LUND UNIVERSITY

Scalable Distributed Kalman Filtering for Mass-Spring Systems

Henningsson, Toivo; Rantzer, Anders

2007

[Link to publication](#)

Citation for published version (APA):

Henningsson, T., & Rantzer, A. (2007). *Scalable Distributed Kalman Filtering for Mass-Spring Systems*. Paper presented at 46th IEEE Conference on Decision and Control, 2007, New Orleans, LA, United States.

Total number of authors:

2

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Fig. 2. Graph of a cycle, with nodes labeled from $\mathbb{Z}_n$. Neighbors are nodes with labels that differ by one (modulo n). The cycle is invariant with respect to translation (rotation).

integers modulo n , and the system looks the same when the nodes are shifted cyclically, see Fig. 2.

Let $x \in \mathbb{R}^n$ contain one scalar for each node (such as a state, input or output). An operator (matrix) A acting on x is *translation invariant* if A looks the same no matter what node is chosen as number 0. This property is equivalent to that A is *circulant* (see [4]), i.e. the columns of A are shifts of the same vector $a \in \mathbb{R}^n$:

$$A = \begin{pmatrix} a_0 & a_{-1} & a_{-2} & \dots & a_1 \\ a_1 & a_0 & a_{-1} & \dots & a_2 \\ a_2 & a_1 & a_0 & \dots & a_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{-1} & a_{-2} & a_{-3} & \dots & a_0 \end{pmatrix}$$

where $a_{-k} = a_{n-k}$. A then acts as a periodic spatial convolution with a as impulse response or *convolution kernel*. If A is symmetric then a is symmetric: $a_{-j} = a_j$ for all j .

B. Modeling Assumptions

Assume that the symmetric M, D and K matrices are circulant. Let the applied force f in (1) be

$$f = u + Mw_p$$

where u is the control input and w_p is unknown process noise. The gain M on w_p is not essential (any translation invariant gain can be used), but will make the presentation cleaner. Position sensors at each node give the measurements

$$y = z + w_m \quad (2)$$

where w_m is measurement noise. The noises w_p and w_m are for simplicity assumed Gaussian, independent identically distributed (i.i.d.) in space, and white noise in time for all frequencies of interest, with incremental variance σ_p^2 and σ_m^2 .

Example 1 (Finite Element Method (FEM) Model):

Beginning with the wave equation on a homogeneous string

$$m\ddot{z} - d\frac{\partial^2 \dot{z}}{\partial x^2} - k\frac{\partial^2 z}{\partial x^2} = f$$

applying a Galerkin discretization with piecewise linear finite elements (see [8], the example in ch. 18, sec. 2, p. 476 contains essentially the same discretization), and choosing appropriate scaling factors we arrive at the model matrices

$$M = \frac{1}{2}\text{circ} \begin{bmatrix} 1 & 4 & 1 \end{bmatrix}, \quad D = \delta K, \quad K = \frac{1}{2}\text{circ} \begin{bmatrix} -1 & 2 & -1 \end{bmatrix}$$

where δ is a damping parameter and $A = \text{circ}[a_{-1} \ a_0 \ a_1]$ is tridiagonal circulant matrix with values a_{-1}, a_0, a_1 . Each node is coupled only to its immediate neighbors.

III. MODAL TRANSFORMATION

To proceed, we shall exploit that M, D and K can be diagonalized by a common unitary matrix S . Any circulant matrix $A \in \mathbb{R}^{n \times n}$ acting on some $x \in \mathbb{R}^n$ can be transformed into a diagonal matrix $D = S^*AS$ by the Discrete Fourier Transform $y = S^*x$, with $S = S(n)$ given by

$$S_{jk}(n) = \frac{1}{\sqrt{n}} e^{i\frac{2\pi kj}{n}} = \frac{1}{\sqrt{n}} e^{i\kappa(k)j}. \quad (3)$$

The modal operator D is completely described by its diagonal $\text{diag}(D) = d = \sqrt{n}S^*a$.

The columns of S make up an orthonormal basis of eigenvectors or *modes* of A . They can be indexed by the *spatial frequency* κ in (3) that determines how fast the mode shape changes as the position j increases:

$$\kappa = 2\pi k/n, \quad k \in \mathbb{Z}, |\kappa| \leq \pi$$

where the spatial Nyquist Frequency at $\kappa = \pi$ is the highest that can be represented across the nodes.

A. Model Transformation

The system will now be decoupled by transformation to modal coordinates. The transformation $z' = S^*z$ applied to the nodal dynamics (1) yield the decoupled modal dynamics

$$M'\ddot{z}' + D'\dot{z}' + K'z' = f' \quad (4)$$

where $M' = S^*MS > 0, D' = S^*DS \geq 0$, and $K' = S^*KS \geq 0$ are modal mass, damping and stiffness matrices that are diagonal, and the applied modal forces are

$$f' = u' + M'w'_p = S^*u + M'S^*w_p$$

where w'_p is the modal process noise. The nodal measurements (2) are transformed by $y' = S^*y$ into the modal

$$y' = z' + w'_m$$

where $w'_m = S^*w_m$ is the modal measurement noise. The modal noises w'_p and w'_m have the same distributions as w_p and w_m since these are i.i.d. Gaussian and S is unitary.

Multiplying by $(M')^{-1}$, the modal dynamics (4) become

$$\ddot{z}' + \Gamma\dot{z}' + \Lambda z' = (M')^{-1}u' + w'_p$$

where $\Gamma = (M')^{-1}D'$ and $\Lambda = (M')^{-1}K'$ are diagonal.

B. State Space Form

Dropping the primes for readability and introducing the state variable $x = (z^T \ v^T)^T = (z^T \ \dot{z}^T)^T$ the modal system can be written in state space form as

$$\begin{aligned} \dot{x} &= \underbrace{\begin{pmatrix} 0 & I \\ -\Lambda & -\Gamma \end{pmatrix}}_A x + \underbrace{\begin{pmatrix} 0 \\ I \end{pmatrix}}_B (M^{-1}u + w_p), \\ y &= \underbrace{\begin{pmatrix} I & 0 \end{pmatrix}}_C x + w_m. \end{aligned} \quad (5)$$

Fig. 3. Frequency response of the operators m , k and $\omega = \sqrt{k/m}$ as a function of spatial frequency κ . The linear fit ω_p of $|\kappa|$ versus ω shows that ω is quite linear in $|\kappa|$ and thus the wave speed is rather constant.

The modes are uncoupled, so with $x \in \mathbb{R}^2$ and u, y, w_p and w_m scalar each mode can be considered separately as

$$\begin{aligned} \dot{x} &= \underbrace{\begin{pmatrix} 0 & 1 \\ -\lambda & -\gamma \end{pmatrix}}_A x + \underbrace{\begin{pmatrix} 0 \\ 1 \end{pmatrix}}_B (m^{-1}u + w_p), \\ y &= \underbrace{\begin{pmatrix} 1 & 0 \end{pmatrix}}_C x + w_m \end{aligned} \quad (6)$$

keeping in mind that the parameter triple (m, γ, λ) will vary over the modes. Kalman Filter design for the entire system reduces to design for a parametrized family of low-order systems, as remarked in [2]. The same goes for controllability/observability analysis, controller design, performance evaluation etc. With a large number of similar modes it might suffice to do all calculations for a representative subset, if one accounts for the minor added uncertainty.

The resonance frequency $\omega = \lambda^{1/2}$ and the damping γ are the essential parameters that differentiate the modes. We may expect that a mode's resonance frequency ω is an increasing function of $|\kappa|$, which is the situation for the wave equation where $\omega = c|\kappa|$ and c is the wave speed.

Example 2 (Modal operators): The modal form of the FEM model operators can be found by letting them act on a mode $x_j = e^{i\kappa j}x_0$ with spatial frequency κ . We get

$$m = 2 + \cos(\kappa), \quad k = 1 - \cos(\kappa), \quad d = \delta(1 - \cos(\kappa)).$$

Fig. 3 shows m, k and $\omega = \lambda^{1/2}$ as a function of κ . Also shown is the best linear fit ω_p of $|\kappa|$ versus ω , to compare with the wave equation. We see that m drops and k rises for increasing $|\kappa|$, and that ω is quite linear with $|\kappa|$.

IV. KALMAN FILTERING

We will investigate Kalman Filtering one mode at a time. The filter estimates the state x through the filter dynamics

$$\dot{\hat{x}} = A\hat{x} + Bm^{-1}u + L(y - C\hat{x}) \quad (7)$$

where $\hat{x}$ is the state estimate, L the filter gain, and A, B, C are as in (6). The filter is designed by letting

$$L = PC^T R_m^{-1} \quad (8)$$

where the covariance matrix of the estimation error

$$P = \begin{pmatrix} p_{zz} & p_{zv} \\ p_{zv} & p_{vv} \end{pmatrix}$$

satisfies the Riccati equation

$$AP + PA^T + R_p - PC^T R_m^{-1} CP = 0 \quad (9)$$

where $R_p = B\sigma_p^2 B^T$, and $R_m = \sigma_m^2$. Note that in this case $L = (l_z \ l_v)^T = PC^T \sigma_m^{-2} = \sigma_m^{-2} (p_{zz} \ p_{zv})^T$.

Let $\tilde{x} = x - \hat{x}$ be the estimation error. From (6) and (7)

$$\dot{\tilde{x}} = (A - LC)\tilde{x} + Bw_p - LCw_m$$

which is a linear filter fed with noise w_p and w_m . The error covariance P can be found from the Lyapunov Equation

$$(A - LC)P + P(A - LC)^T + R_p + LR_m L^T = 0 \quad (10)$$

for arbitrary L . P will be minimal with L according to the Riccati solution; smaller L will not attenuate $\tilde{x}$ fast enough and larger L will feed too much measurement noise into $\tilde{x}$.

A. Modal State Estimation

We look at the Kalman Filter for zero damping γ ; this should be the hardest case since the system is only nominally stable and no modes are damped well enough to be neglected. Only the parameter $\lambda = \omega^2$ will distinguish the modes.

The Riccati equation (9) can be solved analytically, giving

$$\begin{aligned} \frac{p_{zz}}{\sigma_m^2} &= \sqrt{2 \frac{p_{zv}}{\sigma_m^2}} & \frac{p_{zv}}{\sigma_m^2} &= \frac{\lambda_0^2}{\sqrt{\lambda^2 + \lambda_0^2} + \lambda} \\ \lambda_0 &= \omega_0^2 = \frac{\sigma_p}{\sigma_m} & \frac{p_{vv}}{\sigma_m^2} &= \sqrt{\lambda^2 + \lambda_0^2} \frac{p_{zz}}{\sigma_m^2} \end{aligned} \quad (11)$$

where the scaling parameter λ_0 describes the balance between process and measurement noise intensity. It gives the break frequency between slow modes where process noise dominates and fast modes where measurement noise dominates, and the filter convergence rate for slow modes.

Fig. 4 shows P with the choice $\omega_0 = 0.2$, with coordinate scales normalized so that $p_{zz}(\omega = 0) = p_{vv}(\omega = 0) = 1$. The appearance of the plot only depends on ω_0 , which sets the frequency scale. Filter behavior is qualitatively different for modes below and above ω_0 . When $\omega \ll \omega_0$, P and therefore L are almost constant; the covariance tends to

$$P(\omega = 0) = \sqrt{2} \sigma_p \sigma_m \begin{pmatrix} \omega_0^{-1} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \omega_0 \end{pmatrix}$$

and the filter characteristic polynomial tends to $s^2 + \sqrt{2}\omega_0 s + \omega_0^2$ with time constant $T = \sqrt{2}/\omega_0$ and damping $\zeta = 1/\sqrt{2}$. Apparently the value of ω is unimportant as long as it is much slower than the filter.

When $\omega > \omega_0$, p_{zz} and p_{zv} which correspond to l_z and l_v fall off, while p_{vv} rises. The covariance tends to

$$P(\omega \gg \omega_0) = \sigma_p \sigma_m \begin{pmatrix} \omega^{-1} & \frac{\omega_0^2}{2\omega^2} \\ \frac{\omega_0^2}{2\omega^2} & \omega \end{pmatrix}$$

and the characteristic polynomial tends to $s^2 + \frac{\omega_0^2}{\omega} s + \omega^2$ with time constant $T = 2\omega/\omega_0^2 \gg \omega_0^{-1}$ and damping

Fig. 4. Normalized Kalman Filter estimation error covariances as a function of modal resonance frequency ω , with break frequency $\omega_0 = 0.2$. The filter gains $l_z = \sigma_m^{-2} p_{zz}$, $l_v = \sigma_m^{-2} p_{zv}$ are proportional to the two lower curves.

$\zeta = \omega_0^2/(2\omega^2) \ll 1$. For these modes the filter convergence is much slower than the process dynamics. The filter gain L and displacement error p_{zz} fall off since it is hard for a white noise force to excite a high frequency mode. The covariance p_{zv} drops since the states are estimated over many cycles. The velocity error p_{vv} rises since large differences in velocity amplitude are only observable as small differences in displacement. With $\delta > 0$ damping would eventually force all p_{jk} down for large enough ω since it limits the excitation.

If ω is monotonic in $|\kappa|$ as we expect, the filter gains l_z and l_v are spatial low pass filters with cutoff at around $\kappa_0 = \kappa(\omega = \omega_0)$, e.g. $\kappa_0 = \omega_0/c$ for the wave equation.

V. SPATIAL REALIZATION

Once a Kalman Filter is designed in the modal domain, it can be realized in the nodal domain through inverse transformation. The nodal realizations will not be sparse in general, but the coupling between nodes will fall off rapidly with distance allowing good sparse approximations.

A. Exact Realization

The modal filter dynamics (7) for all modes are

$$\dot{\hat{x}}' = A' \hat{x}' + B'(M')^{-1} u' + L'(y' - C' \hat{x}')$$

with A, B and C according to (5), $L = (L_z \ L_v)^T$ the collection of modal filter gains (L_z, L_v diagonal) and the primes put back for clarity. The corresponding nodal dynamics are

$$\dot{\hat{x}} = \underbrace{\begin{pmatrix} 0 & I \\ -\tilde{K} & -\tilde{D} \end{pmatrix}}_A \hat{x} + \underbrace{\begin{pmatrix} 0 \\ I \end{pmatrix}}_B M^{-1} u + \underbrace{\begin{pmatrix} L_z \\ L_v \end{pmatrix}}_L (y - \underbrace{(I \ 0)}_C \hat{x})$$

where $\tilde{K} = M^{-1}K$ and $\tilde{D} = M^{-1}D$ are symmetric and circulant, $L_z = SL'_z S^*$ and $L_v = SL'_v S^*$ are spatial filter gains (convolutions), $\hat{x} = (\hat{z}^T \ \hat{v}^T)^T$ is the vector of nodal state estimates and y the nodal measurements. L_z and L_v are convolutions that feed the measurement error from each node to its neighborhood.

The matrices M^{-1} and L are not sparse, but if the problem data M^{-1} fades exponentially with distance then L will

Fig. 5. Normalized spatial realization of L_z, L_v, M^{-1} with $\kappa_0 = 0.15\pi$.

Fig. 6. Normalized spatial $|l_z|, |l_v|$ and $|m^{-1}|$. The dashed line shows that the filter coefficients fall off exponentially approximately as $e^{-\kappa_1 |j|}$, $\kappa_1 = 0.78\kappa_0$; however there is an interval of quicker falloff in the beginning.

decay exponentially with distance at least asymptotically as shown in [2], and should be amenable to spatial truncation.

Example 3 (Spatial realization of L and M^{-1}): Figs. 5 and 6 show the convolution kernels of L_z, L_v and M^{-1} for the FEM model with $\omega_0 = 0.2$ and $\kappa_0 \approx 0.15\pi$, normalized so that the largest element is 1. All three kernels do in fact decay exponentially in space, so a narrow truncation should suffice to come close to optimal performance.

B. Spatial Truncation

Let $\mathcal{N}_r$ be the space of symmetric convolution kernels with support only on indices $|j| \leq r \pmod{n}$. A convolution kernel a can be approximated by a kernel $b \in \mathcal{N}_r$ by setting all elements with index $|j| > r \pmod{n}$ to zero. This *spatial truncation* is an orthogonal projection onto $\mathcal{N}_r$.

The effect of spatial truncation as seen from the modal domain is to truncate the cosine series expansion of the frequency response a' after $r+1$ terms, which might be acceptable if a' varies only slowly over the modes.

C. Truncation of Filter Gains

To implement an observer with sparse matrices, L_z, L_v , and M^{-1} must be approximated. For brevity we will not consider the approximation of M^{-1} . The potentially most

Fig. 7. Normalized full l_z and l_v compared to truncated. The $\mathcal{N}_2$ approximation of l_v is negative on the lobe around $\omega = 1$.

important influence of errors in L_z and L_v is to perturb the modal observer dynamics with characteristic polynomial

$$\det(sI - (A - LC)) = s^2 + (\gamma + l_z)s + (\lambda + l_v + \gamma l_z).$$

The gain l_v is important for low ω but is dominated by λ when ω is high since it can be seen from (11) that nominally $l_v/\lambda \leq \lambda_0^2/(2\lambda^2)$. The gain l_z provides damping and is needed as long as γ is small — it should never become negative since the filter dynamics could be destabilized. As can be seen from (10) the other effect of changes in L is the change of the measurement noise injection $LR_m L^T$.

Example 4 (Spatial truncation): Fig. 7 shows modal optimal l_z and l_v together with $\mathcal{N}_2'$ and higher order approximations ($\mathcal{N}_5'$ and $\mathcal{N}_9'$ respectively). Fig. 8 compares the filter error covariances for optimal L and truncation to $\mathcal{N}_2'$.

Although Fig. 7 shows that l_z and l_v can be well approximated with modest nodal support, Fig. 8 shows that already with the narrow $\mathcal{N}_2'$ truncation performance is very close to optimal in this case. The most apparent deviations are a small deterioration for low frequencies caused by error in l_v , and a small bump around $\omega = 1$ where a value of l_z of about half the optimal degrades damping.

With this insight into the sensitivities on l_v and l_z , they could be adjusted for improved performance. Too much damping is generally better than too little. We note however that already simple spatial truncation works very well.

While truncation changes the total error covariance only modestly, the source of error can shift so that some estimates are mostly sensitive to process noise w_p and others to measurement noise w_m . The design is scalable in that the same gains with the same sparse structure for each node works just as well independent of the number of nodes n .

VI. SYSTEMS WITH PARTIAL SENSING

Kalman Filtering for systems with sensors placed only periodically on the nodes will now be investigated. Let $n = mq$ with m, q positive integers and let the restriction (downsampling) operator $R \in \mathbb{R}^{m \times n}$ be such that $(Rx)_j = x_{qj \pmod n}$. The restriction operator corresponds to periodic downsampling of a signal from $\mathbb{R}^n$ to $\mathbb{R}^m$. The prolongation operator R^T injects sampled values back to the original grid.

Let $S_m^* = S(m)^*$ according to (3) be the modal transform on $\mathbb{R}^m$. The modal restriction operator R' is given by $R' =$

Fig. 8. Normalized estimation error covariance comparison for optimal and spatially truncated Kalman Filter with 2 nearest neighbors kept in the filter gain. The covariances for the truncated filter lie slightly above the optimal.

Fig. 9. Graph of periodic string with partial sensing. In this example sensors are placed on every other node, colored white in the figure.

$S_m^* R S$, and has the effect of *aliasing* together modes into groups of q . Values for aliased modes are added together and scaled by R' , and the averages injected back by $(R')^T$.

A. Kalman Filter Structure

Let there be sensors only once every q nodes, as in Fig. 9. In the modal domain the dynamics are the same as in (5) while the measurements take the form

$$y_R = Ry = R \begin{pmatrix} I & 0 \end{pmatrix} x + w_m$$

where the (modal) restriction operator R downsamples the measurements to $y_R \in \mathbb{R}^m$, and w_m is now in $\mathbb{R}^m$.

The aliasing introduced by R means that only the sum of displacements can be measured for each group of aliased modes. Since the groups have no coupling between them filters can be designed independently for each. When two aliased modes have the same resonance ω there is a loss of observability, and the filter has to rely on Γ for damping.

Example 5 (Partial sensing): Consider the FEM model with position sensors every $q = 2$ nodes. The modes alias in pairs $(\kappa, \kappa_{\max} - \kappa)$ and there is a loss of observability at $\kappa = \kappa_{\max}/2$ corresponding to a standing wave with zeros at all sensors. The unobservable mode has $\omega = \omega_{ns} = 0.707$.

To ensure stability and limit the filter time constant for unobservable and close to unobservable modes we make the physical damping be $\zeta = 0.01$ at ω_{ns} by an appropriate choice of δ . The measurement noise intensity σ_m^2 is halved to give a fair comparison with halved number of sensors.

Fig. 10. Normalized estimation error covariance comparison for Kalman Filter and Kalman Filter approximated by spatial truncation with 2 step neighbors kept in the filter gain, with sensors on every other node. The covariances for the approximated filter lie slightly above the optimal covariances, except for p_{zv} at $\omega > 0.7$ which is slightly lower.

Fig. 11. Normalized translation invariant spatial realization of l_z, l_v with sensors on every other node. The slowly decaying oscillatory component of l_v is expressed mostly on odd shifts j corresponding to nodes without sensors.

Fig. 10 shows estimation error covariances for partial sensing, with full and truncated filter gains. The cross covariances between coupled modes are not shown but are generally small except for ω close to ω_{ns} . The damping causes the covariances to go down for high frequencies. The only significant difference with partial sensor placement is the small but sharp bumps around ω_{ns} where the unobservable mode contributes a big estimation error.

Fig. 11 shows the spatial realization of the optimal filter gains (which can be made convolutions acting on the virtual measurement errors $R^T(y_R - Rz)$). The bumps in Fig. 10 result in slowly decaying components in the nodal filter gains. This is hardly noticeable in l_z but very visible in l_v , which seemingly should be hard to truncate. If ω_0 is high we see from Fig. 10 that the sharp components will be closer to the break point in P and therefore stronger, with greater risk to cause trouble.

Let us compare the estimation errors for full and truncated filter gains in Fig. 10. The truncation leads to small approximation errors, the most noticeable are the same as in Fig. 8. The slowly decaying spatial component of l_v that

was truncated appears insignificant compared to λ . Behavior around the bumps is unaffected, since the observer loop is broken. The results when using truncated l 's from the fully sensed case are virtually indistinguishable from Fig. 10.

VII. CONCLUSION

This paper investigates distributed Kalman Filtering for mass-spring systems on periodic grids. The Kalman Filter for an undamped system is derived analytically and it is argued that actual filter performance is mildly sensitive to perturbations in filter gain. Performance comparison for optimal filters and filters with spatially truncated filter gains shows that the difference can be small even with narrow truncation. The design is scalable in that the same results are achieved with the same sparse gains in the nodes independent of the size of the system.

A system with only periodically placed sensors is also investigated. The example shows that close to optimal filter performance can be achieved through narrow spatial truncation even though the optimal filter gains do not fall off rapidly with distance in this case; indeed truncating the filter gains for the system with full sensing which do fall off rapidly works just as well.

A number of points deserve greater clarification: How are the observability problems caused by missing sensors best quantified? How should the Kalman Filter (in-)sensitivity to perturbation in filter gains be described? How is uncertainty in the process model best handled?

With the properties exhibited here in mind, some directions for further investigation are indicated:

- What are the implications of a non-diagonal measurement structure on the modal state estimation problem?
- How can we generalize to less regular systems?
- Is there a distributed design procedure for approximate distributed Kalman Filters of this kind?

ACKNOWLEDGMENT

The authors would like to thank P. Hagander, P. Alriksson, M. Karlsson, and A. Gattami for valuable feedback.

REFERENCES

- [1] Peter Alriksson and Anders Rantzer. Distributed Kalman filtering using weighted averaging. In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*, Kyoto, Japan, July 2006.
- [2] B. Bamieh, F. Paganini, and M.A. Dahleh. Distributed control of spatially invariant systems. *IEEE Trans. Autom. Control*, 47(7):1091–1107, July 2002.
- [3] Geir E. Dullerud and Raffaello D'Andrea. Distributed control of heterogeneous systems. *IEEE Trans. Autom. Control*, December 2004.
- [4] R. M. Gray. *Toeplitz And Circulant Matrices: A Review*. Now Publishers, 2006. Also at <http://www-ee.stanford.edu/~gray/>.
- [5] Nader Motee and Ali Jadbabaie. Optimal control of spatially distributed systems. *IEEE Trans. Autom. Control*, September 2006.
- [6] Anders Rantzer. A separation principle for distributed control. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, December 2006.
- [7] D.P. Spanos, R. Olfati-Saber, and R.M. Murray. Approximate distributed Kalman filtering in sensor networks with quantifiable performance. In *Fourth Int. Symp. Inform. Process. Sensor Networks*, 2005.
- [8] O. C. Zienkiewicz and R. L. Taylor. *The Finite Element Method*. Butterworth-Heinemann, fifth edition, 2000.