf{x}}_{n,\ell|n-1} - \mathbf{P}_{n,\ell|n} \sum_{i \in \mathfrak{Y}_\ell} \mathbf{H}_{n,i}^T \mathbf{R}_{n,i}^{-1} \mathbf{v}_{n,i} \right). \quad (1.63)$$

Using $\tilde{\mathbf{x}}_{n,\ell|n-1} = \mathbf{F}_{n-1}\tilde{\mathbf{x}}_{n-1,\ell|n-1} + \mathbf{G}_{n-1}\mathbf{u}_{n-1}$, we arrive at the recursion

$$\begin{aligned} \tilde{\mathbf{x}}_{n,k|n} &= \sum_{\ell \in \mathfrak{Y}_k} a_{\ell k} \left((\mathbf{I}_N - \mathbf{P}_{n,\ell|n}\mathbf{S}_{n,\ell})\mathbf{F}_{n-1}\tilde{\mathbf{x}}_{n-1,\ell|n-1} \right. \\ &\quad \left. + (\mathbf{I}_N - \mathbf{P}_{n,\ell|n}\mathbf{S}_{n,\ell})\mathbf{G}_{n-1}\mathbf{u}_{n-1} - \mathbf{P}_{n,\ell|n} \sum_{i \in \mathfrak{Y}_\ell} \mathbf{H}_{n,i}^T \mathbf{R}_{n,i}^{-1} \mathbf{v}_{n,i} \right). \end{aligned} \quad (1.64)$$

This relation shows that the network error vector evolves according to the following dynamics:

$$\begin{aligned} \tilde{\boldsymbol{\chi}}_{n|n} &= \mathcal{A}^T (\mathbf{I}_N - \mathcal{P}_{n|n}\mathcal{S}_n) (\mathbf{I}_K \otimes \mathbf{F}_{n-1}) \tilde{\boldsymbol{\chi}}_{n-1|n-1} \\ &\quad + \mathcal{A}^T (\mathbf{I}_{KN} - \mathcal{P}_{n|n}\mathcal{S}_n) (\mathbf{I}_K \otimes \mathbf{G}_{n-1}) \mathbf{U}_{n-1} - \mathcal{A}^T \mathcal{P}_{n|n} \mathcal{J}^T \mathcal{H}_n^T \mathcal{R}_n^{-1} \mathbf{V}_n, \end{aligned} \quad (1.65)$$

which we rewrite more compactly as

$$\tilde{\boldsymbol{\chi}}_{n|n} = \mathcal{F}_n \tilde{\boldsymbol{\chi}}_{n-1|n-1} + \mathcal{G}_n \mathbf{U}_{n-1} - \mathcal{D}_n \mathbf{V}_n, \quad (1.66)$$

18 CHAPTER 1 Chapter Title

where

$$\mathcal{F}_n \triangleq \mathcal{A}^T(\mathbf{I}_{KN} - \mathcal{P}_{n|n}\mathcal{S}_n)(\mathbf{I}_K \otimes \mathbf{F}_{n-1}), \quad (1.67)$$

$$\mathcal{G}_n \triangleq \mathcal{A}^T(\mathbf{I}_{KN} - \mathcal{P}_{n|n}\mathcal{S}_n)(\mathbf{I}_K \otimes \mathbf{G}_{n-1}), \quad (1.68)$$

$$\mathcal{D}_n \triangleq \mathcal{A}^T \mathcal{P}_{n|n} \mathcal{J}^T \mathcal{H}_n^T \mathcal{R}_n^{-1}. \quad (1.69)$$

If we now let $\mathcal{P}_{\tilde{\mathcal{X}},n}$ denote the covariance matrix of the network error vector,

$$\mathcal{P}_{\tilde{\mathcal{X}},n} \triangleq \mathbb{E}[\tilde{\mathcal{X}}_{n|n}\tilde{\mathcal{X}}_{n|n}^T], \quad (1.70)$$

it then follows from (1.66) that this matrix satisfies the Lyapunov recursion

$$\mathcal{P}_{\tilde{\mathcal{X}},n} = \mathcal{F}_n \mathcal{P}_{\tilde{\mathcal{X}},n-1} \mathcal{F}_n^T + \mathcal{G}_n (\mathbf{1}_K \mathbf{1}_K^T \otimes \mathbf{Q}_{n-1}) \mathcal{G}_n^T + \mathcal{D}_n \mathcal{R}_n \mathcal{D}_n^T. \quad (1.71)$$

In order to analyze the stability and performance of the diffusion filter, we shall consider here the following conditions – see [28] for definitions of the notions of stabilizability and detectability.

Assumption. *It is assumed that the model parameters $\{\mathbf{F}, \mathbf{G}, \mathbf{H}_k, \mathbf{Q}, \mathbf{R}_k\}$ do not depend on the time index, n . We further collect the measurement matrices in the neighborhood of agent k into the block column matrix:*

$$\mathbf{H}_k^{\text{loc}} \triangleq \text{blkcol}\{\mathbf{H}_{k_1}, \mathbf{H}_{k_2}, \dots, \mathbf{H}_{k_{d_k}}\}, \quad (1.72)$$

and assume that the pair $(\mathbf{F}, \mathbf{G}\mathbf{Q}^{1/2})$ is stabilizable and the pair $(\mathbf{F}, \mathbf{H}_k^{\text{loc}})$ is detectable. \square

Under the stabilizability and detectability conditions, it is known from the convergence properties of the discrete Riccati recursions that each entry $\mathbf{P}_{n,k|n}$ of $\mathcal{P}_{n|n}$ converges to the quantity \mathbf{P}_k^+ defined by [28]:

$$(\mathbf{P}_k^+)^{-1} = (\mathbf{P}_k)^{-1} + (\mathbf{H}_k^{\text{loc}})^T (\mathbf{R}_k^{\text{loc}})^{-1} \mathbf{H}_k^{\text{loc}} = (\mathbf{P}_k)^{-1} + \mathbf{S}_k, \quad (1.73)$$

where

$$\mathbf{R}_k^{\text{loc}} \triangleq \text{blkdiag}\{\mathbf{R}_{k_1}, \mathbf{R}_{k_2}, \dots, \mathbf{R}_{k_{d_k}}\}, \quad (1.74)$$

and \mathbf{P}_k is the unique stabilizing solution of the following discrete algebraic Riccati equation (DARE):

$$\mathbf{P}_k = \mathbf{F}\mathbf{P}_k^+ \mathbf{F}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T = \mathbf{F}\mathbf{P}_k \mathbf{F}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T - \mathbf{K}_{p,k} \mathbf{R}_{e,k} \mathbf{K}_{p,k}^T, \quad (1.75)$$

where

$$\mathbf{K}_{p,k} \triangleq \mathbf{F}\mathbf{P}_k (\mathbf{H}_k^{\text{loc}})^T \mathbf{R}_{e,k}^{-1}, \quad \mathbf{R}_{e,k} \triangleq \mathbf{R}_k^{\text{loc}} + \mathbf{H}_k^{\text{loc}} \mathbf{P}_k (\mathbf{H}_k^{\text{loc}})^T. \quad (1.76)$$

Accordingly, the quantities $\{\mathcal{P}_{n|n}, \mathcal{F}_n, \mathcal{G}_n, \mathcal{D}_n\}$ converge to steady-state values denoted

by [12]:

$$\mathcal{P}^+ \triangleq \lim_{n \rightarrow \infty} \mathcal{P}_{n|n} \text{ (block diagonal),} \quad (1.77)$$

$$\mathcal{F} \triangleq \lim_{n \rightarrow \infty} \mathcal{F}_n = \mathcal{A}^T (\mathbf{I}_{KN} - \mathcal{P}^+ \mathcal{S}) (\mathbf{I}_K \otimes \mathbf{F}), \quad (1.78)$$

$$\mathcal{G} \triangleq \lim_{n \rightarrow \infty} \mathcal{G}_n = \mathcal{A}^T (\mathbf{I}_{KN} - \mathcal{P}^+ \mathcal{S}) (\mathbf{I}_K \otimes \mathbf{G}), \quad (1.79)$$

$$\mathcal{D} \triangleq \lim_{n \rightarrow \infty} \mathcal{D}_n = \mathcal{A}^T \mathcal{P}^+ \mathcal{J}^T \mathcal{H}^T \mathcal{R}^{-1}, \quad (1.80)$$

where \mathcal{S} , \mathcal{H} , and \mathcal{R} are written without the time subscript n because their entries are now time-independent. Furthermore, the stabilizability and detectability conditions ensure that (see Lemma 1 in [12]):

$$\mathcal{X} \triangleq (\mathbf{I}_{KN} - \mathcal{P}^+ \mathcal{S}) (\mathbf{I}_K \otimes \mathbf{F}) \text{ is block diagonal and stable.} \quad (1.81)$$

Note that the limiting matrix \mathcal{F} in (1.78) has the form

$$\mathcal{F} = \mathcal{A}^T \mathcal{X}. \quad (1.82)$$

We will verify soon that, under some conditions, \mathcal{F} is also stable for any left-stochastic matrix \mathcal{A} and, in view of Lemma 1 in the Appendix (Section 1.5), the error covariance matrix $\mathcal{P}_{\bar{\mathcal{X}},n}$ in (1.71) converges to the unique solution of the Lyapunov equation:

$$\mathcal{P}_{\bar{\mathcal{X}}} = \mathcal{F} \mathcal{P}_{\bar{\mathcal{X}}} \mathcal{F}^T + \mathcal{G} (\mathbb{1}_K \mathbb{1}_K^T \otimes \mathbf{Q}) \mathcal{G}^T + \mathcal{D} \mathcal{R} \mathcal{D}^T. \quad (1.83)$$

Using the vec notation (which stacks the columns of a matrix on top of each other), we can solve for $\mathcal{P}_{\bar{\mathcal{X}}}$ and write

$$\text{vec}(\mathcal{P}_{\bar{\mathcal{X}}}) = (\mathbf{I}_{(KN)^2} - \mathcal{F}^T \otimes \mathcal{F})^{-1} \text{vec}(\mathcal{G} (\mathbb{1}_K \mathbb{1}_K^T \otimes \mathbf{Q}) \mathcal{G}^T + \mathcal{D} \mathcal{R} \mathcal{D}^T). \quad (1.84)$$

In this way, the steady-state mean-square error at any agent k will be given by

$$\lim_{n \rightarrow \infty} \mathbb{E} [\|\bar{\mathbf{x}}_{n,k|n}\|^2] = \text{Tr}(\mathcal{P}_{\bar{\mathcal{X}}} \mathcal{I}_k), \quad (1.85)$$

where \mathcal{I}_k is an $K \times K$ block diagonal matrix with blocks of size $N \times N$; it contains the identity matrix at block (k, k) and zeros everywhere else. That is, it holds that $\mathcal{I}_k = \mathbf{e}_k \mathbf{e}_k^T \otimes \mathbf{I}_N$, where \mathbf{e}_k is the basis column vector of size $K \times 1$ with unit entry at location k and zeros everywhere else. Consequently, the average mean-square error across the network is

$$\text{MSE}_{\text{av}} = \frac{1}{K} \text{Tr}(\mathcal{P}_{\bar{\mathcal{X}}}). \quad (1.86)$$

We verify the above claims by showing that the matrix \mathcal{F} is stable and that, in view of this fact, the Lyapunov recursion (1.71) converges to the unique solution of the Lyapunov equation (1.83). These facts follow from the results of two lemmas presented in the Appendix.

1.3 DISTRIBUTED PARTICLE FILTERING

In the second part of this chapter, we discuss distributed particle filtering. We will focus on consensus-based methods in Subsection 1.3.4 and on diffusion-based methods in Subsection 1.3.5. Before that, we present the underlying state-space model, which generalizes (1.2) and (1.3), and we discuss some basics of sequential Bayesian estimation and review the particle filter. We adopt a Bayesian framework for estimation, which means that both the state \mathbf{x}_n and the measurements $\mathbf{z}_{n,k}$ are modeled as random vectors [36].

1.3.1 STATE-SPACE MODEL

Differently from Section 1.2, the state-space model is now allowed to be nonlinear. The temporal dynamics or evolution of the state \mathbf{x}_n is characterized by the state-evolution model (a.k.a. state-transition model or system model)

$$\mathbf{x}_{n+1} = \mathbf{g}_n(\mathbf{x}_n, \mathbf{u}_n), \quad n = 0, 1, \dots, \quad (1.87)$$

which extends (1.2). In (1.87), $\mathbf{g}_n(\cdot, \cdot)$ is a known, generally nonlinear function, and \mathbf{u}_n is state noise that is statistically independent and identically distributed (iid) across time n and also statistically independent of the state sequence. The probability density function (pdf) of \mathbf{u}_n is assumed known. Equation (1.87) then determines the *state-transition pdf* $f(\mathbf{x}_n|\mathbf{x}_{n-1})$ for $n = 1, 2, \dots$.

Furthermore, the statistical dependence of the measurement $\mathbf{z}_{n,k}$ of agent $k \in \{1, 2, \dots, K\}$ on the state \mathbf{x}_n is characterized by a measurement model

$$\mathbf{z}_{n,k} = \mathbf{h}_{n,k}(\mathbf{x}_n, \mathbf{v}_{n,k}), \quad n = 1, 2, \dots; \quad k = 1, 2, \dots, K, \quad (1.88)$$

which extends (1.3). In (1.88), $\mathbf{h}_{n,k}(\cdot, \cdot)$ is a known, generally nonlinear function, and $\mathbf{v}_{n,k}$ is observation noise that is statistically independent across time n and across the agents k , and also statistically independent of the state sequence and of the state noise. The pdf of $\mathbf{v}_{n,k}$ is assumed known. Equation (1.88) then determines the *local likelihood function* of agent k , $f(\mathbf{z}_{n,k}|\mathbf{x}_n)$, for $n = 1, 2, \dots$. Because $\mathbf{v}_{n,k}$ is assumed independent of $\mathbf{v}_{n,k'}$ for $k' \neq k$, the measurements $\mathbf{z}_{n,k}$ at different agents k are conditionally independent given the state \mathbf{x}_n . Thus, the *global (all-agents) likelihood function* $f(\mathbf{z}_n|\mathbf{x}_n)$, with $\mathbf{z}_n \triangleq (\mathbf{z}_{n,1}^T \mathbf{z}_{n,2}^T \cdots \mathbf{z}_{n,K}^T)^T$, factorizes into the local likelihood functions, i.e.,

$$f(\mathbf{z}_n|\mathbf{x}_n) = \prod_{k=1}^K f(\mathbf{z}_{n,k}|\mathbf{x}_n). \quad (1.89)$$

The above independence assumptions can be relaxed; for example, the state noise process \mathbf{u}_n and the observation noise processes $\mathbf{v}_{n,k}$ may be allowed to be dependent [37]. Further, the conditional independence of the observations can be removed too [38, 39]. Also, the noises in the state and observation equations do not have to be

independent across time [40, 41]. However, the independence assumptions stated above are required for the consensus-based DPF methods presented below. We also note that the case where the state-evolution function $\mathbf{g}_n(\mathbf{x}_n, \mathbf{u}_n)$ is linear in both \mathbf{x}_n and \mathbf{u}_n and the measurement function $\mathbf{h}_{n,k}(\mathbf{x}_n, \mathbf{v}_{n,k})$ is linear in both \mathbf{x}_n and $\mathbf{v}_{n,k}$ was considered in Section 1.2 (see (1.2) and (1.3)). Furthermore, when the state noise \mathbf{u}_n in (1.87) and the observation noise \mathbf{v}_n in (1.88) have Gaussian distributions, these distributions are completely characterized by their first- and second-order moments. Finally, we point out that the network may be dynamic in that the set of neighboring nodes may vary with time (i.e., we may have $\mathfrak{N}_{n,k}$ instead of \mathfrak{N}_k). However, in the rest of the chapter, we will continue to consider the network to be static.

The above state-evolution and measurement models along with the associated statistical assumptions imply two further conditional independence properties. First, the current state \mathbf{x}_n is conditionally independent of all the past measurements, $\mathbf{z}_{1:n-1,k} \triangleq (\mathbf{z}_{1,k}^T \ \mathbf{z}_{2,k}^T \ \cdots \ \mathbf{z}_{n-1,k}^T)^T$, given the previous state \mathbf{x}_{n-1} , i.e.,

$$f(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{z}_{1:n-1,k}) = f(\mathbf{x}_n | \mathbf{x}_{n-1}), \quad (1.90)$$

for all $k = 1, 2, \dots, K$. Second, the current measurements $\mathbf{z}_{n,k}$ are conditionally independent of all the past measurements, $\mathbf{z}_{1:n-1,k'}$, given the current state \mathbf{x}_n , i.e.,

$$f(\mathbf{z}_{n,k} | \mathbf{x}_n, \mathbf{z}_{1:n-1,k'}) = f(\mathbf{z}_{n,k} | \mathbf{x}_n), \quad (1.91)$$

for any $k' \in \{1, 2, \dots, K\}$, including in particular $k' = k$.

1.3.2 SEQUENTIAL BAYESIAN ESTIMATION

We address the problem of sequential estimation of the state \mathbf{x}_n from the total (all-agents) measurement sequence $\mathbf{z}_{1:n} \triangleq (\mathbf{z}_1^T \ \mathbf{z}_2^T \ \cdots \ \mathbf{z}_n^T)^T$. In the Bayesian context, this essentially amounts to calculating the *posterior pdf* $f(\mathbf{x}_n | \mathbf{z}_{1:n})$, from which optimal Bayesian estimators and quantities characterizing their performance can be derived. In particular, the MMSE estimator [36] is given by the first moment of $f(\mathbf{x}_n | \mathbf{z}_{1:n})$, i.e.,

$$\widehat{\mathbf{x}}_n^{\text{MMSE}} = \mathbb{E}[\mathbf{x}_n | \mathbf{z}_{1:n}] = \int_{\mathbb{R}^N} \mathbf{x}_n f(\mathbf{x}_n | \mathbf{z}_{1:n}) d\mathbf{x}_n, \quad n = 1, 2, \dots \quad (1.92)$$

In a centralized scenario, based on the conditional independence properties (1.90) and (1.91), the current posterior pdf $f(\mathbf{x}_n | \mathbf{z}_{1:n})$ can be calculated sequentially (recursively) from the previous posterior pdf $f(\mathbf{x}_{n-1} | \mathbf{z}_{1:n-1})$ and the current all-agents measurement vector \mathbf{z}_n . This recursion consists of a prediction step calculating $f(\mathbf{x}_n | \mathbf{z}_{1:n-1})$ from $f(\mathbf{x}_{n-1} | \mathbf{z}_{1:n-1})$ and a subsequent time-update or correction step calculating $f(\mathbf{x}_n | \mathbf{z}_{1:n})$ from $f(\mathbf{x}_n | \mathbf{z}_{1:n-1})$. The prediction step involves the state-transition pdf $f(\mathbf{x}_n | \mathbf{x}_{n-1})$, and the update step involves the global likelihood function $f(\mathbf{z}_n | \mathbf{x}_n)$ and, thus, the all-agents measurement vector \mathbf{z}_n . Unfortunately, the prediction and update steps as well as the calculation of the MMSE estimate in (1.92) are usually computationally infeasible. A prominent exception is the case of the linear-Gaussian

22 CHAPTER 1 Chapter Title

state-space model, where optimal sequential Bayesian state estimation reduces to the Kalman filter considered in Subsection 1.2.3 [19, 28, 29].

Several computationally feasible approximations to optimal sequential Bayesian state estimation have been proposed in the centralized case; these include the extended Kalman filter [19, 42, 36], the Gaussian sum filter [43], the unscented (sigma-point) Kalman filter [44, 45], the cubature Kalman filter [46], and the particle filter [47, 20, 48, 49]. In particular, particle filters are well suited to any nonlinear state-space model and any distributions of the state and observation noises, and they perform well in situations where Kalman filter-based methods perform poorly. This generally comes at the cost of an increased computational complexity. Below, we will discuss distributed methods for particle filtering in decentralized agent networks. We start with a review of the basic particle filter.

1.3.3 REVIEW OF THE PARTICLE FILTER

A particle filter (PF) performs an approximation of optimal sequential Bayesian state estimation that is based on Monte Carlo simulation and importance sampling [47, 20, 48, 49]. In a centralized scenario, the posterior pdf $f(\mathbf{x}_n|\mathbf{z}_{1:n})$ is represented in an approximative manner by M randomly drawn samples or *particles* $\mathbf{x}_n^{(m)}$ and corresponding weights $w_n^{(m)}$, where $m = 1, 2, \dots, M$. Specifically, at each time step n , propagation of the posterior pdf (i.e., $f(\mathbf{x}_{n-1}|\mathbf{z}_{1:n-1}) \rightarrow f(\mathbf{x}_n|\mathbf{z}_{1:n})$) is replaced by propagation of the particles and weights (i.e., $\{(\mathbf{x}_{n-1}^{(m)}, w_{n-1}^{(m)})\}_{m=1}^M \rightarrow \{(\mathbf{x}_n^{(m)}, w_n^{(m)})\}_{m=1}^M$). As in the case of optimal sequential Bayesian estimation considered in Subsection 1.3.2, this propagation consists of a prediction step and an update or correction step.

In the prediction step at time n , for each preceding particle $\mathbf{x}_{n-1}^{(m)}$, a new particle $\mathbf{x}_n^{(m)}$ is sampled from a suitably chosen proposal pdf (a.k.a. importance pdf) $q(\mathbf{x}_n|\mathbf{x}_{n-1}^{(m)})$. In the simplest case, $q(\mathbf{x}_n|\mathbf{x}_{n-1}^{(m)})$ is chosen as $f(\mathbf{x}_n|\mathbf{x}_{n-1}^{(m)})$, i.e., the state-transition pdf $f(\mathbf{x}_n|\mathbf{x}_{n-1})$ conditioned on $\mathbf{x}_{n-1} = \mathbf{x}_{n-1}^{(m)}$; the resulting PF algorithm is known as the sequential importance resampling (SIR) filter. However, more sophisticated “adapted” proposal pdfs that also involve the current measurement z_n can result in improved estimation performance [49, 50] (see Subsection 1.3.4.7).

In the update step at time n , for each particle $\mathbf{x}_n^{(m)}$, a nonnormalized weight is calculated according to

$$\tilde{w}_n^{(m)} = w_{n-1}^{(m)} \frac{f(z_n|\mathbf{x}_n^{(m)})f(\mathbf{x}_n^{(m)}|\mathbf{x}_{n-1}^{(m)})}{q(\mathbf{x}_n^{(m)}|\mathbf{x}_{n-1}^{(m)})}, \quad m = 1, 2, \dots, M. \quad (1.93)$$

For the SIR filter, this simplifies to

$$\tilde{w}_n^{(m)} = w_{n-1}^{(m)} f(z_n|\mathbf{x}_n^{(m)}), \quad m = 1, 2, \dots, M. \quad (1.94)$$

Subsequently, the weights are normalized according to $w_n^{(m)} = \tilde{w}_n^{(m)} / \sum_{m'=1}^M \tilde{w}_n^{(m')}$. The set of particles and normalized weights $\{(\mathbf{x}_n^{(m)}, w_n^{(m)})\}_{m=1}^M$ constitutes a Monte Carlo representation of the posterior pdf $f(\mathbf{x}_n|\mathbf{z}_{1:n})$. From $\{(\mathbf{x}_n^{(m)}, w_n^{(m)})\}_{m=1}^M$, a corresponding

approximation of the MMSE state estimate $\widehat{\mathbf{x}}_n^{\text{MMSE}}$ in (1.92) can be computed as the weighted sample mean, i.e.,

$$\widehat{\mathbf{x}}_n = \sum_{m=1}^M w_n^{(m)} \mathbf{x}_n^{(m)}. \quad (1.95)$$

Finally, if a suitable criterion is satisfied (as discussed in [51, 20]), the set $\{(\mathbf{x}_n^{(m)}, w_n^{(m)})\}_{m=1}^M$ is resampled to avoid an effect known as particle degeneracy. In the simplest case [51], the resampled particles are obtained by sampling with replacement from the set $\{(\mathbf{x}_n^{(m)}, w_n^{(m)})\}_{m=1}^M$, where $\mathbf{x}_n^{(m)}$ is sampled with probability $w_n^{(m)}$. This results in M resampled particles $\mathbf{x}_n^{(m)}$. The weights are redefined as $w_n^{(m)} = 1/M$.

This recursive algorithm is initialized at time $n=0$ by M particles $\mathbf{x}_0^{(m)}$, $m = 1, 2, \dots, M$, which are drawn from a suitable prior pdf $f(\mathbf{x}_0)$. The initial weights are equal, i.e., $w_0^{(m)} = 1/M$ for all m .

In a distributed implementation based on a decentralized network of agents, each agent runs a local instance of a PF, hereafter briefly referred to as “local PF.” The local PF at agent k observes only the local measurements $z_{n,k}$ directly; however, it receives from the neighbor agents indirect information about the measurements of the other agents in the network and also provides to the neighboring agents indirect information about its own measurements. The type of information that is exchanged between neighboring agents depends on the specific method used for distributed particle filtering. A distributed PF (DPF) is based on the PF algorithm summarized above. However, it modifies that algorithm to account for the fact that each agent runs its own local PF, and it employs some networkwide distributed scheme (such as consensus, gossip, or diffusion) to disseminate and fuse local information provided by the agents. Several types of DPF methods have been proposed, including [17, 52, 53, 54, 55, 13, 56, 57, 58, 59, 60, 61, 62, 63, 64]. In what follows, we discuss two classes of DPF methods that rely on consensus and diffusion schemes for distributed information dissemination and fusion.

1.3.4 CONSENSUS-BASED METHODS

In a DPF that emulates the performance of the basic PF, the local PF at agent k attempts to track a particle representation $\{(\mathbf{x}_{n,k}^{(m)}, w_{n,k}^{(m)})\}_{m=1}^M$ of the global posterior pdf $f(\mathbf{x}_n | z_{1:n})$. According to (1.93), this requires the *global* likelihood function (GLF) $f(z_n | \mathbf{x}_n)$ (evaluated at the current local particles, i.e., $\mathbf{x}_n = \mathbf{x}_{n,k}^{(m)}$), rather than merely the *local* likelihood function (LLF) of agent k , $f(z_{n,k} | \mathbf{x}_n)$. Because agent k by itself is only able to calculate the LLF, calculation of the GLF requires information provided by the other agents.

1.3.4.1 DPF Based on Likelihood Consensus

A recently proposed class of DPF algorithms performs a distributed calculation of the GLF using the *likelihood consensus* (LC) scheme [13, 56]. To develop that scheme,

24 CHAPTER 1 Chapter Title

we first take the logarithm of (1.89), which yields

$$\ln f(\mathbf{z}_n|\mathbf{x}_n) = \sum_{k=1}^K \ln f(\mathbf{z}_{n,k}|\mathbf{x}_n), \quad (1.96)$$

where \ln denotes the natural logarithm. We can thus write the GLF as

$$f(\mathbf{z}_n|\mathbf{x}_n) = \exp(\ln f(\mathbf{z}_n|\mathbf{x}_n)) = \exp(K\lambda(\mathbf{x}_n; \mathbf{z}_n)), \quad (1.97)$$

with

$$\lambda(\mathbf{x}_n; \mathbf{z}_n) \triangleq \frac{1}{K} \ln f(\mathbf{z}_n|\mathbf{x}_n) = \frac{1}{K} \sum_{k=1}^K \ln f(\mathbf{z}_{n,k}|\mathbf{x}_n), \quad (1.98)$$

where (1.96) has been used in the last step. Hence, the GLF $f(\mathbf{z}_n|\mathbf{x}_n)$ has been rewritten in terms of the *average* of the K log-LLFs $\ln f(\mathbf{z}_{n,k}|\mathbf{x}_n)$, $k = 1, 2, \dots, K$. Because the log-LLFs $\ln f(\mathbf{z}_{n,k}|\mathbf{x}_n)$ are functions of the unknown state $\mathbf{x}_n \in \mathbb{R}^N$, rather than simply numbers, we cannot directly use a distributed averaging scheme such as the average consensus algorithm to calculate $\lambda(\mathbf{x}_n; \mathbf{z}_n)$.

The solution provided by the LC scheme is based on an approximate (finite-order) function expansion of the log-LLFs $\ln f(\mathbf{z}_{n,k}|\mathbf{x}_n)$. Let $\{\varphi_r(\mathbf{x}_n)\}_{r=1}^R$ denote a system of R functions that is known to all agents. Possible choices of these functions include monomials [13], orthogonal polynomials [65], Fourier basis functions [56], and spline functions [66]. We consider the approximations of the log-LLFs given by

$$\ln f(\mathbf{z}_{n,k}|\mathbf{x}_n) \approx \sum_{r=1}^R \alpha_{n,k,r}(\mathbf{z}_{n,k}) \varphi_r(\mathbf{x}_n), \quad k = 1, 2, \dots, K, \quad (1.99)$$

with suitable expansion coefficients $\alpha_{n,k,r}(\mathbf{z}_{n,k})$, $r = 1, 2, \dots, R$. Note that these expansion coefficients involve the local measurements $\mathbf{z}_{n,k}$, and thus they are different at different agents k . Inserting (1.99) into (1.98) and changing the order of summation yields

$$\lambda(\mathbf{x}_n; \mathbf{z}_n) \approx \tilde{\lambda}(\mathbf{x}_n; \mathbf{z}_n) \triangleq \sum_{r=1}^R a_{n,r}(\mathbf{z}_n) \varphi_r(\mathbf{x}_n), \quad (1.100)$$

where

$$a_{n,r}(\mathbf{z}_n) \triangleq \frac{1}{K} \sum_{k=1}^K \alpha_{n,k,r}(\mathbf{z}_{n,k}). \quad (1.101)$$

Using in (1.97) the approximation $\tilde{\lambda}(\mathbf{x}_n; \mathbf{z}_n)$ instead of $\lambda(\mathbf{x}_n; \mathbf{z}_n)$, we obtain a corresponding approximation of the GLF, namely,

$$f(\mathbf{z}_n|\mathbf{x}_n) \approx \tilde{f}(\mathbf{z}_n|\mathbf{x}_n) \triangleq \exp(K\tilde{\lambda}(\mathbf{x}_n; \mathbf{z}_n)). \quad (1.102)$$

1.3.4.2 Distributed Calculation of the GLF

Within the approximation given by (1.100) and (1.102), the distributed calculation of the GLF $f(\mathbf{z}_n|\mathbf{x}_n)$ at time n amounts to the distributed calculation of the R numbers $a_{n,r}(\mathbf{z}_n)$, $r = 1, 2, \dots, R$, each of which is an average of the respective local expansion coefficients $\alpha_{n,k,r}(\mathbf{z}_{n,k})$, $k = 1, 2, \dots, K$. Note that because the measurements $\mathbf{z}_{n,k}$ have been observed and thus are fixed, we now have to average numbers instead of functions. Thus, the average expansion coefficients $a_{n,r}(\mathbf{z}_n)$, $r = 1, 2, \dots, R$ can be calculated in a distributed manner via R instances of an average consensus algorithm [24, 25]. In addition (see (1.102)), each agent also needs to know the total number of agents, K . Distributed algorithms for counting the number of agents are available (e.g., [67]).

In iteration i of the r th instance of the average consensus algorithm, which is used to calculate $a_{n,r}(\mathbf{z}_n)$, each agent k updates an “internal state” $\zeta_{k,r}^{(i)}$ according to

$$\zeta_{k,r}^{(i)} = \sum_{k' \in \mathfrak{N}_k} \omega_{k,k'}^{(i)} \zeta_{k',r}^{(i-1)}. \quad (1.103)$$

Here, $\{\omega_{k,k'}^{(i)}\}_{k' \in \mathfrak{N}_k}$ is a set of weights whose choice is discussed in [24, 68, 25]. Hence, $\zeta_{k,r}^{(i)}$ is a linear combination of the preceding (i.e., at iteration $i - 1$) internal state of agent k and the preceding internal states of the neighbor agents k' . The iteration is initialized by choosing the initial internal states as $\zeta_{k,r}^{(0)} = \alpha_{n,k,r}(\mathbf{z}_{n,k})$. Then, for a suitable choice of the weights, and using our assumption that the agent network is connected, it can be shown [25] that the internal state $\zeta_{k,r}^{(i)}$ of each agent k converges to the desired average $a_{n,r}(\mathbf{z}_n)$, i.e.,

$$\lim_{i \rightarrow \infty} \zeta_{k,r}^{(i)} = a_{n,r}(\mathbf{z}_n) = \frac{1}{K} \sum_{k'=1}^K \alpha_{n,k',r}(\mathbf{z}_{n,k'}). \quad (1.104)$$

For a finite number i_{\max} of iterations, complete convergence cannot be achieved; this means that the internal states $\zeta_{k,r}^{(i_{\max})}$ will be (slightly) different from $a_{n,r}(\mathbf{z}_n)$, and also (slightly) different across the agents k .

Several standard choices of the weights $\omega_{k,k'}^{(i)}$ are available [68, 24]. A popular choice is given by the Metropolis weights [24]

$$\omega_{k,k'} = \begin{cases} \frac{1}{\max\{|\mathfrak{N}_k|, |\mathfrak{N}_{k'}|\}}, & k' \neq k, \\ 1 - \sum_{k'' \in \mathfrak{N}_k \setminus \{k\}} \omega_{k,k''}, & k' = k. \end{cases} \quad (1.105)$$

Note that these weights do not depend on the iteration index i . Their calculation at agent k requires that agent k knows both $|\mathfrak{N}_k|$ and $|\mathfrak{N}_{k'}|$ for all $k' \in \mathfrak{N}_k$. Certain other choices of the weights require less knowledge [24].

In each iteration of the consensus scheme, agent k has to broadcast its internal states $\zeta_{k,r}^{(i)}$, $r = 1, 2, \dots, R$ to its neighbors $k' \in \mathfrak{N}_k \setminus \{k\}$. Furthermore, agent k receives the internal states $\zeta_{k',r}^{(i)}$, $r = 1, 2, \dots, R$ from its neighbors $k' \in \mathfrak{N}_k \setminus \{k\}$. Thus, during one time step n , each agent k has to broadcast a total of $i_{\max} R$ real numbers to its

neighbors.

1.3.4.3 Calculation of the Local Expansion Coefficients

The local expansion coefficients $\alpha_{n,k,r}(z_{n,k})$, $r = 1, 2, \dots, R$ arising in the log-LLF approximation (1.99) are calculated *locally* at the respective agent k . This can be done by means of a least-squares (LS) fit between the log-LLF $\ln f(z_{n,k}|\mathbf{x}_n)$ and the approximating expansion $\sum_{r=1}^R \alpha_{n,k,r}(z_{n,k}) \varphi_r(\mathbf{x}_n)$. In view of the way the local PF operates, the approximation does not need to be good for all possible states \mathbf{x}_n but only in those regions of the state space where the current particles $\mathbf{x}_{n,k}^{(m)}$ are located. Therefore, the LS fit at agent k minimizes, with respect to the vector of expansion coefficients $\boldsymbol{\alpha}_{n,k} \triangleq (\alpha_{n,k,1}(z_{n,k}) \cdots \alpha_{n,k,R}(z_{n,k}))^T$, the LS approximation error of the expansion (1.99) evaluated at the particles $\mathbf{x}_{n,k}^{(m)}$, $m = 1, 2, \dots, M$, i.e.,

$$\sum_{m=1}^M \left(\ln f(z_{n,k}|\mathbf{x}_{n,k}^{(m)}) - \sum_{r=1}^R \alpha_{n,k,r}(z_{n,k}) \varphi_r(\mathbf{x}_{n,k}^{(m)}) \right)^2 = \|\boldsymbol{\xi}_{n,k} - \mathbf{\Phi}_{n,k} \boldsymbol{\alpha}_{n,k}\|^2 \rightarrow \min_{\boldsymbol{\alpha}_{n,k}}. \quad (1.106)$$

Here, $\boldsymbol{\xi}_{n,k} \triangleq (\ln f(z_{n,k}|\mathbf{x}_{n,k}^{(1)}) \cdots \ln f(z_{n,k}|\mathbf{x}_{n,k}^{(M)}))^T$ and $\mathbf{\Phi}_{n,k}$ is an $M \times R$ matrix with columns $\boldsymbol{\phi}_{n,k,r} \triangleq (\varphi_r(\mathbf{x}_{n,k}^{(1)}) \cdots \varphi_r(\mathbf{x}_{n,k}^{(M)}))^T$. Assuming that $M \geq R$ (i.e., there are at least as many particles as expansion coefficients) and that $\mathbf{\Phi}_{n,k}$ has full rank, the solution to the LS problem (1.106) is given by [69]

$$\widehat{\boldsymbol{\alpha}}_{n,k} = \mathbf{\Phi}_{n,k}^{\#} \boldsymbol{\xi}_{n,k}, \quad \text{with } \mathbf{\Phi}_{n,k}^{\#} \triangleq (\mathbf{\Phi}_{n,k}^T \mathbf{\Phi}_{n,k})^{-1} \mathbf{\Phi}_{n,k}^T. \quad (1.107)$$

Numerical aspects of computing $\widehat{\boldsymbol{\alpha}}_{n,k}$ in (1.107) are discussed in [69, 70]. The elements $\widehat{\alpha}_{n,k,r}$ of $\widehat{\boldsymbol{\alpha}}_{n,k}$ are then used to initialize the LC according to $\zeta_{k,r}^{(0)} = \widehat{\alpha}_{n,k,r}$.

A summary of the LC-based DPF algorithm is provided in listing (1.108).

LC-based DPF

each agent starts with its own set of particles $\mathbf{x}_{0,k}^{(m)} \sim f(\mathbf{x}_0)$ and identical weights $w_{0,k}^{(m)} = 1/M$, where $m = 1, 2, \dots, M$.

repeat at every agent k for $n \geq 1$:

(propose a new set of particles)

$$\mathbf{x}_{n,k}^{(m)} \sim q(\mathbf{x}_n | \mathbf{x}_{n-1,k}^{(m)}), \quad m = 1, 2, \dots, M$$

(calculate an approximation $\tilde{f}_k(\mathbf{z}_n | \mathbf{x}_n)$ of the GLF $f(\mathbf{z}_n | \mathbf{x}_n)$)

- calculate the coefficient vector $\tilde{\alpha}_{n,k}$ according to (1.107)
- for $r = 1, 2, \dots, R$:
 - initialize the internal LC state: $\zeta_{k,r}^{(0)} = \tilde{\alpha}_{n,k,r}$
 - for $i = 1, 2, \dots, i_{\max}$, update the internal LC state $\zeta_{k,r}^{(i)}$ according to (1.103)
 - calculate $\tilde{f}_k(\mathbf{z}_n | \mathbf{x}_n) = \exp(K \sum_{r=1}^R \zeta_{k,r}^{(i_{\max})} \varphi_r(\mathbf{x}_n))$ (cf. (1.102), (1.100))

(1.108)

(calculate the weights of the proposed particles)

- calculate nonnormalized weights:

$$\tilde{w}_{n,k}^{(m)} = w_{n-1,k}^{(m)} \frac{\tilde{f}_k(\mathbf{z}_n | \mathbf{x}_n) f(\mathbf{x}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}{q(\mathbf{x}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}, \quad m = 1, 2, \dots, M$$

- normalize weights:

$$w_n^{(m)} = \tilde{w}_n^{(m)} / \sum_{m'=1}^M \tilde{w}_n^{(m')}$$

(compute the state estimate)

$$\hat{\mathbf{x}}_{n,k} = \sum_{m=1}^M w_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)}$$

(if necessary, perform resampling [51, 20])

end

1.3.4.4 Exponential Family

An alternative LC scheme is possible if the LLFs of all the agents k (viewed as conditional pdfs of $\mathbf{z}_{n,k}$) belong to the exponential family of distributions [71], i.e.,

$$f(\mathbf{z}_{n,k} | \mathbf{x}_n) = c_{n,k}(\mathbf{z}_{n,k}) \exp(\mathbf{b}_{n,k}^T(\mathbf{x}_n) \mathbf{d}_{n,k}(\mathbf{z}_{n,k}) - g_{n,k}(\mathbf{x}_n)), \quad k = 1, 2, \dots, K, \quad (1.109)$$

with some functions $c_{n,k}(\cdot) \in \mathbb{R}_+$, $\mathbf{b}_{n,k}(\cdot) \in \mathbb{R}^q$, $\mathbf{d}_{n,k}(\cdot) \in \mathbb{R}^q$, and $g_{n,k}(\cdot) \in \mathbb{R}$, where $q \geq N$. Inserting this expression into (1.89) yields for the GLF

$$f(\mathbf{z}_n | \mathbf{x}_n) = C_n(\mathbf{z}_n) \exp(K G_n(\mathbf{x}_n; \mathbf{z}_n)), \quad (1.110)$$

where $C_n(\mathbf{z}_n) \triangleq \prod_{k=1}^K c_{n,k}(\mathbf{z}_{n,k})$ and

$$G_n(\mathbf{x}_n; \mathbf{z}_n) \triangleq \frac{1}{K} \sum_{k=1}^K (\mathbf{b}_{n,k}^T(\mathbf{x}_n) \mathbf{d}_{n,k}(\mathbf{z}_{n,k}) - g_{n,k}(\mathbf{x}_n)). \quad (1.111)$$

The normalization factor $C_n(\mathbf{z}_n)$ does not depend on the state \mathbf{x}_n and is hence irrelevant due to weight normalization (see Subsection 1.3.3).

As an alternative to expanding the log-LLFs as in (1.99), we may use separate

finite-order function expansions of the state-dependent functions involved in (1.111) [13], i.e.,

$$\mathbf{b}_{n,k}(\mathbf{x}_n) \approx \sum_{r=1}^{R_b} \boldsymbol{\beta}_{n,k,r} \varphi_r(\mathbf{x}_n), \quad g_{n,k}(\mathbf{x}_n) \approx \sum_{r=1}^{R_g} \gamma_{n,k,r} \psi_r(\mathbf{x}_n), \quad k = 1, 2, \dots, K. \quad (1.112)$$

Here, the coefficients $\boldsymbol{\beta}_{n,k,r}$ are vectors of the same dimension as $\mathbf{b}_{n,k}(\mathbf{x}_n)$. The coefficients $\boldsymbol{\beta}_{n,k,r}$ and $\gamma_{n,k,r}$ can again be calculated by means of LS fitting, similarly to Subsection 1.3.4.3; note that one separate LS fit is required for each component of $\boldsymbol{\beta}_{n,k,r}$. Inserting (1.112) into (1.111) and changing the order of summation yields

$$G_n(\mathbf{x}_n; \mathbf{z}_n) \approx \widetilde{G}_n(\mathbf{x}_n; \mathbf{z}_n) \triangleq \sum_{r=1}^{R_b} B_{n,r}(\mathbf{z}_n) \varphi_r(\mathbf{x}_n) - \sum_{r=1}^{R_g} \Gamma_{n,r} \psi_r(\mathbf{x}_n), \quad (1.113)$$

with

$$B_{n,r}(\mathbf{z}_n) \triangleq \frac{1}{K} \sum_{k=1}^K \boldsymbol{\beta}_{n,k,r}^T \mathbf{d}_{n,k}(\mathbf{z}_{n,k}), \quad \Gamma_{n,r} \triangleq \frac{1}{K} \sum_{k=1}^K \gamma_{n,k,r}. \quad (1.114)$$

Using in (1.110) the approximation $\widetilde{G}_n(\mathbf{x}_n; \mathbf{z}_n)$ instead of $G_n(\mathbf{x}_n; \mathbf{z}_n)$ yields an approximation of the GLF.

The R_b numbers $B_{n,r}(\mathbf{z}_k)$, $r = 1, 2, \dots, R_b$ and the R_g numbers $\Gamma_{n,r}$, $r = 1, 2, \dots, R_g$ are seen in (1.114) to be averages of the local quantities $a_{n,k,r}(\mathbf{z}_{n,k}) \triangleq \boldsymbol{\beta}_{n,k,r}^T \mathbf{d}_{n,k}(\mathbf{z}_{n,k})$ and $\gamma_{n,k,r}$, respectively. Hence, they can be calculated in a distributed manner via $R_b + R_g$ instances of an average consensus algorithm, similarly to Subsection 1.3.4.2. Note that this distributed calculation assumes that each agent k knows its own functions $\mathbf{b}_{n,k}(\cdot)$, $\mathbf{d}_{n,k}(\cdot)$, and $g_{n,k}(\cdot)$, but not the respective functions of the other agents $k' \neq k$.

This distributed calculation of (an approximation of) the GLF $f(\mathbf{z}_n | \mathbf{x}_n)$ may be preferable over the general scheme presented in Subsections 1.3.4.1–1.3.4.3 if it is easier to approximate the functions $\mathbf{b}_{n,k}(\mathbf{x}_n)$ and $g_{n,k}(\mathbf{x}_n)$ than the LLFs $\ln f(\mathbf{z}_{n,k} | \mathbf{x}_n)$. In particular, there is a reduction in the number of real numbers each agent has to broadcast to its neighbors if $R_b + R_g < R$, where R is the order of function expansion used in the general scheme.

1.3.4.5 Gaussian Observation Noise

An important special case of an exponential-family LLF arises with additive Gaussian observation noise [13]. Here, the measurement model of agent k in (1.88) is specialized according to

$$\mathbf{z}_{n,k} = \mathbf{m}_{n,k}(\mathbf{x}_n) + \mathbf{v}_{n,k}, \quad n = 1, 2, \dots; \quad k = 1, 2, \dots, K, \quad (1.115)$$

where $\mathbf{m}_{n,k}(\mathbf{x}_n)$ is a generally nonlinear function and the additive observation noise $\mathbf{v}_{n,k}$ satisfies the independence properties formulated in Section 1.3.1 and, in addition, is Gaussian with zero mean and covariance matrix $\mathbf{C}_{n,k}$. It follows that the LLF of

agent k is given, up to a normalization factor, by

$$f(\mathbf{z}_{n,k}|\mathbf{x}_n) \propto \exp\left(-\frac{1}{2}(\mathbf{z}_{n,k} - \mathbf{m}_{n,k}(\mathbf{x}_n))^T \mathbf{C}_{n,k}^{-1}(\mathbf{z}_{n,k} - \mathbf{m}_{n,k}(\mathbf{x}_n))\right). \quad (1.116)$$

This is easily verified to be an exponential-family LLF (1.109) with

$$\mathbf{b}_{n,k}(\mathbf{x}_n) = \mathbf{m}_{n,k}(\mathbf{x}_n), \quad (1.117)$$

$$\mathbf{d}_{n,k}(\mathbf{z}_{n,k}) = \mathbf{C}_{n,k}^{-1} \mathbf{z}_{n,k}, \quad (1.118)$$

$$g_{n,k}(\mathbf{x}_n) = \frac{1}{2} \mathbf{m}_{n,k}^T(\mathbf{x}_n) \mathbf{C}_{n,k}^{-1} \mathbf{m}_{n,k}(\mathbf{x}_n). \quad (1.119)$$

As in Subsection 1.3.4.4, we may use separate function expansion approximations of $\mathbf{b}_{n,k}(\mathbf{x}_n) = \mathbf{m}_{n,k}(\mathbf{x}_n)$ and $g_{n,k}(\mathbf{x}_n)$. However, an alternative approach is possible because according to (1.119), $g_{n,k}(\mathbf{x}_n)$ is a function of $\mathbf{m}_{n,k}(\mathbf{x}_n)$ [13]. Indeed, we may insert a function expansion approximation of $\mathbf{b}_{n,k}(\mathbf{x}_n) = \mathbf{m}_{n,k}(\mathbf{x}_n)$, i.e.,

$$\mathbf{m}_{n,k}(\mathbf{x}_n) \approx \tilde{\mathbf{m}}_{n,k}(\mathbf{x}_n) \triangleq \sum_{r=1}^{R_b} \boldsymbol{\beta}_{n,k,r} \varphi_r(\mathbf{x}_n), \quad (1.120)$$

into (1.119) to obtain an “induced” function expansion approximation of $g_{n,k}(\mathbf{x}_n)$,

$$\begin{aligned} g_{n,k}(\mathbf{x}_n) \approx \tilde{g}_{n,k}(\mathbf{x}_n) &\triangleq \frac{1}{2} \tilde{\mathbf{m}}_{n,k}^T(\mathbf{x}_n) \mathbf{C}_{n,k}^{-1} \tilde{\mathbf{m}}_{n,k}(\mathbf{x}_n) \\ &= \frac{1}{2} \sum_{r_1=1}^{R_b} \sum_{r_2=1}^{R_b} \boldsymbol{\beta}_{n,k,r_1}^T \mathbf{C}_{n,k}^{-1} \boldsymbol{\beta}_{n,k,r_2} \varphi_{r_1}(\mathbf{x}_n) \varphi_{r_2}(\mathbf{x}_n). \end{aligned} \quad (1.121)$$

Using any one-to-one mapping of the double index $(r_1, r_2) \in \{1, 2, \dots, R_b\} \times \{1, 2, \dots, R_b\}$ to a single index $r \in \{1, 2, \dots, R_g\}$, with $R_g = R_b^2$, we can rewrite (1.121) as (cf. (1.112))

$$\tilde{g}_{n,k}(\mathbf{x}_n) = \sum_{r=1}^{R_g} \gamma_{n,k,r} \psi_r(\mathbf{x}_n), \quad (1.122)$$

where $\gamma_{n,k,r} = \frac{1}{2} \boldsymbol{\beta}_{n,k,r_1}^T \mathbf{C}_{n,k}^{-1} \boldsymbol{\beta}_{n,k,r_2}$ and $\psi_r(\mathbf{x}_n) = \varphi_{r_1}(\mathbf{x}_n) \varphi_{r_2}(\mathbf{x}_n)$. It can be easily verified that the resulting approximate GLF (cf. (1.89)) can be written as $\tilde{f}(\mathbf{z}_k|\mathbf{x}_k) \propto \exp(-\frac{1}{2} \tilde{\mathbf{Q}}_n(\mathbf{z}_k, \mathbf{x}_k))$ with $\tilde{\mathbf{Q}}_n(\mathbf{z}_k, \mathbf{x}_k) \triangleq \sum_{k=1}^K (\mathbf{z}_{n,k} - \tilde{\mathbf{m}}_{n,k}(\mathbf{x}_n))^T \mathbf{C}_{n,k}^{-1} (\mathbf{z}_{n,k} - \tilde{\mathbf{m}}_{n,k}(\mathbf{x}_n))$. This equals the true GLF $f(\mathbf{z}_k|\mathbf{x}_k)$, except that the means $\mathbf{m}_{n,k}(\mathbf{x}_n)$ are replaced by their approximations $\tilde{\mathbf{m}}_{n,k}(\mathbf{x}_n)$.

1.3.4.6 Distributed Gaussian Particle Filter

The Gaussian PF (GPF) [72] uses a Gaussian approximation of the posterior pdf $f(\mathbf{x}_n|\mathbf{z}_{1:n})$. This Gaussian approximation is derived from a weighted particle set, which is calculated in a similar way as in the PF except that no resampling is required. Note that the measurement model is still allowed to be non-Gaussian, as in (1.88).

30 CHAPTER 1 Chapter Title

In a distributed GPF, the local GPF at agent k propagates a Gaussian approximation $\mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}_{n,k}, \boldsymbol{\Sigma}_{n,k})$ of the global posterior pdf $f(\mathbf{x}_n | \mathbf{z}_{1:n})$. At time n , particles $\mathbf{x}_{n,k}^{(m)}$, $m = 1, 2, \dots, M$ are drawn from the preceding Gaussian approximation $\mathcal{N}(\mathbf{x}_{n-1}; \boldsymbol{\mu}_{n-1,k}, \boldsymbol{\Sigma}_{n-1,k})$. Then, the prediction and update steps are performed as described in Section 1.3.3. From the resulting weighted particles $\{(\mathbf{x}_{n,k}^{(m)}, w_{n,k}^{(m)})\}_{m=1}^M$, a new Gaussian approximation $\mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}_{n,k}, \boldsymbol{\Sigma}_{n,k})$ is determined according to

$$\boldsymbol{\mu}_{n,k} = \sum_{m=1}^M w_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)} \quad \text{and} \quad \boldsymbol{\Sigma}_{n,k} = \sum_{m=1}^M w_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)\text{T}} - \boldsymbol{\mu}_{n,k} \boldsymbol{\mu}_{n,k}^{\text{T}}. \quad (1.123)$$

The weighted sample mean $\boldsymbol{\mu}_{n,k}$ also constitutes the local approximation $\widehat{\mathbf{x}}_{n,k}$ of the MMSE state estimate $\widehat{\mathbf{x}}_n^{\text{MMSE}}$ in (1.92), i.e., $\widehat{\mathbf{x}}_{n,k} = \boldsymbol{\mu}_{n,k}$.

Just as in the conventional DPF (cf. (1.93)), the update step of the local GPFs requires the GLF evaluated at the current local particles, i.e., $f(\mathbf{z}_n | \mathbf{x}_n)$ for $\mathbf{x}_n = \mathbf{x}_{n,k}^{(m)}$, $m = 1, 2, \dots, M$. An approximation of $f(\mathbf{z}_n | \mathbf{x}_{n,k}^{(m)})$ can be calculated in a distributed manner by the LC as explained earlier. However, the following variation yields a significant reduction of computational complexity at the cost of some increase in local communications [13]. Inspired by the parallel GPF implementation proposed in [73], the idea is to “distribute” the M particles over the K agents, such that each local GPF uses a significantly reduced set of particles, and to combine the results of the local GPFs via additional average consensus operations. The local GPF at agent k uses only $M_k < M$ particles, where $\sum_{k=1}^K M_k = M$. In the local weight update step at agent k , for each particle $\mathbf{x}_{n,k}^{(m)}$ from the reduced local particle set, $m \in \{1, 2, \dots, M_k\}$, a nonnormalized weight $\widetilde{w}_{n,k}^{(m)}$ is calculated according to (1.93). This requires the GLF evaluated at the M_k particles, i.e., $f(\mathbf{z}_n | \mathbf{x}_{n,k}^{(m)})$, $m = 1, 2, \dots, M_k$, which was previously calculated (approximately) by the LC. Furthermore, the sum of the resulting nonnormalized weights is calculated, i.e., $\widetilde{W}_{n,k} = \sum_{m=1}^{M_k} \widetilde{w}_{n,k}^{(m)}$.

Next, from the weighted particles $\{(\mathbf{x}_{n,k}^{(m)}, \widetilde{w}_{n,k}^{(m)})\}_{m=1}^{M_k}$, a local nonnormalized mean vector and a local nonnormalized correlation matrix are calculated at each agent k as

$$\boldsymbol{\mu}'_{n,k} = \sum_{m=1}^{M_k} \widetilde{w}_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)}, \quad \mathbf{R}'_{n,k} = \sum_{m=1}^{M_k} \widetilde{w}_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)} \mathbf{x}_{n,k}^{(m)\text{T}}. \quad (1.124)$$

Finally, the local means and correlations from all the agents are combined into a global mean and covariance:

$$\bar{\boldsymbol{\mu}}_{n,k} = \frac{1}{W_{n,k}} \sum_{k=1}^K \boldsymbol{\mu}'_{n,k}, \quad \bar{\boldsymbol{\Sigma}}_{n,k} = \frac{1}{W_{n,k}} \sum_{k=1}^K \mathbf{R}'_{n,k} - \bar{\boldsymbol{\mu}}_{n,k} \bar{\boldsymbol{\mu}}_{n,k}^{\text{T}}, \quad (1.125)$$

where $W_{n,k} \triangleq \sum_{k=1}^K \widetilde{W}_{n,k}$. Approximations of the sums $\sum_{k=1}^K \boldsymbol{\mu}'_{n,k}$, $\sum_{k=1}^K \mathbf{R}'_{n,k}$, and $\sum_{k=1}^K \widetilde{W}_{n,k}$ can be calculated in a distributed manner by means of an average consensus algorithm (again assuming that the number of agents, K , is known to each agent). Subsequently, the division by $W_{n,k}$ and the subtraction of $\bar{\boldsymbol{\mu}}_{n,k} \bar{\boldsymbol{\mu}}_{n,k}^{\text{T}}$ in (1.125)

are performed locally at each agent. This results in local approximations $\mu_{n,k}$ and $\Sigma_{n,k}$ of, respectively, $\bar{\mu}_{n,k}$ and $\bar{\Sigma}_{n,k}$ at each agent k . The state estimate at agent k , $\hat{\mathbf{x}}_{n,k}$, is taken to be $\mu_{n,k}$.

1.3.4.7 Distributed Proposal Adaptation

The choice of the proposal pdf $q(\mathbf{x}_n; \mathbf{x}_{n-1}^{(m)})$ used in the weight update step (1.93) strongly affects the performance of a PF. The proposal pdf should be similar to the posterior pdf $f(\mathbf{x}_n | \mathbf{z}_{1:n})$ [20]. Using the state-transition pdf $f(\mathbf{x}_n | \mathbf{x}_{n-1}^{(m)})$ as proposal pdf, as is done in the SIR filter, does not cater to this goal; in particular, it does not take into account the current measurement \mathbf{z}_n . The design of a proposal pdf taking into account the measurements is known as *proposal adaptation* [49, 50]. In a DPF, the proposal pdfs used by the local PFs should be adapted to the total (all-agents) measurement $\mathbf{z}_n = (\mathbf{z}_{n,1}^T \mathbf{z}_{n,2}^T \cdots \mathbf{z}_{n,K}^T)^T$.

A distributed method for this type of “global” proposal adaptation can again be based on the average consensus principle [56]. In this method, each agent first calculates a “predistorted” local posterior pdf. A Gaussian approximation of the global posterior pdf is then obtained by fusing all the predistorted local posterior pdfs; this approximate global posterior pdf is used by the local PFs as a proposal pdf. The method is inspired by [52] but employs a different predistortion that enables the use of a Gaussian filter for calculating the predistorted local posterior pdfs. To derive the method, we note that the global posterior pdf can be developed as $f(\mathbf{x}_n | \mathbf{z}_{1:n}) = f(\mathbf{x}_n | \mathbf{z}_{1:n-1}, \mathbf{z}_n) \propto f(\mathbf{z}_n | \mathbf{x}_n, \mathbf{z}_{1:n-1}) f(\mathbf{x}_n | \mathbf{z}_{1:n-1}) = f(\mathbf{z}_n | \mathbf{x}_n) f(\mathbf{x}_n | \mathbf{z}_{1:n-1})$, where Bayes’ rule and (1.91) have been used. Inserting (1.89), we further obtain

$$f(\mathbf{x}_n | \mathbf{z}_{1:n}) \propto \left(\prod_{k=1}^K f(\mathbf{z}_{n,k} | \mathbf{x}_n) \right) f(\mathbf{x}_n | \mathbf{z}_{1:n-1}). \quad (1.126)$$

We now define a predistorted, nonnormalized “local pseudoposterior pdf” at agent k as

$$\tilde{f}(\mathbf{x}_n | \mathbf{z}_{1:n-1}, \mathbf{z}_{n,k}) \triangleq f(\mathbf{z}_{n,k} | \mathbf{x}_n) (f(\mathbf{x}_n | \mathbf{z}_{1:n-1}))^{1/K}. \quad (1.127)$$

Furthermore, we consider the network-wide product of all the local pseudoposterior pdfs,

$$\prod_{k=1}^K \tilde{f}(\mathbf{x}_n | \mathbf{z}_{1:n-1}, \mathbf{z}_{n,k}) = \left(\prod_{k=1}^K f(\mathbf{z}_{n,k} | \mathbf{x}_n) \right) f(\mathbf{x}_n | \mathbf{z}_{1:n-1}) \propto f(\mathbf{x}_n | \mathbf{z}_{1:n}), \quad (1.128)$$

where (1.126) has been used in the last step. According to (1.128), the product of all the local pseudoposterior pdfs equals the global posterior pdf up to a factor. This fact can be used to calculate a Gaussian approximation of the global posterior pdf $f(\mathbf{x}_n | \mathbf{z}_{1:n})$ —note that $f(\mathbf{x}_n | \mathbf{z}_{1:n})$ involves all the measurements and in fact would be the optimal proposal pdf—by a distributed evaluation of the leftmost side of (1.128). To make this possible, we use Gaussian approximations of the local pseudoposterior pdfs, i.e.,

32 CHAPTER 1 Chapter Title

$$\tilde{f}(\mathbf{x}_n | \mathbf{z}_{1:n-1}, \mathbf{z}_{n,k}) \approx \mathcal{N}(\mathbf{x}_n; \tilde{\boldsymbol{\mu}}_{n,k}, \tilde{\boldsymbol{\Sigma}}_{n,k}), \quad (1.129)$$

and of the global posterior pdf, i.e.,

$$f(\mathbf{x}_n | \mathbf{z}_{1:n}) \approx q(\mathbf{x}_n; \mathbf{z}_n) \triangleq \mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n). \quad (1.130)$$

Note that (1.130) also defines the global proposal pdf $q(\mathbf{x}_n; \mathbf{z}_n)$. Inserting the approximations (1.129) and (1.130) into the factorization $f(\mathbf{x}_n | \mathbf{z}_{1:n}) \propto \prod_{k=1}^K \tilde{f}(\mathbf{x}_n | \mathbf{z}_{1:n-1}, \mathbf{z}_{n,k})$ (see (1.128)) yields

$$\mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n) \propto \prod_{k=1}^K \mathcal{N}(\mathbf{x}_n; \tilde{\boldsymbol{\mu}}_{n,k}, \tilde{\boldsymbol{\Sigma}}_{n,k}). \quad (1.131)$$

Using this relation and the rules for a product of Gaussian densities [52, 74], the global mean $\boldsymbol{\mu}_n$ and global covariance $\boldsymbol{\Sigma}_n$ —which determine $q(\mathbf{x}_n; \mathbf{z}_n) = \mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n)$ —can be calculated from the local means $\tilde{\boldsymbol{\mu}}_{n,k}$ and local covariances $\tilde{\boldsymbol{\Sigma}}_{n,k}$ according to

$$\boldsymbol{\mu}_n = \boldsymbol{\Sigma}_n \sum_{k=1}^K (\tilde{\boldsymbol{\Sigma}}_{n,k})^{-1} \tilde{\boldsymbol{\mu}}_{n,k}, \quad \boldsymbol{\Sigma}_n = \left(\sum_{k=1}^K (\tilde{\boldsymbol{\Sigma}}_{n,k})^{-1} \right)^{-1}. \quad (1.132)$$

Here, the sums $\sum_{k=1}^K (\tilde{\boldsymbol{\Sigma}}_{n,k})^{-1} \tilde{\boldsymbol{\mu}}_{n,k}$ and $\sum_{k=1}^K (\tilde{\boldsymbol{\Sigma}}_{n,k})^{-1}$ can be computed in a distributed manner by means of an average consensus algorithm.

The calculation of the local mean $\tilde{\boldsymbol{\mu}}_{n,k}$ and local covariance $\tilde{\boldsymbol{\Sigma}}_{n,k}$, as defined in (1.129), at agent k is based on the observation that (1.127) can be interpreted as the update step of a Bayesian filter using the predistorted predicted posterior pdf $(f(\mathbf{x}_n | \mathbf{z}_{1:n-1}))^{1/K}$ instead of the true predicted posterior pdf $f(\mathbf{x}_n | \mathbf{z}_{1:n-1})$. Each agent first calculates a Gaussian approximation of the predicted posterior pdf,

$$f(\mathbf{x}_n | \mathbf{z}_{1:n-1}) \approx \mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}'_{n,k}, \boldsymbol{\Sigma}'_{n,k}), \quad (1.133)$$

where $\boldsymbol{\mu}'_{n,k} = \frac{1}{M} \sum_{m=1}^M \mathbf{x}_{n,k}^{(m) \prime}$ and $\boldsymbol{\Sigma}'_{n,k} = \frac{1}{M} \sum_{m=1}^M \mathbf{x}_{n,k}^{(m) \prime} (\mathbf{x}_{n,k}^{(m) \prime})^T - \boldsymbol{\mu}'_{n,k} (\boldsymbol{\mu}'_{n,k})^T$. Here, $\{\mathbf{x}_{n,k}^{(m) \prime}\}_{m=1}^M$ is a set of temporary particles that are randomly drawn from $f(\mathbf{x}_n | \mathbf{x}_{n-1,k}^{(m)})$, where $\{\mathbf{x}_{n-1,k}^{(m)}\}_{m=1}^M$ is the set of particles resulting from the preceding filtering step at time $n-1$. The approximation (1.133) implies

$$(f(\mathbf{x}_n | \mathbf{z}_{1:n-1}))^{1/K} \approx \mathcal{N}(\mathbf{x}_n; \boldsymbol{\mu}'_{n,k}, K\boldsymbol{\Sigma}'_{n,k}). \quad (1.134)$$

Using the Gaussian approximations (1.129) and (1.134) in (1.127) gives

$$\mathcal{N}(\mathbf{x}_k; \tilde{\boldsymbol{\mu}}_{n,k}, \tilde{\boldsymbol{\Sigma}}_{n,k}) = f(\mathbf{z}_{n,k} | \mathbf{x}_k) \mathcal{N}(\mathbf{x}_k; \boldsymbol{\mu}'_{n,k}, K\boldsymbol{\Sigma}'_{n,k}). \quad (1.135)$$

This can be calculated by the update step of a Gaussian filter with input mean $\boldsymbol{\mu}'_{n,k}$, input covariance $K\boldsymbol{\Sigma}'_{n,k}$, and measurement $\mathbf{z}_{n,k}$. This Gaussian filter update step is performed locally at each agent and produces $\tilde{\boldsymbol{\mu}}_{n,k}$ and $\tilde{\boldsymbol{\Sigma}}_{n,k}$. Examples of a Gaussian filter include the extended Kalman filter [19, 42, 36], the unscented Kalman filter [44, 45], the cubature Kalman filter [46], and the filters described in [75].

1.3.5 DIFFUSION-BASED METHODS

In this subsection, we address diffusion cooperation among agents as in Subsection 1.2.7 but implemented with particle filtering [60]. However, now the considered state-space models are nonlinear and the pdfs of \mathbf{u}_n in (1.87) and $\mathbf{v}_{n,k}$ in (1.88) are assumed to be known up to proportionality constants. We reiterate that these pdfs can be of any form.

The diffusion-based DPF presented here follows the spirit of cooperation from Subsection 1.2.7 in the sense that the agents share moments (or parameters of approximating distributions) of the unknown state. We assume that the agents at time $n = 0$ start with the generation of their own particles $\mathbf{x}_{0,k}^{(m)}$, $m = 1, 2, \dots, M$, which are drawn from a prior pdf $f(\mathbf{x}_0)$, i.e.,

$$\mathbf{x}_{0,k}^{(m)} \sim f(\mathbf{x}_0), \quad m = 1, 2, \dots, M; \quad k = 1, 2, \dots, K. \quad (1.136)$$

Thus, before the tracking of the state starts, each agent has a representation of the prior pdf of \mathbf{x}_0 given by $\{\mathbf{x}_{0,k}^{(m)}, w_{0,k}^{(m)} = \frac{1}{M}\}_{m=1}^M$. Similarly, before processing the measurement vector $\mathbf{z}_{n,k}$, the local PF of agent k has a representation of the posterior pdf of $\mathbf{x}_{n-1,k}$ given by $\{\mathbf{x}_{n-1,k}^{(m)}, w_{n-1,k}^{(m)} = \frac{1}{M}\}_{m=1}^M$, which serves as a prior for the state \mathbf{x}_n .

Next, we explain how we obtain $\{\mathbf{x}_{n,k}^{(m)}, w_{n,k}^{(m)} = \frac{1}{M}\}_{m=1}^M$ from $\{\mathbf{x}_{n-1,k}^{(m)}, w_{n-1,k}^{(m)} = \frac{1}{M}\}_{m=1}^M$. The PF of each agent k propagates the particles $\mathbf{x}_{n-1,k}^{(m)}$ to particles that represent possible values of the state at time n . The propagation is carried out by drawing samples from the importance pdf $q(\mathbf{x}_n | \mathbf{x}_{n-1,k}^{(m)})$, i.e.,

$$\tilde{\mathbf{x}}_{n,k}^{(m)} \sim q(\mathbf{x}_n | \mathbf{x}_{n-1,k}^{(m)}), \quad m = 1, 2, \dots, M; \quad k = 1, 2, \dots, K. \quad (1.137)$$

As described in Subsection 1.3.3, the PF subsequently computes the weights of these particles. This is accomplished by

$$\tilde{w}_{n,k}^{(m)} \propto \frac{f(\mathbf{z}_{n,k} | \tilde{\mathbf{x}}_{n,k}^{(m)}) f(\tilde{\mathbf{x}}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}{q(\tilde{\mathbf{x}}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}, \quad (1.138)$$

where $\sum_{m=1}^M \tilde{w}_{n,k}^{(m)} = 1$. With the particles $\tilde{\mathbf{x}}_{n,k}^{(m)}$ and their weights $\tilde{w}_{n,k}^{(m)}$, the agent k has a local approximation of the posterior pdf of \mathbf{x}_n given by $\{\tilde{\mathbf{x}}_{n,k}^{(m)}, \tilde{w}_{n,k}^{(m)}\}_{m=1}^M$. The cooperation among the agents requires that they exchange these approximating local posterior pdfs and fuse them to form new local posteriors. The best approach in terms of accuracy would be that each agent broadcasts all the particles and weights to its neighbors, but that would be quite costly in terms of communication load because the number of particles and weights is usually large. An alternative is to use the particles and weights to construct a parametric distribution, that is, estimate the parameters of an approximating distribution of the weighted particles. Then, the agents would only exchange the parameters of the approximating distribution.

In the sequel, the approximating distributions are multivariate Gaussians. Thus, the agents proceed by computing the means and covariance matrices of the approxi-

34 CHAPTER 1 Chapter Title

mating Gaussians. First, the agents obtain the means by

$$\tilde{\boldsymbol{\mu}}_{n,k} = \sum_{m=1}^M \tilde{w}_{n,k}^{(m)} \tilde{\boldsymbol{x}}_{n,k}^{(m)}, \quad k = 1, 2, \dots, K, \quad (1.139)$$

and then the covariance matrices by

$$\tilde{\boldsymbol{\Sigma}}_{n,k} = \sum_{m=1}^M \tilde{w}_{n,k}^{(m)} (\tilde{\boldsymbol{x}}_{n,k}^{(m)} - \tilde{\boldsymbol{\mu}}_{n,k}) (\tilde{\boldsymbol{x}}_{n,k}^{(m)} - \tilde{\boldsymbol{\mu}}_{n,k})^T, \quad k = 1, 2, \dots, K. \quad (1.140)$$

In the next step, the agents exchange the means and covariance matrices with their neighbors and follow up with merging all the Gaussian posteriors into one Gaussian. To that end, the agents use coefficients $a_{\ell k}$ that quantify how much they trust their neighbors (defined earlier in Subsection 1.2.1, and here with $a_{\ell k}$ representing how much agent k trusts agent ℓ). We note that $\sum_{\ell \in \mathfrak{N}_k} a_{\ell k} = 1$. In the chapter on Bayesian Approach to Inference from this book (see its Subsection 1.2.3), it is explained that the Gaussian after merging has a covariance matrix and mean that can be expressed in terms of the individual covariance matrices and means of the agents by

$$\boldsymbol{\Sigma}_{n,k} = \left(\sum_{\ell \in \mathfrak{N}_k} a_{\ell k} (\tilde{\boldsymbol{\Sigma}}_{n,\ell})^{-1} \right)^{-1} \quad (1.141)$$

and

$$\boldsymbol{\mu}_{n,k} = \boldsymbol{\Sigma}_{n,k} \left(\sum_{\ell \in \mathfrak{N}_k} a_{\ell k} (\tilde{\boldsymbol{\Sigma}}_{n,\ell})^{-1} \tilde{\boldsymbol{\mu}}_{n,\ell} \right), \quad (1.142)$$

respectively. More specifically, this choice of $\boldsymbol{\mu}_{n,k}$ and $\boldsymbol{\Sigma}_{n,k}$ minimizes the weighted sum of Kullback-Leibler distances $\sum_{\ell \in \mathfrak{N}_k} a_{\ell k} \mathcal{D}(\mathcal{N}_{n,k} \| \tilde{\mathcal{N}}_{n,\ell})$, where $\mathcal{N}_{n,k}$ is the resulting Gaussian after merging, and $\tilde{\mathcal{N}}_{n,\ell}$, $\ell \in \mathfrak{N}_k$ are the Gaussians of agent k and its neighbors.

At this point, the agents have more than one way of proceeding. A simple way is that the agents use the newly obtained Gaussian and that each draws a set of particles that represent it, i.e.,

$$\boldsymbol{x}_{n,k}^{(m)} \sim \mathcal{N}(\boldsymbol{x}_n; \boldsymbol{\mu}_{n,k}, \boldsymbol{\Sigma}_{n,k}), \quad m = 1, 2, \dots, M; \quad k = 1, 2, \dots, K. \quad (1.143)$$

After generation of these particles, the agents have a discrete representation, $\{\boldsymbol{x}_{n,k}^{(m)}, w_{n,k}^{(m)} = \frac{1}{M}\}_{m=1}^M$, of the posterior pdf at time n that serves as a prior pdf of \boldsymbol{x}_{n+1} .

Another approach would be to avoid fusion of the local posteriors by (1.141) and (1.142) and instead draw particles directly from the mixture of Gaussians with components $\tilde{\mathcal{N}}(\boldsymbol{x}_n; \tilde{\boldsymbol{\mu}}_{n,\ell}, \tilde{\boldsymbol{\Sigma}}_{n,\ell})$ and weights $a_{\ell k}$, i.e.,

$$\boldsymbol{x}_{n,k}^{(m)} \sim \sum_{\ell \in \mathfrak{N}_k} a_{\ell k} \tilde{\mathcal{N}}(\boldsymbol{x}_n; \tilde{\boldsymbol{\mu}}_{n,\ell}, \tilde{\boldsymbol{\Sigma}}_{n,\ell}), \quad m = 1, 2, \dots, M; \quad k = 1, 2, \dots, K. \quad (1.144)$$

1.3 Distributed Particle Filtering 35

Once done, the posterior pdf is again represented by $\{\mathbf{x}_{n,k}^{(m)}, w_{n,k}^{(m)} = \frac{1}{M}\}_{m=1}^M$.

An alternative direction is to first resample the original local particles $\tilde{\mathbf{x}}_{n,k}^{(m)}$ according to their weights $\tilde{w}_{n,k}^{(m)}$. Let the resampled particles be denoted by $\tilde{\mathbf{x}}_{n,k}^{(m)}$. In the next step, these particles are rescaled and shifted so that they correspond to particles of a Gaussian distribution with mean $\boldsymbol{\mu}_{n,k}$ in (1.142) and covariance $\boldsymbol{\Sigma}_{n,k}$ in (1.141). It is not difficult to show that this can be achieved by

$$\mathbf{x}_{n,k}^{(m)} = \mathbf{Q}_{n,k}^T \tilde{\mathbf{Q}}_{n,k}^{-T} (\tilde{\mathbf{x}}_{n,k}^{(m)} - \tilde{\boldsymbol{\mu}}_{n,k}) + \boldsymbol{\mu}_{n,k}, \quad (1.145)$$

where $\mathbf{Q}_{n,k}$ is defined by $\boldsymbol{\Sigma}_{n,k} = \mathbf{Q}_{n,k}^T \mathbf{Q}_{n,k}$ and $\tilde{\mathbf{Q}}_{n,k}$ by $\tilde{\boldsymbol{\Sigma}}_{n,k} = \tilde{\mathbf{Q}}_{n,k}^T \tilde{\mathbf{Q}}_{n,k}$. The particles obtained by (1.145) have all equal weights, and thus they represent the prior pdf for the state \mathbf{x}_{n+1} . This diffusion-based approach is described in the listing (1.146). Note that the particles generated by (1.143) and (1.144) are all different whereas many of the ones obtained by (1.145) may be replicated.

Diffusion particle filter (adapt and then combine form)

each agent starts with its own set of particles $\mathbf{x}_{0,k}^{(m)} \sim f(\mathbf{x}_0)$, $m = 1, 2, \dots, M$.

repeat at every agent k for $n \geq 0$:

(propose a new set of particles)

$$\tilde{\mathbf{x}}_{n,k}^{(m)} \sim q(\mathbf{x}_n | \mathbf{x}_{n-1,k}^{(m)}), \quad m = 1, 2, \dots, M$$

(compute the weights of the proposed particles)

$$\tilde{w}_{n,k}^{(m)} \propto \frac{f(\mathbf{z}_{n,k} | \tilde{\mathbf{x}}_{n,k}^{(m)}) f(\tilde{\mathbf{x}}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}{q(\tilde{\mathbf{x}}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}$$

(find the parameters of a Gaussian distribution that approximates the posterior)

$$\tilde{\boldsymbol{\mu}}_{n,k} = \sum_{m=1}^M \tilde{w}_{n,k}^{(m)} \tilde{\mathbf{x}}_{n,k}^{(m)}$$

$$\tilde{\boldsymbol{\Sigma}}_{n,k} = \sum_{m=1}^M \tilde{w}_{n,k}^{(m)} (\tilde{\mathbf{x}}_{n,k}^{(m)} - \tilde{\boldsymbol{\mu}}_{n,k}) (\tilde{\mathbf{x}}_{n,k}^{(m)} - \tilde{\boldsymbol{\mu}}_{n,k})^T \quad (1.146)$$

(combine own posterior with the posteriors of the neighbors)

$$\boldsymbol{\Sigma}_{n,k} = \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} (\tilde{\boldsymbol{\Sigma}}_{n,\ell})^{-1} \right)^{-1}$$

$$\boldsymbol{\mu}_{n,k} = \boldsymbol{\Sigma}_{n,k} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} (\tilde{\boldsymbol{\Sigma}}_{n,\ell})^{-1} \tilde{\boldsymbol{\mu}}_{n,\ell} \right)$$

(resample the local particles according to their weights)

$$\left\{ (\tilde{\mathbf{x}}_{n,k}^{(m)}, \tilde{w}_{n,k}^{(m)}) \right\}_{m=1}^M \xrightarrow{\text{resampling}} \left\{ (\mathbf{x}_{n,k}^{(m)}, w_{n,k}^{(m)} = \frac{1}{M}) \right\}_{m=1}^M$$

(rescale and shift the resampled local particles)

$$\mathbf{x}_{n,k}^{(m)} = \mathbf{Q}_{n,k}^T \tilde{\mathbf{Q}}_{n,k}^{-T} (\tilde{\mathbf{x}}_{n,k}^{(m)} - \tilde{\boldsymbol{\mu}}_{n,k}) + \boldsymbol{\mu}_{n,k}, \text{ where } \mathbf{Q}_{n,k} \text{ and } \tilde{\mathbf{Q}}_{n,k} \text{ are defined by}$$

$$\boldsymbol{\Sigma}_{n,k} = \mathbf{Q}_{n,k}^T \mathbf{Q}_{n,k} \text{ and } \tilde{\boldsymbol{\Sigma}}_{n,k} = \tilde{\mathbf{Q}}_{n,k}^T \tilde{\mathbf{Q}}_{n,k}, \text{ respectively}$$

end

The performance of the diffusion cooperation can be improved if the calculation of the particle weights of each agent involves the measurements of the respective neighbors. This requires that the agents engage in two rounds of communication, one when they exchange their measurements with their neighbors, and the other when they exchange the parameters of their posterior pdfs. In this case, the calculation of

the weights in (1.138) is replaced by

$$\tilde{w}_{n,k}^{(m)} \propto \frac{\left(\prod_{\ell \in \mathcal{N}_k} f(z_{n,\ell} | \tilde{\mathbf{x}}_{n,k}^{(m)})\right) f(\tilde{\mathbf{x}}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}{q(\tilde{\mathbf{x}}_{n,k}^{(m)} | \mathbf{x}_{n-1,k}^{(m)})}. \quad (1.147)$$

Everything else in the described process remains the same.

In summary, the diffusion-based particle filtering algorithm described in this subsection relies on approximating the local posterior pdfs by Gaussians whose parameters are then exchanged with neighbors. Upon receiving the Gaussian parameters from its neighbors, each agent either uses them to create a new Gaussian that serves as the final posterior pdf or simply exploits them directly as in (1.143) or (1.144) to generate particles that represent the support of the final posterior pdf and the prior pdf of \mathbf{x}_{n+1} . However, we reiterate that the underlying approximating distributions of $\{\tilde{\mathbf{x}}_{n,k}^{(m)}, \tilde{w}_{n,k}^{(m)}\}_{m=1}^M$ used by the agents do not have to be Gaussians.

1.4 CONCLUSIONS

In this chapter, we discussed the problem of distributed filtering in a network of agents. The filtering amounts to estimating a hidden state process based on measurements that are acquired by the agents and contain information about the state process. The agents cooperate by communicating relevant information with their neighbors. Two main types of filtering methods were addressed. The first is distributed Kalman filtering where the state-space models are linear and the state and observation noise processes are zero-mean and white vector processes with known covariance matrices. The second type is distributed particle filtering where the models are nonlinear and the pdfs characterizing the state and observation noise processes are known up to proportionality constants. The distributed Kalman filtering method was based on the diffusion strategy, with a discussion on the relation and differences relative to a consensus-type implementation. The distributed particle filtering methods were addressed by both consensus and diffusion strategies.

ACKNOWLEDGMENTS

The work of A. H. Sayed was supported by the NSF under grants ECCS-1407712 and CCF-1524250. The author is grateful to the IEEE for allowing reproduction of material from [12] in this book chapter. The work of P. M. Djurić was supported by the NSF under grant CCF-1618999. The work of F. Hawatsch was supported by the Austrian Science Fund (FWF) under grant P27370-N30 and by the Czech Science Foundation (GAČR) under grant 17-19638S.

1.5 APPENDIX

Lemma 1 (Convergence of Lyapunov recursion). *Consider a general Lyapunov recursion of the form*

$$\mathbf{Z}_{n+1} = \mathbf{C}_n \mathbf{Z}_n \mathbf{C}_n^T + \mathbf{B}_n \quad (1.148)$$

where the matrix sequences $\{\mathbf{B}_n, \mathbf{C}_n\}$ converge uniformly to $\{\mathbf{B}, \mathbf{C}\}$ as $n \rightarrow \infty$ with \mathbf{C} being a stable matrix (all its eigenvalues lie strictly inside the unit circle). Then, the sequence \mathbf{Z}_n converges to the unique solution \mathbf{Z} of the Lyapunov equation

$$\mathbf{Z} = \mathbf{C} \mathbf{Z} \mathbf{C}^T + \mathbf{B}. \quad (1.149)$$

Proof: Using the vec notation we have

$$\text{vec}(\mathbf{Z}_{n+1}) = (\mathbf{C}_n \otimes \mathbf{C}_n) \text{vec}(\mathbf{Z}_n) + \text{vec}(\mathbf{B}_n) \triangleq \mathbf{T}_n \text{vec}(\mathbf{Z}_n) + \mathbf{d}_n, \quad (1.150)$$

where we introduced the quantities $\mathbf{T}_n = (\mathbf{C}_n \otimes \mathbf{C}_n)$ and $\mathbf{d}_n = \text{vec}(\mathbf{B}_n)$. We know from the assumptions in the lemma that \mathbf{d}_n converges to $\mathbf{d} = \text{vec}(\mathbf{B})$ and \mathbf{T}_n converges to $\mathbf{T} = (\mathbf{C} \otimes \mathbf{C})$, which is a stable matrix since \mathbf{C} is stable. It follows that $\text{vec}(\mathbf{Z}_n)$ converges to $\text{vec}(\mathbf{Z}) = (\mathbf{I} - \mathbf{C} \otimes \mathbf{C})^{-1} \mathbf{d}$, which establishes that the limiting \mathbf{Z} is the solution to the Lyapunov equation (1.149). We can establish this convergence result more formally as follows by adjusting the proof given for Theorem 1 in App. E of [12]. Let $\mathbf{z}_n = \text{vec}(\mathbf{Z}_n)$, $\mathbf{z} = \text{vec}(\mathbf{Z})$, $\mathbf{T}_n - \mathbf{T} = \epsilon \Delta_n$, and $\mathbf{d}_n - \mathbf{d} = \epsilon \delta_n$, for some arbitrary scalar $\epsilon > 0$ and quantities $\{\delta_n, \Delta_n\}$. Using $\mathbf{z} = \mathbf{T} \mathbf{z} + \mathbf{d}$ and $\mathbf{z}_{n+1} = \mathbf{T}_n \mathbf{z}_n + \mathbf{d}_n$, we get

$$\mathbf{z}_{n+1} - \mathbf{z} = \mathbf{T}(\mathbf{z}_n - \mathbf{z}) + \epsilon \Delta_n(\mathbf{z}_n - \mathbf{z}) + \epsilon \Delta_n \mathbf{z} + \epsilon \delta_n. \quad (1.151)$$

Now, since \mathbf{T} is a stable matrix, there exists a sub-multiplicative matrix norm, denoted by $\|\cdot\|_\rho$, such that $\|\mathbf{T}\|_\rho = c < 1$ [76]. Using this norm, and the triangle inequality, we have

$$\|\mathbf{z}_{n+1} - \mathbf{z}\|_\rho \leq \|\mathbf{T}\|_\rho \|\mathbf{z}_n - \mathbf{z}\|_\rho + \epsilon \|\Delta_n\|_\rho \|\mathbf{z}_n - \mathbf{z}\|_\rho + \epsilon \|\Delta_n\|_\rho \|\mathbf{z}\|_\rho + \epsilon \|\delta_n\|_\rho. \quad (1.152)$$

In the limit, as $n \rightarrow \infty$, we can choose $\|\Delta_n\|_\rho \leq 1$ and $\|\delta_n\|_\rho \leq 1$, which implies that

$$\|\mathbf{z}_{n+1} - \mathbf{z}\|_\rho \leq (\|\mathbf{T}\|_\rho + \epsilon) \|\mathbf{z}_n - \mathbf{z}\|_\rho + \epsilon(\|\mathbf{z}\|_\rho + 1), \text{ as } n \rightarrow \infty. \quad (1.153)$$

But since $\|\mathbf{T}\|_\rho < 1$, we can select ϵ small enough such that $\|\mathbf{T}\|_\rho + \epsilon < 1$ and, hence,

$$\lim_{n \rightarrow \infty} \|\mathbf{z}_{n+1} - \mathbf{z}\|_\rho = \epsilon(\|\mathbf{z}\|_\rho + 1)/(1 - \|\mathbf{T}\|_\rho - \epsilon). \quad (1.154)$$

Since ϵ can be chosen arbitrarily small, it follows that $\|\mathbf{z}_{n+1} - \mathbf{z}\|_\rho \rightarrow 0$ as $n \rightarrow \infty$. ■

Lemma 2 (Stability of \mathcal{F}). *For any left-stochastic matrix \mathcal{A} and block diagonal matrix \mathcal{D} with the 2-induced norms of its blocks strictly bounded by one, it holds that the matrix product $\mathcal{B} = \mathcal{A}^T \mathcal{D}$ is stable. Consequently, when $\|(\mathbf{I}_N - \mathbf{P}_k^+ \mathbf{S}_k) \mathbf{F}\|_2 < 1$,*

38 CHAPTER 1 Chapter Title

the matrix \mathcal{F} defined below is stable:

$$\mathcal{F} \triangleq \mathcal{A}^T \mathcal{X}, \quad \mathcal{X} \triangleq (\mathbf{I}_{KN} - \mathcal{P}^+ \mathcal{S})(\mathbf{I}_K \otimes \mathbf{F}). \quad (1.155)$$

Proof: This argument adjusts the proof given for Lemma 2 in [12], where the ρ -norm should be replaced by the block-maximum norm defined as follows [77]. Let $\mathbf{x} = \text{col}\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}$ denote a $K \times 1$ block column vector whose individual entries are themselves vectors of size $N \times 1$ each. The block maximum norm of \mathbf{x} is denoted by $\|\mathbf{x}\|_{b,\infty}$ and is defined as $\|\mathbf{x}\|_{b,\infty} = \max_{1 \leq k \leq K} \|\mathbf{x}_k\|$. This vector norm induces a block-maximum matrix norm. Let \mathcal{A} denote an arbitrary $K \times K$ block matrix with individual block entries of size $N \times N$ each. Then, the block-maximum norm of \mathcal{A} is defined as

$$\|\mathcal{A}\|_{b,\infty} \triangleq \max_{\mathbf{x} \neq 0} \frac{\|\mathcal{A}\mathbf{x}\|_{b,\infty}}{\|\mathbf{x}\|_{b,\infty}}. \quad (1.156)$$

The block-maximum norm has several useful properties — see [77]. In particular, when \mathbf{A} is $K \times K$ left-stochastic and $\mathcal{A} = \mathbf{A} \otimes \mathbf{I}_N$, then it can be verified that $\|\mathcal{A}^T\|_{b,\infty} = 1$. Likewise, when \mathcal{D} is block diagonal with the 2-induced norm of its diagonal blocks bounded by one, it is easy to check that $\|\mathcal{D}\|_{b,\infty} < 1$. Consequently, for $\mathcal{B} = \mathcal{A}^T \mathcal{D}$ we get

$$\rho(\mathcal{B}) \leq \|\mathcal{B}\|_{b,\infty} \leq \|\mathcal{A}^T\|_{b,\infty} \|\mathcal{D}\|_{b,\infty} < 1, \quad (1.157)$$

and we conclude that \mathcal{B} is stable. Alternatively, we note that

$$\|\mathcal{B}^n\|_{b,\infty} \leq (\|\mathcal{A}^T\|_{b,\infty})^n (\|\mathcal{D}\|_{b,\infty})^n \rightarrow 0, \text{ as } n \rightarrow \infty. \quad (1.158)$$

The matrix \mathcal{F} in (1.155) has a form that fits into this formulation with $\mathcal{D} = (\mathbf{I}_{KN} - \mathcal{P}^+ \mathcal{S})(\mathbf{I}_K \otimes \mathbf{F}) = \mathcal{X}$, which is block diagonal. Moreover, since by assumption $\|(\mathbf{I}_{KN} - \mathcal{P}_k^+ \mathcal{S}_k) \mathbf{F}\|_2 < 1$, it holds that $\|\mathcal{D}\|_{b,\infty} = \|\mathcal{X}\|_{b,\infty} < 1$. ■

REFERENCES

1. Zhu Y, You Z, Zhao J, Zhang K, Li XR, The optimality for the distributed Kalman filtering fusion with feedback. *Automatica* 2001; 37(9):1489–1493.
2. Coates M, Distributed particle filters for sensor networks. In: *Proceedings of the 3rd ACM/IEEE International Symposium on Information Processing in Sensor Networks*, 2004, pp. 99–107.
3. Spanos DP, Olfati-Saber R, Murray RM, Approximate distributed Kalman filtering in sensor networks with quantifiable performance. In: *Proceedings of the 4th ACM/IEEE International Symposium on Information Processing in Sensor Networks*, 2005, pp. 133–139.
4. Ribeiro A, Giannakis GB, Roumeliotis SI, SOI-KF: Distributed Kalman filtering with low-cost communications using the sign of innovations. *IEEE Transactions on Signal Processing* 2006; 54(12):4782–4795.

5. Olfati-Saber R, Distributed Kalman filtering for sensor networks. In: 46th IEEE Conference on Decision and Control, 2007, pp. 5492–5498.
6. Khan UA, Moura JM, Distributing the Kalman filter for large-scale systems. *IEEE Transactions on Signal Processing* 2008; 56(10):4919–4935.
7. Cattivelli FS, Lopes CG, Sayed AH, Diffusion strategies for distributed Kalman filtering: Formulation and performance analysis. In: *Proceedings of the IEEE IAPR Workshop on Cognitive Information Processing*, 2008, pp. 36–41.
8. Cattivelli FS, Sayed AH, Diffusion mechanisms for fixed-point distributed Kalman smoothing. In: *16th European Signal Processing Conference*, 2008, pp. 1–5.
9. Carli R, Chiuso A, Schenato L, Zampieri S, Distributed Kalman filtering based on consensus strategies. *IEEE Journal on Selected Areas in Communications* 2008; 26(4):622–633.
10. Cattivelli F, Sayed AH, Diffusion distributed Kalman filtering with adaptive weights. In: *Proceedings of Asilomar Conference on Signals, Systems and Computers*, 2009, pp. 908–912.
11. Olfati-Saber R, Kalman-consensus filter: Optimality, stability, and performance. In: *Proceedings of the 48th IEEE Conference on Decision and Control*, held jointly with the 28th Chinese Control Conference, 2009, pp. 7036–7042.
12. Cattivelli FS, Sayed AH, Diffusion strategies for distributed Kalman filtering and smoothing. *IEEE Transactions on Automatic Control* 2010; 55(9):2069–2084.
13. Hlinka O, Slučiak O, Hlawatsch F, Djurić PM, Rupp M, Likelihood consensus and its application to distributed particle filtering. *IEEE Transactions on Signal Processing* 2012; 60(8):4334–4349.
14. Djurić PM, Wang Y, Distributed Bayesian learning in multiagent systems: Improving our understanding of its capabilities and limitations. *IEEE Signal Processing Magazine* 2012; 29(2):65–76.
15. Sayed AH, Adaptation, learning, and optimization over networks. *Foundations and Trends® in Machine Learning* 2014; 7(4-5):311–801.
16. Zhao F, Guibas LJ, *Wireless Sensor Networks: An Information Processing Approach*. Amsterdam, The Netherlands: Morgan Kaufmann, 2004.
17. Hlinka O, Hlawatsch F, Djurić PM, Distributed particle filtering in agent networks: A survey, classification, and comparison. *IEEE Signal Processing Magazine* 2013; 30(1):61–81.
18. Sayed AH, Adaptive networks. *Proceedings of the IEEE* 2014; 102(4):460–497.
19. Anderson BDO, Moore JB, *Optimal Filtering*. Englewood Cliffs, NJ: Prentice Hall, 1979.
20. Arulampalam MS, Maskell S, Gordon N, Clapp T, A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing* 2002; 50(2):174–188.
21. Li XR, Jilkov VP, Survey of maneuvering target tracking. Part I. Dynamic models. *IEEE Transactions on Aerospace and Electronic Systems* 2003; 39(4):1333–1364.
22. Tu SY, Sayed AH, Diffusion strategies outperform consensus strategies for distributed estimation over adaptive networks. *IEEE Transactions on Signal Processing* 2012; 60(12):6217–6234.
23. Towfic ZJ, Chen J, Sayed AH, Excess-risk of distributed stochastic learners. *IEEE Transactions on Information Theory* 2016; 62(10):5753–5785.
24. Xiao L, Boyd S, Lall S, A scheme for robust distributed sensor fusion based on average consensus. In: *Proceedings of the 4th ACM/IEEE International Symposium on Information Processing in Sensor Networks*, Los Angeles, CA, 2005, pp. 63–70.
25. Olfati-Saber R, Fax JA, Murray RM, Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE* 2007; 95(1):215–233.
26. Aysal TC, Yildiz ME, Sarwate AD, Scaglione A, Broadcast gossip algorithms for consensus.

40 CHAPTER 1 Chapter Title

- IEEE Transactions on Signal Processing 2009; 57(7):2748–2761.
27. Dimakis AG, Kar S, Moura JMF, Rabbat MG, Scaglione A, Gossip algorithms for distributed signal processing. *Proceedings of the IEEE* 2010; 98(11):1847–1864.
 28. Kailath T, Sayed AH, Hassibi B, *Linear Estimation*. Prentice Hall Upper Saddle River, NJ, 2000.
 29. Sayed AH, *Adaptive Filters*. Wiley, 2008.
 30. Chang KC, Saha RK, Bar-Shalom Y, On optimal track-to-track fusion. *IEEE Transactions on Aerospace and Electronic Systems* 1997; 33(4):1271–1276.
 31. Atherton DP, Bather JA, Briggs AJ, Data fusion for several Kalman filters tracking a single target. *IEE Proceedings – Radar, Sonar and Navigation* 2005; 152(5):372–376.
 32. Mitchell HB, *Multi-sensor Data Fusion: An Introduction*. Springer, 2007.
 33. Julier SJ, Uhlmann JK, A non-divergent estimation algorithm in the presence of unknown correlations. In: *Proceedings of the American Control Conference*, vol. 4, vol. 4, 1997, pp. 2369–2373.
 34. Uhlmann JK, Covariance consistency methods for fault-tolerant distributed data fusion. *Information Fusion* 2003; 4(3):201–215.
 35. Hu J, Xie L, Zhang C, Diffusion Kalman filtering based on covariance intersection. *IEEE Transactions on Signal Processing* 2012; 60(2):891–902.
 36. Kay SM, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.
 37. Djurić PM, Khan M, Johnston DE, Particle filtering of stochastic volatility modeled with leverage. *IEEE Journal of Selected Topics in Signal Processing* 2012; 6(4):327–336.
 38. Hlinka O, Hlawatsch F, Distributed particle filtering in the presence of mutually correlated sensor noises. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 6269–6273.
 39. Moldaschl M, Gansterer WN, Hlinka O, Meyer F, Hlawatsch F, Distributed decorrelation in sensor networks with application to distributed particle filtering. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE, 2014, pp. 6117–6121.
 40. Urteaga I, Djurić PM, Sequential estimation of hidden ARMA processes by particle filtering: Part I. *IEEE Transactions on Signal Processing* 2017; 65(2):482–493.
 41. Urteaga I, Djurić PM, Sequential estimation of hidden ARMA processes by particle filtering: Part II. *IEEE Transactions on Signal Processing* 2017; 65(2):494–504.
 42. Tanizaki H, *Nonlinear Filters: Estimation and Applications*. Berlin, Germany: Springer, 1996.
 43. Alspach D, Sorenson H, Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Transactions on Automatic Control* 1972; 17(4):439–448.
 44. Julier SJ, Uhlmann JK, Unscented filtering and nonlinear estimation. *Proceedings of the IEEE* 2004; 92(3):401–422.
 45. van der Merwe R, *Sigma-point Kalman filters for probabilistic inference in dynamic state-space models*. Ph.D. thesis, OGI School of Science and Engineering, Oregon Health and Science University, Hillsboro, OR, 2004.
 46. Arasaratnam I, Haykin S, Cubature Kalman filters. *IEEE Transactions on Automatic Control* 2009; 54(6):1254–1269.
 47. Gordon NJ, Salmond DJ, Smith AFM, Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)* 1993; 140(2):107–113.
 48. Djurić PM, Kotecha JH, Zhang J, Huang Y, Ghirmai T, Bugallo MF, et al., Particle filtering. *IEEE Signal Processing Magazine* 2003; 20(5):19–38.

49. Cappé O, Godsill SJ, Moulines E, An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE* 2007; 95(5):899–924.
50. Bugallo MF, Elvira V, Martino L, Luengo D, Míguez J, Djurić PM, Adaptive importance sampling: The past, the present, and the future. *IEEE Signal Processing Magazine* 2017; 34(4):60–79.
51. Li T, Bolic M, Djurić PM, Resampling methods for particle filtering: Classification, implementation, and strategies. *IEEE Signal Processing Magazine* 2015; 32(3):70–86.
52. Oreshkin BN, Coates MJ, Asynchronous distributed particle filter via decentralized evaluation of Gaussian products. In: *Proceedings of FUSION*, 2010.
53. Farahmand S, Roumeliotis SI, Giannakis GB, Set-membership constrained particle filter: Distributed adaptation for sensor networks. *IEEE Transactions on Signal Processing* 2011; 59(9):4122–4138.
54. Üstebay D, Coates M, Rabbat M, Distributed auxiliary particle filters using selective gossip. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, 2011, pp. 3296–3299.
55. Mohammadi A, Asif A, Consensus-based distributed unscented particle filter. In: *Proceedings of IEEE Statistical Signal Processing Workshop*, 2011, pp. 237–240.
56. Hlinka O, Hlawatsch F, Djurić PM, Consensus-based distributed particle filtering with distributed proposal adaptation. *IEEE Transactions on Signal Processing* 2014; 62(12):3029–3041.
57. Bordin CJ, Bruno MGS, Consensus-based distributed particle filtering algorithms for cooperative blind equalization in receiver networks. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, 2011, pp. 3968–3971.
58. Bandyopadhyay S, Chung SJ, Distributed estimation using Bayesian consensus filtering. In: *Proceedings of American Control Conference*, 2014, pp. 634–641.
59. Mohammadi A, Asif A, Distributed consensus + innovation particle filtering for bearing/range tracking with communication constraints. *IEEE Transactions on Signal Processing* 2015; 63(3):620–635.
60. Wang H, Djurić PM, Diffusion in networks by cooperative particle filtering. In: *Proceedings of the IEEE Workshop on Computational Advances in Multi-sensor Adaptive Processing*, 2017.
61. Yu JY, Coates MJ, Rabbat MG, Blouin S, A distributed particle filter for bearings-only tracking on spherical surfaces. *IEEE Signal Processing Letters* 2016; 23(3):326–330.
62. Li J, Nehorai A, Distributed particle filtering via optimal fusion of Gaussian mixtures. *IEEE Transactions on Signal and Information Processing over Networks* 2017; To be published.
63. Vázquez MA, Míguez J, A robust scheme for distributed particle filtering in wireless sensors networks. *Signal Processing* 2017; 131:190–201.
64. Savic V, Wymeersch H, Zazo S, Belief consensus algorithms for fast distributed target tracking in wireless sensor networks. *Signal Processing* 2014; 95:149–160.
65. Gautschi W, *Orthogonal Polynomials: Computation and Approximation*. Oxford, U.K.: Oxford University Press, 2004.
66. Unser M, Splines: A perfect fit for signal and image processing. *IEEE Signal Processing Magazine* 1999; 16(6):22–38.
67. Mosk-Aoyama D, Shah D, Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory* 2008; 54(7):2997–3007.
68. Xiao L, Boyd S, Fast linear iterations for distributed averaging. *Systems & Control Letters* 2004; 53(1):65–78.
69. Björck Å, *Numerical Methods for Least Squares Problems*. Philadelphia, PA: SIAM, 1996.

42 **CHAPTER 1** Chapter Title

70. Lawson CL, Hanson RJ, Solving Least Squares Problems. Philadelphia, PA: SIAM, 1995.
71. Bishop CM, Pattern Recognition and Machine Learning. New York, NY: Springer, 2006.
72. Kotecha JH, Djurić PM, Gaussian particle filtering. *IEEE Transactions on Signal Processing* 2003; 51(10):2592–2601.
73. Bolić M, Athalye A, Hong S, Djurić PM, Study of algorithmic and architectural characteristics of Gaussian particle filters. *Journal of Signal Processing Systems* 2010; 61(2):205–218.
74. Gales MJF, Airey SS, Product of Gaussians for speech recognition. *Computer Speech & Language* 2006; 20(1):22–40.
75. Ito K, Xiong K, Gaussian filters for nonlinear filtering problems. *IEEE Transactions on Automatic Control* 2000; 45(5):910–927.
76. Horn RA, Johnson CR, *Matrix Analysis*. Cambridge University Press, 1990.
77. Sayed AH, Diffusion adaptation over networks. In: Chellapa R, Theodoridis S, editors, *Academic Press Library in Signal Processing*, vol. 3, vol. 3, Academic Press, Elsevier, 2014; pp. 323–454.