

Photoacoustic tomography with a single detector in a reverberant cavity

B. T. Cox^{a)} and P. C. Beard

Department of Medical Physics and Bioengineering, University College London, Gower Street, London WC1E 6BT, United Kingdom

(Received 27 August 2008; revised 3 December 2008; accepted 11 December 2008)

In conventional biomedical photoacoustic tomography (PAT), ultrasonic pulses generated through the absorption of nanosecond pulses of near-infrared light are recorded over an array of detectors and used to recover an image of the initial acoustic pressure distribution within soft tissue. This image is related to the tissue optical coefficients and therefore carries information about the tissue physiology. For high resolution imaging, a large-area detector array with a high density of small, sensitive elements is required. Such arrays can be expensive, so reverberant-field PAT has been suggested as a means of obtaining PAT images using arrays with a smaller number of detectors. By recording the reflections from an acoustically reverberant cavity surrounding the sample, in addition to the primary acoustic pulse, sufficient information may be captured to allow an image to be reconstructed without the need for a large-area array. An initial study using two-dimensional simulations was performed to assess the feasibility of using a *single* detector for PAT. It is shown that reverberant-field data recorded at a single detector are sufficient to reconstruct the initial pressure distribution accurately, so long as the shape of the reverberant cavity makes it ray-chaotic. The practicalities of such an approach to photoacoustic imaging are discussed.

© 2009 Acoustical Society of America. [DOI: 10.1121/1.3068445]

PACS number(s): 43.35.Ud, 43.20.Ks, 43.60.Pt, 43.28.We [TDM]

Pages: 1426–1436

I. INTRODUCTION

In the past decade or so, the potential of photoacoustic tomography¹ (PAT) to become a useful biomedical imaging tool has been clearly demonstrated.^{2–6} For imaging soft-tissue to depths where the effect of optical scattering becomes too great for “ballistic” optical imaging modalities, e.g., optical microscopy or optical coherence tomography, PAT provides a means of obtaining an image based on optical contrast but still with good resolution (typically $<100\ \mu\text{m}$ at depths of 5–10 mm). It is a hybrid technique in the sense that *optical* pulses are used to generate *ultrasonic* pulses within the tissue. The absorption of short (nanosecond) laser pulses by tissue chromophores results in a small but rapid increase in the local pressure, which then propagates as a high frequency ultrasonic pulse (tens of megahertz). These ultrasonic pulses are recorded by a detector array and, via a numerical reconstruction algorithm, used to form an image.

The spatially varying pressure rise caused when the light energy is absorbed is called the “initial pressure distribution.” It is this that a PAT image represents, and it is proportional to the optical energy per unit volume absorbed by the tissue. As this is closely related to the optical absorption coefficient, PAT images are related to the distribution of the chromophores contained in the tissue and carry information about the tissue structure and function. Because of this reliance on optical absorption, PAT has the potential to be used spectroscopically, so could be used in applications such as

molecular imaging, in which an image of the distribution of endogenous or exogenous chromophores with known optical absorption spectra is obtained.^{7–10}

Current experimental PAT systems use arrays of detectors to collect data over a sufficiently large measurement surface to form an image. This becomes expensive as the demand for better resolution, and therefore larger arrays, increases. In addition, for a real time imaging system, each of the detector array elements requires dedicated electronics, further increasing the cost. For targets that can be enclosed, or partially enclosed, by an acoustically reverberant cavity, it has been shown that by using the information stored in the reverberant field, the size of the measurement surface (and therefore the number of detection points) can be reduced without losing image quality.¹¹ This paper suggests taking this approach one step further: using a single point measurement of the reverberation of a photoacoustically generated signal within an acoustically reverberant cavity to form a PAT image. This is applicable only to targets that can be enclosed by the reverberant cavity, but avoids both the use of a large and expensive array and point-by-point scanning. Two-dimensional numerical examples are used to demonstrate how good quality PAT images could be obtained from measurements made using just a single ultrasound detector in a chaotic cavity, and the advantages, practicalities, and limitations of such a scheme are discussed.

II. PHOTOACOUSTIC IMAGE RECONSTRUCTION

Consider a short laser pulse incident on some optically heterogeneous biological tissue (such as skin). A region of slightly increased pressure, $p_0(\mathbf{x})$, is generated locally wher-

^{a)}Electronic mail: bencox@mpb.ucl.ac.uk

ever the optical energy is absorbed (such as in a blood vessel). When the laser pulse duration is much shorter than the acoustic travel time across a characteristic distance of p_0 (such as the vessel diameter), the subsequent acoustic propagation can be modeled as an initial value problem for the wave equation with the initial conditions $p(\mathbf{x})|_{t=0}=p_0(\mathbf{x})$ and $dp/dt|_{t=0}=0$. In conventional PAT, an image of $p_0(\mathbf{x})$ is reconstructed from measurements of the acoustic pressure waves $p(\mathbf{x}_s, t)$ made by detectors placed at points \mathbf{x}_s on a measurement surface S enclosing V , where $\text{supp}(p_0) \in V$. (Usually the detectors are assumed to be much smaller than the shortest acoustic wavelength in the pulse, and omnidirectional, which is never quite the case in practice.)

In the past decade or so, numerous algorithms to recover $p_0(\mathbf{x})$ from $p(\mathbf{x}_s, t)$ have been proposed. For recent reviews see Refs. 5 and 12. The algorithms differ in the measurement geometries assumed (e.g., spherical, cylindrical, planar, and arbitrary), in the types of solution (e.g., closed form, infinite series, and numerical), and in their computational efficiency. However, all the algorithms make the simplifying assumption that the measurement surface is acoustically transparent and does not affect the free-field propagation of the photoacoustically generated waves. (Experimentally this assumption has been accommodated by truncating the measurements before any reflections, e.g., from other parts of the measurement apparatus, are received, i.e., by time-gating out any reflections.) The image reconstruction described in this paper is quite different in the sense that the free-field assumption is not made, and the reflected or reverberant field is included in the formulation.

A. Limited-view problem

Experimental imaging studies have been carried out using circular, cylindrical, spherical, and planar measurement surfaces. With a spherical measurement surface enclosing V it is possible to record the photoacoustic waves traveling away from the source in all directions. With a planar measurement surface at most half of the waves can be recorded, as a plane cannot surround the source region. Nevertheless, if the plane is infinite in extent this is still sufficient information to recover the source exactly, as shown by Anastasio *et al.*¹³ In practice, an actual planar measurement surface will have a finite size, and not even half of the emitted waves can be recorded. This loss of information means that there are insufficient data to reconstruct the image exactly. The approximation of an infinite or closed measurement surface by one of limited extents—the “limited view,” “partial scan,” or “finite aperture” problem—therefore leads to artifacts and distortion in the reconstructed image. [It has been shown^{14–16} that the edge of a source region (an abrupt change in absorbed energy density) can be reconstructed stably only if the normal to the edge crosses the measurement surface. For a spherical measurement surface enclosing V this is readily achievable, but, as hinted at above, for a single planar measurement surface it is necessary for it to extend to infinity to fulfill this requirement.] In practice, image reconstruction algorithms designed with a complete measurement surface in mind, e.g., a sphere or an infinite plane, are applied to partial




FIG. 1. A rectangular reverberant cavity with a single detector at point (x_d, y_d) and a point source.

data sets with the missing information implicitly replaced by zeros, even though such an extension of the data is clearly unphysical. The use of range conditions to constrain the extension of the data to be more realistic has been proposed,¹⁷ but this does not overcome the fundamental lack of data.

B. Reverberant-field PAT

When approximating an infinite planar measurement surface by one of finite extent, one way to ameliorate the limited-view problem is to reflect the sound back onto the measurement surface by placing acoustically reflecting walls perpendicular to, and at the ends of, the planar measurement aperture. The effect of this is to introduce an infinite, periodically repeating array of acoustic image sources. The data can now be extended periodically so that $p(\mathbf{x}_s, t)$ is known over the entire infinite measurement plane, even though it was only recorded over a finite region. An efficient reconstruction algorithm that exploits the periodicity exists for this case.^{11,18} In effect, the information that would have been lost is retained in the reverberant field set up between the reflecting walls.

This use of a reverberant acoustic field for imaging is in contrast to most image reconstruction algorithms, which, as mentioned above, suppose that the measurement surface sits within an acoustic free-field, i.e., the waves emanating from the photoacoustic source p_0 travel outward through the measurement surface, unimpeded and unaffected by either the measurement surface itself or any other obstacles. The use of acoustic reflectors, in addition to being one way to tackle the limited-view problem, suggests that smaller detector arrays could replace larger ones without sacrificing the quality of the images, if reverberation can be used to replace the missing data. The question then naturally arises as to how far this idea can be taken. For instance, if p_0 is restricted to a box with reflecting walls, as shown in Fig. 1, is it possible to reconstruct p_0 from measurements of the reverberant field at a single point?¹⁹

III. RECTANGULAR CAVITY BACKPROJECTION

As a first step, consider the simple scenario: can the position of a single point source (which emits one impulse) be reconstructed from the times at which the pulse and its reflections arrive at a single detector? For a rectangular cav-




FIG. 2. Image sources and detectors. The shaded region is the actual rectangular cavity; the other rectangles are images. (a) Image sources. Those within a circle of radius ct have contributed to the signal measured at the detector by time t , where c is the sound speed. Rays are shown connecting the image sources to the omnidirectional detector, which records the times of arrivals of the pulses but not their directions. (b) Image detectors. These can be useful when reconstructing the position of a point source. The circles illustrate the backprojection of three impulses received at the detector for three of the image sources, showing that they coincide at the source point.

ity, the reflections from the walls can be represented by image sources, which make the calculations of the arrival times straightforward and provide a simple geometrical way to understand the problem. [Two-dimensional examples are used throughout this paper for ease of presentation and in order to limit the calculation times, but the same principles apply in three dimensions (3D).]

Figure 1 shows a $15\text{ mm} \times 20\text{ mm}$ cavity with a single point source and a single point omnidirectional detector, and Fig. 2(a) illustrates the idea of image sources due to the single point source within the reverberant cavity. The image sources within a circle of radius ct , where c is the speed of sound, contribute to the signal measured at the detector up to time t . If the source emits a pulse at time $t=0$, then the times of arrival of the pulse and its reflections can be straightforwardly calculated by using this notion of image sources: each image source supplies one pulse at time (r/c) , where r is the distance between the detector and the image source. The effect of geometrical spreading on the waves was included by reducing the amplitude by a factor of $1/\sqrt{r}$ (in 3D this would be $1/r$). At this initial stage no attempt has been made to model the photoacoustic wave propagation more realistically than this, although it is fully modeled in the other examples given below and in Sec. IV. Figure 3 shows the impulse train $p(t)$ that would be recorded by a detector positioned in the corner, $(x_d, y_d) = (0\text{ mm}, 20\text{ mm})$, due to a point source at $(8\text{ mm}, 7\text{ mm})$. (There are three times at which two reflections have coincided exactly and therefore have twice the amplitude of the neighboring pulses.)

In Fig. 2(a), rays connecting the image sources to the detector have been drawn. Of course, the image sources emit rays in all directions, but it is clear from Fig. 2(a) that the direction in which the *detected* ray leaves its image source can, and usually will, be different for different image sources. This implies that many of the normals to the edge of the source will reach a single measurement point sooner or later. Given that the normals to the edges of a source are

required to cross the measurement surface for that edge to be stably reconstructed, as mentioned in Sec. II A, this suggests that it might be possible to reconstruct the edges of the source accurately from a measurement at a single detector.

By considering the symmetries of the cavity in Fig. 1, it is clear that for some detector positions it is not possible to determine the point source uniquely, as point sources at two different positions can give identical pulse trains. For example, if the detector lies on any axis of reflection or rotation symmetry then it is not possible to associate the pulse train with a unique source. However, if the detector is placed away from these symmetry axes, then it is possible to locate the point source uniquely. To demonstrate this, the concept of image detectors—the detector equivalents of image sources—will be introduced.

The pulse train in Fig. 3 was calculated by determining the distance (and therefore propagation time) between the detector and a set of image sources, as described above. Each image source represents one reflection, and therefore contributes one pulse to the pulse train. It could equally well have been generated by calculating the distance between the actual

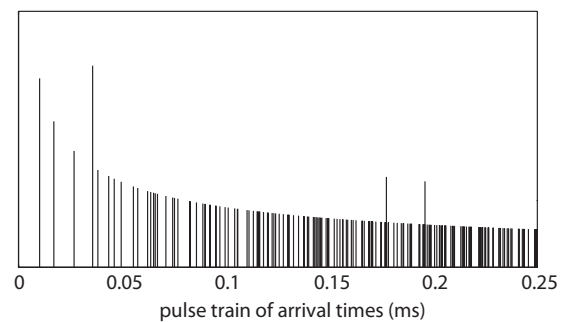


FIG. 3. The first 0.25 ms of the train of pulses received by the detector as calculated using image sources. The amplitude shows cylindrical spreading \sqrt{ct} dependence (sound speed $c=1500\text{ m/s}$). On three occasions two pulses arrive at the detector simultaneously, hence the three pulses with larger amplitude.




FIG. 4. Backprojecting the train of pulses from the nearest 500 image detectors shows that the point source can be located. The detector is located at the corner (0 mm, 20 mm) and the source is at (8 mm, 7 mm).

source position and a set of image *detectors*. The measured signal at each hypothetical image detector would consist of a single pulse, indicating the distance from the source. If these measured signals are denoted $p_i(t)$ for image detectors $i = 1, \dots$, then the pulse train for the actual detector would be the summation of the individual signals, $p(t) = \sum_i p_i(t)$. Each image detector represents one reflection and, as with the image sources, contributes one pulse to the measured pulse train. If it were possible to measure the signals $p_i(t)$ individually, then the position and amplitude of the point source could be determined by summing the backprojections of each $p_i(t)$

$$p_0(\mathbf{x}) \approx \sum_{i=1}^N \left(\frac{\sqrt{|\mathbf{x} - \mathbf{d}_i|}}{N} \right) p_i(|\mathbf{x} - \mathbf{d}_i|/c), \quad (1)$$

where \mathbf{d}_i is the position of the image detector i , and the term $\sqrt{|\mathbf{x} - \mathbf{d}_i|}$ was included to account for the cylindrical spreading. This will give the position and amplitude of the point source exactly in the limit $N \rightarrow \infty$. In practice, the set of signals $\{p_i(t); i=1, \dots\}$ is not measured, but only its summation $p(t)$. Replacing $p_i(t)$ with $p(t)$ in Eq. (1) gives an approximate backprojection formula

$$p_0(\mathbf{x}) \approx \sum_{i=1}^N \left(\frac{\sqrt{|\mathbf{x} - \mathbf{d}_i|}}{N} \right) p(|\mathbf{x} - \mathbf{d}_i|/c). \quad (2)$$

Figure 2(b) illustrates the principle by showing the backprojection of three pulse arrival times from the actual detector position and two image detectors, and the three coincide only at the source point. The backprojection of the time series in Fig. 3, from the 25 image detectors nearest to the actual detector, gives the image in Fig. 4. It is clear that with only this single measurement, the source position can nevertheless be located accurately and uniquely. This is similar to single channel time-reversal imaging to locate point sources, which is discussed further in Sec. V.




FIG. 5. The first part of a photoacoustic time series simulated from the initial pressure distribution in Fig. 6(A) for the point (0.03 mm, 19.96 mm).

Next, the same backprojection was used to reconstruct an estimate of the initial pressure distribution from a simulated *photoacoustic* time series “measurement” rather than just pulse arrival times. The “measured” signal in Fig. 5 was simulated using a *k*-space method²⁰ from the initial pressure distribution shown in Fig. 6(A). The detector was positioned close to the corner at (0.03 mm, 19.96 mm), and 1% Gaussian noise was added. Figure 6(B) shows the results of backprojecting the 0.5 ms duration time series from the 500 image detector positions nearest to the actual detector. It is apparent that the main features of the initial pressure distributions have been recovered, but also that the image is contaminated by artifacts and noise. A second example, which will be compared to the modal reconstruction approach described in Sec. IV, is also shown in Fig. 7. (These images are placed later in the text to allow the comparison with the images generated using other techniques to be made easily.) The initial pressure distribution, Fig. 7(A), consists of nine small circular sources with Gaussian profiles which are just about located accurately in the image, Fig. 7(B), but the errors due to artifacts are so significant that the image quality is poor.

IV. MODAL RECONSTRUCTION

The backprojection reconstruction described in Sec. III can give at best a crude approximation to the true initial pressure distribution. Furthermore, it is only simple to implement when the reverberant cavity is a regular shape, such as a rectangle, because the positions of the image detectors can then be calculated straightforwardly. For a more accurate reconstruction, and one that could be applied to cavities of more general shape, a reconstruction based on the acoustic modes of a cavity was used. This has the advantage that it is based more closely on the physics of the system, it is exact (within the limits imposed by ill-posedness) and is generally applicable, in that mode shapes and eigenvalues can be calculated for a cavity of any shape.

A. General case

A photoacoustically generated acoustic pressure field at a point (x, y) in a reverberant cavity, at a time t , can be written as a sum of normal modes




FIG. 6. (A) Circular initial pressure distribution used to generate the time series in Fig. 5. (B) Image formed by backprojecting the reverberant time series from the 500 image detector positions nearest to the actual detector. The circular initial pressure distribution has been recovered, but the image suffers from artifacts and noise.

$$p(x, y, t) = \sum_m A_m \phi_m(x, y) \cos(\omega_m t), \quad (3)$$

where subscript m indicates the mode with amplitude A_m , modal frequency ω_m , and normalized mode shape $\phi_m(x, y)$. In this notation, the initial pressure distribution is

$$p_0(x, y) = \sum_m A_m \phi_m(x, y), \quad (4)$$

and it is clear that $\partial p / \partial t|_{t=0} = 0$, i.e., the particle velocity is initially zero, which is true for photoacoustic waves. By discretizing the continuous variables x , y , and t as x_n , y_n , and t_i , where $n=0, \dots, N$ and $i=0, \dots, T$, Eqs. (3) and (4) can be written in matrix-vector form

$$\mathbf{p} = M_t \mathbf{A}, \quad \mathbf{p}_0 = M_0 \mathbf{A}, \quad (5)$$

where $\mathbf{p} = [p(x_d, y_d, t_0), \dots, p(x_d, y_d, t_T)]^T$, $\mathbf{p}_0 = [p_0(x_0, y_0), \dots, p_0(x_N, y_N)]^T$, and $\mathbf{A} = [A_0, \dots, A_M]^T$. M is the total number of modes that are included in the reconstruction. Each column of the $N \times M$ matrix M_0 contains the shape of a single mode m at all the points (x_n, y_n) , $n=0, \dots, N$

$$M_0(n, m) = \phi_m(x_n, y_n), \quad (6)$$

and each column of the $T \times M$ matrix M_t contains the values of a single mode m at the detector position (x_d, y_d) at all the time points t_i , $i=0, \dots, T$

$$M_t(i, m) = \phi_m(x_d, y_d) \cos(\omega_m t_i). \quad (7)$$

The matrices M_0 and M_t can be generated in advance from knowledge of the mode shapes and frequencies of the cavity. For simple geometries these may be known analytically; otherwise it might be necessary to calculate them numerically using the finite element (FE) method, for instance. (In which case it would also be possible to take into account known

variations in the acoustic parameters within the cavity, such as sound speed heterogeneities.)

From Eq. (5) it is clear that an approximation to \mathbf{p}_0 can be found from the measured time series, \mathbf{p} , in two steps: first, estimate the modal amplitudes \mathbf{A} by inverting the first equation,

$$\mathbf{A} = M_t^{-1} \mathbf{p}, \quad (8)$$

and then use the second equation to obtain an estimate for \mathbf{p}_0 . Written as one equation these two operations are

$$\mathbf{p}_0 = M_0 M_t^{-1} \mathbf{p}, \quad (9)$$

which could also be written as

$$\mathbf{p}_0 = G_{\text{reverb}}^{-1} \mathbf{p}, \quad (10)$$

where the columns of the Green's function matrix G_{reverb} contain the time series (impulse response functions) that would be recorded at (x_d, y_d) due to impulses leaving the points (x_n, y_n) , $n=0, \dots, N$, at time $t=0$. For simple geometries G_{reverb} can be found analytically; when the boundary is more complex, G_{reverb} could in principle be constructed column by column from time domain calculations of the impulse response functions. Clearly, $G_{\text{reverb}}^{-1} = M_0 M_t^{-1}$.

Equations (9) and (10) show that the success of this imaging approach depends on the invertibility of the matrix M_t (or G_{reverb}). This in turn depends on the geometry of the cavity, as will be shown below by comparing the reconstruction using a rectangular cavity with that in a chaotic cavity (a quarter Sinai billiard).

B. Rectangular cavity

For a rectangular cavity, the mode shapes are separable in x and y and take on simple forms, so Eq. (3) becomes




FIG. 7. (Color online) (A) Initial pressure distribution, p_0 , consisting of nine small circles with Gaussian profiles. (B) Image of p_0 obtained by backprojection using Eq. (2). (C) Image of p_0 obtained by a modal reconstruction in a rectangular cavity. (D) Image of p_0 obtained by a modal reconstruction in a chaotic cavity. (E) Profiles at $y=10$ mm through (A), (B), (C), and (D), corresponding to the exact p_0 (solid line), backprojection (dashed), modal reconstruction with rectangular cavity (dot-dashed), and chaotic cavity (dotted). (F) Mean squared error as a function of truncation order for the images from rectangular (+) and chaotic (○) cavities.

$$p(x,y,t) = \sum_{k,l} A_{kl} \cos(k\pi x/L_x) \cos(l\pi y/L_y) \cos(\omega_{kl}t), \quad (11)$$

where $k=0, \dots, K$, $l=0, \dots, L$, and each pair $\{k, l\}$ corresponds to a single mode with amplitude A_{kl} , modal frequency

$$\omega_{kl} = c \sqrt{(k\pi/L_x)^2 + (l\pi/L_y)^2}, \quad (12)$$

and mode shape $\cos(k\pi x/L_x) \cos(l\pi y/L_y)$. Similarly, Eq. (4) becomes

$$p_0(x,y) = \sum_{k,l} A_{kl} \cos(k\pi x/L_x) \cos(l\pi y/L_y). \quad (13)$$

Each column of the $N \times KL$ matrix M_0 contains, therefore, the shape of a single mode $\{k, l\}$ at all the points (x_n, y_n) :

$$M_0(n, \{k, l\}) = \cos\left(\frac{k\pi x_n}{L_x}\right) \cos\left(\frac{l\pi y_n}{L_y}\right), \quad (14)$$

and each column of the $T \times KL$ matrix M_t contains the values of a single mode $\{k, l\}$ at the detector position (x_d, y_d) at the time points $[t_0, \dots, t_T]$:




FIG. 8. Singular values of matrix M_t for the rectangular cavity.

$$M_t(i, \{k, l\}) = \cos\left(\frac{k\pi x_d}{L_x}\right) \cos\left(\frac{l\pi y_d}{L_y}\right) \cos(\omega_{kl} t_i). \quad (15)$$

This modal reconstruction approach was used to estimate the initial pressure distribution shown in Fig. 7(A) from the same single time series measurement (not shown) that was used to reconstruct Fig. 7(B). The $T \times 1$ vector of measured data, \mathbf{p} , included 1% Gaussian noise and consisted of $T=1600$ equally spaced samples with $t_{\max}=0.16$ ms, and the $15 \text{ mm} \times 20 \text{ mm}$ reconstruction mesh consisted of $N=90 \times 90=8100$ rectangular pixels. Matrix M_t was constructed with the 1000 lowest frequency modes, and its singular value decomposition (SVD) was calculated using MATLAB's SVD function, which uses LAPACK's DGESVD routine. The singular value spectrum is shown in Fig. 8. The physical significance of the very small singular values will be discussed below, but numerically they make M_t ill-conditioned, and therefore its (pseudo)inverse is very sensitive to noise in the data. Regularization must be included to suppress the noise, and here a truncated-SVD inversion was used. The mean squared error between the estimated and true images is shown, again with + symbols, in Fig. 7(F) as a function of the truncation order. Figure 7(C) shows the image with the truncation order set to 770 singular values. This reconstruction is much more accurate than the backprojection, Fig.

7(B), but it still used less than 80% of the available modes, as many had to be discarded in order to regularize the solution. If a way could be found to improve the conditioning of the matrix M_t , so that more modes can be included in the inversion, it should be possible to improve the accuracy of the image further.

C. Modal degeneracy

The columns of matrix M_t are given by Eq. (7), and each one is a scaled cosine oscillating with modal frequency ω_m . For M_t to be invertible it is necessary that the columns are linearly independent. This can only be the case when no two modal frequencies are the same, i.e., $\omega_m \neq \omega_n \forall m, n$. In the case of a rectangular cavity, the modal frequencies are given by Eq. (12), so many modes have the same modal frequency, and the condition does not hold. This modal degeneracy gives rise to the practically zero singular values in the spectrum shown in Fig. 8 and thus to the noninvertibility of M_t . In order to improve the conditioning of M_t , therefore, it is necessary to choose a cavity shape that minimizes the number of modes with the same modal frequency. Cavity shapes with this property are already known from the study of waves in reverberant cavities known as quantum or wave chaos.²¹⁻²³ Indeed, the eigenvalues for cavities of different shapes have been studied in some depth, and it has been shown that the distributions of the spacings between the eigenvalues (or modes) take on quite precise and universal forms which depend on the shape of the cavity. Figure 9(a) is a histogram showing the distribution of the spaces between the lowest 4000 modal frequencies in a $20 \text{ mm} \times 15.1 \text{ mm}$ rectangular cavity as a function of normalized modal spacing, s (modal spacing/number of modes included). This distribution agrees closely with Poisson's distribution $\exp(-s)$, also shown, as is expected for integrable cavities such as this. (An integrable cavity is one in which Hamilton's equations describing the cavity's ray dynamics form an integrable system.) It is notable that this distribution, which applies universally to any cavity of this type, has its peak at a modal spacing of zero. This is of practical significance because it means that there will be many degenerate modal frequencies, and therefore imperfect image reconstruction, for any inte-




FIG. 9. Histograms showing the distributions of the spaces between the lowest 4000 modal frequencies as a function of normalized modal spacing, s (modal spacing/number of modes included). (a) For a $20 \times 15.1 \text{ mm}^2$ rectangular cavity, the distribution agrees closely with Poisson's distribution $\exp(-s)$ (solid line). (b) For a chaotic cavity (a quarter Sinai billiard), the distribution agrees closely with Wigner's distribution $(\pi/2)s \exp(-\pi s^2/4)$ (solid line).

grable cavity. (In fact, it is even worse than this for the case investigated in Sec. IV B. When the sides of a rectangular cavity are in simple ratios there is an even larger peak close to zero spacing.²³ The cavity used in Sec. IV B was 20 mm × 15 mm, not 20 mm × 15.1 mm, and the resulting increase in degeneracy is shown in Fig. 9(a) with dotted lines.) If the modal frequencies were spaced so that no two modes share the same modal frequency, and degeneracy is thereby reduced, then the columns of M_t would be linearly independent and more of the modes can be used in the reconstruction, leading to more accurate imaging. That this is true for chaotic cavities is shown in Fig. 9(b) and discussed below.

D. Chaotic cavity

Although the description given for the acoustic field in Sec. IV A was in terms of modes, ray descriptions can also be helpful and intuitive. According to this model, rays are emitted from a source, reflect specularly from the walls of the reverberant cavity, and contribute to the detected signal when they pass through the measurement point. Consider rays emitted in all directions from a source. For some cavity geometries some of these rays will never reach the detector. They may, for instance, follow a periodic orbit around the cavity never encountering the measurement point if it is not on their orbit. The rectangular cavity used here is one example of a type of cavity in which this can happen. It seems intuitive that if some of the rays from a point do not encounter the detector, then that point cannot be reconstructed exactly from the measured data. It seems clear, when thinking in terms of rays, that what is required is a cavity in which every ray travels through every point in the cavity, without getting trapped into periodic orbits. Such a cavity is termed a chaotic (or ray-chaotic) cavity.

To convert the 20 mm × 15 mm cavity used so far into a chaotic cavity, a quadrant of a circle of radius 7 mm, centered at (0,0), was removed to form a “quarter Sinai billiard.” (A Sinai billiard is a rectangular cavity with a circle removed from the center.) The modal spacing distribution for all chaotic cavities follows Wigner’s distribution, $(\pi/2)s \exp(-(\pi s^2/4))$, shown as the solid line in Fig. 9(b), and the distribution of the spacings of the first 4000 modes of the quarter Sinai billiard clearly follows this trend. The principal point to note is that in a Wigner distribution, there are no modal spacings of zero, i.e., no two modes have the same modal frequency. This is exactly what was required to improve the conditioning of M_t and to lead to better images.

The modal frequencies of the chaotic cavity cannot be calculated using simple algebraic expressions, as they can for the rectangular case, and so a FE model was used. A FE simulation was constructed in MATLAB and the function `sptarn` was used to calculate the eigenvalues (Arnoldi estimates) and eigenvectors of the cavity, and thereby construct matrix M_t . A different FE mesh and the MATLAB function `hyperbolic` were used to calculate the measured photoacoustic time series for a detector at (0.03 mm, 19.96 mm), and 1% Gaussian noise was added. The reconstruction was performed, as above, using a truncated SVD. The singular




FIG. 10. Singular values of matrix M_t for rectangular (+) and chaotic (.) cavities.

value spectrum is shown in Fig. 10 (dots) alongside the SVD spectrum from the rectangular case (+), and it shows that the modal degeneracy has been removed so almost all the modes can now be used in the reconstruction without causing ill-conditioning.

Figure 7(A) shows an initial pressure distribution used to compare the three reconstructions: simple backprojection [Fig. 7(B)], modal inversion in a rectangular cavity [Fig. 7(C)], and modal inversion with the chaotic cavity [Fig. 7(D)].

Figure 7(D) shows the image recovered when 920 modes were included, which when compared to Fig. 7(C) shows better recovery of both the amplitudes and the shapes of the circles. The improvement in the reconstruction of the amplitudes is confirmed by Fig. 7(E), which shows profiles at $y=10$ mm through the images in Figs. 7(A)–7(D), corresponding to the exact p_0 (solid line), p_0 reconstructed using backprojection (dashed), modal reconstruction with rectangular cavity (dot-dashed), and reconstruction with chaotic cavity (dotted). Figure 7(F) shows the mean squared error as a function of the truncation order for the rectangular (cross) and chaotic (open circle) cavities. A much larger percentage of the modes can be used in the chaotic case, and the error is consequently significantly lower than in the rectangular case.

A second example comparing reverberant PAT imaging in the rectangular and chaotic cavities was also calculated. The inversions were performed exactly as before but for the initial pressure distribution shown in Fig. 11(A). The images for the rectangular and chaotic cavities are given in Figs. 11(B) and 11(C), and profiles through these images at $y=15$ mm are shown in Fig. 11(D). Again it is clear that the chaotic cavity gives better results in terms of both amplitude and shape.

V. DISCUSSION

A. Matrix inversion

The aim of this paper is to suggest and try to demonstrate that it might be possible to use a single point measurement of reverberation to do photoacoustic imaging. To this end, simple two-dimensional examples have been given, in

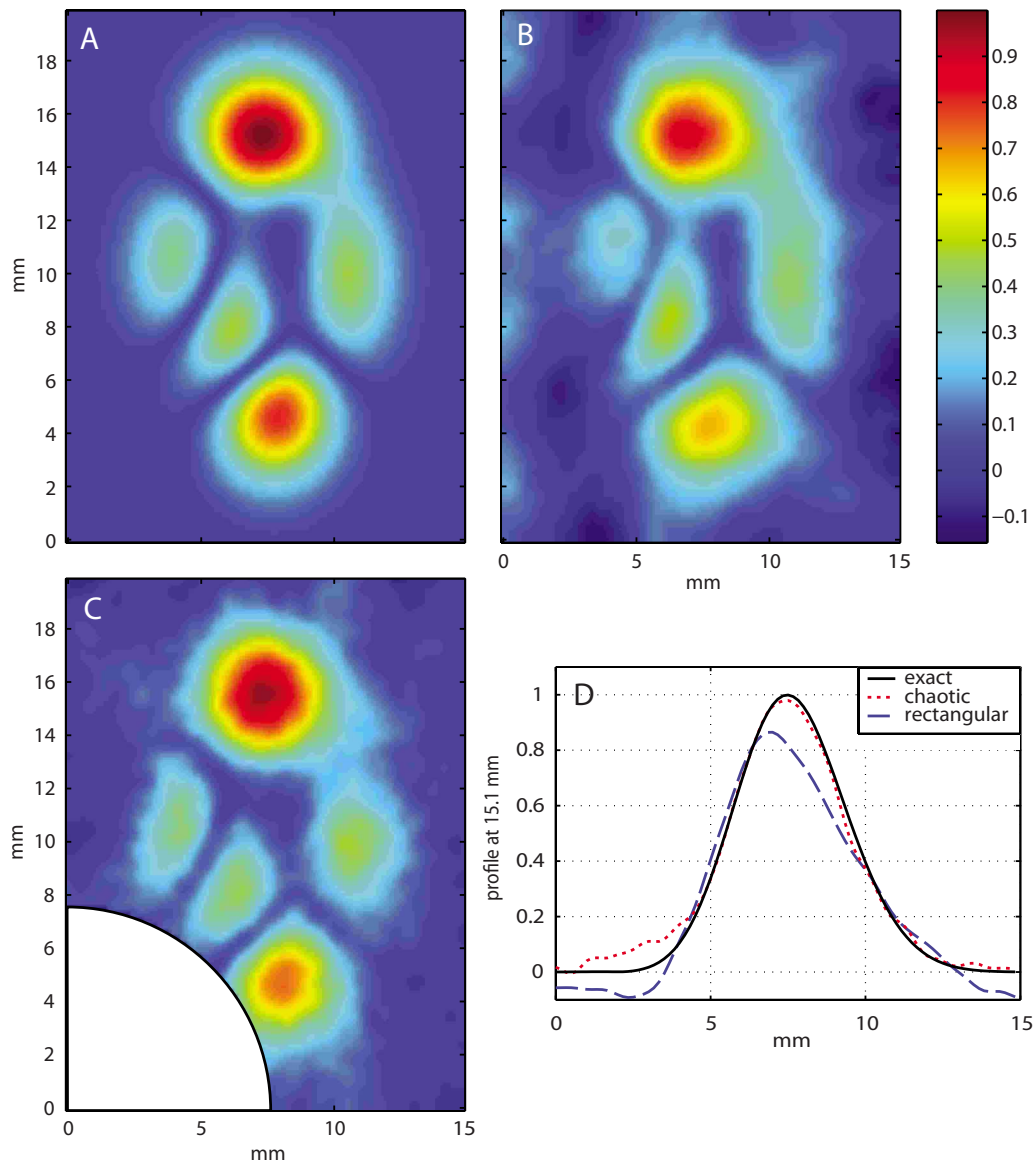


FIG. 11. (Color online) (A) Example initial pressure distribution, p_0 . (B) Image of p_0 obtained using a rectangular cavity. (C) Image of p_0 obtained using a chaotic cavity. (D) Profiles at $y=15$ mm through (A), (B), and (C), corresponding to the exact p_0 (solid line), p_0 reconstructed using modal reconstruction with rectangular cavity (dashed), and reconstruction with chaotic cavity (dotted).

which a SVD was used to calculate the matrix inversions. The SVD provides a fundamental decomposition of the matrices, and as such helps to provide insight into the reasons why a chaotic cavity is preferable to an integrable one. (It is worth noting that the inversion of M_t only needs to be performed once for each geometry, i.e., cavity shape and detector position, so for a particular experimental arrangement M_t^{-1} could be precalculated, and each image formed with a single matrix-vector multiplication. It does not need to be calculated for every new initial pressure distribution.) The SVD is only one way to calculate the inversions, and the truncated SVD the only one way to regularize the solution. The literature on algorithms to estimate matrix inversions and regularization is substantial, and there is insufficient space here for a study of which algorithms are best suited to this particular problem. In some cases, perhaps for large-scale problems, a different method altogether may be advantageous. For instance, an optimization approach, in which a

norm of the error between the measured data and a model is minimized, may have computational benefits as well as the advantage that a constraint could be added to the cost function to include the prior information that p_0 must be non-negative everywhere.

B. Model assumptions

In the analysis above, a number of idealizations were made to simplify the situation. For instance, it was assumed that the sound speed is the same everywhere within the cavity, that there is no bulk acoustic absorption, and that the walls are perfectly reflecting (have a reflection coefficient of 1). In practice these assumptions will not quite be true, and the sound field will differ from the model, especially at higher frequencies where the absorption may be significant. The time series used in the examples was 0.5 ms duration, over which time the sound waves will have traveled about

0.75 m, or almost 40 times the length of the cavity. The high frequency part of the signal will suffer in two ways over this distance: it will be susceptible to being scattered by localized sound speed heterogeneities and it may be strongly absorbed through the bulk absorption of the medium. Both of these effects will act to reduce the spatial resolution achievable in the image. To counter these effects, both the spatial variation in sound speed and the absorption could be included in the model, if they are known beforehand, and the mode shapes, frequencies, and amplitudes could then be calculated taking these variations into account. When the high frequency absorption is significant, however, the acoustic pressure signal may fall well below the noise floor, and regularization will be required in the inversion to ensure a stable solution. The frequency dependence of absorption varies, so the frequency at which this will start to be a problem will depend on the type of tissue.

The assumption of perfectly reflecting walls will be difficult to achieve in practice, but materials with a large acoustic impedance can achieve high pressure reflection coefficients for sound traveling from water, e.g., for steel $R \approx 0.94$ and for glass $R \approx 0.8$. As above, if the reflection coefficients are known, they can be included in the model. (This will add to the damping of the modes.) It may be possible to do calibration experiments in the cavity containing a known medium, such as distilled water, to ensure that the model of the geometry and boundary conditions are well modeled. Another possibility is to replace one of the walls with an air (pressure-release) boundary and take into account the $R \approx -1$ reflection coefficient when calculating the modal information. This sort of open-topped cavity would have the practical advantage of making the cavity more accessible for inserting and removing samples.

C. Detectors

The question studied in this paper has been whether an image could be obtained from reverberant measurements using *one* detector. Intuitively there would be advantages in using data from more than one measurement point. First, if a single detector happens to be positioned close to the node of a mode, then the contribution of that mode to the measured pressure will be below the noise floor, and it will not be included in the image reconstruction. With two detectors, it is less likely that a mode will be missed in this way, so in some circumstances could improve the image. Second, if the reverberant signal is attenuated through absorption, a particular ray may be too weak to detect by the time it reaches the detector, but a second detector, some distance from the first, may encounter the ray earlier in its “lifetime” and thus with a higher signal-to-noise ratio. However, the conditioning of the matrix M_t , which has been the chief concern of much of this paper, would not be improved, as its conditioning depends on the columns having different modal frequencies, and two different detector points would still see the same modal frequencies. Measurements recorded in different parts of the cavity will also be affected by absorption differently, so it may be that when the absorption is significant, multiple measurements can be used to improve the image.

The detectors in this analysis have been assumed to be omnidirectional, broadband, and located at a single point. “Real” ultrasound detectors are often directional, bandlimited, and of finite size. There is a class of ultrasound detectors, however, which may be suitable for this application. Optical interferometric ultrasound detectors^{6,24} can have a very wide, flat, sensitivity range (from dc to >50 MHz), with small element sizes ($<50 \mu\text{m}$) and could be incorporated into the wall of the cavity, so not affecting the reverberant field.

D. Time-reversal imaging

When using time-reversal to refocus a signal through a heterogeneous or reflecting medium, the measurements are made at detector points and then time-reversed versions are transmitted back into the medium from the same positions. When time-reversal is used for imaging, the retransmission into the medium is replaced by a simulated retransmission in a numerical model. Experiments on time-reversing array measurements of photoacoustic signals back to their origin have been performed²⁵ and time-reversal has been suggested as a way to reconstruct a photoacoustic image^{26,27} when measurements have been made over an array of detectors.

The time-reversal of a single photoacoustic measurement in a reverberant cavity in order to form an image has yet to be investigated, although the idea of using single channel time-reversal in a reverberant cavity to detect the location of a point source has been studied and applied to communication²⁸ and touch sensitive interactive surfaces.²⁹ In three interesting papers, Draeger and co-workers^{30–32} described and analyzed experimental and numerical studies of the time-reversal of a single point measurement of the reverberation generated by a point source back to its origin. These studies were conducted using flexural waves on a chaotically shaped silicon plate, but are analogous to the acoustic case, and demonstrated that the signal could indeed be refocused. A finite difference model was used to show the waves converging on the point source,³⁰ perhaps the first example of time-reversal imaging using single channel reverberation data. As an imaging technique, single point time-reversal differs in a number of ways to the modal inversion given here, e.g., it requires the chaoticity of the cavity and so does not work in a rectangular cavity,^{21,32} and it is inherently approximate, so can never give an exact image.

It is instructive to understand where the approximation in single channel time-reversal arises, and why it does not apply to the modal inversion described here. This can be explained both in terms of modes or rays. The ray explanation is that while each pulse arrives at the detector from one specific direction, it is retransmitted in all directions, as the omnidirectional detector does not record the direction of arrival.³² This is similar to the approximation made in the backprojection in Sec. III, which introduced significant artifacts while still roughly locating the regions of large initial pressure. In terms of modes, the measured reverberation signal consists of a contribution from each mode weighted by the modal amplitudes at the detector point, $A_m \phi_m(x_d, y_d)$. The time-reversed signal that is retransmitted into the medium (or

the model in the case of imaging) will only excite the modes at the detector (now transmitter) point with the same amplitude. Therefore, unless the detector is at a maximum of a mode, its contribution to the overall acoustic field will be underestimated by $(A_m \phi_m(x_d, y_d))^2$. Draeger and Fink³¹ suggested that the solution to this problem would be to know the amplitude of each eigenmode at the detector point and compensate for it, which is precisely the role of $\phi_m(x_d, y_d)$ in Eq. (7). This amplitude effect only now remains a problem for modes that fall below the noise floor at the detection point. To avoid this, the detector was placed close to the corner of the cavity where the boundary conditions enforce a maximum for all modes.

VI. SUMMARY

Photoacoustic tomography using reverberation recorded by a single detector was proposed, and two-dimensional simulations using a simple imaging algorithm based on a modal inversion were described. The technique was analyzed using a SVD, and it was shown that better quality images can be obtained when the reverberant cavity is chaotic. The practicalities of such an imaging technique, and its advantages over single channel time-reversal imaging, were discussed.

ACKNOWLEDGMENTS

The authors would like to thank Professor Simon Arridge for useful and interesting discussions. This work was supported by the Engineering and Physical Sciences Research Council (UK).

¹Photoacoustic tomography (PAT) is also known as photoacoustic imaging or optoacoustic tomography. When the optical excitation pulse is replaced by microwave or rf radiation, the technique is called thermoacoustic tomography. The image reconstruction procedure is the same in both cases.

²R. A. Kruger, P. Liu, Y. R. Fang, and C. R. Appledorn, "Photoacoustic ultrasound (PAUS)—Reconstruction tomography," *Med. Phys.* **22**, 1605–1609 (1995).

³C. G. A. Hoelen, F. F. M. de Mul, R. Pongers, and A. Dekker, "Three-dimensional photoacoustic imaging of blood vessels in tissue," *Opt. Lett.* **23**, 648–650 (1998).

⁴X. Wang, Y. Pang, G. Ku, X. Xie, G. Stoica, and L. V. Wang, "Noninvasive laser-induced photoacoustic tomography for structural and functional in vivo imaging of the brain," *Nat. Biotechnol.* **21**, 803–806 (2003).

⁵M. Xu and L. V. Wang, "Photoacoustic imaging in biomedicine," *Rev. Sci. Instrum.* **77**, 041101 (2006).

⁶E. Z. Zhang, J. G. Laufer, and P. C. Beard, "Backward-mode multiwavelength photoacoustic scanner using a planar Fabry-Perot polymer film ultrasound sensor for high-resolution three-dimensional imaging of biological tissues," *Appl. Opt.* **47**, 561–577 (2008).

⁷R. Kruger, W. Kiser, D. Reinecke, G. Kruger, and K. Miller, "Thermoacoustic molecular imaging of small animals," *Mol. Imaging* **2**, 113–123 (2003).

⁸J. G. Laufer, C. Elwell, D. Delpy, and P. Beard, "In vitro measurements of absolute blood oxygen saturation using pulsed near-infrared photoacoustic spectroscopy: Accuracy and resolution," *Phys. Med. Biol.* **50**, 4409–4428 (2005).

⁹J. G. Laufer, C. Elwell, D. Delpy, and P. Beard, "Quantitative spatially

resolved measurement of tissue chromophore concentrations using photoacoustic spectroscopy: Application to the measurement of blood oxygenation and haemoglobin concentration," *Phys. Med. Biol.* **52**, 141–168 (2007).

¹⁰B. T. Cox, S. R. Arridge, and P. C. Beard, "Simultaneous estimation of chromophore concentration and scattering distributions from multiwavelength photoacoustic images," *Proc. SPIE* **6856**, 68560Y (2008).

¹¹B. T. Cox, S. R. Arridge, and P. C. Beard, "Photoacoustic tomography with a limited-aperture planar sensor and a reverberant cavity," *Inverse Probl.* **23**, S95–S112 (2007).

¹²P. Kuchment and L. Kunyansky, "Mathematics of thermoacoustic tomography," *Eur. J. Appl. Math.* **19**, 191–224 (2008).

¹³M. A. Anastasio, J. Zhang, X. Pan, Y. Zou, G. Ku, and L. V. Wang, "Half-time image reconstruction in thermoacoustic tomography," *IEEE Trans. Med. Imaging* **24**, 199–210 (2005).

¹⁴A. K. Louis and E. T. Quinto, *Surveys on Solution Methods for Inverse Problems* (Springer-Verlag, Vienna, 2000), pp. 147–154.

¹⁵Y. Xu, L. V. Wang, G. Ambartsoumian, and P. Kuchment, "Reconstructions in limited-view thermoacoustic tomography," *Med. Phys.* **31**, 724–733 (2004).

¹⁶M. A. Anastasio and J. Zhang, "Image reconstruction in photoacoustic tomography with truncated cylindrical measurement apertures," *Proc. SPIE* **6086**, 608610 (2006).

¹⁷S. K. Patch, "Thermoacoustic tomography—consistency conditions and the partial scan problem," *Phys. Med. Biol.* **49**, 2305–2315 (2004).

¹⁸B. T. Cox and P. C. Beard, "Exact photoacoustic image reconstruction using a planar sensor array and image sources," *Proc. SPIE* **6437**, 64371H (2007).

¹⁹B. T. Cox, S. R. Arridge, and P. C. Beard, "Photoacoustic tomography using reverberant field data from a single detector," *Proc. SPIE* **6856**, 68561X (2008).

²⁰B. T. Cox, S. Kara, S. R. Arridge, and P. C. Beard, "k-space propagation models for acoustically heterogeneous media: Application to biomedical photoacoustics," *J. Acoust. Soc. Am.* **121**, 3453–3464 (2007).

²¹G. Tanner and N. Søndergaard, "Wave chaos in acoustics and elasticity," *J. Phys. A: Math. Theor.* **40**, R443–R509 (2007).

²²M. C. Gutzwiller, *Chaos in Classical and Quantum Mechanics*, Interdisciplinary Applied Mathematics (Springer-Verlag, New York, 1990).

²³H.-J. Stöckmann, *Quantum Chaos: An Introduction* (Cambridge University Press, Cambridge, 1999).

²⁴Y. Hou, J. S. Kim, H. Sheng-Wen, S. Ashkenazi, L. J. Guo, and M. O'Donnell, "Characterization of a broadband all-optical ultrasound transducer—From optical and acoustical properties to imaging," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **55**, 1867–1877 (2008).

²⁵E. Bossy, K. Daoudi, A.-C. Boccara, M. Tanter, J.-F. Aubry, G. Montaldo, and M. Fink, "Time reversal of photoacoustic waves," *Appl. Phys. Lett.* **89**, 184108 (2006).

²⁶Y. Xu and L. V. Wang, "Time reversal and its application to tomography with diffracting sources," *Phys. Rev. Lett.* **92**, 033902 (2004).

²⁷P. Burgholzer, G. J. Matt, M. Haltmeier, and G. Palttauf, "Exact and approximate imaging methods for photoacoustic tomography using an arbitrary detection surface," *Phys. Rev. E* **75**, 046706 (2007).

²⁸M. G. Heinemann, A. Larrazza, and K. B. Smith, "Acoustic communications in an enclosure using single-channel time-reversed acoustics," *Appl. Phys. Lett.* **80**, 694–696 (2002).

²⁹R. K. Ing, N. Quieffin, S. Catheline, and M. Fink, "In solid localization of finger impacts using acoustic time-reversal process," *Appl. Phys. Lett.* **87**, 204104 (2005).

³⁰C. Draeger and M. Fink, "One-channel time reversal of elastic waves in a chaotic 2-D cavity," *Phys. Rev. Lett.* **79**, 407–410 (1997).

³¹C. Draeger and M. Fink, "One-channel time reversal in chaotic cavities: Theoretical limits," *J. Acoust. Soc. Am.* **105**, 611–617 (1999).

³²C. Draeger, J.-C. Aime, and M. Fink, "One-channel time reversal in chaotic cavities: Experimental results," *J. Acoust. Soc. Am.* **105**, 618–625 (1999).