

Optimal Estimation of Deterioration from Diagnostic Image Sequence

Dimitry Gorinevsky, *Fellow, IEEE*, Seung-Jean Kim, *Member, IEEE*, Shawn Beard, Stephen Boyd, *Fellow, IEEE*, and Grant Gordon

Abstract—Estimation of mechanical structure damage can greatly benefit from the knowledge that the damage accumulates irreversibly over time. This paper formulates a problem of estimation of a pixel-wise monotonic increasing (or decreasing) time series of images from noisy blurred image data. Our formulation includes temporal monotonicity constraints and a spatial regularization penalty. We cast the estimation problem as a large-scale quadratic programming (QP) optimization and describe an efficient interior-point method for solving this problem. The method exploits the special structure of the QP and scales well to problems with more than a million of decision variables and constraints. The proposed estimation approach performs well for simulated data. We demonstrate an application of the approach to diagnostic images obtained in structural health monitoring experiments and show that it provides a good estimate of the damage accumulation trend while suppressing spatial and temporal noises.

Index Terms—Optimal estimation, spatio-temporal filtering, regularization, interior-point methods, isotonic regression, structural health monitoring, damage, monotonic

I. INTRODUCTION

We consider a time series of noisy blurred images where the underlying images are constrained to be pixel-wise monotonically increasing (or decreasing). Such problem arises when estimating mechanical structure damage. It is known that mechanical damage accumulates irreversibly over time; this can be described through the monotonicity constraints. Despite its apparent simplicity and usefulness, such formulation has not been studied earlier, to the best of the authors' knowledge.

The main application considered in the paper is structural health monitoring (SHM). This work was initially motivated by interest of commercial aircraft industry in automating, using SHM, mandatory periodic inspections of aircraft structure, see [1], [25], [59]. SHM has been attracting much attention recently; an initial overview of work in this area can be found, for instance, in [2], [55], [56], [57], [58]. Most of the published SHM work is on component technologies; a small portion is on integrated SHM systems producing diagnostic images. The SHM literature considers detection of defects or damage locations from a diagnostic image. Most papers consider *pattern recognition* and *detection* problems: a gray-scale image is the input, a binary damage image or damage

location and size are possible outputs. In contrast, this paper considers damage *estimation* problems: a *time series* of SHM diagnostic images is an input, a series of gray-scale estimate images is the output. The authors are unaware of any prior work, apart from theirs, considering such formulation.

The health monitoring literature distinguishes between diagnostics (detecting and identifying faults) and prognostics (forecasting of remaining useful life). The prognostics requires estimation of the damage state. An overview of current work on structural prognostics can be found in [11], [15]. The existing prognostics work is focused on estimation and prediction (forecasting) of lumped parameters. The formulation herein is related to prognostic estimation for SHM images.

The problems of estimating underlying monotonic changes from noisy images can be also formulated in geophysics (e.g., analysis of earthquake precursors, petroleum extraction), in medical imaging (e.g., a growing tumor imaged at different times), in environmental sciences (e.g., irreversible changes caused by global warming trends), and other areas.

Deblurring of noisy images is covered in many textbooks, e.g., [14], [18], [35]. Particularly relevant to this work are the papers [20], [22], [24], [25], where linear deblurring filters are designed in spatial frequency domain. Taking into account the monotonicity of damage accumulation leads to nonlinear filtering problems.

There is substantial earlier work on lumped nonlinear estimation with monotonicity constraints. It includes 'isotonic regression' work that was driven by applications in statistics and operations research and is summarized in the books [4], [49]. Monotonicity constraints in signal processing problems are also considered in [19], [21], [48], [54]. Maximum a posteriori probability (MAP) estimation assuming a monotonic walk prior model leads to (convex) QP problems, which can be solved efficiently. For monotonic signals, the filtering approach based on constrained optimization provides a substantial improvement over standard linear filtering methods.

Most of the prior work on estimation with the monotonicity constraints considers univariate data, some considers multivariate data, and none deals with image data. This paper considers a series of images monotonic in time. The MAP problem is formulated as a large scale QP problem with the linear monotonicity constraints. The novelty of our formulation is that it includes the spatial dimensions of the problem and the monotonicity constraints in the time dimension.

The image processing problem in this paper can be solved using standard QP solvers when the total number of decision variables (i.e., the pixels in the image series) is modest, say, under 10,000. The SHM application in Section VI of this

Copyright (c) 2009 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org

D. Gorinevsky, S.-J. Kim, and S. Boyd are with Information Systems Lab, Dept. of Electrical Engineering, Stanford University, Stanford, CA 94305

S. Beard is with Acellent Technologies Inc., Sunnyvale, CA 94085

G. Gordon is with Honeywell Laboratories, Phoenix, AZ 85027

Manuscript received March 14, 2008. Revised on July 14, 2008 and on October 06, 2008. This work was supported by the NSF Grant ECS-0529426

paper has around a million decision variables, well beyond the capacity of standard QP solvers.

A number of prior papers formulate image processing problems as QPs and other constrained convex optimization problems, e.g., see [16], [17], [29], [37], [39]. Most relevant to the solution method described in this paper is the primal-dual interior-point method described in [31], which was successfully applied to medical image estimation problems of similar size (about a million variables). In this paper we use a specialized interior-point method for solving the QP problems with l_1 regularization penalties and monotonicity constraints. The method customizes a truncated Newton interior-point method described in [32] to estimation problems involving a time series of two-dimensional data. Similar to [31], we compute the search step using a preconditioned conjugate gradient (PCG) method.

The solution method described in this paper improves on [31], [32] by exploiting the special structure of the problem. We do not explicitly form Hessian and other large scale sparse matrices. Instead, we use convolutions with spatial kernels for efficient application of these operators. This leads to faster computations and greatly reduced memory requirements. A simple Matlab implementation of the method, described in more detail in Section III, can solve the QP problems, with more than two million variables and two million constraints that arise in SHM applications, on a PC in a few tens of minutes.

The main contribution of this paper is in demonstrating a practical solution to the problem of estimating monotonically accumulating damage from a series of images. The contribution spans several knowledge areas. Our secondary contributions in these areas are summarized below.

First, we introduce a new type of estimation problems for a time series of noisy images where the underlying signal is pixel-wise monotonic in time. Apparently, despite its usefulness, such problem formulation has not been considered earlier.

Second, we propose an interior-point PCG method for solving large QP problems of the formulated type. We demonstrate that by exploiting the problem structure (which is described by convolution operators), the method scales almost linearly with the problem size both in computational effort and in memory requirements.

Third, we develop an approach to selecting (tuning) a sparse spatial regularization penalty in the optimization problem such that the solution satisfies engineering specification for noise rejection and estimation error. Tuning the regularization penalty in frequency domain is new, though, in spirit, the approach is related to earlier work on frequency-domain tuning of spatial filters.

Finally, we demonstrate an important application to structural health monitoring. The problem setup includes collecting a series of diagnostic image data and estimating spatio-temporal signal of the underlying damage from this noisy data. Such problem formulation in itself is novel in SHM, even apart from the new estimation approach used. We demonstrate estimation of the underlying damage from experimental SHM data.

Section II gives a mathematical formulation of the estimation problem. Section III describes the specialized interior-point method for this problem. Section IV discusses selection of the regularization penalty. Section V verifies the efficiency and performance of the approach for simulated data with known ground truth. Section VI validates the approach in SHM application with experimental data.

II. OPTIMIZATION-BASED FORMULATION

This section introduces the problem formulation.

A. Estimation problem

Consider an observed data set Y comprising a sequence of N_t diagnostic images $Y(t) \in \mathfrak{R}^{N_1, N_2}$:

$$Y = \{Y(1), \dots, Y(N_t)\}. \quad (1)$$

We assume that damage intensity at each spatial location (pixel) is described by a real number. A truth data set X comprises a corresponding sequence of underlying damage maps $X(t) \in \mathfrak{R}^{N_1, N_2}$:

$$X = \{X(1), \dots, X(N_t)\}. \quad (2)$$

We assume a linear model of the diagnostic imaging system

$$Y(t) = \mathbf{B}X(t) + e(t), \quad (3)$$

where \mathbf{B} is a linear operator $\mathfrak{R}^{N_1, N_2} \rightarrow \mathfrak{R}^{N_1, N_2}$ and $e(t) \in \mathfrak{R}^{N_1, N_2}$ is the observation noise.

We consider the problem of estimating X from Y . Considering (3) as an inverse problem and directly solving it for $X(t)$ while assuming $e(t) = 0$ does not work for an ill-conditioned blur operator \mathbf{B} . A standard approach is to regularize the problem with a penalty on the size of nonsmoothness of X . Such a regularization approach is commonly used in image processing. We explicitly incorporate the knowledge about the damage being irreversible, by introducing pixel-wise monotonicity constraints, $X(t) \geq X(t-1)$, where $A \geq B$ means that $A_{ij} \geq B_{ij}$ for all i and j .

The truth data set X (2) can be estimated by solving the following constrained optimization problem:

$$L = \frac{1}{2} \sum_{t=1}^{N_t} \|Y(t) - \mathbf{B}X(t)\|_F^2 + \frac{1}{2} \sum_{t=1}^{N_t} (X(t), \mathbf{R}X(t)) + \rho \sum_{t=2}^{N_t} \|X(t) - X(t-1)\|_1 \rightarrow \min \quad (4)$$

$$\text{subject to } X(t) - X(t-1) \geq 0, \quad t = 2, \dots, N_t \quad (5)$$

where (U, V) denotes a dot product of the two images U and V considered as flat vectors, $\|U\|_F^2 = (U, U)$, and $\|U\|_1$ is the ℓ_1 -norm of U (i.e., the sum of the absolute values of the pixels).

The first sum in (4) corresponds to the inverse problem of finding $X(t)$ such that $\mathbf{B}X(t) = Y(t)$. This inverse problem is ill-conditioned (sensitive to noise with high spatial frequency) and has to be regularized. The regularized problem including the first two terms in (4) corresponds to spatial Wiener filtering for each of the images $Y(t)$ separately. The choice of spatial

regularization operator $\mathbf{R} : \mathfrak{R}^{N_1, N_2} \rightarrow \mathfrak{R}^{N_1, N_2}$ in the second sum is important; it is discussed in Section IV. The positive scalar ρ in the third sum scales temporal regularization.

The problem (4)–(5) is a convex QP problem, which has a single global optimum. One of advantages of such formulation of the estimation problem is that it works with incomplete data. (The missing data pixels can be just dropped from the sum of the squared errors, the first sum in (4)). Even for a series of images of a moderate size, the QP problem (4)–(5) becomes very large. The SHM data example in Section VI has a series of 24 images 171×171 pixels each, which leads to a QP with about one million of variables and one million constraints.

In a special case, where B and R are identity (or scaled identity) operators, the problem (4)–(5) separates into $N_1 \cdot N_2$ independent problems of pixel-wise time series regressions with the monotonicity constraint. Each monotonic regression problem can be solved independently and efficiently, see [19], [21]. Such solution does not use any spatial information and will not be spatially smooth. In this paper we are interested in the problem that combines the spatial smoothness and temporal monotonicity.

B. Bayesian estimation interpretation

The QP problem (4)–(5) can be interpreted as a Bayesian MAP estimate of the form

$$L = -\log P(Y|X) - \log P(X) \rightarrow \min. \quad (6)$$

The first sum in (4) is the data fit error corresponding to the observation likelihood $-\log P(Y|X)$ in (6). The observation model implied by the first sum in (4) is that the noise $e(t)$ in (3) is zero-mean identically distributed Gaussian noise (white noise) independent for different spatial locations and for different t . The observation model can be easily generalized to include known spatial or temporal correlations of the noise $e(t)$.

The last two sums in (4) express the prior likelihood $-\log P(X)$ in (6); they correspond to spatial and temporal regularization terms. The constraints come from the temporal part of the prior model. It is assumed that the prior probabilities are independent: $P(X) = P_S(X) \cdot P_T(X)$, where the subscripts indicate the spatial and temporal priors, respectively. Such separable stochastic models are commonly used in multidimensional signal processing.

Spatial prior: For the spatial probability structure we use a Gaussian Random Field (GRF) model which leads to a quadratic log-likelihood in the spatial prior

$$-\log P_S(X(t)) = -\frac{1}{2} (X(t), \mathbf{R}X(t)) + C_R \quad (7)$$

where \mathbf{R} is a positive definite operator and C_R is a normalization constant. One special case of the regularization operator is

$$\mathbf{R} = r_0 \mathbf{I}, \quad (8)$$

where r_0 is a scalar and \mathbf{I} is the identity operator. Another special case is

$$\mathbf{R} = r_1 \Delta, \quad (9)$$

where Δ is the Laplace operator. Both identity and Laplace regularization operators are often used in image processing problems.

A generalization of identity and Laplace operators is an operator \mathbf{R} with a symmetric positive definite $(2K + 1)$ -tap finite impulse response (FIR) convolution kernel $R = \{r_{ij}\}_{i,j=-K}^K$. Such a regularization operator \mathbf{R} can be introduced through a Gaussian Markov Random Field (GMRF) prior probability structure for the underlying damage. A GMRF probability structure specifies the conditional probability for a pixel through its neighborhood and can be expressed in the form

$$x_{jk} = r_{00}e_{jk} + \sum_{l,m=-K; l,m \neq 0}^K r_{lm}x_{j+l, k+m} \quad (10)$$

where e_{jk} are white noise variables. In the context of damage estimation, the model (10) expresses that a damage at a given spatial location is correlated with the neighboring location damages.

Temporal prior: To explain the temporal prior $P_T(X)$, we consider the underlying damage for a single image pixel. We omit the pixel index in the following discussion. We consider a first-order monotonic random walk model

$$x(t+1) = x(t) + \xi(t), \quad (11)$$

where $\xi(t)$ is independent process noise with an identical exponential distribution

$$p_\xi(x) = \begin{cases} e^{-x/\rho}, & x \geq 0 \\ 0, & x < 0. \end{cases} \quad (12)$$

The increments $\xi(t)$ have zero probability of being negative. This model reflects the fact that the damage is accumulating irreversibly. Such a monotonic damage accumulation model corresponds to the Palmgren-Miner rule used in mechanical damage analysis. The reader is referred to [19], [21] for more details and references.

The temporal prior probability for the pixel can be expressed as a product of the independent probabilities of the increments (we assume that no prior for $x(1)$ is available)

$$-\log P_T(\{x(1), \dots, x(N)\}) = -\sum_{t=2}^N \log p_\xi(x(t) - x(t-1)) = \rho \sum_{t=2}^N |x(t) - x(t-1)|, \quad (13)$$

In accordance with (12), the second equality holds subject to $x(t) - x(t-1) \geq 0$. By adding up the prior terms of the form (13) for all pixels in the image, we obtain the last summation term in (4).

C. Problem structure

In what follows, we consider the filter described by the QP problem (4)–(5). We consider \mathbf{R} and ρ as tuning parameters of the filter without resorting to the MAP interpretation. This paper considers blur operator \mathbf{B} and the regularization operator \mathbf{R} that are FIR convolution operators. This problem structure allows for efficient solution.

If an image U is considered as a flat vector, the FIR operators \mathbf{B} and \mathbf{R} correspond to sparse square matrices of compatible size. Multiplying by such matrix (even in a sparse matrix form) is inefficient for a large image size. The imaging blur operator \mathbf{B} can be efficiently applied as a 2-D convolution

$$\mathbf{B}X(t) = B**X(t), \quad (14)$$

where B is a spatially-invariant FIR (finite impulse response) PSF (point spread function) kernel of the blur. The notation $**$ in (14) stands for two-dimensional convolution. We assume that one of the standard image processing approaches to boundary condition handling in convolution is used; see, e.g., [14], [18], [35] for more details on these approaches.

For the model (10), the spatial regularization operator \mathbf{R} in (4) has the form

$$\mathbf{R}X(t) = R**X(t), \quad (15)$$

where R is a non-causal 2-D FIR convolution kernel with a maximum tap delay K and entries r_{lm} . A necessary and sufficient condition for (7) to represent a valid probability density function is that the 2-D FIR convolution kernel R is symmetric and positive-definite [34]. Section IV considers the selection of the kernel R (filter tuning) in more detail. The scaled identity operator (8) is a convolution operator with $K = 0$. The Laplace operator (9) corresponds to $K = 1$, $r_{0,0} = r_1$, $r_{1,0} = r_{-1,0} = r_{0,1} = r_{0,-1} = -r_1/4$ and other entries $r_{ij} = 0$.

III. A SPECIALIZED LARGE-SCALE QP SOLVER

This section describes an interior-point method for solving the problem (4)–(5). Our approach is based on truncated Newton method. Such methods were earlier used for image enhancement, deconvolution, and deblurring, e.g., see [40]. The success of the method critically depends on finding a preconditioner that gives an effective trade-off between the computational complexity and the accelerated convergence. The new contribution in this section is in establishing the preconditioner and in demonstrating how the specific structure of the problem (4)–(5) can be otherwise exploited to reduce the computational complexity and the memory requirements of the solution method.

With some abuse of notation we consider the images $X(t)$ and $Y(t)$ as flat vectors in \mathfrak{R}^M obtained by stacking all $M = N_1N_2$ image elements (pixels). Instead of the decision vector $X = (X(1), \dots, X(N_t)) \in \mathfrak{R}^N$, where $N = N_tN_1N_2$, in (4)–(5), we introduce a new decision vector $Z = (Z(1), \dots, Z(N_t)) \in \mathfrak{R}^N$, where

$$Z(1) = X(1), \quad Z(t) = X(t) - X(t-1), \quad t = 2, \dots, N_t.$$

The inverse variable transformation is $X(t) = \sum_{k=1}^t Z(k)$. This can be represented as $X = \mathbf{L}Z$, where \mathbf{L} is a sparse lower block-triangular matrix.

In terms of these new variables, (4) is equivalent to finding $Z(1), \dots, Z(N_t) \in \mathfrak{R}^M$ that solve

$$G(Z) = \Psi(\mathbf{L}Z) + \rho \sum_{t=2}^{N_t} Z(t) \rightarrow \min \quad (16)$$

subject to $Z(t) \geq 0, \quad t = 2, \dots, N_t, \quad (17)$

where

$$\Psi(X) = \frac{1}{2} \sum_{t=1}^{N_t} (\|Y(t) - \mathbf{B}X(t)\|^2 + X(t)^T \mathbf{R}X(t))$$

If Z^* solves the problem (16)–(17), then $X^* = \mathbf{L}Z^*$ solves the original problem (4)–(5).

A. The barrier method

The *logarithmic barrier* for the nonnegativity constraints (17) has the form

$$\Phi(Z) = - \sum_{t=2}^{N_t} \sum_{i=1}^M \log Z_i(t), \quad (18)$$

The logarithmic barrier function (18) is smooth and convex in its domain. We augment the objective function (16) with the barrier function (18) to obtain

$$\phi_\tau(Z) = \tau G(Z) + \Phi(Z), \quad (19)$$

where $\tau > 0$ is a barrier weight. The augmented function $\phi_\tau(Z)$ is smooth, strictly convex, and bounded below, and so has a unique minimizer $Z^*(\tau) \in \mathfrak{R}^N$. The set $\{Z_\tau^* \mid \tau > 0\}$ defines a curve in \mathfrak{R}^N , parameterized by τ , which is called the *central path*. The minimizer of (19) is no more than N/τ -suboptimal, so the central path leads to an optimal solution. See [8, §11] for more on the central path and its properties.

In a classic primal barrier method, the barrier subproblem that finds the minimizer of (19) is solved for an increasing sequence of values of τ until N/τ is smaller than the required tolerance. Standard references on interior-point methods include [43], [44], [62], [63].

B. A truncated Newton interior-point method

For solving the large-scale QP problem (16)–(17) we use a *truncated Newton method* (also known as a conjugate gradient Newton method) [51], [12]. This is a modification of the barrier method with the search direction computed approximately, using a PCG method.

The most computationally expensive part of the method is related to operation with the Hessian $\mathbf{H} = \nabla^2 \phi_\tau(Z) \in \mathfrak{R}^{N,N}$ of the augmented objective function (19). The Hessian has the form

$$\mathbf{H} = \text{block diag}(\text{diag}(Z(1)), \dots, \text{diag}(Z(N_t))) \\ + \tau \mathbf{L}^T \text{block diag}(\mathbf{B}^T \mathbf{B} + \mathbf{R}, \dots, \mathbf{B}^T \mathbf{B} + \mathbf{R}) \mathbf{L} \quad (20)$$

We will not go into the details of the PCG algorithm, and, instead, refer the reader to [52], [44]. An important part of the PCG is choosing a preconditioner $\mathbf{P} \in \mathfrak{R}^{N,N}$, a symmetric positive definite linear operator that approximates the Hessian \mathbf{H} . We use a diagonal preconditioner \mathbf{P} , which retains all diagonal entries of \mathbf{H} . Such a preconditioner retains the Hessian of the logarithmic barrier (the first matrix in the sum (20)).

The PCG algorithm needs a good initial search direction and an effective truncation rule. As an initial search direction, we use the search direction from the previous step. The truncation rule for the PCG algorithm gives the condition for terminating

the algorithm when either the cumulative number of PCG steps exceeds the given limit N_{pcg} , or the gradient is less than the relative tolerance ϵ_{pcg} . We change the relative tolerance adaptively as

$$\epsilon_{\text{pcg}} = \min \{0.1, \xi\eta/\|g\|_2\}, \quad (21)$$

where η is the duality gap at the current iteration and ξ is an algorithm parameter. The choice of $\xi = 0.01$ appears to work well for a wide range of problems. In other words, we solve the Newton system with low accuracy at early iterations, and increase the accuracy as the duality gap decreases.

C. Complexity and performance

Each iteration of the PCG algorithm involves a handful of inner products, the matrix-vector product $\mathbf{H}p$ with $p \in \mathbb{R}^N$ and a solve step with the preconditioner \mathbf{P} in computing $\mathbf{P}^{-1}r$ with $r \in \mathbb{R}^N$. The solve step $\mathbf{P}^{-1}r$ can be computed in $O(N)$ flops, since \mathbf{P} is diagonal.

The most computationally expensive operation for a PCG step is the matrix-vector product $\mathbf{H}p$ with $p \in \mathbb{R}^N$. In accordance with (20), this can be done without actually forming matrix \mathbf{H} . Operators \mathbf{L} and \mathbf{L}^T can be applied by computing cumulative sums; operators \mathbf{B} , \mathbf{B}^T , and \mathbf{R} , by computing two-dimensional convolutions of each image $Z(t)$ in the series with the FIR kernels B and R . Each of these operations takes $O(N)$ flops.

The memory requirement of the truncated Newton interior-point method is modest, so the method is able to solve very large problems, for which forming the Hessian \mathbf{H} , let alone computing the search direction, would be prohibitively expensive. The runtime of the truncated Newton interior-point method is determined by the product of s , the total number of PCG steps required over all iterations, and the cost of a PCG step. In extensive testing, we found that the total number of PCG steps ranges between a few hundred and several thousand to compute a solution with a relative tolerance of 0.01.

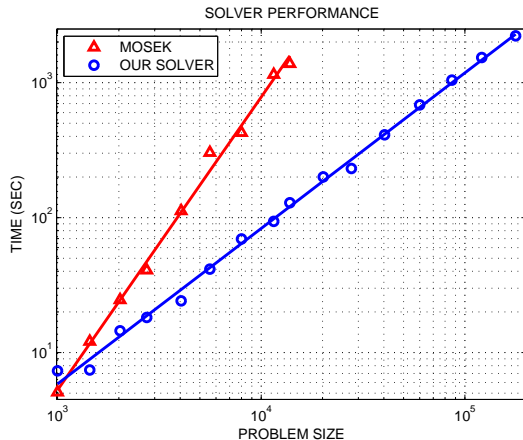


Fig. 1: Computation time (in seconds) depending on the problem size for the proposed solver and Mosek

The performance of a Matlab implementation of the truncated Newton interior-point method described above is illustrated in Figure 1. A series of the QP problems (4)–(5) were

solved for different N_1 , N_2 , N_t . The detailed formulation of the problem is the same as discussed in Section V. The logarithmic plot shows dependence of the solution time (in seconds) on the problem size $N = N_1 \cdot N_2 \cdot N_t$. The shallow curve describes the near-linear computational complexity of the developed method, while the steeper one a near-quadratic complexity of obtaining the solution using the Mosek QP solver [41]. Note that for very small problems Mosek, with its optimized C/C++ implementation, performs faster than the simple Matlab implementation of the proposed method. For large problems, the Matlab implementation is much faster. On our test PC with 2 Gb RAM, Mosek runs out of memory for $N \approx 20 \cdot 10^3$. The proposed method was tested and works well for up to $N \approx 20 \cdot 10^6$.

IV. TUNING REGULARIZATION OPERATOR

In the optimization problem (4)–(5) for a given model (3) we need to tune the spatial filtering (the operator \mathbf{R}) and the time-domain filtering (the scalar ρ). As with the time-domain monotonic regression problem discussed in [19], [21], the parameter ρ can be set, based on (11)–(12), to a ratio of the standard deviations of the noise and the signal.

We consider the FIR convolution kernel R in (15) as a filter design parameter. Our approach to choosing R is related to passband equalization by Wiener filtering; an added requirement is that R has a fixed FIR structure. Linear multidimensional control and filtering problems, which are related in spirit, are considered in [20], [22], [23], [24]. These prior papers design *weights* of linear multidimensional filters. Herein, we consider a different problem of designing *regularization penalty* in the optimization formulation.

A. Tuning requirements

Consider the spatial response of the proposed filter, steady-state in time to a steady state input $X(t) = X_*$. Substituting $X(t) = X_*$, $Y(t) = Y_*$, and $N \gg 1$ into (4) and removing the inactive monotonicity constraints (5) leads to a steady-state optimization problem

$$\frac{1}{2}\|Y_* - \mathbf{B}X\|_F^2 + \frac{1}{2}(X, \mathbf{R}X) \rightarrow \min \quad (22)$$

To obtain an optimal estimate X , we assume that $Y_* = \mathbf{B}X_* + e_*$, where e_* is the steady state (spatial) noise. Substituting this into (22) and solving for X yields

$$X = [\mathbf{I} - (\mathbf{B}^T\mathbf{B} + \mathbf{R})^{-1}\mathbf{R}] X_* + (\mathbf{B}^T\mathbf{B} + \mathbf{R})^{-1}\mathbf{B}^T e_*, \quad (23)$$

The first term in (23) describes recovery of the steady state signal and the second term describes the noise amplification. The design goal is to find an optimized tradeoff between the two goals: the recovery error gain $(\mathbf{B}^T\mathbf{B} + \mathbf{R})^{-1}\mathbf{R}$ being small (which requires \mathbf{R} to be small) and the noise amplification gain $(\mathbf{B}^T\mathbf{B} + \mathbf{R})^{-1}\mathbf{B}^T$ being small (which requires \mathbf{R} to be large).

For an operator \mathbf{R} of the form (15) the tradeoff given by (23) can be conveniently analyzed in spatial frequency domain, at the cost of neglecting the boundary effects; *e.g.*, see [23],

[35]). The convolution operator R in (4), (15) can be expressed through a 2-D optical transfer function

$$r(v_1, v_2) = (2\pi)^{-2} \sum_{-L \leq k, l \leq L} r_{ij} e^{-2\pi i k v_1} e^{-2\pi i l v_2}. \quad (24)$$

Similarly, the blur convolution operator \mathbf{B} in (14) can be expressed through a transfer function $b(v_1, v_2)$. The operator \mathbf{B}^T in (23) can also be represented in the convolution form and corresponds to a complex conjugate transfer function $b(-v_1, -v_2) = \bar{b}(v_1, v_2)$.

In the spatial frequency domain, (23) can be expressed as

$$x_e(v_1, v_2) = \left(1 - \frac{r(v_1, v_2)}{|b(v_1, v_2)|^2 + r(v_1, v_2)}\right) x_*(v_1, v_2) + \frac{\bar{b}(v_1, v_2)}{|b(v_1, v_2)|^2 + r(v_1, v_2)} e(v_1, v_2) \quad (25)$$

We require that noise amplification gain in (25) is bounded. The bound e_0 has a meaning of signal to noise ratio (the amplified noise should be below the signal)

$$\left| \frac{b(v_1, v_2)}{|b(v_1, v_2)|^2 + r(v_1, v_2)} \right| \leq e_0. \quad (26)$$

Spatial filter (23) has the same form as the steady state spatial filter considered in [23], with \mathbf{R} taking place of the smoothing operator and \mathbf{B}^T of the feedback operator in [23]. A robust stability condition for the filter is derived [23] and has the form (26). The robust stability is guaranteed for $|\delta b(v_1, v_2)| \leq 1/e_0$, where $\delta b(v_1, v_2)$ is the transfer function of the uncertainty in \mathbf{B} .

To ensure signal recovery performance in (25), we define the in-band frequency set as $\Omega_{In} \equiv \{(v_1, v_2) : b(v_1, v_2) > h_0\}$. The design parameter h_0 roughly describes the noise level; in-band, the system response is above the noise level. In this band we strive to invert the blur operator and recover the truth signal x_* with a minimal distortion level e_1 :

$$\left| \frac{r(v_1, v_2)}{|b(v_1, v_2)|^2 + r(v_1, v_2)} \right| \leq e_1, \quad \{v_1, v_2\} \in \Omega_{In}, \quad (27)$$

Outside the in-band set, the blur operator gain is low, the noise overwhelms the signal, and the filter has degraded performance.

The design specifications (26), (27) can be used for tuning the scalar weights r_0 or r_1 for the regularization operators (8) or (9). The tuning can be done by trial and error, using the plots shown in Figure 3. A more systematic approach to choosing the regularization operator R is detailed in Appendix. The approach is to fix e_0 (26) and then minimize e_1 in (27). For a symmetric kernel R , the transfer function $r(v_1, v_2)$ is real. It is also linear in the design parameters: r_0 in (8), r_1 in (9), or r_{ij} in (15). We multiply (26) and (27) through by the denominators, which are positive, to obtain linear inequalities. After gridding spatial frequencies we obtain an LP problem, which can be solved efficiently. The described design method yields the kernel R and the achievable filter performance parameter e_1 for a given blur kernel B . There are two tuning parameters: noise amplification e_0 , which roughly corresponds to the signal to noise ratio, and bandwidth parameter h_0 , which roughly corresponds to the noise level. More details on the design can be found in Appendix.

B. Tuning example




Fig. 2: Spatial operators in the optimization problem statement.

We illustrate the tuning of the regularization operator, with a Gaussian blur B the half-width σ of which is 2 pixels. This example corresponds to the simulation study in Section V and experimental data in Section VI. The Gaussian blur was truncated to yield a point-spread function (PSF) operator B illustrated on the left plot in Figure 2. The noncausal FIR operator B had a maximal ± 6 pixels tap delay along each spatial coordinate. By assuming a 128×128 spatial frequency grid, the LP (33), (34) was solved to obtain a central symmetric operator R with maximal $L = 2$ tap delays (5×5 FIR convolution kernel), which is shown in the right plot of Figure 2. The in-band frequency set Ω_{In} was chosen by considering a set of grid frequencies where the blur operator gain exceeds $h_0 = 0.55$ of the maximal (zero-frequency) gain. The maximal noise amplification gain in (26) was chosen as $e_0 = 1.2$. The design yields the in-band signal recovery distortion factor $e_1 = 0.3393$ in (27). The two tuning parameters h_0 and e_0 were chosen by trials and errors such that the problem is feasible and a reasonably looking solution can be found.

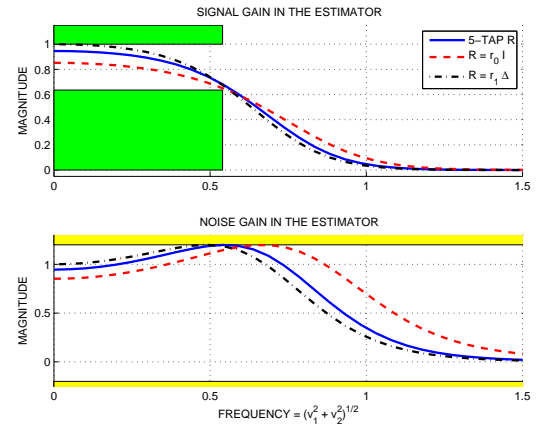


Fig. 3: Frequency response specifications for signal recovery (upper plot) and noise amplification (lower plot). Solid lines - the 5-tap operator R , dashed lines - $\mathbf{R} = r_0 \mathbf{I}$, dash-dotted lines - $\mathbf{R} = r_1 \Delta$.

We also solved the design problem for the regularization operators (8) and (9). The optimal scaling factors are $r_0 = 0.1736$ for the scaled identity operator $\mathbf{R} = r_0 \mathbf{I}$ and $r_1 = 2.1846$ for the scaled Laplace operator $\mathbf{R} = r_1 \Delta$. These designs yield in-band signal recovery distortion factors $e_1 = 0.3641$ and $e_1 = 0.3542$ respectively.

The three frequency domain designs are illustrated in Figure 3. The designs are central symmetric and the 2-D transfer

functions primarily depend on the magnitude $\sqrt{v_1^2 + v_2^2}$ of the spatial frequency vector. The upper plot in Figure 3 shows the estimator signal gain, the magnitude of the first-term transfer function in (25) for the three designs. The lower plot shows the noise gain (the magnitude of the second-term transfer function in (25)).

V. SIMULATION RESULTS

The performance and practical usefulness of the proposed signal processing approach were verified in simulation. For the simulated data, the ground truth data (underlying damage) is available. This allows verifying the performance of the filter. (Section VI will describe an application to experimentally collected SHM data.)

A. Simulated data

The simulated data was a series of $N_t = 20$ images with $N_1 = 32$ by $N_2 = 30$ pixels. The blur operator \mathbf{B} was modeled as the FIR convolution operator B shown on the left in Figure 2 and discussed in the example of Subsection IV-B.

We simulated the underlying damage as a series of N_t images $U(t) \in \mathbb{R}^{N_1, N_2}$ which are zero outside of an elliptic domain 6×7 pixels in the center of the image and constant inside this domain. The damage $U(t)$ was pixel-wise ramped up from zero to the end value $U(N_t)$ over the middle third of the time interval. The observed data $Y(t)$ (1) were simulated by blurring $U(t)$ and adding noise $e(t)$ to get $Y(t) = B**U(t) + e(t)$. The noise $e(t)$ had standard deviation e_0 , which were varied in the simulation runs. The last image $Y(t = 20)$ of a raw data set Y obtained for $e_0 = 0.4$ is illustrated in Figure 4 (the upper left plot).

The experiments described in the next section involve similar damage pattern and similar blur.

B. Filtering results

The designed optimization-based filter was implemented using the interior-point method described in Section III. We solved the problem with relative accuracy 1%. The solver allows achieving much better relative accuracy, but this accuracy is more than adequate for practical use. The algorithm was implemented in Matlab and run on a 2GHz PC. We used $\eta = 1$ in the adaptive rule (21).

We compared the results for the following five filters:

- F1** This filter uses the regularization operator \mathbf{R} in (4)–(5) designed as in Subsection IV-B. The 5×5 FIR kernel R is illustrated on the right plot in Figure 2.
- F2** This filter uses the scaled identity regularization $\mathbf{R} = r_0 \mathbf{I}$, where $r_0 = 0.1736$ was designed as in Subsection IV-B. The same $\rho = 0.2$ was used.
- F3** This filter uses the Laplace regularization operator $\mathbf{R} = r_1 \mathbf{\Delta}$, where $r_1 = 2.1846$ as described in Subsection IV-B.
- F4** This is a simple spatial low-pass filter; it convolves the noisy images with the blur operator B . This operator has just the right spatial bandwidth for the filtering since, since the spatial harmonics of the observed signal outside of B band mostly contain noise.

F5 A simple spatio-temporal filter. The output of filter is passed through a pixel-wise EWMA (exponentially weighted moving average) filter with 0.7 filter factor.

The same time domain regularization parameter $\rho = 0.2$ was used in filters **F1**, **F2**, and **F3**. The filtering results are not very sensitive to this parameter. Filters **F1**, **F2**, and **F3** were tuned as discussed in Subsections IV-A and IV-B.

The filter performance was evaluated by considering the following two error metrics:

- E1** The detection error metric is defined as the ratio of the recovered signal maximum absolute value outside of the damage area (where the ground truth signal is zero) and the mean value of the recovered signal in the damage area. The ‘outside of the damage area’ is defined as a complement to the damage domain with the linear size increased by a factor of $\sqrt{3}$. A simple practical approach to detecting the damage from the estimate is by thresholding. This metric describes the extent to which the recovered damage signal stands above the clutter. This determines applicability of the thresholding approach.
- E2** The tracking error metric is defined as the mean square error of restoring the underlying damage signal. The error is computed for the blurred signals as $\|\mathbf{B}X - \mathbf{B}U\|$. This metric quantifies filter performance for damage trending and prognostics. Though the stated problem is to remove the noise and blur from the underlying signal U , we compare $\mathbf{B}U$ and $\mathbf{B}X$. This is because the minimized loss index (4) includes $\|Y - \mathbf{B}X\|$, where \mathbf{B} is a spatial low pass filter; the inverse problem is ill-defined. The filter would make $\mathbf{B}U - \mathbf{B}X$ small, but would not necessarily make $U - X$ small.

We studied the dependence of the filtering results on the signal to noise ratio, by varying the noise standard deviation e_0 in the range from 0.1 to 1; all other parameters of the simulation and of the estimation filter were fixed. The results of the study are shown in two tables. Table I shows the detection error metric **E1** depending on the standard deviation e_0 of the observation noise. Smaller **E1** indicates better performance, values above unity indicate that damage cannot be detected. The four rows correspond to the filters **F1**, **F2**, **F3**, and **F4**. All rows show a noticeable error for the small signal-to-noise ratios (small e_0). The error is related to the blurring of the input signal, and is present irrespectively of the noise. In practice, damage is commonly detected by thresholding the signal at 0.5 of the maximum value (at 6 dB). Therefore, the detection error metric **E1** should be below the threshold level 0.5. In Table I, this holds for filter **F1** if $e_0 \leq 0.1$. For the optimization-based filters **F1**, **F2**, and **F3**, this holds for $e_0 \leq 0.4$.

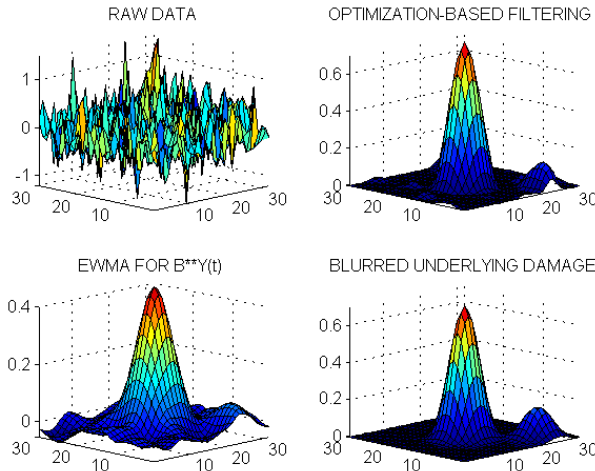
Table II shows the tracking error metric **E2** computed for the same simulation as in Table I. Filter **F1** with 5-tap R is close to but slightly worse than filter **F3** (Laplacian regularization); filter **F2** (scaled identity regularization) has about 50% larger error than **F1** or **F2**; and the performance of filters **F4** and **F5** based on the simple spatial convolution is about twice worse.

Overall, filter **F1** provides the best balance of the two metrics; filter **F2** yields the best tracking error metric **E2**; and filter **F3** provides the best detection error metric **E1**. The three

Noise standard deviation e_0	0.1	0.2	0.4	0.6	0.8	1
Filter F1 (5-tap R)	0.2579	0.2754	0.2722	0.5252	0.5614	0.7787
Filter F2 ($R = r_0 I$)	0.2260	0.2445	0.2439	0.5996	0.7012	0.8254
Filter F3 ($R = r_1 \Delta$)	0.2857	0.3020	0.2926	0.5007	0.5021	0.7636
Filter F4 (convolution with B)	0.4632	0.5952	0.8496	1.0861	1.9305	2.0393
Filter F5 (F4 plus EWMA)	0.4311	0.5064	0.6394	1.2435	1.1961	1.7261

TABLE I: Detection error metric comparison results


Noise standard deviation e_0	0.1	0.2	0.4	0.6	0.8	1
Filter F1 (5-tap R)	3.0598	3.3861	3.7380	4.1272	4.6145	5.2035
Filter F2 ($R = r_0 I$)	4.0542	4.3209	4.6503	4.8576	5.3247	5.7649
Filter F3 ($R = r_1 \Delta$)	2.8182	3.1618	3.4735	3.9690	4.4600	5.1363
Filter F4 (convolution with B)	4.1964	5.5598	8.6145	11.8626	15.7361	19.5784
Filter F5 (F4 plus EWMA)	6.2828	6.6329	7.4085	8.6174	9.2635	12.0161

TABLE II: Tracking error metric comparison results

Fig. 4: Last image in raw data (upper left plot), for **F1** filter (upper right), for **F5** (lower left), $B**U(N_t)$ (lower right). The surfaces show the damage estimates depending on the spatial location coordinates (in inches).

optimization-based filters **F1**, **F2**, and **F3** perform much better than the simple filters **F4** and **F5**. The EWMA filtering in **F5** somewhat improves performance compared to filter **F4** for measurement noise with $e_0 = 0.6$ and larger. Adding EWMA to filter **F4** is unhelpful for smaller noise levels.

The Laplacian regularization (filter **F3**) might provide the best balance of the complexity and performance. Two notes are in order in that regard. First, the above results were obtained for an optimally tuned Laplacian penalty. Selecting an arbitrary penalty weight does not guarantee the performance. Second, the results of this section were obtained for a specific Gaussian blur kernel. For a more demanding blur operator, a higher order regularization operator R might be necessary.

Filter performance is illustrated in Figures 4 and 5. The simulated signal $U(t)$ included two features. The first feature had 31 central pixels with value 1. The second feature had 12 pixels with value 0.5 near the corner of the domain. Features were ramped up at different times (see Figure 5). The signal $U(t)$ was blurred with a Gaussian kernel B with the width of 2 pixels; a white observation noise with standard deviation 0.4 was added. With the peak-to-peak magnitude of signal $U(t)$ being unity, this correspond to a signal-to-noise (S/N) ratio


Fig. 5: Temporal responses for the middle pixel of the first feature (upper plot) and the middle pixel of the second feature (lower plot). Filtering with **F1** (solid line), with **F5** (dash line), and the blurred signal $B**U(t)$ (dash-dotted line).

close to one.

The last image $Y(N_t)$ for simulated raw data is shown in the upper left plot on Figure 4. The last image corresponds to fully ramped up signal $U(t)$. The upper right plot shows the last estimate $X(N_t)$ for filter **F1**; the lower left shows similar results for **F5**; and the lower right plot shows the blurred signal $B**U(N_t)$ with no noise added. We show $B**U(N_t)$, rather than $U(N_t)$ because the high-frequency information beyond the bandwidth of B is lost when recovering X . Since B is a good low pass filter, the blurred signal $B**U(N_t)$ should have about the same frequency content as the recovered X . Note that metric **E1** in Table I uses the mean value of the signal in the damage area, which is about three times smaller than the peak value. It also uses the maximum value of the noise signal over all scans, which is larger than the value seen in the plots on Figure 4.

One can see that the optimization-based filter **F1** recovers well both features in the underlying (blurred) signal. For the simple spatio-temporal smoothing filter **F5**, the second feature is below the noise level of the recovered signal (defined as 150% of maximum noise value). Figure 5 shows signal and estimate time series at the middle pixels of the first and second features. By design, nonlinear filter **F5** suppresses the smaller

second feature more than the larger first feature.

VI. SHM APPLICATION RESULTS

This section discusses an application of the proposed approach to an SHM system in a laboratory experiment setting. The first goal of the section is to illustrate how the approach can be applied in practice, what steps are involved. We also put the approach in perspective of the existing SHM literature. The second goal of the section is to show that the approach provides practically useful damage estimates. There is no certainty that the SHM system output increases monotonically with damage. However, we assume that the underlying damage is monotonic (irreversible) and attribute any deviation from monotonicity to measurement errors.

A. SHM test data

The experimental data was collected with a commercial SHM sensing system developed by Acellent Technologies and previously demonstrated in a variety of structural health monitoring applications [5], [6], [7], [46]. The overall setup of the experimental data collection and processing is reflected in Figure 6. The two upper blocks correspond to transducers, electronics, and software of a commercial SHM system used in the experiments. We considered this SHM system as a black box capable of producing a snapshot diagnostic map. In experiments, we collected and accumulated a series of the diagnostic maps (lower left block) and then processed the collected series using the algorithms described in the earlier sections of the paper (the lower right block).




Fig. 6: Data processing in the experimental system.

A variety of SHM and NDE (nondestructive evaluation) systems exist. These systems are based on different physical principles and employ a variety of proprietary signal processing algorithms. The main envisioned practical application for the proposed approach is for estimating structural damage from an accumulated *time series* of the diagnostic maps produced by SHM or NDE systems.

B. Ultrasonic SHM systems

The Acellent SHM system detects and locates structural defects using piezoelectric ceramic disc to transmitted and

received ultrasonic signals in a so called pitch-catch configuration [30]. The transducers are distributed on the surface of the structure and connected to a portable diagnostic unit comprised of sensor/actuator amplifiers, filters, a function generator, data acquisition card, and a laptop computer with diagnostic software. Each transducer is driven with a pre-selected signal, typically a modulated sinusoidal tone burst, so that elastic waves are generated and propagate through the structure to be recorded by neighboring transducers acting as sensors. The received signals are compared with a previously recorded baseline signal to identify the locations and extents of structure properties changes. The changes are considered indicative of damages or other structural anomalies; the results are reported out as a smoothed damage map estimate.



Fig. 7: Flat composite panel with 49 sensors.

The system uses elastic waves, primarily in the form of Lamb waves, see for example [50], [60], for integrated structural health monitoring. A deeper discussion or various aspects and operation principles of such systems can be found in [3], [10], [13], [28], [30], [47]. Very few such integrated systems are available commercially. The reader is referred to [5], [6], [7], [46] for more technical details on the Acellent SHM system. Consistent with the formulation of aircraft monitoring problem in [25], [59], we used the output of the system as a black box.

C. SHM test data

To collect the data, tests were performed on a $1.2 \text{ m} \times 1.2 \text{ m}$ flat composite panel 13 mm thick. The panel was instrumented with 49 transducers placed in a 7×7 grid with 178 mm spacing. Damage was induced in the panel through consecutive impacts, using equal blows that that were calibrated to produce barely visible impact damage (BVID) by adjusting the drop height of the dead load. The blows were repeated nine times at the same location; see Figure 7. Using an environmental chamber to control the temperature of the panel, the data was collected between the impacts at two different temperatures 20°C and 40°C .

Environmental effects, such as temperature differences, cause additional changes in the sensed signals, and can confound damage detection schemes. The effects of temperature

variations on guided wave SHM systems has been considered by others, see, *e.g.*, [33], [36], [47]. Thermal calibration techniques utilizing multiple baselines collected at various temperatures can be employed to mitigate these effects, but still leave a residual error. Outdoor structures and vehicles, especially aircrafts, are subject to a range of temperatures many times larger than could be enacted in the study. To emulate the temperature-induced error, only a partial thermal compensation was applied in the data processing (over a range of operating temperatures from 25°C - 35°C). The diagnostic images were collected outside the range and therefore are influenced by environmental variation. The images are illustrated in Figure 8. The horizontal axes show panel coordinates in inches, the vertical axis is damage indication in relative intensity units. (Unit damage is well visible). There is a visible temperature-caused variation between the left (20°C) and right (40°C) images.




Fig. 8: Diagnostic images collected at 20°C and 40°C after 3, 6 and 9 impacts. The surfaces show damage estimates (in relative units) depending on the spatial coordinates (in inches).

The initial data set obtained in the experiments contains 8 pairs of images with $171 \times 171 = 29,241$ pixels each. We used a bootstrapping-like method to increase the number of images in the sequence and create a more realistic emulation of the random temperature swings. From a single pair of the images $\{Y_{20}(t), Y_{40}(t)\}$ obtained for the same panel damage at two different temperatures we created $n_b > 2$ samples. We then computed linear interpolations of the two images to approximate data for n_b in-between temperatures:

$$Y(\tau) = a_\tau Y_{20}(t) + (1 - a_\tau) Y_{40}(t), \quad (k = 1, \dots, n_b), \quad (28)$$

where $\tau = n_b t + k$, ($k = 1, \dots, n_b$) is the time index of the generated data set and a_τ are random variables uniformly distributed on the interval $[0, 1]$. We assumed $n_b = 3$: three scans were generated according to (28) for each damage state.

As Figure 8 illustrates, the environment variation is about 25% of the maximum signal (compare the size of the side lobe in the lower right plot with the central peak).

D. SHM filtering results

The underlying damage maps $X(t)$ were estimated from the observed data $Y(t)$ by solving the QP problem (4)–(5) for optimization-based filtering. Blur kernel B was not known exactly; it was approximated by a Gaussian kernel of width 2. The tuning parameters of the filter were the same as described in Subsection IV-B. For the noise amplification gain tuning $e_0 = 1.2$, the design is robust to $1/e_0$ (about 80%) uncertainty in the blur operator (see Subsection IV-B).

The filter (4)–(5) implicitly assumes white observation noise. Figure 8 indicates that the noise is in-fact a low-pass noise. Thus, the designed filter could be overly conservative on high spatial frequencies. Such design has an increased robustness to high-frequency uncertainty of the blur operator, which is desirable. A discussion on choosing the temporal regularization parameter ρ can be found in [19], [21]. We used $\rho = 0.01$.

The proposed filtering approach was applied to the test data set illustrated in Figure 8. The raw images in Figure 8 show the SHM system response. In experiments we know that, in fact, the damage is concentrated in a single spot. The filtering results are shown in Figure 9. The displayed images are $X(t) - X(1)$ for the filter (4)–(5). (We assume that initially there is no damage and subtract the baseline.) The right plots in Figure 9 show the filtered data $X(t) - X(1)$. The left plots shows the test data $Y(t)$. The right plots in Figure 9 have a single peak, which accurately recovers the damage location. The restored signal looks blurred because the filter cannot fully restore the high-frequency harmonics of the underlying signal with B being a good low pass filter. In fact, the underlying damage might have sharper boundaries.

The SHM system measures the damage indirectly, by observing the changes in ultrasonic wave propagation compared to an initial baseline state. It would be desirable to compare the filtered SHM data with the ground truth data. Though the ground truth is not fully known, it is partially known. In the experiments, the damage of the composite panel is concentrated at the panel center. The damage includes surface indentation and some debonding/delamination of the surface layers of the composite. The damage increases after each blow and remains constant between the blows. The filter output matches this prior knowledge. The proposed nonlinear filtering scheme substantially improves the quality of the damage estimate by removing the phantom damage at multiple locations on the plate.

VII. CONCLUSIONS

We have considered estimation of a time series of images pixel-wise monotonic in time. The problem is motivated by structural health monitoring; the damage accumulating in a structure needs to be distinguished from the noise (scatter). We have formulated the estimation problem as optimization




Fig. 9: Filtering results for the test data set. The surfaces show damage depending on the spatial coordinates (in inches). The left plot shows original data at several time instances, the right plots shows the corresponding filtered signal data.

of a regularized loss index and proposed a method for tuning the spatial regularization operator.

The estimation with monotonicity constraints leads to a large-scale structured QP problem. We developed an interior-point method for solving large-scale QP problems of this type. Our simple Matlab implementation can handle quadratic programs with several million variables and constraints in a few tens of minutes on a PC.

We have verified the approach in a simulation study and showed that it performs well for signal to noise ratio above unity. We have also validated the approach by applying it to diagnostic images of structural damage obtained in experiments.

ACKNOWLEDGEMENTS

Dr. Bao Liu and Dr. Tom Chang from Acellent conducted the impact tests on the panel and collected the raw sensor data; the authors are grateful for their help and would like to recognize their contribution. The authors thank Kwangmoo Koh, ISL, Stanford EE Department, for helpful comments and suggestions in the work on the large-scale sparse QP solver. The authors appreciate the contribution of Prof. Fu-Kuo Chang, Stanford AA Department, in initiating the interaction between the theoretical and experimental parts of this work.

APPENDIX

In designing the convolution regularization operator R , we look for a symmetric solution such that $r(v_1, v_2)$ is real and positive (this ensures the operator \mathbf{R} is positive semidefinite). Since $r(v_1, v_2) \geq 0$, the denominators in (26), (27) are

real positive. Thus, the design specifications (27), (26) can be multiplied by the denominator and presented as linear in $r(v_1, v_2)$ inequalities:

$$0 \leq r(v_1, v_2) \leq s|b(v_1, v_2)|^2, \quad (v_1, v_2) \in \Omega_{In}, \quad (29)$$

$$|b(v_1, v_2)| \leq e_0|b(v_1, v_2)|^2 + e_0r(v_1, v_2) \quad (30)$$

where $s = e_1/(1 - e_1)$ will be considered as a slack variable.

Following [22], [23], a general symmetric kernel R can be presented as a linear combination of elementary symmetric kernels to yield a real transfer function of the form

$$r(v_1, v_2) = p^T(v_1, v_2)\bar{r}, \quad (31)$$

where $\bar{r} \in \mathfrak{R}^K$ is the vector of the design parameters and the components of the vector $p(v_1, v_2) \in \mathfrak{R}^K$ are transfer functions of the elementary symmetric kernels. The specific form of $p(v_1, v_2)$ in the linear expansion (31) depends on the chosen family of the parametric kernels.

The symmetry types for 2-D operators are discussed, e.g., in [35]. In the example below, we assume a 8-fold symmetry: $r_{m,n} = r_{-m,-n} = r_{-m,n} = r_{m,-n} = r_{n,m} = r_{-n,-m} = r_{-n,m} = r_{n,-m}$. The elementary symmetric kernels of a convolution operator with a given maximum tap delay are described in [22], [23]. The transfer function can be presented in the form (31) with $K = 1$ for the simple regularization operators (8) and (9). For the scaled identity regularization operator (8) $\bar{r} = r_0$ and $p(v_1, v_2) = (2\pi)^{-2}$. For the scaled identity regularization operator (9) $\bar{r} = r_1$ and $p(v_1, v_2) = (2\pi)^{-2}(\sin^2 \pi v_1 + \sin^2 \pi v_2)$, the optical transfer function of the Laplace operator.

Introduce the decision vector $x \in \mathfrak{R}^{K+1}$ combining \bar{r} and the slack variables

$$x = (s, \bar{r}), \quad (32)$$

By using (29), (30) and (32) we obtain the following constrained optimization problem

$$s \rightarrow \min, \quad (33)$$

$$\text{subject to } c(v_1, v_2) + D^T(v_1, v_2)x \geq 0, \quad (34)$$

where the vector $c(v_1, v_2)$ and the matrix $D^T(v_1, v_2)$ collect the linear inequalities expressing the problem (29)–(32). One can introduce a grid of the frequency points and consider the inequality constraints (34) on the grid only. The problem (33), (34) then becomes a linear program with a large number of constraints and $K + 1$ decision variables (32) that can be efficiently solved by an off-the-shelf LP solver. The optimal decision vector x (32) defines the regularization operator R in accordance with (31).

REFERENCES

- [1] J.D. Achenbach, "On the road from schedule-based nondestructive inspection to structural health monitoring," *6th Int. Workshop on Structural Health Monitoring*, pp. 16-28, DEStech Pub., Lancaster, PA, 2007.
- [2] D.E. Adams, *Health Monitoring of Structural Materials and Components: Methods with Applications*, J. Wiley and Sons, Ltd., West Sussex, England, 2007.
- [3] D.N. Alleyne and P. Cawley, "The interaction of Lamb waves with defects," *IEEE Trans. on Ultrason. Ferroelectr. Freq. Control*, Vol. 39, 1992, pp. 381-397.

- [4] R. Barlow, D. Bartholomew, J. Bremner, and H. Brunk, *Statistical Inference under Order Restrictions; the Theory and Application of Isotonic Regression*, New York: Wiley, 1972.
- [5] S. Beard, A. Kumar, X. Qing, H. Chan, C. Zhang, and T. Ooi, "Practical issues in real-world implementation of Structural Health Monitoring systems," *SPIE Smart Structures and Material Systems*, March 2005.
- [6] S. Beard, B. Liu, P. Qing and D. Zhang, "Challenges in implementation of SHM," *6th Internat. Workshop on Structural Health Monitoring (IWSHM)*, Stanford, September 2007
- [7] S. Beard, P. Qing, M. Hamilton, and D. Zhang, "Multifunctional software suite for Structural Health Monitoring using SMART technology," *Proc. 2nd European Workshop on Structural Health Monitoring*, July 2004, Germany.
- [8] S. Boyd, and L. Vandenberg, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004
- [9] A. Buades, B. Coll, and J. Morel, "A review of image denoising algorithms, with a new one," *SIAM Journal on Multiscale Modeling and Simulation*, Vol. 4, No. 2, 2005, pp. 490–530.
- [10] A.J. Croxford, Wilcox, PD, Drinkwater, BW, and Konstantinidis, G "Strategies for guided-wave structural health monitoring," *Proc. R. Soc. A: Math., Phys. and Eng.Sciences*, Vol. 463, 2007, pp. 2961–2981.
- [11] *Damage Prognosis*, Ed.: D. J. Inman, C. R. Farrar, V. Lopes, V. Steffen, Wiley 2005
- [12] R. Dembo and T. Steihaug, "Truncated-Newton algorithms for large-scale unconstrained optimization," *Mathematical Programming*, Vol. 26, 1983, pp. 190–212.
- [13] O. Diligent, and M.J.S. Lowe, "Reflection of the s0 Lamb mode from a flat bottom circular hole," *J. Acoust. Soc. Am.*, Vol. 118, 2005, pp. 2869–2879.
- [14] D. Dudgeon and R. Mersereau, *Multidimensional Digital Signal Processing*, Prentice-Hall, 1984.
- [15] C.R. Farrar and N.A.J. Lieven "Damage prognosis: the future of structural health monitoring," *Phil. Trans. R. Soc. A.*, Vol. 365, 2007, pp. 623–632
- [16] H. Fu, M. Ng, M. Nikolova, and J. Barlow, "Efficient minimization methods of mixed ℓ_1 - ℓ_1 and ℓ_2 - ℓ_1 norms for image restoration," *SIAM J. on Scientific Computing*, Vol. 27, No. 6, 2006, pp. 1881–1902.
- [17] D. Goldfarb and W. Yin, "Second-order cone programming based methods for total variation image restoration," *SIAM J. on Scientific Computing*, Vol. 27, No. 2, 2005, pp. 622–645.
- [18] R. Gonzalez and R. Woods, *Digital Image Processing*, 2nd ed., Upper Saddle River, N.J.: Prentice Hall, 2002.
- [19] D. Gorinevsky, "Monotonic regression filters for trending gradual deterioration faults," *American Control Conference*, pp. 5394–5399, Boston, MA, June 2004.
- [20] D. Gorinevsky, "Feedback loop design and analysis for iterative localized image deblurring," *44th IEEE CDC and ECC'05*, Seville, Spain, December 2005.
- [21] D. Gorinevsky, "Efficient filtering using monotonic walk model," *American Control Conference*, Seattle, WA, June 2008.
- [22] D. Gorinevsky and S. Boyd, "Optimization-based design and implementation of multi-dimensional zero-phase IIR filters," *IEEE Transactions on Circuits and Systems - I*, Vol. 53, No. 2, 2006, pp. 372–383.
- [23] D. Gorinevsky, S. Boyd, and G. Stein, "Design of low-bandwidth spatially distributed feedback," *IEEE Transactions on Automatic Control*, Vol. 53, No. 2, 2008, pp. 257–272.
- [24] D. Gorinevsky and G. Gordon, "Spatio-temporal filter for structural health monitoring," *American Control Conference*, Minneapolis, MN, June 2006
- [25] D. Gorinevsky, G. Gordon, S. Beard, A. Kumar, and F.-K. Chang, "Design of integrated SHM system for commercial aircraft applications," *5th Internat. Workshop on Structural Health Monitoring*, Stanford, CA, September 2005.
- [26] M. Hanke and J. Nagy, "Restoration of atmospherically blurred images by symmetric indefinite conjugate gradient techniques," *Inverse Problems*, Vol. 12, 1996, pp. 157–173.
- [27] L. He, M. Burger, and S. Osher, "Iterative total variation regularization with non-quadratic fidelity," *Journal of Mathematical Imaging and Vision*, Vol. 26, No. 1–2, 2006, pp. 167–184.
- [28] R. Hedl, R. Hamza, G.A. Gordon, "Automated corrosion detection using ultrasound Lamb waves," *6th International Workshop on Structural Health Monitoring*, pp. 1315-1323, DEStech Pub., Lancaster, PA, 2007.
- [29] D. Hochbaum, "An efficient algorithm for image segmentation, Markov random fields and related problems," *Journal of the ACM*, Vol. 48, No. 4, 2001, pp. 686–701.
- [30] J.-B. Ihn and F.-K. Chang "Pitch-catch Active Sensing Methods in Structural Health Monitoring for Aircraft Structures," *Structural Health Monitoring*, Vol 7, No. 1, 2008, pp. 5–15.
- [31] C.A. Johnson, J. Seidel, A. Sofer, "Interior-point methodology for 3-D PET reconstruction," *IEEE Tran. on Medical Imaging*, Vol. 19, No. 4, 2000, pp. 271–285.
- [32] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky, "A method for large-scale ℓ_1 -regularized least squares problems with applications in signal processing and statistics," *IEEE Journal of Selected Topics in Signal Processing*, Vol. 1, No. 4, 2007, pp. 606–617.
- [33] G. Konstantinidis, B.W. Drinkwater and P.D. Wilcox, "The temperature stability of guided wave structural health monitoring systems," *Smart Materials and Structures*, Vol. 15, 2006, pp. 967–976.
- [34] S. Li, *Markov Random Field Modelling in Computer Vision*, Springer-Verlag, 1995.
- [35] J. Lim, *Two-dimensional Signal and Image processing*, Englewood Cliffs, N.J.: Prentice Hall, 1990.
- [36] Y. Lu, J.E. Michaels, "A methodology for structural health monitoring with diffuse ultrasonic waves in the presence of temperature variations," *Ultrasonics*, Vol. 43, 2005, pp. 717–731.
- [37] B. Hunt, "The application of constrained least squares estimation to image restoration by digital computer," *IEEE Transactions on Computers*, Vol. C-22, No. 9, pp. 805–812, 1973.
- [38] K. Koh, S.-J. Kim, and S. Boyd, "An interior-point method for large-scale ℓ_1 -regularized logistic regression," *Journal of Machine Learning Research*, Vol. 8, pp. 1519-1555, 2007.
- [39] V. Kolmogorov, *Primal-dual Algorithm for Convex Markov Random Fields*, 2005. Microsoft Technical Report MSR-TR-2005-117.
- [40] R. Molina, J. Mateos, and A. Katsaggelos, "Blind deconvolution using a variational approach to parameter, image, and blur estimation," *IEEE Tran. on Image Processing*, Vol. 15, No. 12, 2006, pp. 3715–3727.
- [41] MOSEK ApS, *The MOSEK Optimization Tools Version 5.0. Optimization Tool Manual*, 2007. Available: www.mosek.com.
- [42] J. Nagy, R. Plemmons, and T. Torgersen, "Iterative image restoration using approximate inverse preconditioning," *IEEE Tr. on Image Processing*, Vol. 5, No. 7, 1996, pp. 1151–1162.
- [43] Y. Nesterov and A. Nemirovsky, *Interior-Point Polynomial Methods in Convex Programming*, Volume 13 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, 1994.
- [44] J. Nocedal and S. Wright, *Numerical Optimization*. Springer Series in Operations Research. Springer, 1999.
- [45] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, "An iterative regularization for total variation based image restoration method," *SIAM Journal on Multiscale Modeling and Simulation*, Vol. 4, No. 2, 2005, pp. 460–489.
- [46] X. P. Qing, H.-L. Chan, S. Beard, and A. Kumar, "An active diagnostic system for structural health monitoring of rocket engines," *J. of Intelligent Material Systems and Structures*, Vol. 17, No. 6, 2006, pp. 619-628
- [47] A. Raghavan and C. Cesnik, "Review of guided-wave structural health monitoring," *The Shock and Vibration Digest*, Vol. 39, No. 2, 2007, pp. 91–114.
- [48] A. Restrepo and A. Bovik, "On the statistical optimality of locally monotonic regression," *IEEE Transactions on Signal Processing*, Vol. 42, No. 6, 1994, pp. 1548–1550.
- [49] T. Robertson, F. Wright, and R. Dykstra, *Order Restricted Statistical Inference*, New York: Wiley, 1988.
- [50] J.L. Rose, *Ultrasonic Waves in Solid Media*, Cambridge University Press, Cambridge, 1999.
- [51] A. Ruszczyński, *Nonlinear Optimization*. Princeton University Press, 2006.
- [52] Y. Saad, *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, 2nd edition, 2003.
- [53] S. Samar, D. Gorinevsky, and S. Boyd, "Moving horizon filter for monotonic trends," *IEEE Conference on Decision and Control*, Paradise Island, Bahamas, December 2004.
- [54] N. Sidiropoulos and R. Bro, "Mathematical programming algorithms for regression-based nonlinear filtering in R^N ," *IEEE Tran. on Signal Processing*, Vol. 47, No. 3, 1999, pp. 771–782.
- [55] *Health Monitoring of Aircraft Structures*, Staszewski, W. J., Boller, C.

and Tomlinson, G. R. (eds), J. Wiley and Sons, West Sussex, England, 2003.

- [56] H. Sohn, C.R. Farrar, F.M. Hemez, D.D. Shunk, D.W. Stinemat, B.R. Nadler, and J.J. Czarnecki, *A Review of Structural Health Monitoring Literature: 1996–2001* Los Alamos National Laboratory Report, LA-13976-MS, 2004
- [57] *Structural Health Monitoring-From Diagnostics, Prognostics to Structural Health Management*, F.-K. Chang (Ed.), Destech Publishing, Proc. 4th Int. Workshop on SHM, Stanford, CA, Sept. 2003.
- [58] *Structural Health Monitoring-Advancements and Challenges for Implementation*, F.-K. Chang (Ed.), Destech Publishing, Proc. 5th Int. Workshop on SHM, Stanford, CA, Sept. 2005.
- [59] A. Trego, E. Haugse, A. Akdeniz, "Structural Health Management (SHM) technology implementation on commercial airplanes," *2nd European Workshop on Structural Health Monitoring*, Munich, Germany, July 7-9, 2004
- [60] I.A. Viktorov, *Rayleigh and Lamb Waves*, Plenum Press, New York, 1967.
- [61] C. Vogel and M. Oman, "Fast, robust total variation-based reconstruction of noisy, blurred images," *IEEE Transactions on Image Processing*, Vol. 7, No. 6, 1998, pp. 813–824.
- [62] S. Wright, *Primal-dual Interior-point Methods*. SIAM, Philadelphia, PA, USA, 1997.
- [63] Y. Ye, *Interior Point Algorithms: Theory and Analysis*. John Wiley & Sons, 1997.



PLACE PHOTO HERE

Shawn Beard received his B.S. from the University of Washington in 1990, M.S. from Caltech in 1991, and Ph.D. from Stanford University in 2001, all in Aeronautics and Astronautics.

From 1991 to 1996, he served as a research engineer/specialist at McDonnell Douglas where he was responsible for aerodynamic sizing of missile interceptors, structural design of space station components, and was a member of the C-17 team that won the Collier Trophy symbolizing the nation's top aeronautical achievement of 1994. In 2000, Shawn joined Acellent Technologies, Inc. and currently serves as the Chief Technology Officer. He has developed numerous simulations and failure analysis methodologies including crash energy absorption in braided composites, aircraft battle damage assessment, and active damage detection in metal and composite structures, and his team was awarded the 2003 SPIE Smart Structures Award for advanced nondestructive inspection technology. Shawn currently holds three patents and has twelve patents pending.



PLACE PHOTO HERE

Dimitry Gorinevsky (M'91–SM'98–F'06) received a Ph.D. degree from department of mechanics and mathematics of Moscow (Lomonosov) University, in 1986, and a M.Sc. in aerospace engineering from the Moscow Institute of Physics and Technology, in 1982.

He is a Consulting Professor of Electrical Engineering at Stanford University and heads a consultancy Mitek Analytics LLC working with Government and industry. He worked for Honeywell for 10 years. Prior to that, he held research, engineering,

and academic positions in Moscow, Russia; Munich, Germany; Toronto and Vancouver, Canada. His interests are in decision and control systems applications across many industries. He has authored a book, more than 140 reviewed technical papers and a dozen patents. Dr. Gorinevsky is a former Associate Editor of IEEE Transactions on Control Systems Technology. He is a recipient of Control Systems Technology Award, 2002, and Transactions on Control Systems Technology Outstanding Paper Award, 2004, of the IEEE Control Systems Society.



PLACE PHOTO HERE

Stephen Boyd (S'82–M'85–SM'97–F'98) received the A.B. degree in Mathematics from Harvard University in 1980, and the Ph.D. in Electrical Engineering and Computer Science from the University of California, Berkeley, in 1985.

He is currently the Samsung Professor of Engineering, and Professor of Electrical Engineering in the Information Systems Laboratory at Stanford University. His current research focus is on convex optimization applications in control, signal processing, and circuit design.



PLACE PHOTO HERE

Seung-Jean Kim (M02) received the Ph.D. degree in electrical engineering from Seoul National University, Seoul, Korea. Since October 2008, he has been with Quantitative Strategies, Citi Alternative Investments, where he is a vice president. From April 2002 to September 2008, he held a consulting assistant professor position at the Information Systems Laboratory (ISL), Department of Electrical Engineering, Stanford University. From October 2004 to March 2005, he worked for a startup in electronic design automation. From July 2002 to September

2004, he was a postdoctoral scholar at the ISL. His current research interests include convex and large-scale numerical optimization, computational finance, and computational methods for machine learning, statistics, and time series analysis.



PLACE PHOTO HERE

Grant Gordon (M'02) received his Ph.D. degree from the department of Engineering Science and Mechanics at the Pennsylvania State University in 1995, and a M.Sc. in Engineering Mechanics also from Penn State in 1990.

He is the technical lead for research and application of structural health monitoring at the Honeywell Technology Research Center developing prognostic and diagnostic tools for predictive maintenance of airframes. Prior to joining Honeywell, he was an assistant professor of Acoustics at Penn State University and a research scientist with Babcock and Wilcox. He has published over 40 technical articles, holds seven patents, and has received various industry and academic awards including Aviation Week's Technology Innovation Award for 2003. Dr. Gordon is a certified Program Management Professional (PMP) and serves on various Aerospace Industry Committees.