

Estimating the sources of global sea level rise with data assimilation techniques

Carling C. Hay^a, Eric Morrow^b, Robert E. Kopp^c, and Jerry X. Mitrovica^{b,1}

^aDepartment of Physics, University of Toronto, Toronto, ON, M5S 1A7, Canada; ^bDepartment of Earth and Planetary Sciences, Harvard University, Cambridge, MA 02138; and ^cDepartment of Earth and Planetary Sciences and Rutgers Energy Institute, Rutgers University, Piscataway, NJ 08854

Edited by Jeffrey Shaman, Columbia University, New York, NY, and accepted by the Editorial Board March 22, 2012 (received for review October 26, 2011)

A rapidly melting ice sheet produces a distinctive geometry, or fingerprint, of sea level (SL) change. Thus, a network of SL observations may, in principle, be used to infer sources of meltwater flux. We outline a formalism, based on a modified Kalman smoother, for using tide gauge observations to estimate the individual sources of global SL change. We also report on a series of detection experiments based on synthetic SL data that explore the feasibility of extracting source information from SL records. The Kalman smoother technique iteratively calculates the maximum-likelihood estimate of Greenland ice sheet (GIS) and West Antarctic ice sheet (WAIS) melt at each time step, and it accommodates data gaps while also permitting the estimation of nonlinear trends. Our synthetic tests indicate that when all tide gauge records are used in the analysis, it should be possible to estimate GIS and WAIS melt rates greater than ~ 0.3 and ~ 0.4 mm of equivalent eustatic sea level rise per year, respectively. We have also implemented a multimodel Kalman filter that allows us to account rigorously for additional contributions to SL changes and their associated uncertainty. The multimodel filter uses 72 glacial isostatic adjustment models and 3 ocean dynamic models to estimate the most likely models for these processes given the synthetic observations. We conclude that our modified Kalman smoother procedure provides a powerful method for inferring melt rates in a warming world.

climate | Kalman filter

Sea level (SL) changes occur on timescales ranging from seconds to millions of years. At the longest of these scales, over millions of years, changes are driven by processes such as continental collision, sedimentation, variations in the mean spreading rate of the sea floor, and dynamic topography fluctuations due to mantle convection (1, 2). At intermediate timescales, from hundreds to tens of thousands of years, glacial isostatic adjustment (GIA) in response to the Late Pleistocene glaciation yields a global pattern of SL change that varies systematically from the near field to the far field of ancient ice cover (3). Finally, at shorter timescales, a suite of processes contributes to the complex spatiotemporal pattern of modern SL change (4). These latter processes can be separated into “dynamic effects,” which include ocean tides, ocean–atmosphere interactions, salinity and temperature variations, and ocean circulation, and “static equilibrium effects,” which are the gravitational, deformational, and rotational signatures of mass flux from polar ice sheets and mountain glaciers (5). In this paper we explore the connection between the processes and sources that give rise to SL change and the geographic variability inherent in the associated observational constraints to investigate whether—and to what accuracy—the signal from recent ice mass fluxes may be estimated from modern tide gauge records of SL.

Twentieth century rates of SL change have been estimated using data from tide gauges: devices that measure relative sea level (i.e., the difference in height between the sea surface and the sea floor) at a specific geographic location. The Permanent Service for Mean Sea Level (PSMSL) compiles tide gauge records for >1,000 sites distributed around the globe. Estimates of globally averaged 20th century SL change rates based on various subsets of these tide gauge records (6–11) fall within the range of 1–2 mm/y. In the last decade of the 20th century and in the current century, tide gauges have been complemented by satellite altimetry measurements of sea surface height and ocean float measurements of temperature

and salinity (12). These records will be critical for understanding recent and future SL rise, but they are too short to inform our understanding of 20th century sea levels.

Estimating mean global SL alone ignores the information content inherent to geographic variability. As an example, it is now well understood that the rapid melting of any individual grounded ice reservoir gives rise to a distinct and highly nonuniform pattern—or fingerprint—of static equilibrium sea level change (13–17). In particular, in the near field of a melting ice sheet (within $\sim 2,000$ km of the margin), SL will fall due to both crustal uplift and the reduction of the gravitational pull on the ocean from the ice sheet. This fall can be an order of magnitude greater than the equivalent rise in mean global SL associated with the meltwater addition to the ocean. In contrast, in the far field of the ice sheet, SL will rise with (generally) greater amplitude as the distance from the ice sheet increases; this rise can exceed the global mean value by $\sim 30\%$. These fingerprints provide a framework for moving beyond inferences of globally averaged SL rise to estimate the contributions from individual meltwater sources. As an example, Mitrovica et al. (16) argued that the lower rates of SL change inferred from tide gauge records at European sites relative to the global average were consistent with 20th century melting from Greenland at a level equivalent to a eustatic SL rise of ~ 0.5 mm/y.

Efforts to fingerprint modern records of SL change face potential contamination from dynamic processes, and thus robust estimates of ice mass flux must address the extent to which this contamination can be accounted for. Dynamic SL processes can be subdivided into three categories, due respectively to (i) internal, unforced dynamics (e.g., ocean tides and some air–sea interactions), (ii) atmospheric and ocean warming, and (iii) freshwater addition.

In this study, we outline and evaluate a statistical formalism for estimating global sea level change and the individual contributions to this change in the presence of the three dynamic SL effects listed above. Most previous work has estimated mean global SL change by computing the linear least-squares fit to an assumed constant, long-term trend in select tide gauge records (8, 10). In contrast, our formulation makes no assumption in regard to the stationarity of the trends. This generalization is accomplished by using a Kalman filter (18) and smoother (19) to recursively compute the maximum-likelihood estimate of the total SL change as well as an estimate of its associated uncertainty.

There are several advantages in using this approach. Due to the recursive form of the algorithm, the Kalman smoother naturally accommodates measurements with data gaps. This is an important feature, given that tide gauge records extend back into the 19th

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Fostering Advances in Interdisciplinary Climate Science” held March 31–April 2, 2011, at the AAAS Auditorium in Washington, DC. The complete program and audio files of most presentations are available on the NAS Web site at www.nasonline.org/climate-science.html.

Author contributions: C.C.H., E.M., R.E.K., and J.X.M. designed research; C.C.H. and E.M. performed research; C.C.H., E.M., and R.E.K. contributed new reagents/analytic tools; C.C.H. and E.M. analyzed data; and C.C.H., E.M., R.E.K., and J.X.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. J.S. is a guest editor invited by the Editorial Board.

¹To whom correspondence should be addressed. E-mail: jxm@eps.harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1117683109/-DCSupplemental.

century but are characterized by pervasive and extensive data gaps. Accommodating data gaps permits the use of a larger selection of data than would be available when selecting and truncating records on the basis of, among other criteria, record continuity. The Kalman smoother assimilates observations when they are available, and in the absence of observations the smoother relies on the model dynamics. The recursive nature of the estimator approach is also useful because it reduces the storage requirements necessary when a large volume of observations are available. Additionally, although in essence a least-squares estimator, the Kalman smoother permits the determination of time-varying nonlinear trends in the measurement data that are overlooked using a simple linear least-squares approach. Finally, as our assumed model dynamics are linear, the Kalman smoother computes the optimal estimate of the state at every time step after convergence. We note that all of these advantages hold when the Kalman smoother is applied in applications that move beyond the determination of the total SL change to estimate time-varying contributions from individual sources of meltwater, such as the Greenland ice sheet (GIS) and the West Antarctic ice sheet (WAIS). Finally, we extend the Kalman smoother by including a multimodel filter (20), which provides a straightforward methodology for including uncertain contributions to the SL changes from sources other than meltwater, including ongoing GIA and thermosteric effects.

This study has two overarching goals. The first is to provide a rigorous and computationally efficient framework for analyzing the growing and disparate database of SL records. The second is to assess the accuracy with which individual sources of meltwater are detectable in synthetic SL records that have the same temporal and geographic data extent as presently available SL observations and are also characterized by realistic levels of dynamic variability. Our analysis is motivated by the following interrelated questions: For a suite of ice sheet melting scenarios, can we robustly identify individual ice sheets' static equilibrium SL signals in the presence of dynamic variability over the last century given the current spatio-temporal sampling of the SL field? Are there specific areas in the world where observations are particularly effective in constraining these contributions? To what extent are these estimates degraded by uncertainties associated with other contributors to the SL change, including GIA and thermosteric signals? Is it possible to use 20th century SL measurements to estimate GIA and thermosteric signals?

This paper focuses on a description of the modified Kalman smoother procedure and a discussion of its performance when applied to synthetic data derived from tide gauge observations. *Statistical Formalism* describes in detail the formalism and the *Implementation* section describes the coverage and noise characteristics of the existing tide gauge dataset, the sea level fingerprint predictions that are adopted in the construction of the synthetics, the models used to account for dynamic SL and GIA, and the implementation of our statistical procedure as applied to tide gauge data. Finally, *Results* discusses the performance of the Kalman smoother in a large suite of synthetic tests.

Statistical Formalism

Kalman Smoother. The smoothing algorithm implemented in this study can be divided into two main steps: a forward filtering pass and a backward smoothing pass. In the forward pass, the filter uses all past information to obtain an optimal estimate of the state vector, whereas the smoothing step runs backward in time, computing a smoothed state estimate using information from both the forward and the backward passes.

In the forward pass, the objective of the Kalman filter (21) is to use observations (or measurements), z_k , to estimate the state vector, x_k , at every time step (k). We model the observations with five different contributions to SL in the state vector: (i) SL changes due to melting of the GIS at a rate given by α_k ; (ii) SL changes due WAIS melting at a rate of β_k ; (iii) a geographically uniform SL rise, U_k , that accounts for any unmodeled sources of SL rise such as mountain glaciers; (iv) SL trends due to GIA (G_i); and (v) the SL contribution, S_i , due to ocean dynamics including thermal expansion.

In the filter, the observation model \hat{z}_k is defined as

$$\hat{z}_k = \mathbf{H}_k \mathbf{x}_k + \varepsilon_k, \quad [1]$$

where \mathbf{H}_k is the observation matrix, \mathbf{x}_k is the state vector, and ε_k is the white observation noise, arising primarily from ocean dynamics (*Sea Level Predictions* section), defined to have a mean of zero and a covariance \mathbf{R} .

At every time step we define our state vector as

$$\mathbf{x}_k = [\mathbf{h}_k \quad U_k \quad \alpha_k \quad \beta_k \quad \mathbf{v}_k]^T, \quad [2]$$

where U_k , α_k , and β_k are the scalar weightings of the first three contributors to sea level in our synthetic datasets, as listed above; the vector \mathbf{h}_k is the SL at each site i that includes contributions from GIS and WAIS melting, uniform SL rise, GIA, and ocean dynamics; and \mathbf{v}_k accounts for unforced dynamic SL variability. We account for any temporal correlation in the dynamic variability by modeling \mathbf{v}_k as an autoregressive (AR) process of order r (*SI Text*). The sum of \mathbf{h}_k and \mathbf{v}_k at every time step are the model's estimate of the observations (\hat{z}_k). Thus, \mathbf{h}_k is a row vector whose length is equal to the number of tide gauge sites, n_{sites} , used in the analysis, and \mathbf{v}_k is a row vector whose length is equal to $n_{\text{sites}} \cdot r$.

At every time step we construct the observation matrix \mathbf{H}_k that maps the model state vector into the observation space. \mathbf{H}_k , defined as

$$\mathbf{H}_k = \frac{\partial \hat{z}_k}{\partial \mathbf{x}_k}, \quad [3]$$

generates the observation matrix

$$\mathbf{H}_k = [I_{n_k \times n_k} \quad \mathbf{0}_{n_k \times 3} \quad \mathbf{H}_v], \quad [4]$$

where \mathbf{H}_v is the following $\langle n_k \times (n_k \cdot r) \rangle$ matrix,

$$\mathbf{H}_v = \begin{bmatrix} \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} & \dots & \begin{bmatrix} 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \end{bmatrix} \end{bmatrix}, \quad [5]$$

where each block has the dimension $\langle n_k \times r \rangle$. The sizes of both matrices depend on the number of observations, n_k , available at time step k .

Each time step, k , contains a forecast step and an analysis step. In the forecast step, predictions of a new estimate of the state vector, $\hat{\mathbf{x}}_k^f$, and its covariance, \mathbf{P}_k^f , are based on the analysis estimates from the previous time step using the following equations:

$$\hat{\mathbf{x}}_k^f = \Phi \hat{\mathbf{x}}_{k-1}^a + \mathbf{B} \mathbf{u}_k + \mathbf{w} \quad [6]$$

$$\mathbf{P}_k^f = \Phi \mathbf{P}_{k-1}^a \Phi^T + \mathbf{Q}. \quad [7]$$

The propagation from time $k - 1$ to time k depends on the state transition matrix Φ , the model noise \mathbf{w} , defined to have zero mean and covariance \mathbf{Q} , the input control parameter \mathbf{u} , and the matrix \mathbf{B} that maps \mathbf{u} into the state space. The input control matrix is determined a priori and does not depend on the filter's estimate of the state vector at any time step.

The state transition matrix, Φ , predicts how the state variables evolve between time steps,

$$\Phi = \begin{bmatrix} \mathbf{I}_{n_{\text{sites}} \times n_{\text{sites}}} & \Delta t [\mathbf{I}_{n_{\text{sites}} \times 1} \quad \mathbf{y}_{\text{GIS}} \quad \mathbf{y}_{\text{WAIS}}] & \mathbf{0}_{n_{\text{sites}} \times (r \cdot n_{\text{sites}})} \\ \mathbf{0}_{3 \times n_{\text{sites}}} & \mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times (r \cdot n_{\text{sites}})} \\ \mathbf{0}_{(r \cdot n_{\text{sites}}) \times n_{\text{sites}}} & \mathbf{0}_{(r \cdot n_{\text{sites}}) \times 3} & \mathbf{p} \end{bmatrix}, \quad [8]$$

which assumes that the globally uniform SL change and both ice sheet melt rates remain constant between time steps. The matrix

ρ contains the AR coefficients that characterize the noise at every site (*SI Text*).

The input control parameter \mathbf{u}_k governs how SL will vary due to modeled GIA and ocean dynamics:

$$\mathbf{u}_k = \begin{bmatrix} \mathbf{G} \\ \dot{\mathbf{S}}_k \end{bmatrix}. \quad [9]$$

Here \mathbf{G} is a vector of GIA rates at every site i , whereas $\dot{\mathbf{S}}_k$ is the instantaneous rate of SL change due to ocean dynamics. \mathbf{G} is set to be one GIA model (*Sea Level Predictions*) and $\dot{\mathbf{S}}_k$ is derived from one historical climate model (*Dynamic Sea Level*). Although one GIA model and one climate model must be chosen in the implementation of the filter, these models can be varied in repeated tests to examine the sensitivity of the filter on model choice. This procedure is discussed in detail in *Multimodel Kalman Filter*.

With the input control parameter introduced, we define the matrix \mathbf{B} that relates \mathbf{u}_k to the state space:

$$\mathbf{B} = \Delta t \begin{bmatrix} \mathbf{I}_{n_{\text{sites}} \times n_{\text{sites}}} & \mathbf{I}_{n_{\text{sites}} \times n_{\text{sites}}} \\ \mathbf{0}_{3 \times n_{\text{sites}}} & \mathbf{0}_{3 \times n_{\text{sites}}} \\ \mathbf{0}_{(n_{\text{sites}} \cdot \tau) \times n_{\text{sites}}} & \mathbf{0}_{(n_{\text{sites}} \cdot \tau) \times n_{\text{sites}}} \end{bmatrix}. \quad [10]$$

Substituting Φ , \mathbf{u}_k , and \mathbf{B} into Eq. 6 results in the following recursive equations for our five state variables:

$$\begin{aligned} \mathbf{h}_k^f &= \mathbf{h}_{k-1}^a + \Delta t (U_{k-1}^a + \alpha_{k-1}^a \mathbf{y}_{\text{GIS}} + \beta_{k-1}^a \mathbf{y}_{\text{WAIS}} + \mathbf{G} + \dot{\mathbf{S}}_k) + \mathbf{w}_h \\ U_k^f &= U_{k-1}^a + w_U \\ \alpha_k^f &= \alpha_{k-1}^a + w_\alpha \\ \beta_k^f &= \beta_{k-1}^a + w_\beta \\ \mathbf{v}_k^f &= \rho \mathbf{v}_{k-1}^a + \mathbf{w}_v. \end{aligned} \quad [11]$$

The SL height, \mathbf{h}_k^f , thus depends on the SL at the previous time step, \mathbf{h}_{k-1}^a ; the estimated melt rates, U_{k-1}^a , α_{k-1}^a , and β_{k-1}^a from the previous time step; and the control inputs \mathbf{G} and $\dot{\mathbf{S}}_k$. The predictions for the melt rates U_k^f , α_k^f , and β_k^f are modeled as random processes with AR1 coefficients equal to 1. As previously noted, the temporally correlated measurement noise, \mathbf{v}_k^f , is also modeled as an AR process, of order r , with coefficients in the matrix ρ . We acknowledge that aspects of this state propagation model may be simplistic, particularly the time-invariant melt rates, but we rely on the available observations to correct the state propagation.

The model noise for all state variables is characterized by \mathbf{w} , which has zero mean and covariance \mathbf{Q} . Although \mathbf{Q} can take on a time-varying structure, for simplicity we choose the following constant form:

$$\mathbf{Q} = \begin{bmatrix} \mathbf{V}_h & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \sigma^2 \mathbf{I}_{3 \times 3} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{R}' \end{bmatrix}. \quad [12]$$

The first upper diagonal is an $n_{\text{sites}} \times n_{\text{sites}}$ matrix corresponding to the noise in the modeled heights. The \mathbf{V}_h structure is an exponential covariance matrix:

$$\mathbf{V}_h = \sigma_h^2 e^{-\frac{\|\mathbf{D}\|}{\tau}}. \quad [13]$$

Here τ is the decorrelation length scale and \mathbf{D} is the distance between tide gauge sites (21). The parameter σ_h^2 gives the variance of each SL height, and σ^2 is the variance of the three melt rates (which is assumed to be the same for all three). The final element of \mathbf{Q} is the covariance of \mathbf{w}_v that is associated with the correlated noise. This element of \mathbf{Q} is set to \mathbf{R}' , which contains \mathbf{R} (the covariance of the detrended tide gauge observations—see *SI Text* for a description of \mathbf{R}). The elements of \mathbf{R} are assigned within \mathbf{R}' to the positions corresponding to the first lagged correlated noise state for each site.

Following the forecast step is the analysis step when the observations are assimilated. The analysis state vector, $\hat{\mathbf{x}}_k^a$, and the analysis error covariance, \mathbf{P}_k^a , can be updated using Eqs. 14 and 15, respectively:

$$\hat{\mathbf{x}}_k^a = \hat{\mathbf{x}}_k^f + \mathbf{K}_k (\mathbf{z}_k - \mathbf{H}_k \hat{\mathbf{x}}_k^f) \quad [14]$$

$$\mathbf{P}_k^a = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k^f. \quad [15]$$

In Eq. 14, the analysis state vector is a linear combination of the forecast estimate of the state vector, $\hat{\mathbf{x}}_k^f$, and the weighted difference between the observations and the measurement prediction, $\mathbf{K}_k (\mathbf{z}_k - \mathbf{H}_k \hat{\mathbf{x}}_k^f)$. The weighting \mathbf{K}_k , known as the Kalman gain matrix, minimizes the *posteriori* error covariance in Eq. 15 and is calculated following Eq. 16:

$$\mathbf{K}_k = \mathbf{P}_k^f \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_k^f \mathbf{H}_k^T + \mathbf{R}_k)^{-1}. \quad [16]$$

If observations are unavailable, then $\hat{\mathbf{x}}_k^a = \hat{\mathbf{x}}_k^f$, and $\mathbf{P}_k^a = \mathbf{P}_k^f$. Once the forward pass through the observations is complete, the backward smoothing pass is implemented. Whereas the Kalman filter pass uses all observations at and before the time at which the state vector is estimated, the smoothing pass runs backward in time, using intermediate results from the forward pass so that the final state estimate depends on all available observations over time. This type of smoothing algorithm is known as a fixed-interval Rauch–Tung–Striebel (RTS) two-pass smoother (19).

The smoothed estimate at the final time step t_f is initialized in the following manner:

$$\hat{\mathbf{x}}_{k=t_f}^s = \hat{\mathbf{x}}_{k=t_f}^a. \quad [17]$$

$\hat{\mathbf{x}}_k^f$ is then recursively calculated using Eq. 18,

$$\hat{\mathbf{x}}_k^s = \hat{\mathbf{x}}_k^a + \mathbf{A}_k (\hat{\mathbf{x}}_{k+1}^s - \hat{\mathbf{x}}_{k+1}^f), \quad [18]$$

where

$$\mathbf{A}_k = \mathbf{P}_k^a \Phi^T \mathbf{P}_{k+1}^{f^{-1}}. \quad [19]$$

The associated uncertainty on the smoothed estimate, \mathbf{P}_k^s , is also calculated on the backward pass:

$$\mathbf{P}_k^s = \mathbf{P}_k^a + \mathbf{A}_k (\mathbf{P}_{k+1}^s - \mathbf{P}_{k+1}^f) \mathbf{A}_k^T. \quad [20]$$

It is important to emphasize that although the backward pass does not involve the processing of actual observations, it does use the full forward filtering solution. As we are estimating historical melt rates rather than predicting future ones, the smoothing pass ensures that at every point in time our estimated state vector and its associated uncertainty include all information available between the first and the final time steps.

Multimodel Kalman Filter. As evident in the description of the forward filter pass, the Kalman filter forecast step requires the construction of the input control parameter \mathbf{u}_k (Eq. 6), which depends on the choice of models both for GIA and the ocean dynamic signal. Although there is only one correct GIA model and one correct ocean dynamics model, we are, in practice, unsure which models those are. Filter convergence can occur in the case of a poorly modeled \mathbf{u}_k ; however, this situation will produce residuals between the state estimate and the observations that can grow over time (22). This problem can be addressed by either estimating the unknown model parameters by adding them to the state vector or extending our Kalman smoother to include and test multiple models (20). We implement a multimodel filter where we define each model, $\boldsymbol{\mu}_j$, to be a combination of one GIA model and one ocean dynamics model, resulting in a total of N

different model pairings. The objective of the multimodel filter is to use the residuals between each filtered-model estimate and the observations to determine the most likely combination of GIA and ocean dynamics models.

For each model, μ_j , we define the likelihood of one model given the observations at time step k as the conditional probability density function,

$$f(\hat{\mathbf{z}}_k | \mu_j) = \frac{1}{(2\pi)^{n_{\text{states}}/2} \sqrt{|\mathbf{C}_{k, \mu_j}|}} e^{-\frac{1}{2} (\hat{\mathbf{z}}_k - \mathbf{H}\hat{\mathbf{x}}_{k, \mu_j}^f)^T \mathbf{C}_{k, \mu_j}^{-1} (\hat{\mathbf{z}}_k - \mathbf{H}\hat{\mathbf{x}}_{k, \mu_j}^f)}; \quad [21]$$

i.e., $\hat{\mathbf{z}}_k$ is normally distributed with mean $\mathbf{H}\hat{\mathbf{x}}_{k, \mu_j}^f$ and covariance \mathbf{C}_{k, μ_j} , where \mathbf{C}_{k, μ_j} is defined as

$$\mathbf{C}_{k, \mu_j} = \mathbf{H}\mathbf{P}_{k, \mu_j}^f \mathbf{H}^T + \mathbf{R}. \quad [22]$$

All variables are consistent with their previous definitions and n_{states} is the size of the state vector. The calculation of f requires the forecasted estimate and covariance of the state vector because these estimates depend on the chosen GIA and ocean dynamics models. We compute the likelihood of each model at each time step, using the filtered state estimates $\hat{\mathbf{x}}_k^f$ and \mathbf{P}_k^f from Eqs. 6 and 7. Calculating f also requires computing the inverse and determinant of the sometimes ill-conditioned matrix \mathbf{C} . To estimate the inverse when the covariance matrix has a condition number exceeding a threshold value, we compute the eigenvalues of \mathbf{C} and then regularize by retaining only the largest eigenvalues and eigenvectors that contain 99.5% of the variance of the covariance matrix. We then reconstruct \mathbf{C} using the truncated eigenvectors and eigenvalues.

At each time step k , we calculate the probability, p_j , that the model, μ_j , is the correct model at that time:

$$p_j(k) = \frac{f(\hat{\mathbf{z}}_k | \mu_j) p_j(k-1)}{\sum_{i=1}^N f(\hat{\mathbf{z}}_k | \mu_i) p_i(k-1)}. \quad [23]$$

The recursive computation of Eq. 23 requires initialization of the probabilities at the first time step. While prior knowledge of the likelihood of individual models can be used to initialize p , in the absence of this information we can consider all N models to be equally likely; i.e., $p_j(0) = 1/N$. Although the likelihood and probabilities could be computed in one batch, we calculate them recursively to avoid large matrix inversions and storage. We then construct the model-conditioned estimate of the state vector, $\hat{\mathbf{x}}_k^*$, and the associated error covariance \mathbf{P}_k^* using Eqs. 24 and 25,

respectively. These quantities are defined to be the weighted combination of each model's smoothed state and covariance estimates with the weightings set to the model probability estimates at the final time step, $p(t_f)$:

$$\hat{\mathbf{x}}_k^* = \sum_{j=1}^N p_j(t_f) \hat{\mathbf{x}}_{\mu_j, k}^s \quad [24]$$

$$\mathbf{P}_k^* = \sum_{j=1}^N p_j(t_f) \left[\mathbf{P}_{\mu_j, k}^s + (\hat{\mathbf{x}}_k - \hat{\mathbf{x}}_{\mu_j, k}^s) (\hat{\mathbf{x}}_k - \hat{\mathbf{x}}_{\mu_j, k}^s)^T \right]. \quad [25]$$

Implementation

Datasets. There are 1,166 tide gauges in the PSMSL revised local reference (RLR) database (www.psmsl.org). Although tide gauge records exist for sites in all continents (Fig. 1), there is a significant sampling bias, with the vast majority of tide gauges found along the coasts of North America and Eurasia. Examination of the time series at individual sites reveals that, in addition to spatial biases, the length of the datasets varies significantly throughout the database, with the longest records located in the Northern Hemisphere, particularly in the United Kingdom and Scandinavia (Fig. S1).

For the present study, we use mean annual tide gauge records to avoid seasonal variability. The largest interannual variability (Fig. 1) occurs along the northern coast of Russia, in closed basins (e.g., the Mediterranean and Black Seas), and in the western Pacific. We have chosen to use tide gauges that have at least 20 y of data in the last 50 y, leaving 596 records for analysis. We have binned the 596 tide gauges into the five regions illustrated in Fig. 1. Region 1 includes all tide gauges located poleward of 60°N, regions 2–4 move southward in 20° latitudinal bins, and region 5 includes all sites south of the equator.

Dynamic Sea level. To estimate long-term, warming-forced dynamic sea level trends, we use three climate models that participated in the World Climate Research Programme's (WCRP's) Coupled Model Intercomparison Project phase 3 (CMIP3) multi-model dataset (23): the Goddard Institute for Space Studies Model-E with the Russell Ocean Model (GISS-ER), the Model for Interdisciplinary Research on Climate Medium Resolution Model 3.2 (MIROC-3.2 medres), and the Meteorological Institute of the University of Bonn ECHAM4 HOPE-G Model (MIUB ECHO-G). The climate models are solely used to determine long-term trends in SL due to changing ocean and atmosphere dynamics; decadal-scale and interannual variability is included in the noise term.

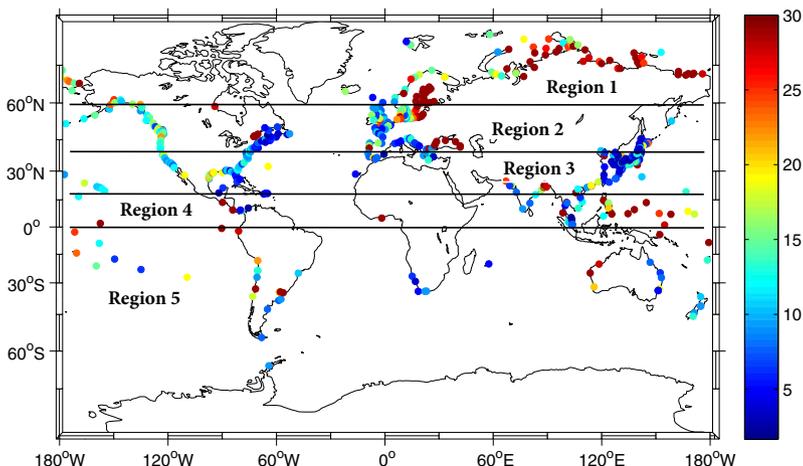


Fig. 1. The variance (cm^2) of detrended sea level observations at 596 tide gauge sites in the PSMSL RLR database subdivided into five regions. Region 1 includes all tide gauges poleward of 60°N, and each successive region (2–4) moves southward in 20° latitudinal bins. All southern hemisphere sites are binned into region 5.

Although the climate model outputs are provided on a global grid, we require coastal points for our tide gauge analysis. We estimate the model output at each site by using the model cell that is coincident with the tide gauge position if it exists. If the tide gauge location lies between multiple model grid points, we use an inverse distance interpolation on all available surrounding grid cells that are within a 1,000-km radius of the site. The 1,000-km radius is appropriate given that the correlation length scale of the tide gauge observations is $\sim 1,500$ km (*SI Text*) and the southern hemisphere is data sparse, thereby requiring a large radius to ensure sufficient observations for interpolation. Despite the apparently large radius, the inverse distance relationship results in significantly higher weightings being given to sites close to the interpolated location.

Sea Level Predictions. The nonuniform patterns of SL change that arise due to rapidly melting ice sheets are commonly referred to as SL fingerprints. The fingerprints we adopt in this study are calculated by assuming geographically uniform ice melting across both GIS and WAIS (24). They are based on a gravitationally self-consistent SL theory (25) that incorporates shoreline migration due to local transgressions and regressions and changes in grounded, marine-based ice cover, as well as feedback into SL of contemporaneous perturbations in the Earth’s rotation vector (26). Our calculations adopt a 1D elastic Earth model with structure prescribed by the seismically inferred Preliminary Reference Earth Model (PREM) (27). This elastic (i.e., instantaneous) treatment of load-induced deformation and gravity changes is valid for melting events with timescales of several centuries or less (see Fig. S2 for GIS and WAIS SL fingerprints).

The effect of ongoing GIA due to the last ice age on relative SL observations is modeled using a more general version of the SL theory described above that incorporates viscoelastic adjustments. GIA predictions are sensitive to both the space–time geometry of the Late Pleistocene ice cover and the Earth’s viscoelastic structure. Most previous analyses of modern SL datasets were corrected using a GIA calculation based on a radial viscosity profile that increases only moderately with depth from the base of the lithosphere to values of $\sim 2\text{--}3 \times 10^{21}$ Pa-s in the deep mantle (6–8, 28, 29). However, Mitrovica and Davis (30) showed that relatively minor variations in the viscoelastic Earth structure map into uncertainties in estimates of globally averaged SL rise of order of 0.5 mm/y. Furthermore, an increase in the lower mantle viscosity to values of $5\text{--}10 \times 10^{21}$ Pa-s, preferred on the basis of several analyses of GIA-related datasets (31, 32), significantly reduces systematic geographic trends in GIA-corrected tide gauge rates along the US East Coast (9).

In the formulation described above, uncertainties in the GIA signal are explored by performing the analysis using a suite of GIA models. We have run 72 GIA models, all of which are based on the global ice sheet reconstruction model ICE-5G (33). The models are distinguished on the basis of the choice for the lithospheric thickness and the viscosities in both the upper and the lower mantle. Specifically, we use lithosphere thicknesses in the range 72–125 km, upper mantle viscosities between 0.3 and 1×10^{21} Pa-s, and lower mantle viscosities between 2 and 20×10^{21} Pa-s. Illustrations of the SL prediction for one GIA model and the variance between the 72 models are presented in *SI Text* (Fig. S3).

Synthetic Datasets. To evaluate the performance of the Kalman smoother with known signals, we construct a suite of synthetic tide gauge data in the following manner. First, we remove long-term trends from the 596 RLR time series. For tide gauges containing <50 y of data we remove a linear trend, and for time series spanning >50 y we remove both a linear and a quadratic signal. Removing only a linear trend from the shorter time series was necessary to avoid the unrealistically large quadratic terms that occasionally resulted from estimating higher-order polynomial coefficients from insufficient data. The intent of the detrending is to remove all long-term trends associated with melting ice sheets, thermosteric and halosteric signals, GIA, and tectonics, leaving primarily decadal and interannual variability,

which we treat as the observation noise. The detrended time series thus capture the shorter-term dynamic variability evident in the observations that is due to ocean circulation changes, long period tides, air–sea interactions, temperature and salinity variations, etc., as well as any regional variability. Second, we add SL fingerprints computed for various scenarios of GIS and WAIS melting. The SL fingerprints are normalized using the equivalent eustatic SL change and therefore the construction of the synthetic observations requires only that a weighting (in units of mm/y) be applied to each of the two fingerprints. A geographically uniform trend is also added to the synthetics to account for any additional, unmodeled, sources of SL rise such as mountain glacier and small ice sheet melting. Finally, we also add at each site SL contributions due to GIA and ocean dynamics, which include thermal expansion of the ocean.

The general expression for the synthetic tide gauge record at a given site, i , is given by

$$z_i(t) = \sum_{k=1}^{t_f} \left[\left(U_k + \alpha_k y_{i,\text{GIS}} + \beta_k y_{i,\text{WAIS}} + G_i \right) \Delta t \right] + S_i(t) + v_i(t), \quad [26]$$

where z_i is the time series of relative SL at site i ; U_k is a geographically uniform SL rise at time $t = k$; $y_{i,\text{GIS}}$ and $y_{i,\text{WAIS}}$ are the normalized GIS and WAIS SL fingerprints, sampled at site i , with prescribed melt rates at time $t = k$ given by α_k and β_k , respectively; G_i is the sea level trend due to GIA at site i ; Δt is time step; S_i is the SL contribution from ocean dynamics at each site i ; and v_i is the site-specific dynamic variability (considered to be noise) in the observations, with covariance \mathbf{R} . Because the noise $v(t)$ is derived from PSMSL data, the synthetic measurements contain the same irregular and pervasive data gaps as the PSMSL data. Although the longest tide gauge record begins in 1807 and all records end by the year 2008, resulting in the full dataset spanning 202 y, we begin our synthetic dataset in 1861 to coincide with the climate model runs. The GIA model added to the synthetics is arbitrarily chosen to have the Earth’s structure characterized by a lithospheric thickness of 125 km and upper and lower mantle viscosities of 0.5×10^{21} Pa-s and 5×10^{21} Pa-s, respectively. As discussed above, the dynamic SL, S_i , is taken from 140 y of a randomly selected historical climate simulation. Specifically, we set the heights from 1861 to 2000 to be the heights from the MIROC3.2 medres historical climate model realization 1 and all dynamic heights after this time as constant at the year 2000 values. From each tide gauge site we use only the linear and quadratic trends in the model-derived dynamic SL time series, which ensures that all short term, interannual, and decadal variability in the synthetic time series arises from the detrended tide gauge observations.

With the synthetic datasets constructed, we can define the covariance matrices \mathbf{R} and \mathbf{Q} and initialize the state vector and its associated uncertainty (*SI Text*). We then recursively apply the algorithm to the full synthetic dataset, computed using one GIA–ocean dynamics model pairing, and run it forward and backward in time. The multimodel filter is implemented using the 72 GIA models (*Sea Level Predictions*) and 3 ocean dynamics models (Dynamic Sea Level), resulting in 216 GIA–ocean dynamic model pairings. For all tests, we calculate and plot the probabilities of each model as a function of time and present the final melt rate estimates and their associated uncertainties.

Results

Results based on applying our Kalman smoother methodology to synthetic tide gauge data are presented in Figs. 2 and 3. The detection analysis is performed for each individual geographic region (Fig. 1) and for the combination of all regions. We begin by presenting the probabilities, calculated using Eq. 23, of each of the 216 GIA–dynamic ocean pairing models as a function of time (Fig. 2). Each frame of Fig. 2 corresponds to results associated with a specific subset of tide gauges, beginning with

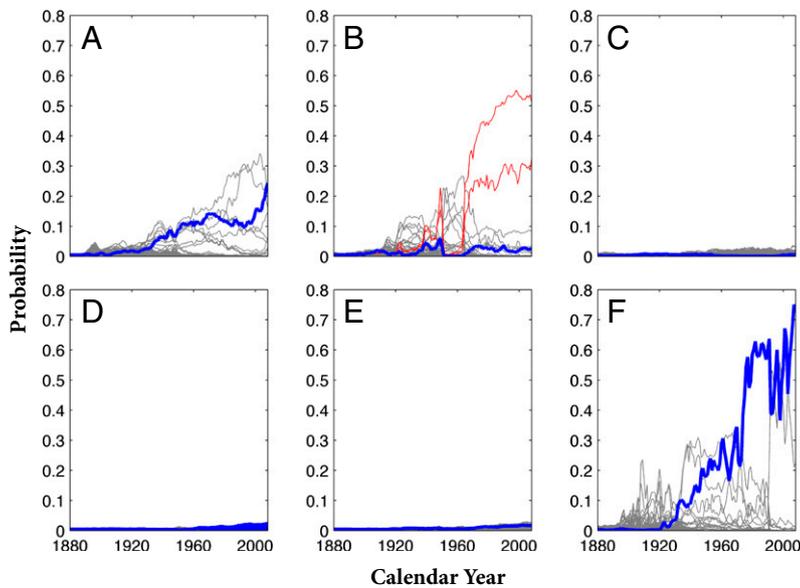


Fig. 2. The probabilities for each of the 216 GIA and ocean dynamics model combinations computed using the multimodel Kalman filter (main text). The thick blue line corresponds to the model that includes the correct GIA and ocean dynamics model pairing, and the gray lines are associated with the remaining 215 GIA–dynamic ocean pairings. Results are shown for each of the five regions in Fig. 1 (A–E), as well as for the case when all 596 tide gauges are used in the analysis (F). The thin red lines in B are the results for 2 model pairings that include the correct ocean dynamics model and a GIA model with the correct lithospheric thickness and correct upper mantle viscosity, but incorrect lower mantle viscosity.

region 1 in Fig. 2A and ending in Fig. 2F, where all 596 tide gauges are considered. In this scenario, the synthetics are computed using constant GIS and WAIS melt rates of 0.3 and 0.5 mm eustatic sea level (esl)/y, respectively, and a uniform rise in SL of 0.8 mm/y. The thick blue line in Fig. 2A–F corresponds to the computed probability of the correct GIA and ocean dynamics pairing, and the gray lines correspond to the remaining 215 pairings. In addition to these probabilities, the melt estimates at the final time step along with their associated uncertainty ($\pm 1\sigma$) are provided in Table 1.

To begin, consider the results for region 1 (Fig. 2A and Table 1). This region is in the near field of the GIS and is therefore a location where the GIS SL fingerprint has a large signal (and signal gradient) due to both the decrease in the gravitational attraction of the GIS on the surrounding ocean and the crustal uplift associated with the ice and water unloading. The probabilities indicate that although the correct GIA–dynamic ocean model pairing becomes more likely over time, this particular region has a peak probability of only 0.2 by the final time step. The spread in probabilities associated with the remaining 215 models suggests that this region would not be the optimal subset to adopt in a detection effort. Indeed, although the WAIS melt rate is inferred with high accuracy, it has very low precision. The

opposite is true for the detection of the GIS melt rate. The GIS melt rate is inferred with greater precision than the WAIS melt rate (± 0.33 compared with ± 0.79) because the latter fingerprint is nearly uniform within region 1. This geometry makes it difficult to detect the signal in the presence of dynamic variability or to separate it from the geographically uniform component (estimated uncertainty of ± 0.92) of the SL synthetics. Detection of the two fingerprints is clearly challenging in this region, even in the case (as in this study) where local dynamical effects associated with the influx of fresh water are not accounted for.

There is no improvement in the accuracy of the GIS detection when data from region 2 are considered (Table 1). In fact, whereas the precision remains almost constant between regions 1 and 2, the estimated GIS melt rate moves farther from the true value. The significant uncertainty in the WAIS result is, once again, due to the difficulty in separating the WAIS fingerprint from a globally uniform SL trend. There is large GIA model variability in region 2 (Fig. S3). This variability makes the analysis of SL data in this region particularly sensitive to the GIA signal, and therefore data from this region provide a potentially strong constraint in estimating the GIA signal. We can explore this in more detail by considering the results from the multimodel filter (Fig. 2B). Examination of Fig. 2B reveals that whereas the

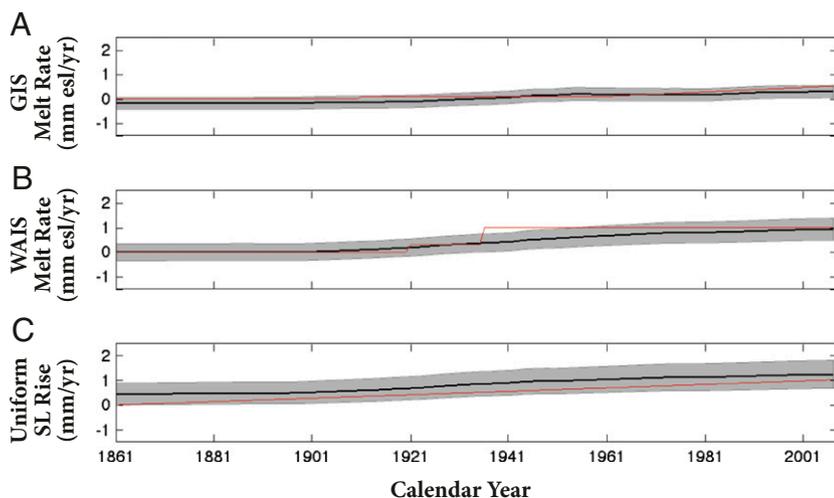


Fig. 3. Analysis of synthetic tide gauge data generated using all 596 tide gauges. Estimates of state variables, i.e., melt rates, after applying the multimodel smoother (main text) are plotted as a function of observation time (black lines), along with the $\pm 1\sigma$ range of uncertainty (gray shading), for estimates of (A) the GIS melt rate, (B) the WAIS melt rate, and (C) the globally uniform SL rise. The correct values (i.e., the values adopted in constructing the synthetics) are plotted in red. A–C show the results for the multimodel Kalman filter using all 216 model combinations and using synthetics constructed assuming accelerations in both ice sheet melt rates and in the uniform component of SL (main text).

Table 1. Final estimates and uncertainty ($\pm 1\sigma$) of the GIS melt rate, the WAIS melt rate, and the rate of globally uniform SL rise (mm/y)

	Greenland	West Antarctic	Uniform
Correct values	0.3	0.5	0.8
Region 1	0.09 \pm 0.33	0.40 \pm 0.79	0.73 \pm 0.92
Region 2	-0.30 \pm 0.34	0.65 \pm 0.73	1.10 \pm 0.88
Region 3	0.58 \pm 0.69	0.55 \pm 0.85	0.56 \pm 0.95
Region 4	0.56 \pm 0.96	0.52 \pm 0.91	0.87 \pm 0.99
Region 5	0.46 \pm 0.55	0.37 \pm 0.52	0.84 \pm 0.87
All sites	0.21 \pm 0.25	0.38 \pm 0.36	1.13 \pm 0.53

The listed estimates are taken from the last time step using the multimodel smoother. The values used in the synthetics ("correct values") are presented along with the results for each subset of tide gauges (Fig. 1) and for all tide gauges.

correct model combination (thick blue line) does not emerge as the most probable, the two most likely model combinations pair the correct ocean dynamics model with GIA models having the correct lithospheric thickness and upper mantle viscosity but incorrect lower mantle viscosity (10×10^{21} Pa·s and 8×10^{21} Pa·s for the most and second most probable models, respectively, instead of the correct value of 5×10^{21} Pa·s) (thin red lines). The filter's bias against estimating the lower mantle viscosity of the GIA model may be due to the fact that region 2 includes a large number of tide gauges situated close to the location of the former Fennoscandian ice sheet (Fig. 1). The depth of the post-glacial adjustments in the near field of an ice sheet depends on the size and weight of the ice sheet. Compared with the large Laurentide ice sheet, the solid earth deformation associated with the smaller Fennoscandian ice sheet was relatively shallow (limited to the top half of the mantle); therefore, GIA corrections in this region will be more sensitive to variations in the viscosity of the upper mantle than the lower. It is important to note that these sensitivities are highly dependent on the GIA model used in the synthetics and also on the binning of tide gauges.

As we consider tide gauge data in the vicinity of the equator (regions 3 and 4), it becomes increasingly difficult to distinguish between the SL fingerprints of GIS and WAIS melting and the uniform SL change, making detection of the fingerprints using data from this region alone highly problematic. This difficulty is reflected in the large errors in the inferred melt rates (Table 1). In addition, it also becomes difficult to distinguish between the 216 GIA–dynamic ocean model pairings using this data subset; indeed the probabilities of the model pairings remain nearly constant over time (Fig. 2 C and D). As we move farther south we find that detection of WAIS melting has a smaller estimated error in region 5 than in any other region (Table 1). This difference is due to the large gradient in the static WAIS fingerprint that this region samples. However, the computed probabilities in this region (Fig. 2E) indicate that we cannot distinguish among the 216 GIA–dynamic ocean model pairings.

When all tide gauges are used in the analysis, the increase in precision of both GIS and WAIS melt rates is significant (Table 1). Because the Kalman smoother computes the state estimate by weighting the variance from the least-squares estimator, the best detection occurs when all available data are included in the analysis. Using the full, global dataset limits contamination associated with long-term regional periodic signals such as the El Niño–Southern Oscillation and also makes it possible to constrain the GIA and ocean dynamics signal (Fig. 2F). In Fig. 2F, the correct model emerges out of the ensemble with a final probability value of ~ 0.75 . The second most probable model pairs the correct ocean dynamics model with a GIA model that has the correct lithosphere thickness and upper mantle viscosity, but an incorrect lower mantle viscosity (8×10^{21} Pa·s instead of the correct value of 5×10^{21} Pa·s).

To illustrate the time-evolving performance of the multimodel Kalman smoother, we present in Fig. 3 the results for one specific melting scenario. Fig. 3 shows the melt rates (black lines) and error ($\pm 1\sigma$, gray shading) estimates as a function of time.

To test the filter's ability to track nonlinear trends, we carried out analyses using synthetics that include accelerations in both the GIS and the WAIS melt rates (Fig. 3 A–C). The initial GIS and WAIS melt rates are both set to 0. In 1910, the GIS melt rates increase to 0.1 mm esl/y. This increase is followed 50 y later by an acceleration of 0.009 mm esl/y². For the WAIS we impose two jumps in the melt rate: one in 1920 and one in 1935, to values 0.3 mm esl/y and 1 mm esl/y, respectively. We also linearly increase the uniform SL rise from an initial value of 0 mm/y to a final value of 1 mm/y. The synthetics computed using these accelerations are not meant to be a recreation of historical tide gauge records, but rather serve as a test of our methodology.

In Fig. 3 A–C, we note that the estimated errors are nearly constant over time. Note that the results in Fig. 3 are the multimodel smoothed estimates and uncertainties (Eqs. 25 and 26), which include the full forward and backward solutions. Therefore, the estimates at each time step include all assimilated observations from 1861 until 2008. In this case, we found that the Kalman smoother was able to track the accelerating rates of all three contributors to SL. Note that the estimate of the WAIS melt rate smooths gradients in the melt rates adopted in the synthetics; once again this result is because the Kalman smoother includes both a forward and a backward component (Fig. 3B). Our estimated uncertainties ($\pm 1\sigma$) for the GIS and WAIS melt rates suggest that historical tide gauge records have information sufficient to robustly infer 20th century GIS and WAIS melt rates exceeding ~ 0.3 and ~ 0.4 mm esl/y, respectively.

Conclusions

Most past studies of recent SL change have focused on determining mean SL trends, generally disregarding the distinct geographic patterns in SL that rapidly melting ice sheets produce. Our tide gauge analysis has provided a series of detection thresholds for inferring individual contributions of the GIS and the WAIS to SL change during the last century.

We have implemented a multimodel Kalman smoother data assimilation technique that iteratively performs a least-squares analysis whenever observations are available and in the absence of observations relies on the model dynamics to compute the best estimate of the state variables. In this study, we apply the technique to estimate three contributions to synthetic SL trends: melting from the Greenland and West Antarctic ice sheets and a globally uniform rise representing unmodeled components. The Kalman smoother is well suited for this problem because it accommodates data gaps and generates minimum mean square estimates of the SL changes. Also, in contrast to ordinary least-squares techniques, the use of the Kalman smoother permits the estimate of a nonstationary signal, which is clearly advantageous for efforts to detect changes in ice sheet melt rates. Because the Kalman smoother is a recursive technique, it also avoids storage constraints when dealing with large datasets. Additionally, the multimodel filter component of our formalism is not only able to account for uncertainties in SL signals due to GIA and ocean dynamics, but also provides the possibility of using SL observations to constrain models of GIA and ocean dynamics probabilistically.

In our tide gauge analysis we used 596 detrended tide gauge time series from the PSMML database to construct a series of synthetic datasets for combined GIS and WAIS melt scenarios. Constructing our synthetic datasets using detrended tide gauge data ensures that our synthetics have the same data gaps as the historical tide gauge records. Having applied the Kalman smoother to subsets of tide gauges as well as the full set of synthetics, we have determined that when all available tide gauges are included in the analysis, it is possible, in principle, to rigorously estimate stationary and accelerating GIS and WAIS melt rates with 1σ uncertainty values of ± 0.3 and ± 0.4 mm esl/y, respectively. The size of these uncertainty values depends on the adopted model

error covariance matrix \mathbf{Q} and the observation noise covariance matrix \mathbf{R} . Initializing these matrices establishes the trade-off between the relative weighting of a priori model constraints and the observations in the inference procedure. Therefore, additional tuning of these covariances (i.e., of the trade-off) may be necessary to optimize filter performance.

The Kalman smoother explicitly includes information on dynamic SL and GIA, thereby providing more robust estimates of melt rates than previous studies. The results from the multimodel filter indicate, perhaps not surprisingly, that data from some regions of the ocean will be particularly useful for constraining models of GIA. Although we somewhat arbitrarily divided the 596 tide gauges into latitudinal regions, more careful binning could be used to select tide gauges that have the greatest sensitivities to the Earth structure governing GIA. Constraining the dynamic SL model was also possible, particularly when the full set of tide gauges was used. Expanding this suite of dynamic ocean models to include a broader group of climate models will be important in the application of this methodology on 20th century tide gauge observations. In addition, we chose to implement a fixed multimodel filter that assumes that there is a single correct model for GIA and ocean dynamics over the full time span of the data. This assumption is certainly valid when considering GIA because the Earth's structure will not change over centennial timescales. However, it may not hold for the ocean dynamics model when the methodology is applied to historical tide gauge records; that is, one of the suite of models for ocean dynamics may best capture the dynamics over one portion of the time history, whereas a second model may be most accurate in a different time window. One option is to include a dynamic multimodel filter, which would allow for the preferred ocean model to vary continuously over time.

In the methodology presented here we have included a uniform component of SL, which we attribute to unmodeled sources of SL change including mountain glaciers. In reality, each melting glacier produces a unique SL fingerprint. The Kalman smoother could easily be extended to estimate these additional source contributions; however, our ability to isolate these individual contributions from a uniform SL rise will be limited unless observations exist in the near field of the glaciers. Where such observations do exist, our methodology should be extended to incorporate the associated SL fingerprint.

This study has focused on the application of the multimodel Kalman smoother to tide gauge data; however, it is also possible to apply our statistical formalism to analyze satellite altimetry records of sea surface height variations. Although the current time span of altimetry data, ~ 20 y, is a limiting factor, the near-global coverage of the dataset can provide potentially powerful constraints on SL fingerprints. The analysis of historical SL will ultimately combine the sparse tide gauge observations with the shorter time window, but nearly global, altimetry observations.

ACKNOWLEDGMENTS. Tide gauge data were provided by the Permanent Service for Mean Sea Level (www.psmsl.org). The code for the regularized expectation-maximization technique was provided by Tapio Schneider, and the code for computing the autoregressive coefficients was provided by Alois Schlögl. We thank the modeling groups, the Program for Climate Model Diagnosis and Intercomparison, and the World Climate Research Programme's (WCRP's) Working Group on Coupled Modeling for their roles in making available the WCRP Coupled Model Intercomparison Project phase 3 (CMIP3) multimodel dataset. Support for this dataset is provided by the Office of Science, US Department of Energy. We also thank the anonymous reviewers for their helpful comments. R.E.K. was supported in part by an appointment to the US Department of Energy American Association for the Advancement of Science Fellowship Program administered by Oak Ridge Institute for Science and Education. J.X.M. acknowledges funding from Harvard University and the Canadian Institute for Advanced Research.

- Moucha R, Forte AM, Mitrovica JX (2008) Dynamic topography and long-term sea-level variations: There is no such thing as a stable continental platform. *Earth Planet Sci Lett* 271:101–108.
- Müller RD, Sdrolias M, Gaina C, Steinberger B, Heine C (2008) Long-term sea-level fluctuations driven by ocean basin dynamics. *Science* 319:1357–1362.
- Milne GA, Mitrovica JX (2008) Searching for eustasy in deglacial sea-level histories. *Quat Sci Rev* 27:2292–2302.
- Cazenave A, Nerem RS (2004) Present-day sea level change: Observations and causes. *Rev Geophys* 42:RG3001.
- Kopp ER, et al. (2010) The impact of Greenland melt on local sea levels: A partially coupled analysis of dynamic and static equilibrium effects in idealized water-hosing experiments. *Climate Change Letters* 103:619–625.
- Peltier WR, Tushingham AM (1989) Global sea level rise and the greenhouse effect: Might they be connected? *Science* 244:806–810.
- Douglas BC (1991) Global sea level rise. *J Geophys Res* 96:6981–6992.
- Douglas BC (1997) Global sea rise: A redetermination. *Surv Geophys* 18:279–292.
- Davis JL, Mitrovica JX (1996) Glacial isostatic adjustment and the anomalous tide gauge records of eastern North America. *Nature* 379:331–333.
- Holgate SJ, Woodworth PL (2004) Evidence for enhanced coastal sea level rise during the 1990s. *Geophys Res Lett*, 31:L07305:1–4.
- Church JA, White NJ, Coleman R, Lambeck K, Mitrovica JX (2004) Estimates of the regional distribution of sea level rise over the 1950–2000 period. *J Clim* 17:2609–2625.
- Cazenave A, Llovel W (2010) Contemporary sea level rise. *Annu Rev Mar Sci* 2:145–173.
- Woodward R (1888) On the form of position of mean sea level. *US Geol Surv Bull* 48: 87–170.
- Clark JA, Lingle CS (1977) Future sea level changes due to West Antarctic ice sheet fluctuations. *Nature* 269:206–209.
- Conrad CP, Hager BH (1997) Spatial variations in the rate of sea level rise caused by the present-day melting of glaciers and ice sheets. *Geophys Res Lett* 24:1503–1506.
- Mitrovica JX, Tamisiea ME, Davis JL, Milne GA (2001) Recent mass balance of polar ice sheets inferred from patterns of global sea-level change. *Nature* 409:1026–1029.
- Tamisiea MJ, Mitrovica JX, Milne GA, Davis JL (2001) Global geoid and sea level changes due to present-day ice mass fluctuations. *J Geophys Res* 106:30849–30863.
- Kalman RE (1960) A new approach to linear filtering and prediction problems. *Trans ASME J Basic Eng* 82(Series D):35–45.
- Rauch HE, Tung F, Striebel CT (1965) Maximum likelihood estimates of linear dynamic systems. *AIAA J* 3:1445–1450.
- Athans M, Chang CB (1976) *Adaptive Estimation and Parameter Identification Using Multiple Model Estimation Algorithm* (MIT Lincoln Lab, Lexington, MA), Tech Rep 1976-28.
- Sherman M (2011) *Spatial Statistics and Spatio-Temporal Data: Covariance Functions and Directional Properties* (Wiley, West Sussex, UK).
- Grewal MS, Andrews AP (2008) *Kalman Filtering: Theory and Practice Using Matlab* (Wiley, Hoboken, NJ), 3rd Ed.
- Meehl GA, et al. (2007) The WCRP CMIP3 multi-model dataset: A new era in climate change research. *Bull Am Meteorol Soc* 88:1383–1394.
- Mitrovica JX, et al. (2011) On the robustness of predictions of sea-level fingerprints. *Geophys J Int* 187:729–742.
- Kendall RA, Mitrovica JX, Milne GA (2005) On post-glacial sea level - 11. Numerical formulation and comparative results on spherically symmetric models. *Geophys J Int* 161:679–706.
- Mitrovica JX, Wahr J, Matsuyama I, Paulson A (2005) The rotational stability of an ice-age earth. *Geophys J Int* 161:491–506.
- Dziewonski AM, Anderson DL (1981) Preliminary reference Earth model. *Phys Earth Planet Inter* 25:297–356.
- Peltier WR, Tushingham AM (1991) Influence of glacial isostatic adjustment on tide gauge measurements of secular sea level. *J Geophys Res* 96:6779–6796.
- Trupin A, Wahr J (1990) Spectroscopic analysis of global tide-gauge sea level data. *Geophys J Int* 100:441–453.
- Mitrovica JX, Davis JL (1995) Present-day post-glacial sea level change far from the Late Pleistocene ice sheets: Implications for recent analyses of tide gauge records. *Geophys Res Lett* 22:2529–2532.
- Lambeck K, Smither C, Johnston P (1998) Sea-level change, glacial rebound and mantle viscosity for northern Europe. *Geophys J Int* 134(1):102–144.
- Mitrovica JX, Forte AM (2004) A new inference of mantle viscosity based upon joint inversion of convection and glacial isostatic adjustment data. *Earth Planet Sci Lett* 225:177–189.
- Peltier WR (2004) Global glacial isostasy and the surface of the ice-age earth: The ice-5G (VM2) model and GRACE. *Annu Rev Earth Planet Sci* 32:111–149.