Deconvolution and Model-Based Restoration of Clipped Ultrasonic Signals

Tomas Olofsson

***Abstract*—This paper presents a simple and general approach for deconvolving ultrasonic signals for which some of the samples have been clipped at the maximum and minimum saturation levels of the analog-to-digital converter. Furthermore, it shows how the deconvolution results can be used to restore the clipped amplitudes. By using the presented methods, the artifacts that typically arise when applying standard deconvolution methods on clipped data can be avoided. The deconvolution problem is stated as maximum a posteriori estimation of the reflection sequence under an assumption of uncorrelated Gaussian measurement noise and with the signal clipping explicitly taken into account in the signal generation model. Apart from the exact solution, two simplified approximate solutions are considered. The first approximation leads to solving a quadratic programming problem with inequality constraints and the second yields a simple closed form linear solution. A comparison under varying noise and clipping distortion conditions shows that the exact solution consistently yields the best performance, but the accuracies of both the approximative solutions are almost as good as the exact solution for low clipping distortion levels. At larger distortion levels, only the first approximative solution can compete in accuracy with the exact solution. Signal restoration results using real ultrasonic data further verify the above conclusions.**

***Index Terms*—Deconvolution, restoration, signal clipping, ultrasonic testing.**

I. Introduction

In ultrasonic testing as well as in many other measurement applications, the signals are acquired by means of analog-to-digital (A/D) converters that quantize the signals into a limited number of equispaced amplitude levels. Since the signal-to-noise ratio (SNR) partly is determined by the quantization errors, it can be improved by increasing the gain, thereby reducing the relative influence of the quantization noise. However, since the dynamic range of ultrasonic signals can often be quite high, an increase in gain comes at the cost of an increased risk of having some of the largest amplitudes in the signal clipped at the saturation levels of the A/D converter. Thus, there is typically a tradeoff between SNR and the risk of signal clipping.

To some extent, the problem can be treated by using a so-called distance amplitude correction (DAC). However, the DAC does not treat the problem of having received echoes from objects of significantly different sizes. In many applications, the defects of interest are small and may occur close to relatively large reflectors. Then, to acquire defect signals with satisfactory amplitude resolution, the gain may need to be increased so much that the signal from the larger object is clipped.

Deconvolution, which is used mainly for improving the resolution of the signals, can be particularly useful for revealing a small defect that has been masked by the echo from a larger nearby reflector. Standard deconvolution methods, such as the Wiener filter, are all designed under the assumption of a linear convolutional signal generation model. Clipping, however, acts nonlinearly on the signal, and due to this mismatch between the model and the reality, the results obtained by standard deconvolution algorithms typically contain artifacts when the signals have been partly clipped. These artifacts may sometimes completely dominate the time interval in which the weaker echo occurs, thus rendering useless deconvolution results.

The purpose of this paper is twofold. First, it presents an approach for solving the problem of the deconvolution of clipped signals. The approach is based on maximum *a posteriori* (MAP) estimation of the so-called reflection sequence under a signal generation model that explicitly takes the signal clipping into account. The optimization criterion associated with the MAP estimation problem unfortunately does not allow for the use of any simple and fast optimization algorithm. Therefore, suitable approximations of this criterion are also given, and the validity of the associated simplified deconvolution algorithms is examined experimentally. Second, this paper examines the possibility of restoring the clipped amplitudes by a direct use of the deconvolution results, leading to a model-based approach for the restoration of clipped signals.

The two basic problems mentioned above, deconvolution and restoration of nonlinearly distorted signals, have been treated earlier in the literature but have only been briefly mentioned in the literature on ultrasonic testing [1]. In the context of digital image restoration, deconvolution of images that have been subject to a nonlinear pointwise distortion, which in that case is due to the saturation curve of the photographic film, has been treated in [2] and [3]. There, the problem was formulated as MAP estimation of a signal which is observed as a noise corrupted, pointwise nonlinearly distorted convolution. Since the models used in [2] and [3] assumed that the noise was superimposed *after* the nonlinearity, these models are not optimally suited for treating clipped signals.

The deconvolution of clipped signals can also be treated by first approximately restoring the clipped samples and then performing deconvolution using any standard method. This approach was suggested, e.g., in [1]. The basic idea that is common to a large number of signal restoration algorithms is to oversample the signal, thereby allowing one to treat the signal as being bandlimited, and then to find a signal that matches

Manuscript received February 1, 2003; revised November 1, 2004. This work was supported by the SPIQNAR project under Contract FIKS-CT-2000-00065.

The author is with the Signals and Systems Group, Department of Material Science, Uppsala University, 751 20 Uppsala, Sweden.

Digital Object Identifier 10.1109/TIM.2005.847222

either exactly or approximately to the undistorted data while at the same time fulfilling the bandlimitation restriction. It should be noted that the problem usually leads to solving equation systems that are ill-conditioned so when noise is present in the signal, solutions obtained without any other regularization than the one that is imposed by the bandlimitation constraints may sometimes be very poor.

The regularization is typically performed by minimizing the energy in the restored signal [4], [5]. For the particular problem of restoring clipped amplitudes, constraints that enforce the solution to exceed the upper and lower clipping limits at the times of interest can also be added to further improve the results [6].

Although this combination of such so-called bandlimited restoration and deconvolution will in general improve the deconvolution results, it has the disadvantage of not fully utilizing all available information in all steps of the processing. In the restoration, only the information that the system's transfer function is bandlimited is used, whereas in the deconvolution, the full knowledge about the transfer function is explicitly utilized. If such information is available, it should be used throughout all calculations. The method presented in this paper has this property.

This paper is organized as follows. In Section II, the problem formulation is presented, the MAP criterion is derived, and the simplifying approximations and their associated algorithms are presented. In Section III, the restoration method is given. The exact and the approximative MAP deconvolution results are compared quantatively using simulated data in Section IV. In this section, the restoration method is also illustrated using real data. Finally, conclusions and a discussion are given in Section V.

II. MAP Deconvolution of Clipped Signals Assuming Gaussian Noise Sources

In the following, we consider time-discrete signals that are sampled with a period T. The integer k will denote the time index which is related to the continuous time t through the relation $t = kT$. We assume that the acquired ultrasonic signal $x(k)$ is a clipped version of a partially unobserved signal $\tilde{x}(k)$ that, in turn, can be modeled accurately as a noise-corrupted convolution between the so-called reflection sequence $r(k)$ and an impulse response $h(k)$ that captures the relevant electroacoustical and acoustoelectrical properties of the transducer. We write this as

$$x(k) = c\left(\tilde{x}(k)\right) = c\left(h(k) * r(k) + e(k)\right) \tag{1}$$

where $e(k)$ is the noise and $c(\cdot)$ denotes the clipping nonlinearity

$$c(x) = \begin{cases} \overline{c} & \text{if } x \geq \overline{c} \\ x & \text{if } \underline{c} < x < \overline{c} \\ \underline{c} & \text{if } x \leq \underline{c}. \end{cases} \tag{2}$$

$\overline{c}$ and $\underline{c}$ are the upper and lower clipping limits, respectively. For example, an ordinary 8-bit A/D converter that represents the amplitudes as signed integers has $\overline{c} = 127$ and $\underline{c} = -128$.

Let N and M denote the number of samples in $x(k)$ and $r(k)$, respectively. By using vector notation with $\mathbf{r} = (r(1), \ldots, r(M))'$, $\mathbf{x} = (x(1), \ldots, x(N))'$, $\tilde{\mathbf{x}} = (\tilde{x}(1), \ldots, \tilde{x}(N))'$, and $\mathbf{e} = (e(1), \ldots, e(N))'$, the model in (1) can be written as

$$\mathbf{x} = \mathbf{c}(\tilde{\mathbf{x}}) = \mathbf{c}(\mathbf{H}\mathbf{r} + \mathbf{e}) \tag{3}$$

where vector $\mathbf{c}(\cdot)$ is defined as $\mathbf{c}(\tilde{\mathbf{x}}) = (c(\tilde{x}(1), \ldots, c(\tilde{x}(N))'$ and $\mathbf{H}$ is an $N \times M$ finite impulse response matrix with elements $(\mathbf{H})_{p,q} = h(p - q + N - M)$. For future use, let $\mathbf{h}_i$ denote the ith row vector in $\mathbf{H}$.

Let $\overline{\mathcal{C}}$ and $\underline{\mathcal{C}}$ denote the sets of vector indices for which the signal is clipped at the upper and lower limits, respectively, and let $\mathcal{L}$ denote the set of indices for the remaining nonclipped components. Furthermore, let us from the (undistorted) vector $\tilde{\mathbf{x}}$ construct three vectors $\tilde{\mathbf{x}}_{\mathcal{L}}$, $\tilde{\mathbf{x}}_{\overline{\mathcal{C}}}$, and $\tilde{\mathbf{x}}_{\overline{\mathcal{C}}}$ that contain only the components in $\tilde{\mathbf{x}}$ with indices in the sets given by the respective subscripts. The composite vector $(\tilde{\mathbf{x}}'_{\mathcal{L}}, \tilde{\mathbf{x}}'_{\overline{\mathcal{C}}}, \tilde{\mathbf{x}}'_{\underline{\mathcal{C}}})'$ can then be written as

$$\begin{pmatrix} \tilde{\mathbf{x}}_{\mathcal{L}} \\ \tilde{\mathbf{x}}_{\overline{\mathcal{C}}} \\ \tilde{\mathbf{x}}_{\underline{\mathcal{C}}} \end{pmatrix} = \begin{pmatrix} \mathbf{H}_{\mathcal{L}} \\ \mathbf{H}_{\overline{\mathcal{C}}} \\ \mathbf{H}_{\underline{\mathcal{C}}} \end{pmatrix} \mathbf{r} + \mathbf{e} \tag{4}$$

where $\mathbf{H}_{\mathcal{L}}$, $\mathbf{H}_{\overline{\mathcal{C}}}$, and $\mathbf{H}_{\underline{\mathcal{C}}}$ are matrices that consist of the rows $\mathbf{h}_i$ with i belonging to the disjoint sets $\mathcal{L}$, $\overline{\mathcal{C}}$, and $\underline{\mathcal{C}}$, respectively. Note that the number of rows in the three matrices is in general different.

In the following, the noise components in $\mathbf{e}$ are assumed to be independent and identically zero-mean Gaussian distributed with variance σ^2, i.e., $\mathbf{e} \sim N(0, \sigma^2\mathbf{I})$, where $\mathbf{I}$ is the identity matrix. The above variance can often be readily estimated from a separate signal in which there are no responses from scatterers present. For instance, in immersion testing, the noise variance can be directly estimated from a signal in the time interval corresponding to the water path. Furthermore, the reflection sequence vector is assumed to be zero-mean Gaussian $\mathbf{r} \sim N(0, \mathbf{C}_r)$, where $\mathbf{C}_r$ is the covariance matrix of $\mathbf{r}$. This distribution should reflect our incomplete prior knowledge about the true reflection sequence. The better informed we are, the better we should be able to adopt, for instance, the covariance matrix to a given application. As a start, we can make the relatively noncommital assumption that the components in $\mathbf{r}$ are independent and identically distributed, leading to $\mathbf{C}_r = \sigma_r^2\mathbf{I}$, where σ_r^2 is a common scalar variance for all components. A suitable value for σ_r^2 can be obtained by calculating the average of the squared amplitudes of a set of representative signals and dividing this with the squared norm of impulse response vector $\mathbf{h}$. In this way, we will match the expected signal energy to the average observed signal energy for this set of signals.

The MAP estimate of $\mathbf{r}$ is defined as

$$\hat{\mathbf{r}} = \arg\max_{\mathbf{r}} p(\mathbf{r} \mid \mathbf{x}, \mathbf{H}) = \arg\max_{\mathbf{r}} p(\mathbf{x} \mid \mathbf{r}, \mathbf{H}) p(\mathbf{r}) \tag{5}$$

where the likelihood $p(\mathbf{x} \mid \mathbf{r}, \mathbf{H})$ can be factorized as

$$p(\mathbf{x} \mid \mathbf{r}, \mathbf{H}) = p(\mathbf{x}_{\mathcal{L}} \mid \mathbf{r}, \mathbf{H}_{\mathcal{L}}) P(\mathbf{x}_{\overline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\overline{\mathcal{C}}}) P(\mathbf{x}_{\underline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\underline{\mathcal{C}}}). \tag{6}$$

The probabilities $P(\mathbf{x}_{\overline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\overline{\mathcal{C}}})$ and $P(\mathbf{x}_{\underline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\underline{\mathcal{C}}})$ should be interpreted as

$$P(\mathbf{x}_{\overline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\overline{\mathcal{C}}}) = P(\mathbf{h}_i\mathbf{r} + e(i) \geq \overline{c} \text{ for } i \in \overline{\mathcal{C}}) \tag{7}$$

and

$$P(\mathbf{x}_{\underline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\underline{\mathcal{C}}}) = P(\mathbf{h}_i\mathbf{r} + e(i) \le \underline{c} \text{ for } i \in \underline{\mathcal{C}}) \tag{8}$$

respectively.

Let us define $\overline{z_i}$ and $\underline{z_i}$ as

$$\overline{z_i} = \mathbf{h}_i\mathbf{r} - \overline{c} \text{ and } \underline{z_i} = \underline{c} - \mathbf{h}_i\mathbf{r}. \tag{9}$$

By using that the components in $\mathbf{e}$ are independent, we can write $P(\mathbf{x}_{\overline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\overline{\mathcal{C}}})$ and $P(\mathbf{x}_{\underline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\underline{\mathcal{C}}})$ as

$$\begin{aligned} P(\mathbf{x}_{\overline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\overline{\mathcal{C}}}) &= \prod_{i\in\overline{\mathcal{C}}} \Phi\left(\frac{\overline{z_i}}{\sigma}\right) \\ P(\mathbf{x}_{\underline{\mathcal{C}}} \mid \mathbf{r}, \mathbf{H}_{\underline{\mathcal{C}}}) &= \prod_{i\in\underline{\mathcal{C}}} \Phi\left(\frac{\underline{z_i}}{\sigma}\right) \end{aligned} \tag{10}$$

where $\Phi(\cdot)$ denotes the normal cumulative distribution function.

We take as our criterion the negative logarithm of the expression in (5), and by using (10) and using that $\mathbf{r}$ and $\mathbf{e}$ are normal distributed, we finally obtain

$$\begin{aligned} J(\mathbf{r}) = \frac{1}{2\sigma^2}\|\mathbf{x}_{\mathcal{L}} - \mathbf{H}_{\mathcal{L}}\mathbf{r}\|^2 + \frac{1}{2}\mathbf{r}'\mathbf{C}_r^{-1}\mathbf{r} \\ - \sum_{i\in\overline{\mathcal{C}}} \ln\Phi\left(\frac{\overline{z_i}}{\sigma}\right) - \sum_{i\in\underline{\mathcal{C}}} \ln\Phi\left(\frac{\underline{z_i}}{\sigma}\right) \end{aligned} \tag{11}$$

where all terms that are constant with respect to $\mathbf{r}$ have been removed.

It is important to note that the above expression will retain its simple structure also when the measurement noise and reflection sequences are non-Gaussian. For instance, if the measurement noise is double sided exponential, the first term should be replaced by a sum of absolute values of the residuals and the last two sums in $J(\mathbf{r})$ should involve the cumulative distribution function of the double sided exponential density instead of the normal density.

No closed-form solution that minimizes $J(\mathbf{r})$ exists and iterative methods must be used for the optimization. Unfortunately, this leads to a relatively slow data processing which limits the practical use of the method. However, an algorithm that minimizes the criterion in (11) has a value in that it provides a reference against which other methods can be compared.

More practical methods can be developed by making approximations of $J(\mathbf{r})$ that are better suited for efficient optimization schemes. Of course, many such approximations are possible and in this paper we restrict our attention to two that are particularly simple.

A. Approximation I

For decreasing values of σ, the slope of $\Phi(z/\sigma)$ increases around $z = 0$ and in the limit $\Phi(z/\sigma)$ becomes a step function which yields

$$\lim_{\sigma\to 0} \ln\Phi\left(\frac{z}{\sigma}\right) = \begin{cases} 0 & z \ge 0 \\ -\infty & z < 0. \end{cases} \tag{12}$$

Thus, for small σ, we should find a good approximation of the minimizing solution to $J(\mathbf{r})$ by solving the optimization problem

$$\hat{\mathbf{r}} = \arg\min_{\mathbf{r}} \frac{1}{2\sigma^2}\|\mathbf{x}_{\mathcal{L}} - \mathbf{H}_{\mathcal{L}}\mathbf{r}\|^2 + \frac{1}{2}\mathbf{r}'\mathbf{C}_r^{-1}\mathbf{r} \tag{13}$$

subject to the constraints

$$\mathbf{h}_i\mathbf{r} > \overline{c} \text{ for } i \in \overline{\mathcal{C}} \text{ and } \mathbf{h}_i\mathbf{r} < -\underline{c} \text{ for } i \in \underline{\mathcal{C}}. \tag{14}$$

This quadratic optimization problem with inequality constraints can be solved using a number of standard numerical packages—see, e.g., [7]. At each iteration in the optimization algorithm, a system of linear equations is solved. For a signal with only a few clipped samples, the number of unknowns in this equation system is approximately M so each iteration consumes on the order of M^3 calculations. However, the number of iterations is difficult to predict. This number depends not just on the number of constraints but also on how the constraints are related. Thus, the overall computation time is difficult to predict.

B. Approximation II

Although the above approximation simplified the problem somewhat, finding $\hat{\mathbf{r}}$ in (13) still requires iterative algorithms for which the calculation time can be difficult to predict. We can, however, obtain a closed-form solution for which the calculation time is predictable by simply ignoring the clipped components completely, i.e., we treat these samples as *missing* data instead of *censored* data. One interpretation of the solution obtained in this way is that we implicitly assume σ to be large. Then the approximation $\Phi(z/\sigma) \approx \Phi(0) = 1/2$ can be used. Now the terms associated with the likelihood of the clipped samples do not influence the solution and the simplified criterion becomes

$$\hat{\mathbf{r}} = \arg\min_{\mathbf{r}} \frac{1}{2\sigma^2}\|\mathbf{x}_{\mathcal{L}} - \mathbf{H}_{\mathcal{L}}\mathbf{r}\|^2 + \frac{1}{2}\mathbf{r}'\mathbf{C}_r^{-1}\mathbf{r}. \tag{15}$$

The minimizing solution is given in closed form as

$$\hat{\mathbf{r}} = \left(\frac{1}{\sigma^2}\mathbf{H}_{\mathcal{L}}'\mathbf{H}_{\mathcal{L}} + \mathbf{C}_r^{-1}\right)^{-1} \mathbf{H}_{\mathcal{L}}'\frac{1}{\sigma^2}\mathbf{x}_{\mathcal{L}}. \tag{16}$$

In practice the solution is obtained by solving a corresponding system of M linear equation so the number of calculations for this solution is on the order of M^3. Since the solution is obtained after one iteration, this will consistently yield an overall computational time that is less than that for the above Approximation II solution.

Of course, both the Approximation I and II solutions treat the available data somewhat inconsistently; The noise variance is only approximated as zero or infinity at the clipped samples, not at the samples in $\mathbf{x}_{\mathcal{L}}$. We should therefore expect the results to be less accurate than the exact MAP solution. Furthermore, since the solution in (16) makes use of less data than the two other solutions, (16) should be expected to have the worst performance. Thus, this solution is more motivated by its practical simplicity and less by optimality considerations.

III. Restoration of the Clipped Samples

The MAP estimation of the clipped sample is somewhat more complicated than the MAP estimation of $\mathbf{r}$. So instead of trying to find strict MAP estimates of these samples, we use a simpler approach. It is clear that if we for some reason should know $\mathbf{r}$, then the optimal restoration of the clipped samples would be trivial, since then $\hat{\mathbf{x}}_{\overline{\mathcal{C}}} = \mathbf{H}_{\overline{\mathcal{C}}}\mathbf{r}$ and $\hat{\mathbf{x}}_{\underline{\mathcal{C}}} = \mathbf{H}_{\underline{\mathcal{C}}}\mathbf{r}$ must be the best possible approximations of $\tilde{\mathbf{x}}_{\overline{\mathcal{C}}} = \mathbf{H}_{\overline{\mathcal{C}}}\mathbf{r} + \mathbf{e}$ and $\tilde{\mathbf{x}}_{\underline{\mathcal{C}}} = \mathbf{H}_{\underline{\mathcal{C}}}\mathbf{r} + \mathbf{e}$. This motivates the straightforward use of the linear estimates

$$\hat{\mathbf{x}}_{\overline{\mathcal{C}}} = \mathbf{H}_{\overline{\mathcal{C}}}\hat{\mathbf{r}} \quad \text{and} \quad \hat{\mathbf{x}}_{\underline{\mathcal{C}}} = \mathbf{H}_{\underline{\mathcal{C}}}\hat{\mathbf{r}}. \tag{17}$$

Note that it is in general not true that the above simple estimates are the correct MAP estimates of $\tilde{\mathbf{x}}_{\overline{\mathcal{C}}}$ and $\tilde{\mathbf{x}}_{\underline{\mathcal{C}}}$.

IV. Experiments

The purpose of the experiments presented in this section was twofold: 1) to quantitatively compare the performance of the different deconvolution algorithms under varying levels of noise and various levels of clipping distortion and 2) to qualitatively verify that the clipped amplitudes could be satisfactorily restored using the technique in Section III.

For real ultrasonic signals, we do not have direct access to the true reflection sequence $r(k)$, which makes a quantitative comparison practically impossible. Therefore we used simulated signals for those comparisons. For the illustration of the restoration, real ultrasonic data were used.

The exact MAP estimate defined as the minimizer to (11) was found by means of a quasi-Newton algorithm and the Approximation I solution defined by (13) and (14) was found using a standard solver for for quadratic programming with inequality constraints. The Matlab optimization toolbox implementations were used to obtain both the above estimates. The solution corresponding to Approximation II was calculated using (16).

For comparison, two Wiener filter solutions were also considered. The first, which we in the following refer to as Wiener I, was calculated mainly for the purpose of illustrating the model error artifacts mentioned in the introduction. It was obtained by means of (16) but with replacing $\mathbf{x}_{\mathcal{L}}$ and $\mathbf{H}_{\mathcal{L}}$ by $\mathbf{x}$ and $\mathbf{H}$. The second, refered to as Wiener II, was used as a reference providing a lower bound on the estimation error. In this solution, $\mathbf{x}_{\mathcal{L}}$ and $\mathbf{H}_{\mathcal{L}}$ were replaced, respectively, by $\tilde{\mathbf{x}}$ and $\mathbf{H}$ in (16). Note that the provided lower bound will be somewhat optimistic since this solution optimally utilized also the information in the clipped amplitudes.

A. Comparison Between Exact and Approximative MAP Deconvolution Algorithms

In the simulations, both $e(k)$ and $r(k)$ were generated as zero-mean white Gaussian sequences with scalar variances σ^2 and σ_r^2, respectively, and thus we had the corresponding covariance matrix $\mathbf{C}_r = \sigma_r^2\mathbf{I}$. To obtain realistic simulation models, the two impulse responses used, $h_1(k)$ and $h_2(k)$, were measured from two real ultrasonic transducers. The impulse responses are shown in Fig. 1.

The simulated signals were clipped at levels $\overline{c} = c$ and $\underline{c} = -c$ for varying c. These c-levels were chosen so that a specified fraction p of the signal samples was clipped on average. The levels were determined by, for different p, solving the equation $p = 1 - 2\Phi(\|\mathbf{h}\|^2/(c\sigma_r))$ for c.

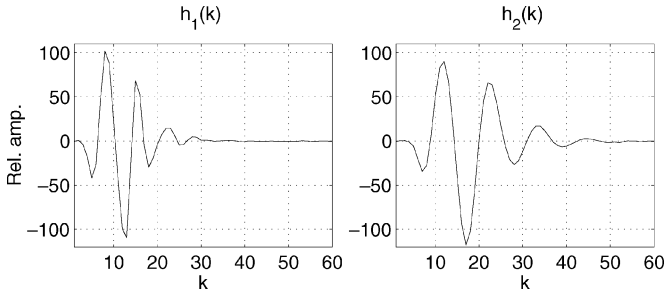

Fig. 1. The impulse responses $h_1(k)$ and $h_2(k)$ used in the simulations. They were acquired from two different ultrasonic transducers as the immersion far-field pulse-echo responses from planar reflectors parallel to the transducer surfaces. Note that the transducer with impulse response $h_2(k)$ was used in the final real data experiments.

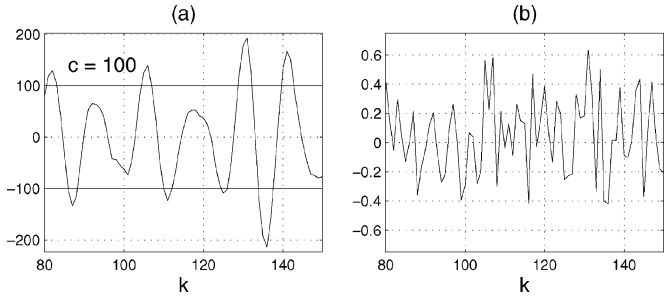

Fig. 2. (a) Part of a signal used in the simulations obtained using $h_2(k)$ as the impulse response. The signal is clipped at the levels indicated by the solid lines. (b) The corresponding reflection sequence $r(k)$.

One example of the signals used in the simulation is shown in Fig. 2 together with its corresponding simulated reflection sequence. The solid lines indicate the upper and lower clipping levels.

The result from deconvolving the signal shown in Fig. 2 is presented in Fig. 3. Note in particular the poor results obtained from an uncritical use of the Wiener filter (Wiener I). Large oscillations occur in this solution due to the clipping. Note that the errors can be reduced by increasing the noise variance parameter σ^2, thereby making the solution more robust to the model errors. However, such a solution would still make no difference between the clipped and the nonclipped samples, and this must lead to a nonoptimal solution. Note also the close resemblance between the exact and the two approximative MAP solutions.

The quantitative comparison was performed by means of Monte Carlo simulations. During the simulations, the reflection sequences were of length $M = 300$, having a variance fixed at $\sigma_r^2 = 0.1$, and the noise variance and the fraction of clipped samples were varied in the ranges $\sigma^2 \in [1{,}1000]$ and $p \in [0.01, 0.50]$, respectively. For each combination of σ^2 and p, $N_{MC} = 50$ signals were generated, each having a length of $N = 300$ samples.

The performance measure used was the normalized mean squared estimation error, i.e.,

$$\epsilon = \frac{E\left[\left(r(k) - \hat{r}(k)\right)^2\right]}{\sigma_r^2}. \tag{18}$$

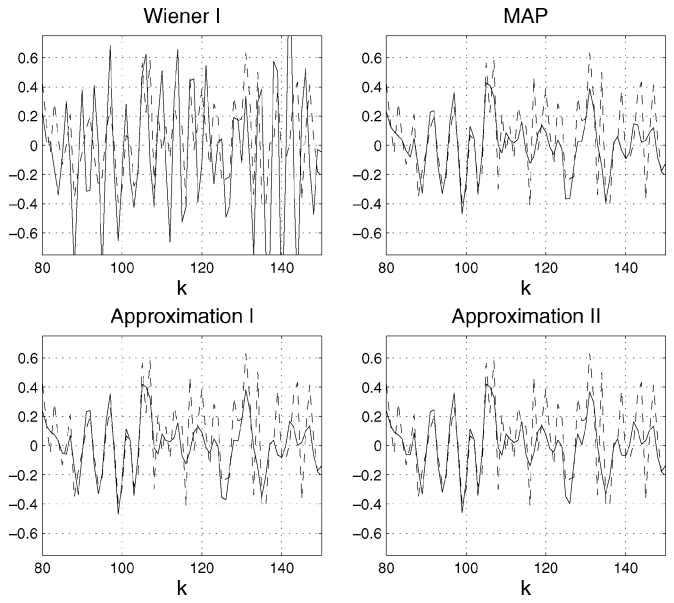

Fig. 3. Results obtained from the Wiener I solution and the exact MAP and two approximate MAP solutions on the signal in Fig. 2. For comparison, the correct reflection sequence $r(k)$ is shown dashed in each plot.

Note that for the trivial estimate $\hat{r}(k) = 0$, we have $\epsilon = 1$. Thus, $\epsilon = 1$ should be interpreted as the worst possible performance for any sensible approach. In the simulations ϵ was straighforwardly estimated as the average

$$\hat{\epsilon} = \frac{1}{N_{MC}} \sum_{n=1}^{N_{MC}} \frac{1}{N-L} \sum_{k=L/2+1}^{N-L/2} \frac{(r(k) - \hat{r}(k))^2}{\sigma_r^2} \qquad (19)$$

where the time interval $[L/2+1, N-L/2]$ was chosen to avoid nonstationary effects at the beginning and the end of the signals. In these simulations, we had $L = 50$. In total, the calculation of $\hat{\epsilon}$ involved an average over $N_{MC}(N-L) = 12\,500$ samples for each combination of σ^2 and p.

The results obtained using impulse response $h_1(k)$ are presented in Fig. 4. The error ϵ and the variance σ^2 are both given in dB. The corresponding results for $h_2(k)$ are shown in Fig. 5.

Both Figs. 4 and 5 illustrate once again the poor performance of the Wiener I solution. Note that the performance for this estimator has an optimum for an intermediate noise level. At this noise level there is a balance between the effects of model errors and the errors due to the measurement noise.

For small levels of clipping distortion, the different MAP versions all behave similarily, with the exact MAP solution having slightly better performance followed, in turn, by the Approximation I and Approximation II solutions. At larger distortion levels, the difference in performance becomes more prominent. Since the relative importance of the information contained in the clipped samples then becomes larger, the optimal treatment of the information in these samples becomes more vital.

Note that the performance curve of the Approximation I solution is in direct accordance with the assumptions made in the approximation. For small σ^2, i.e., when following the assumptions, the difference in performance between the exact MAP and the Approximation I solutions is very small but this difference increases as the noise level increases.

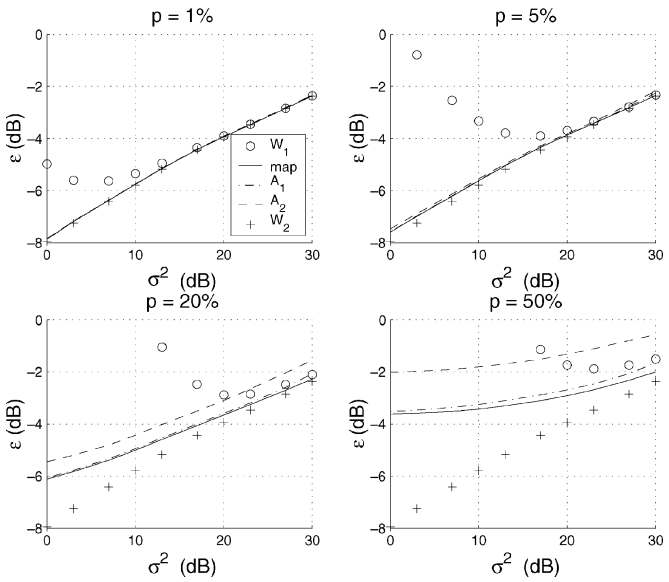

Fig. 4. Performance of the different methods at different noise levels and at different clipping levels. Here the impulse response $h_1(k)$ was used. The normalized estimation error was obtained using (19) with $N_{MC} = 50$, $N = 300$, and $L = 50$.

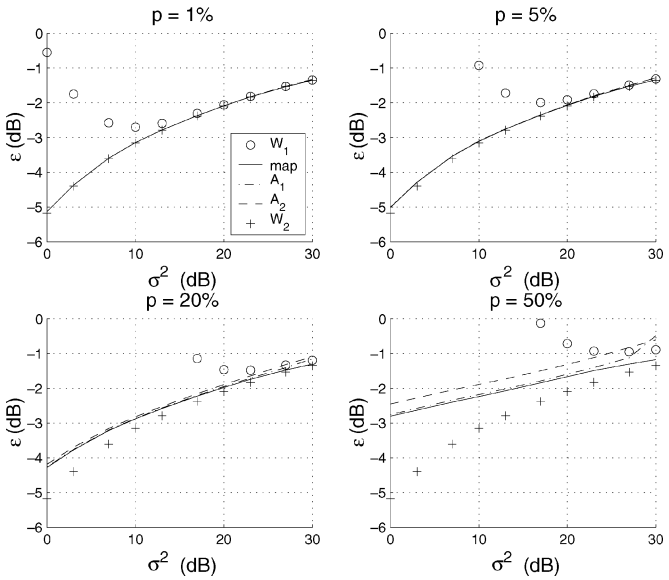

Fig. 5. Performance of the different methods at different noise levels and at different clipping levels. Here the impulse response $h_2(k)$ was used.

Finally, by comparing the difference between the Wiener II solution performance and the MAP solution peformance for the simulations using $h_1(k)$ and $h_2(k)$, we see that the difference is smaller for $h_2(k)$, particularly at large clipping distortion levels. This is because $h_2(k)$ represents a more narrow-band system so the corresponding received signal is more redundant. In terms of information removal, the clipping distortion is of less relative importance than for the more wide-band signal.

B. Signal Restoration

To qualitatively illustrate the signal restoration performance, real ultrasonic signals were obtained using the transducer having the impulse response $h_2(k)$ shown in Fig. 1. The signals were clipped at two levels $c = 25$ and $c = 15$, as shown in Figs. 6 and

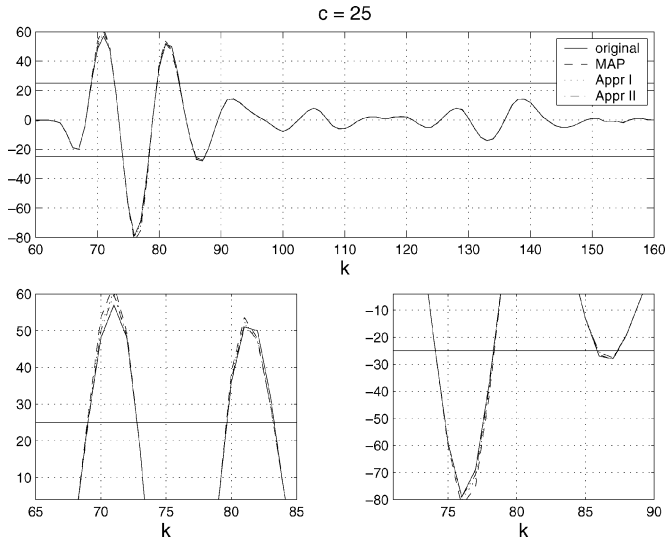

Fig. 6. Example of results from restoring of a real signal clipped at level $c = 25$.

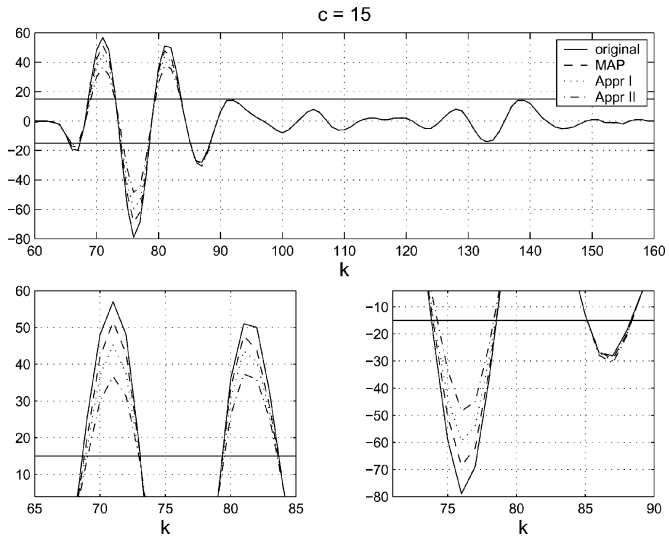

Fig. 7. Example of results from restoring of a real signal clipped at level $c = 15$.

7. The noise and reflection sequence variances were set to $\sigma^2 = 2$ and $\sigma_r^2 = 0.1$, respectively, where the former was estimated from a signal-free part of the signal and the latter was set to approximately match to the energy in the signal.

The signals were deconvolved using the exact MAP and the two approximative MAP solutions and the clipped samples were restored by applying (17) as described in Section III. The restoration results are presented along with signals in the respective figures, and the magnified plots show the intervals of interest.

The signal restoration results further verify the conclusions made from the quantitative experiments presented earlier. Only for a relatively large clipping distortion, the choice of method seems to be critical. For $c = 25$, all solutions yield similar results, but for a larger distortion ($c = 15$), the difference between the methods becomes more significant. Again, the exact MAP solution yields the best results, followed in turn by Approximation I and II.

V. Conclusion and Discussion

This paper has presented an approach for deconvolving and restoring ultrasonic signals that have been clipped at the saturation levels of the A/D converter, thus avoiding artifacts that are generally caused when applying classical deconvolution algorithms to such data. An exact MAP solution and two corresponding approximative solutions have been given for the deconvolution problem, and the performance of the solutions has been evaluated using simulated data.

The results show that the exact solution yields the best performance in general, but that the difference between the methods is relatively small at low distortion levels. Thus, for small distortion levels, the Approximation II solution is recommended because of its computational simplicity. At larger distortion levels the computationally more intensive, but more accurate, Approximation I solution should be considered. The signal restoration results obtained with real data further verified the above conclusions.

Note that although ultrasonic testing has been the application considered in this paper, the same method will be applicable whenever the signal model in (1) is valid. Furthermore, the approach can easily be generalized in a number of ways. For instance, assuming non-Gaussian prior distributions for the reflection sequence will only require the modification of the second term found in the MAP criterion in (11). This means, for instance, that the so-called sparse deconvolution problem can also be straightforwardly treated using the approach presented in the paper. Furthermore, non-Gaussian measurement noise can be treated by modifying the first and the last terms. Note in particular that, as long as the samples $e(k)$ are independent with unimodal probability density functions peaking at zero, the two approximations given in this paper will still be valid.

References

[1] T. J. M. Jeurens, J. C. Somer, F. A. M. Smeets, and A. P. G. Hoeks, "The practical significance of two-dimensional deconvolution in echography," *Ultrason. Imag.*, vol. 9, no. 2, pp. 106–116, Apr. 1987.

[2] B. R. Hunt, "Bayesian methods in nonlinear digital image restoration," *IEEE Trans. Comput.*, vol. C-26, pp. 219–229, Mar. 1977.

[3] H. J. Trussel and B. R. Hunt, "Improved methods of maximum a posteriori restoration," *IEEE Trans. Comput.*, vol. C-27, pp. 57–62, Jan. 1979.

[4] A. Papoulis, "A new algorithm in spectral analysis and band-limited extrapolation," *IEEE Trans. Circuits Syst.*, vol. CAS-22, pp. 735–742, Sep. 1975.

[5] A. K. Jain and S. Ranganath, "Extrapolation algorithms for discrete signals with application in spectral estimation," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, pp. 830–845, Aug. 1981.

[6] J. S. Abel and J. O. Smith, "Restoring a clipped signal," in *Proc. ICASSP 1991*, vol. 3, 1991, pp. 1745–1748.

[7] R. Fletcher, *Practical Methods of Optimization*. New York: Wiley, 1987.

Tomas Olofsson was born in Sandviken, Sweden, on May 2, 1968. He received the Bachelor's degree in engineering physics and the Ph.D degree in signal processing from Uppsala University, Sweden, in 1994 and 2000, respectively.

His main scientific interests are inference problems in ultrasonic nondestructive testing, in particular deconvolution of signals and images.