Chapter 3

Near-surface oscillations

3.1 Introduction†

For almost two decades, methods of local helioseismology (e.g, Hill, 1988; Braun, Duvall, & Labonte, 1987) have been applied to infer properties of the solar interior with varied degrees of success (for a comprehensive review, see Gizon & Birch (2005)). The predominant approach is to construct and subsequently invert models that relate observations to interior properties. As observations have become increasingly sophisticated, the need for refined forward modeling has become apparent. One reason the forward approach is crucial is that although the resonant mechanical modes of the Sun (the diagnostic agents of helioseismology) have been studied carefully, there are still many curious wave properties neglected in models that may prove significant. For example, finite wavelength effects cast doubt on the validity of the ray approximation in some situations (e.g., Hung, Dahlen, & Nolet, 2001; Couvidat et al., 2004); magnetic fields in the case of sunspots are potentially non-trivial contributors to the wavefield. Although these implications have been known for a while now, a systematic means of investigating such factors has only recently been constructed. Such studies

† *This chapter are reproduced from Hanasoge, Duvall, & Couvidat (2007). I was the first author and principally involved in the work. I constructed the code and performed all the simulations. Tom Duvall, Jr. computed the kernels and supplied data for Figures 3.1, 3.5, 3.11, and 3.12. S. Couvidat generated the following Figures 3.3, 3.6, 3.7, and 3.8.*

are difficult to conduct by purely analytical means, requiring the introduction of numerical methods to solve the constituent governing equations of wave motion (e.g., Tong et al., 2003; Hanasoge et al., 2006; Shelyag, 2006; Parchevsky & Kosovichev, 2006).

The effects of sound-speed perturbations on acoustic waves have been investigated in the past (e.g., Jensen et al., 2003; Jensen & Pijpers, 2003; Birch, Kosovichev & Duvall, 2004). In fact, Jensen et al. (2003) constructed a forward model to compute the wavefield associated with a sunspot-type sound-speed profile containing a near-surface decrease in the sound speed and an increase in the deeper layers.

Sensitivity kernels are mathematical structures that relate shifts in observational metrics such as travel times, resonant frequencies, etc. to the anomalies that cause them. Birch & Kosovichev (2000) and Gizon & Birch (2002) introduced finite frequency sensitivity kernels for helioseismology (note: we will use the terms 'finite wavelength' and 'finite frequency' interchangeably in future discussions) that include wave effects in the Born limit, to invert for p-mode interior sound speed, f-mode damping and source perturbations. The computation of these kernels is a non-trivial affair. Moreover, these kernels are obtained in the single scattering limit and the Born approximation itself may not be applicable for problems such as inversions of sunspots (Birch, private communication). In this chapter, we measure kernels at the photospheric level from simulations in two cases for which the Born limit is valid and published kernels are available for comparison (Gizon & Birch, 2002; Birch, Kosovichev & Duvall, 2004).

We also introduce the idea of noise subtraction, based on which we can obtain large signal to noise ratio (SNR) improvement in our simulations. Applying this concept, for certain problems, with as little as 40 hours of computing time, the signal can be boosted by so much that the eventual SNR is at the same level as that of two years of solar data. In §2 we discuss the numerical procedure applied to compute the wavefield. The idea of noise subtraction is introduced in §3.3 along with results that demonstrate its effectiveness. Travel times for the simulated wavefield are estimated with the technique of surface focusing in §3.4. We show how noise subtraction affects the travel times and also validate the results from the simulations in this section. An alternative method to compute kernels for helioseismology is shown to be effective in §3.5. We summarize our results and draw conclusions in §3.6.

3.2 The Simulation

The acoustic wavefield is simulated by numerically solving the linearized 3D Euler equations in Cartesian geometry:

$$\partial_t \rho = -\boldsymbol{\nabla}\boldsymbol{\cdot}(\rho_0 \mathbf{v}) - \Gamma\rho, \tag{3.1}$$

$$\partial_t \mathbf{v} = -\frac{1}{\rho_0}\boldsymbol{\nabla} p - \frac{\rho}{\rho_0} g\hat{\boldsymbol{z}} + \mathbf{S} - \Gamma\mathbf{v}, \tag{3.2}$$

$$\partial_t p = \rho_0 g v_z - \rho_0 c^2 \boldsymbol{\nabla}\boldsymbol{\cdot}\mathbf{v} - \Gamma p, \tag{3.3}$$

where ρ and ρ_0 are the fluctuating and time-stationary background density respectively, p and p_0 are the fluctuating and time-stationary background pressure respectively, $\mathbf{v}$ is the fluctuating vector velocity, $g = g(z)$ is gravity with direction vector $-\hat{\boldsymbol{z}}$, $c = c(x, y, z)$ is the sound speed, $\Gamma = \Gamma(x, y, z) > 0$ is a damping sponge that enhances wave absorption at the boundaries, and $\mathbf{S}$ is the source term. We employ a Cartesian coordinate system (x, y, z) with $\hat{\boldsymbol{z}}$ denoting the unit vector along the vertical or z axis and t, time. In sequential order, equations (3.1) through (3.3) enforce mass, momentum, and energy conservation respectively. In interior regions of the computational box (away from the boundaries), solutions to the above equations are adiabatic since equation (3.3) is applicable only in the case of adiabatic oscillations (e.g., Hanasoge et al., 2006) and there are no other damping terms.

In our computations, waves are excited by a vertically directed dipolar source function, $\mathbf{S} = S(x, y, z, t)\ \hat{\boldsymbol{z}}$, similar to the function described in Hanasoge et al. (2006). The function $S(x, y, z, t)$ is highly localized along the z axis, described by a Gaussian with full width at half maximum (FWHM) of 200 km. Adopting the approach described in Hanasoge et al. (2006), we start our analysis in the frequency

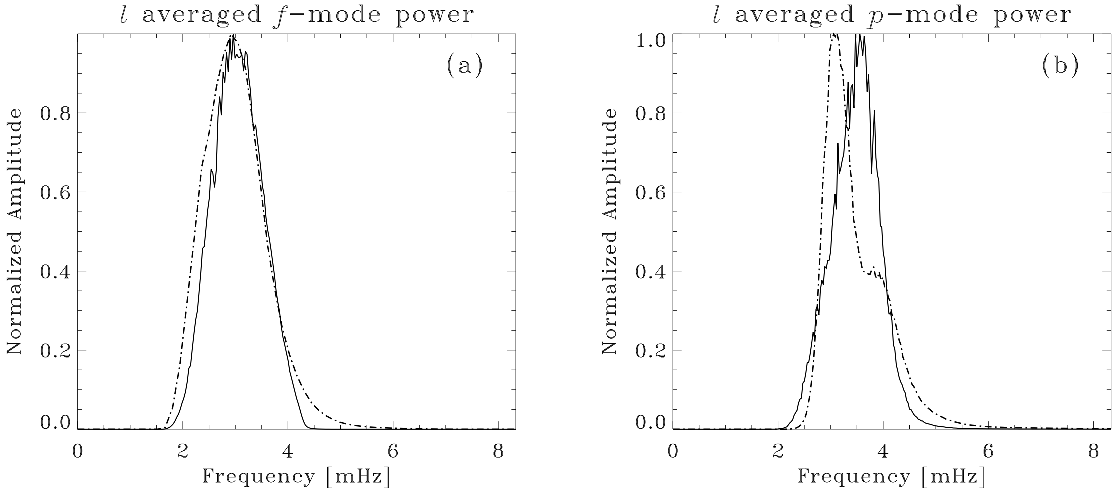

Figure 3.1 Comparison of wavenumber averaged power profiles of MDI high-resolution data (dot-dash line) and simulations (solid line) as a function of frequency. The f-mode power distribution in panel (a) is quite representative of the observations. Because of the lack of damping, the p-mode power distribution in panel (b) does not do quite so well. A phase-speed filter (Filter 'B' of Birch, Kosovichev & Duvall (2004)) was applied to obtain panel (b).

- horizontal wavenumber Fourier space. Because scattering processes are sensitive to the frequency and wavelengths of interacting waves, we attempt to mimic the solar acoustic power spectral distribution as closely as possible (see Figure 3.1). In the Sun, the competing effects of damping that has a super-linear frequency dependence and mode mass that decreases with frequency create a power maximum at $\nu = 3$mHz. However, one immediately realizes that solar damping rates are by no means simple functions of frequency with the consequence that fully including these dissipation rates in our time domain calculations is made all but impossible. Moreover, the aim of these calculations is to separate the influences of perturbative forces, one of which undoubtedly is damping (e.g., Gizon & Birch, 2002). Therefore, although approximate representations of the solar damping function could be incorporated into the equations above, we make the active choice to leave them out at present, with the full intention of pursuing a greater understanding of the effects of dissipation at a future occasion.

As for the dependence of the excitation function on horizontal wavenumbers, each coefficient in Fourier space is assigned a value from the output of a Gaussian distributed random number generator (the Ziggurat algorithm, available online from NETLIB). This results in uniform power across wavenumbers, which is somewhat different in the Sun because the sources are at granular scales and there is some dependence of power on wavenumber. To achieve the demands placed on the frequency dependence of the spectrum, we multiply these coefficients by an a priori specified frequency envelope. Lastly, we prescribe the excitation function so that it possesses no power beyond (to the right of) the f-mode ridge where modal power in the Sun is practically non-existent anyway.

Revisiting the background state, we use time-stationary properties p_0, ρ_0, g and c given by the convectively stable model of Hanasoge et al. (2006) based on model S (Christensen-Dalsgaard et al., 1996). Without additional artificial stabilization, simulations of wave activity in the near surface layers tend to blow up very rapidly due to exponentially growing linear convective instabilities. This stabilization results in a modified dispersion relation for p modes, not quite the same as it is for the Sun, but mercifully, close enough that existing methods of helioseismology may be applied with few alterations. The functional dependence of c presented in the model of Hanasoge et al. (2006) is entirely radial (vertical); for experiments with sound-speed perturbations, we alter c so that it becomes a 3D function of space while keeping p_0 and ρ_0 constant. The latter variables are kept constant so that the delicate hydrostatic balance remains undisturbed, which if tampered with, results in a Kelvin-Helmholtz instability. As seen in equation (3.4) this means that when we alter the sound speed and not the pressure or density, the first adiabatic index Γ_1 changes in the same sense, since

$$c = \sqrt{\frac{\Gamma_1 p_0}{\rho_0}}. \tag{3.4}$$

It can be shown that the degree of convective instability, characterized by the Brünt-Väisälä frequency N ($N^2 < 0$ means instability) in equation (3.5), increases when Γ_1 (or c) is reduced:

$$N^2 = g\left(\frac{1}{\Gamma_1}\frac{\partial \ln p_0}{\partial z} - \frac{\partial \ln \rho_0}{\partial z}\right). \tag{3.5}$$

Consequently, all our theoretical experiments are with local increases in sound speed. Our expectation is that decreases in sound speed affect helioseismic metrics in the opposite sense as corresponding increases would.

3.2.1 Numerical Algorithm

The computational domain is a cuboid that straddles the solar surface, extending from approximately 30 Mm below the photosphere to 2 Mm into the atmosphere. The vertical (radial) grid spacing is such that acoustic travel time between adjacent grid points is constant while the horizontal grid points are equally spaced. The damping sponge described in the previous section is placed adjacent to the six boundaries of the computational box. It is important that the function describing the damping layer decays sufficiently smoothly away from the boundary so that no wave reflections occur at the interface of the absorbent region. Absorbing boundary conditions (Thompson, 1990) are enforced at all boundaries. We choose absorbing over periodic boundaries because we want to avoid the issue of dealing with periodicities associated with the presence of a perturbation.

Derivatives are calculated using sixth-order compact finite differences (Lele, 1992). The solution is evolved in time using an optimized five-stage fourth-order Runge-Kutta scheme (Berland et al., 2006). To avoid aliasing, we apply the three-halves rule (Orszag, 1970), requiring that the number of grid points be at least three-halves the maximum captured wavenumber. In order to avoid vertical (radial) aliasing and the subsequent spectral blocking, we apply the de-aliasing procedure described in Hanasoge & Duvall (2007) every minute in solar time. We also de-alias variables in the horizontal directions by applying a smooth filter that diminishes the upper third of the spectrum and leaves the important lower two-thirds untouched (also at the rate of once per minute). All FFTs are performed using the freely distributed set of routines, the FFTW.

In order to achieve a comparable computation to wall clock time ratio, the code was parallelized according to the Message Passing Interface (MPI) Standard. The domain distribution algorithm is similar to the method described in Hanasoge &

Duvall (2006). The computational box is distributed along the y axis; all points on the x and z axes for a given point on the y axis are located in-processor. The data is transposed and redistributed between processors when the solution has to be filtered and when derivatives along the y axis need to be computed. It is probably true that greater parallel efficiencies may be achieved if the domain distribution is along all axes. However, incumbent to determining the optimal parallel configuration is the investment of considerable time, no doubt entailing many tedious hours of programming and debugging. At present, we only treat relatively small problems ($200 \times 200 \times 40$ Mm^3 is the biggest box) and achieve ratios of 1:1 (usually better) computing to real time. Seeing little merit in further code optimization, we have relegated such investigations to the future.

3.2.2 Power Spectrum

The modal distribution of power we recover upon putting together all the ideas discussed above and subsequently performing the computation is shown in Figure 3.2. We do not have a convincing explanation for the oddly shaped inter-ridge mode tails. Beyond the demarcation of the excitation profile, seen just below the f-mode ridge, excitation levels drop to zero. Also, in accordance with the three-halves rule, the excitation is Fourier limited, not extending beyond two-thirds of the spatial Nyquist wavenumber. The root mean squared (RMS) spatially averaged velocity as a function of time is shown in Figure 3.5; the simulation in this case is seen to achieve statistical stationarity in 4.5 hours. This timescale corresponds to at least a few bounces of the largest wavelength acoustic waves and presumably, full stationarity is reached when the balance between the forcing and the absorbent layers (at all boundaries) is struck.

3.3 Noise Subtraction

Given that we have full control over the excitation mechanism and source function, we can achieve high SNR in the artificial data by *subtracting* the noise out. In other words, having computed a source function, we perform two simulations, one with

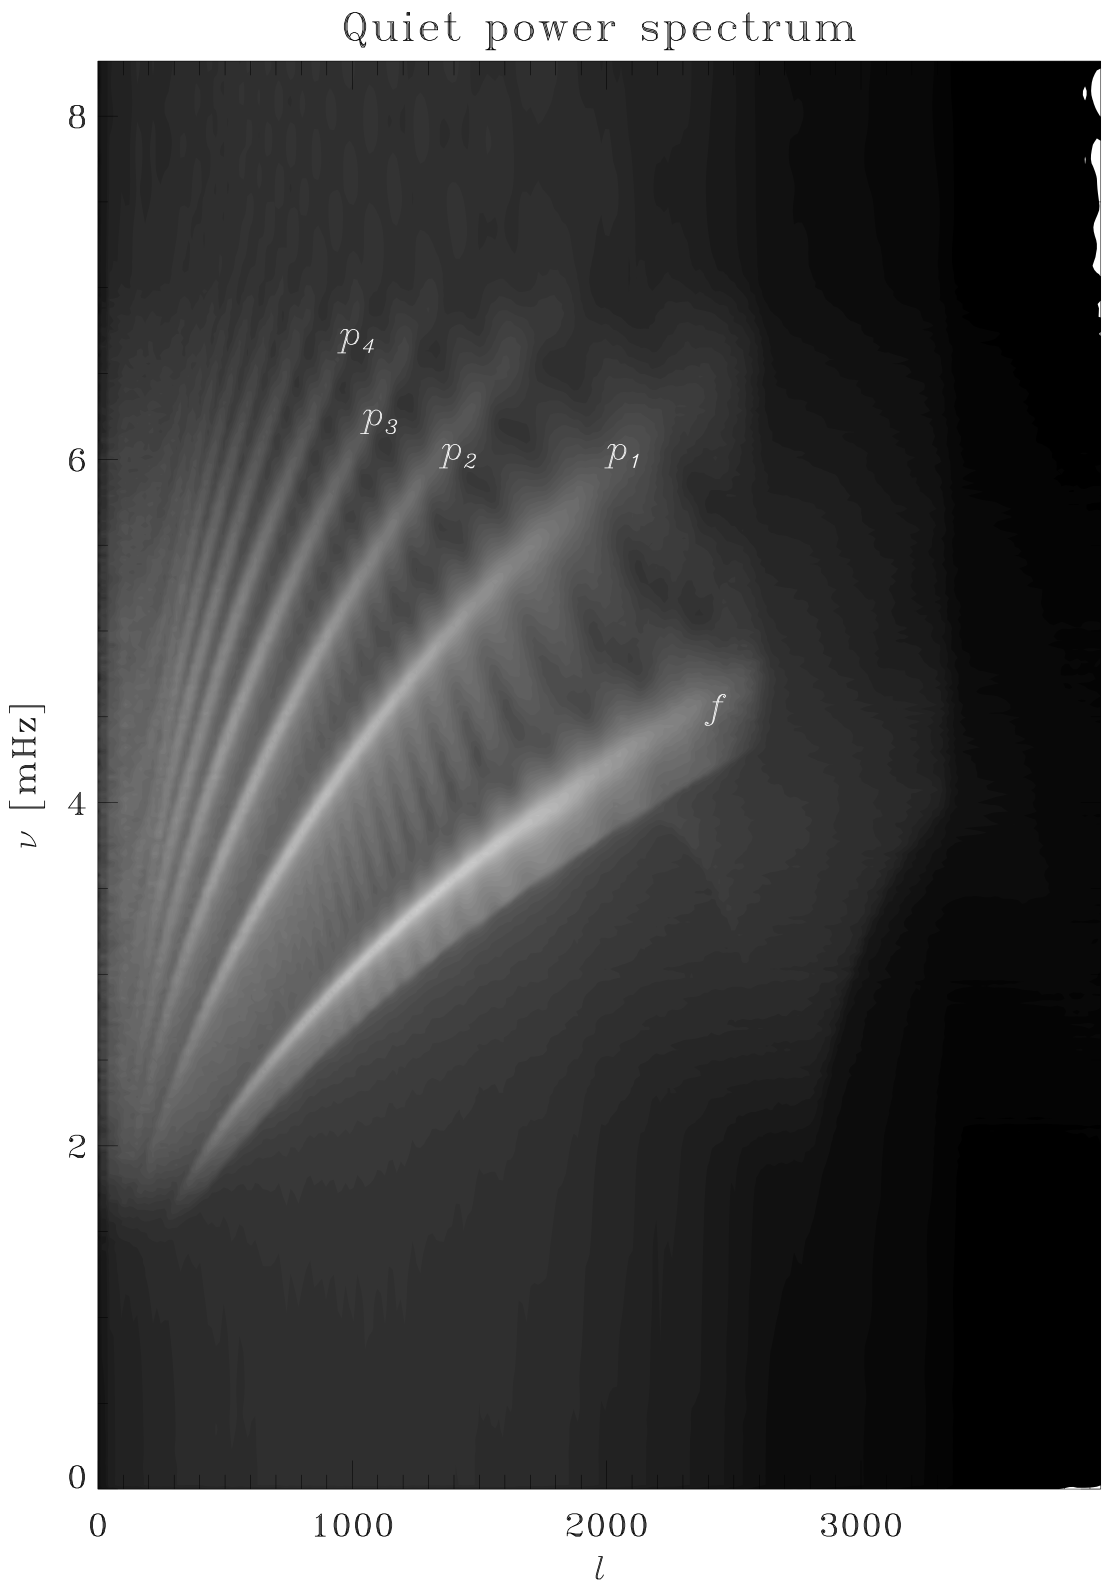

Figure 3.2 'Quiet' power spectrum obtained from a simulation in a $200 \times 200 \times 30$ Mm^3 box. The horizontal axis is spherical harmonic degree, l and the vertical axis is the frequency in milli-Hertz.

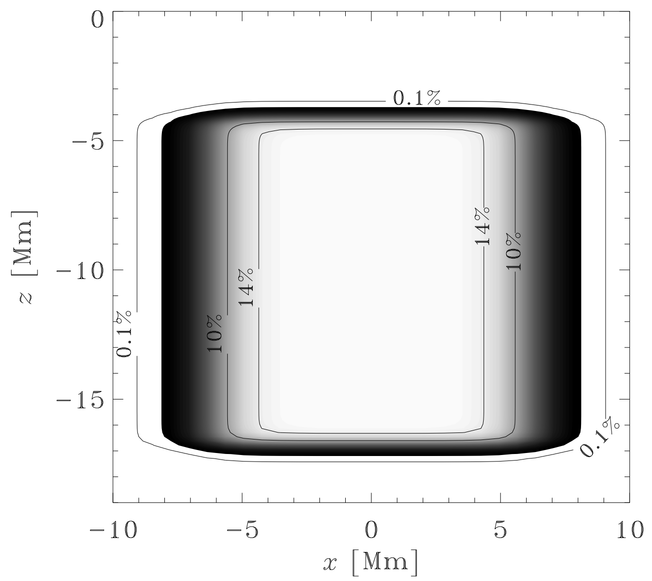

Figure 3.3 Vertical cut in the sound-speed perturbation discussed in Figure 3.4 and section 3.4.1. The labels on the isocontours show the amplitude of $\delta c^2/c^2$.

no perturbations ('quiet' simulation) and another with the perturbation of interest (shown in Figure 3.3). Evidence of the ability of this method is demonstrated by displaying the time-averaged RMS of the velocity differences between a quiet and perturbed model in Figure 3.4(b). Compare this to Figure 3.4(a) which is the time-averaged RMS velocity of the perturbed model. Subsequently, we can subtract the travel times of the quiet data from its perturbed cousin and depending on the size of the perturbation in comparison to the wavelength, excellent signal-to-noise properties can be achieved. The instantaneous difference is shown in panel (c) of Figure 3.4. In some cases where the perturbation is highly sub-wavelength in size, the SNR after noise subtraction is as much as 1000:1. In subsequent sections, we shall elaborate further upon the applications of this technique.

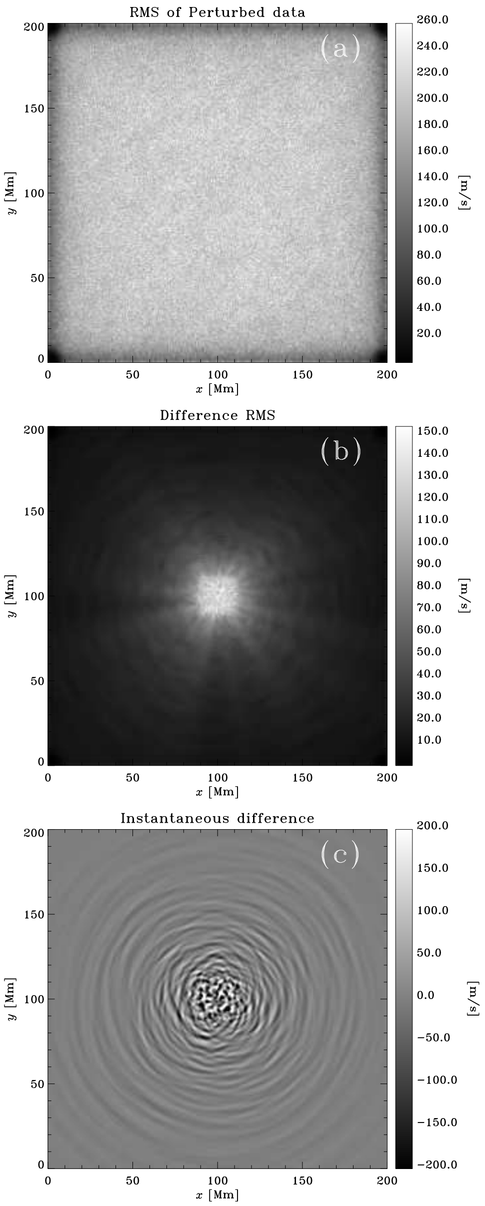

Figure 3.4 Panel (a): the time averaged RMS velocity of the wavefield interacting with a 12 Mm-sized sound-speed increase (shown in Figure 3.3), centered around $(x, y, z) = (100\ \mathrm{Mm}, 100\ \mathrm{Mm}, -10\ \mathrm{Mm})$. The sound-speed anomaly is entirely invisible in panel (a). The darkening towards the spatial edges of the frame is due to the absorbent sponge at work. Panel (b): the time averaged RMS velocity difference of the perturbed datacube and its quiet counterpart. The difference between the quiet and perturbed datacubes is greatest at the location of the perturbation due to enhanced scattering. Panel (c): the instantaneous difference between the perturbed and related quiet data 100 (solar) minutes into the simulation. The anomaly is a scatterer, creating ripples in the wavefield just like a pebble dropped onto the surface of still water. Fine wave structure is visible at the location of perturbation.

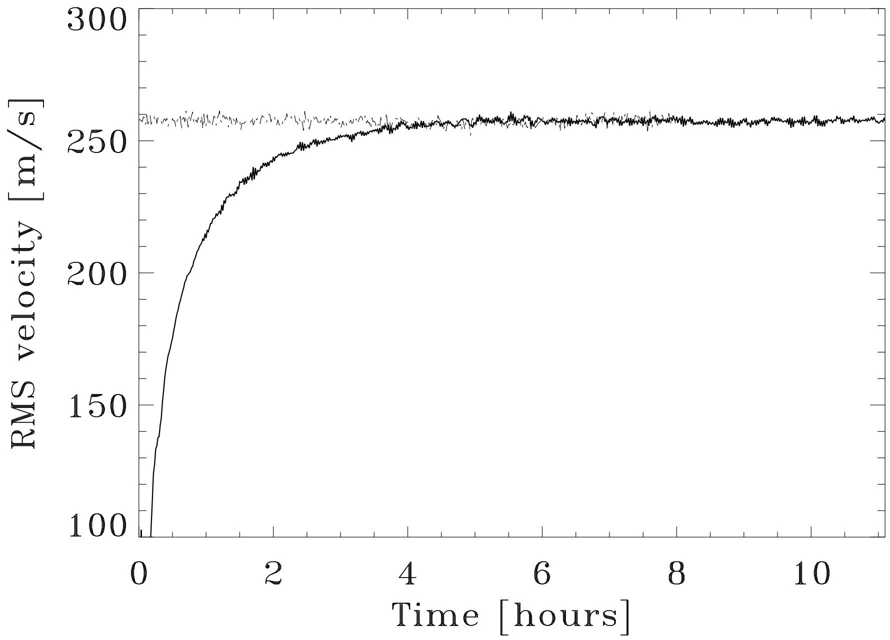

Figure 3.5 The spatial RMS velocity at the photosphere of a simulation (thick line) as compared to MDI high-resolution data (thin line) as a function of time. The system appears to achieve statistical stationarity 4.5 hours into the simulation. Because the system is linear, we can scale velocities by an arbitrary factor; in this case, velocities have been scaled so as to allow comparison with solar values.

3.4 Travel times with Surface Focusing

The time-distance helioseismic formalism introduced by Duvall et al. (1993) is based on the computation of cross-covariances between solar oscillation signals at two locations $\mathbf{r}_1$ and $\mathbf{r}_2$ on the solar surface ($z = 0$). Due to the stochastic excitation of acoustic waves (by convective motions in the Sun) and to the superposition of a large number of waves of different horizontal phase velocities $v_{ph} = \omega/k$, where k is the horizontal wavenumber and ω is the temporal angular frequency, the cross-covariances are very noisy and need to be phase-speed filtered and averaged (Duvall et al., 1997). The Doppler velocity datacube $\phi(\mathbf{r}, t)$ is phase-speed filtered in the Fourier domain using a Gaussian filter $F(k, \omega; \Delta)$ for each travel distance $\Delta = |\mathbf{r}_2 - \mathbf{r}_1|$:

$$F(k, \omega; \Delta) = \exp\left[-\frac{(\omega/k - v)^2}{2\delta v^2}\right], \tag{3.6}$$

where the central phase-speed v is derived from the solar model describing the oscillation power spectrum. For p modes in the ray approximation, v corresponds to ω/k at the lower turning point of a ray that traverses a horizontal distance Δ between successive reflections. The filter width δv is chosen empirically.

The standard method is then to average point-to-point cross-covariances over an annulus centered on $\mathbf{r_1}$ and with a radius Δ. Such point-to-annulus cross-covariances are computed for several distances Δ (55 in this chapter), and then averaged by groups of 5 distances to further increase their SNR. A detailed explanation of all the steps in the analysis process can be found in, e.g., Couvidat, Birch, & Kosovichev (2006). Table C.1 in appendix C.2 lists the distances Δ and phase-speed filter characteristics used here. These values are slightly different from the solar case because of the modified description of the background model.

The point-to-annulus cross-covariances are fitted by two Gabor wavelets (Kosovichev & Duvall, 1997): one each for the positive and negative times. To select the first-bounce ridge, we multiply the temporal cross-covariances by 14-minute wide rectangular window functions prior to the fit. The center t_0 of these windows is listed in Table C.1. The fitting procedure returns the ingoing (subscript i) and outgoing (subscript o) phase travel times $\tau_{\mathrm{i/o}}(\mathbf{r}, \Delta)$. The average of these two travel times, $\tau_{\mathrm{mean}}(\mathbf{r}, \Delta)$ is at first approximation sensitive only to the sound speed $c(\mathbf{r})$ in the region traversed by the wavepacket (see Eq. [3.7]), and similarly the difference $\tau_{\mathrm{diff}}(\mathbf{r}, \Delta)$ only to material flows. Note, by wavepacket, we mean a collection of waves of different frequencies and wavelengths that satisfy a specific dispersion relation, $\omega = \omega(k)$.

3.4.1 Results

We separately compute the acoustic wavefield in the presence of a sound-speed anomaly and the unperturbed counterpart with the same forcing function. Subsequently, we derive travel-time maps related to these datacubes by applying the method described

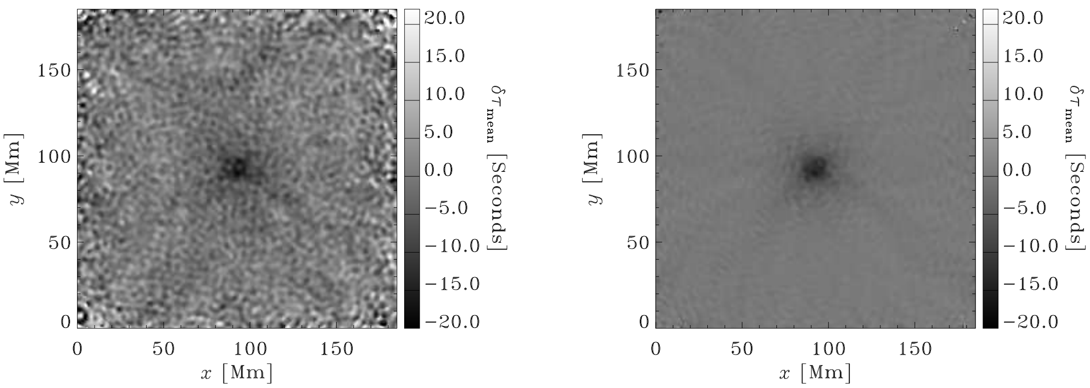

Figure 3.6 Example of mean travel-time perturbation map $\delta\tau_{\text{mean}}(\mathbf{r},\Delta)$ for $\Delta = 30.55$ Mm. Left panel: before quiet map subtraction. Right panel: after subtraction.

in §3.4. Here we shall focus on results related to a cylindrical sound-speed perturbation (see Figure 3.3) with a horizontal radius of about 6 Mm, a height of 12 Mm, centered at a depth of 10 Mm, and with a maximum amplitude of $\delta c^2/c^2 = 15\%$. The travel-time map corresponding to the distance $\Delta = 30.55$ Mm is shown on the left panel of Figure 3.6. On the right panel, the quiet travel-time map has been subtracted: the noise level is considerably reduced, as mentioned in §3.3. The travel-time map on the left panel looks very similar to maps computed for datacubes obtained from the Michelson Doppler Imager (MDI) instrument onboard the Solar and Heliospheric Observatory (SOHO) spacecraft, and because the noise levels between the artificial and real data are comparable, the choice of the excitation mechanism is somewhat justified. Figure 3.7 further shows the impact of the noise subtraction method on the travel-time maps.

In Figure 3.8, the dependence of the azimuthally averaged mean travel-time perturbations, $\delta\tau_{\text{mean}}(\mathbf{r},\Delta)$, on the radial distance from the center of the perturbation is shown by the thick solid line. The RMS variation σ of $\delta\tau_{\text{mean}}(\mathbf{r},\Delta)$ is an estimate of the uncertainty on this value (the error bars on Figure 3.8 are $\pm\sigma$). The thin solid line is the average of the difference travel-time perturbation $\delta\tau_{\text{diff}}(\mathbf{r},\Delta)$ and is close to zero, as expected from the absence of flows in our simulation.

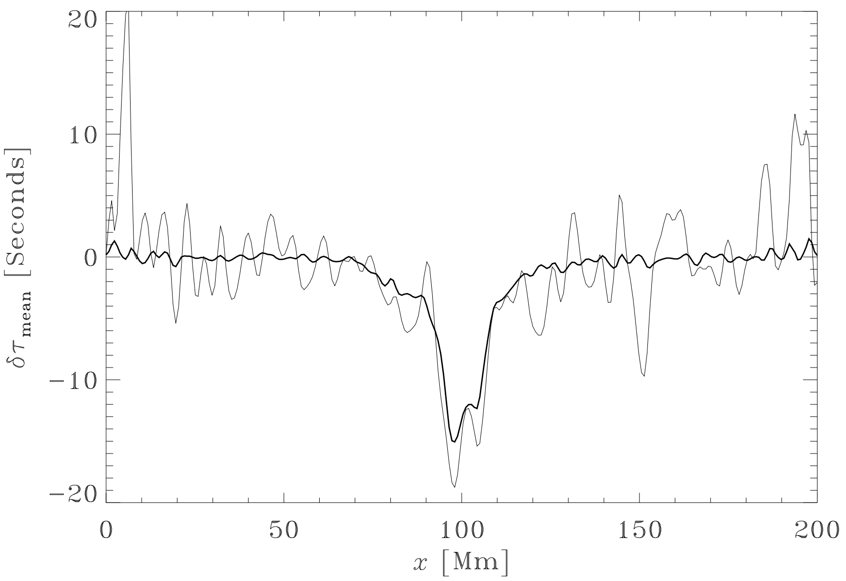

Figure 3.7 Cut across mean travel-time perturbation map $\delta\tau_{\mathrm{mean}}(\mathbf{r},\Delta)$ from Figure 3.6 at $y = 50$ Mm. Thin line: before quiet map subtraction, thick line: after subtraction.

3.4.2 Validation

A first test of the effectiveness of the numerical algorithm is to see if the simulated power spectrum looks reasonable in comparison to the solar modal spectrum. The spectrum shown in Figure 3.4(a) seems to satisfy this basic criterion. Secondly, through ray calculations, we can estimate the mean travel times associated with the background solar model. Although we do not show this here, comparisons between the ray predicted travel times and those obtained from the quiet simulations further fortify our confidence in the computational method. Next, using Born sensitivity kernels for sound-speed perturbations (Birch, Kosovichev & Duvall, 2004), we can derive the expected $\delta\tau_{\mathrm{mean}}(\mathbf{r},\Delta)$ (or $\delta\tau_{\mathrm{diff}}$) by evaluating the right-hand side of the following equation (forward problem):

$$\delta\tau_{\mathrm{mean/diff}}(\mathbf{r},\Delta) = \iint_S d\mathbf{r}' \int_{-d}^{0} dz\, K_{\mathrm{mean/diff}}(\mathbf{r}-\mathbf{r}', z; \Delta)\, \frac{\delta c^2}{c^2}(\mathbf{r}', z) \tag{3.7}$$

where S is the area of the Doppler velocity datacube, d is its depth, $K_{\text{mean/diff}}(\mathbf{r}-\mathbf{r}', z; \Delta)$ are the sensitivity kernels for distance Δ, and $\delta\tau(\mathbf{r}; \Delta)$ is defined as the travel-time shift at the center of the annulus (whose radius is Δ). The kernels we use to calculate the right hand side of equation (3.7) were computed for the standard solar model S (Christensen-Dalsgaard et al., 1996), implying that the comparison is mostly qualitative, since we use a slightly altered description of the near-surface layers (Hanasoge et al., 2006). These are point-to-point kernels, azimuthally averaged the same way as the cross-covariances to produce point-to-annulus kernels. The ingoing and outgoing point-to-annulus kernels are then either averaged or subtracted in relation to $\delta\tau_{\text{mean}}$ and $\delta\tau_{\text{diff}}$.

In Figure 3.8, the dashed line shows the azimuthal average of $\delta\tau_{\text{mean}}(\mathbf{r}, \Delta)$ obtained from equation (3.7) while the solid line is an azimuthal average of the mean travel-times obtained from the simulation. As expected, the two lines are not quite identical; this is perhaps due to small differences between phase-speed filter parameters and the background models in the two cases. Two independent methods were used in this comparison: the Born approximation to solve the forward problem, and the wavefield simulation in conjunction with a code to extract travel times from these calculations. Next, we check to see if the mean travel-time shifts computed with the ray approximation, shown as the dot-dash line in Figure 3.8, are comparable. As expected, ray theory, not accounting for wavefront healing (e.g., Hung, Dahlen, & Nolet, 2001), over-estimates these shifts but is still relatively accurate. It is gratifying to see that although the approaches in these methods differ greatly, there is still close agreement between the travel times. At this juncture, we consider the numerical method validated for the cases discussed here.

3.5 Kernels

Extending equation (3.7) to account for arbitrary perturbations, $\delta q(\mathbf{r}')$, we have:

$$\delta\tau_{\text{mean/diff}}(\mathbf{r}; \Delta) = \int\int_S d\mathbf{r}' \int_{-d}^{0} dz K_{\text{mean/diff}}(\mathbf{r}-\mathbf{r}', z; \Delta)\delta q(\mathbf{r}', z), \tag{3.8}$$

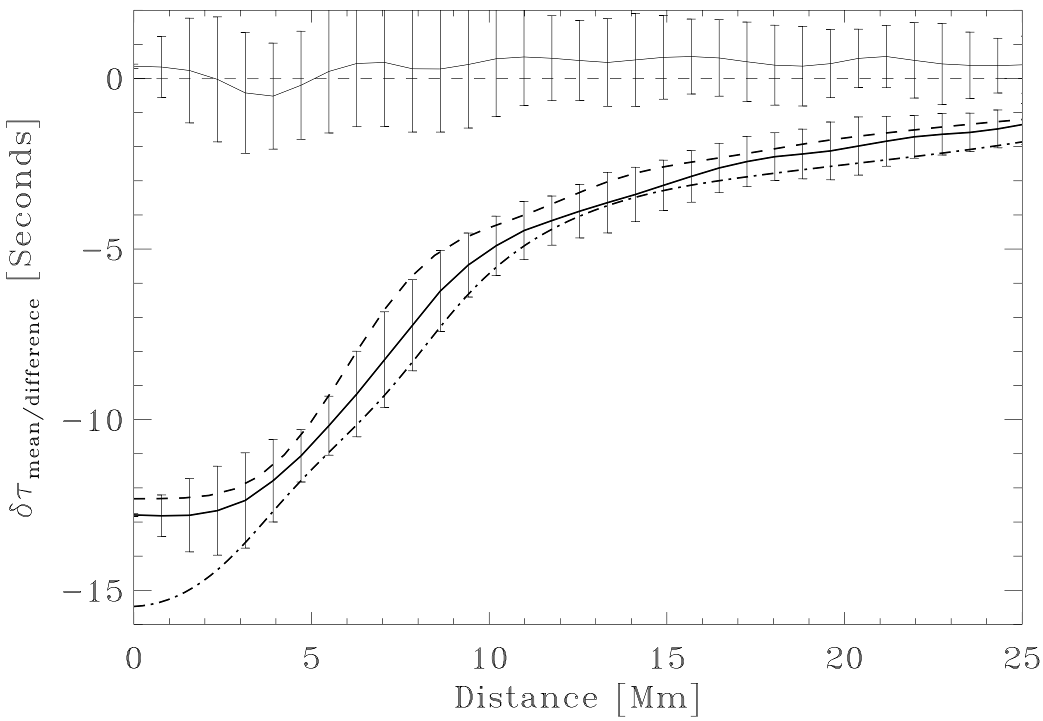

Figure 3.8 Azimuthal average of $\delta\tau_{\mathrm{mean}}(\mathbf{r},\Delta)$ (thick solid line) and $\delta\tau_{\mathrm{diff}}(\mathbf{r},\Delta)$ (thin solid line) for $\Delta = 30.55$ Mm, as a function of the radial distance to the sound-speed perturbation center (simulation). The dot-dash line is the mean travel-time shift computed with ray kernels for the background model in the simulation. The dashed lines are the solution to the forward problem using Born sensitivity kernels. The solar model in the simulation is slightly different from the model used to compute the Born kernels, contributing to the difference between the two travel-time curves.

where for future discussions, we adopt the notation of Duvall, Gizon & Birch (2006), namely that $\delta\tau(\mathbf{r};\Delta)$ is defined as the travel-time shift at position $\mathbf{r}$, located at the center of two observation points spaced distance Δ apart. With this change in definition of travel-times, we move from the center-to-annulus geometry of the previous section to a point-to-point description. There are numerous kinds of anomalies (changes in density, sound speed, pressure, source amplitude, magnetic fields to name a few) each of which is associated with a specific kernel. For perturbations that are spatial delta functions with magnitude M, of the form $\delta q(\mathbf{r}', z) = M\delta(\mathbf{r}')\delta(z - z_0)$, it may be verified from equation (3.8) that the convolution on the right hand side reduces to $(2\pi)^3 MK(\mathbf{r}, z_0; \Delta)$ (the 2π factors are dependent on the definition of the δ function), the conclusion being that the travel-time shifts and the kernel are identical to within a proportionality constant in this limit. Duvall, Gizon & Birch (2006) incorporated this approximation to derive a kernel from observations of thousands of thin magnetic elements on the Sun, justifying it through the argument that the magnetic flux tubes were all much smaller than the wavelength of the f modes used in the analysis. We apply the same technique to derive (1) a p-mode sound-speed kernel and (2) an f-mode source kernel from a thousand randomly placed small (1 Mm in size) source suppressions (see Figure 3.9) for the solar model and setup in our computations.

We choose this set of kernels to study because Gizon & Birch (2002) and Birch, Kosovichev & Duvall (2004) have constructed f-mode source disturbance and p-mode sound-speed perturbation kernels respectively, allowing us the luxury of comparison. In the source perturbation kernel of Gizon & Birch (2002), elliptical features are absent and only hyperbolic features associated with multiple sources are seen; the sound-speed kernels of Birch, Kosovichev & Duvall (2004) contain a mixture of both elliptical and hyperbolic ridges. The ability to recover a source kernel from the simulation of Figure 3.9 is important because we demonstrate proof of concept of the kernel measurement method devised by Duvall, Gizon & Birch (2006). The balance between the randomness of the locations of these elements and the determinism of the sense of the perturbation (all sources are reduced) seems to result in an average structure in the travel times that looks remarkably like the source perturbation kernel from Gizon & Birch (2002), as seen in panels (a) through (d) of Figure 3.5.

The technique through which the SNR of data is improved by enough that it becomes possible to see the travel times associated with these small features is described in detail by Duvall, Gizon & Birch (2006). Firstly, the data is phase-speed filtered to either restrict the acoustic spectrum to waves that travel a given distance, Δ, between bounces or to isolate the f-mode ridge. For all points $\mathbf{r} = (x, y)$ in the domain of interest (Eq. [3.8]), the temporal cross-correlation is obtained by inverse Fourier transforming $\tilde{C}$, where

$$\tilde{C}(\omega, x + \delta x/2, y + \delta y/2) = \tilde{v}^*(\omega, x, y)\tilde{v}(\omega, x + \delta x, y + \delta y), \tag{3.9}$$

which is subsequently fitted to obtain travel times. In equation (3.9), ω is the frequency, $v(t, x, y)$ is the velocity signal at time t and spatial location (x, y), with $\tilde{v}(\omega, x, y)$ representing the Fourier transform of $v(t, x, y)$, $\tilde{v}^*$ the complex conjugate of $\tilde{v}$, and $\arctan(\delta \mathrm{y}/\delta \mathrm{x})$ the orientation of the travel-time map. Within the limits of spatial resolution, there were 77 possible orientations for the simulation that contained the sound-speed perturbation (120×120 Mm^2 in the (x, y) plane, resolved with 144^2 points) and 83 orientations for the source perturbation case (200×200 Mm^2 in the (x, y) plane, resolved with 512^2 points). Subsequently, the resulting two-point travel times were de-rotated by an angle $\arctan(\delta \mathrm{y}/\delta \mathrm{x})$ and averaged. For the simulation with the 1000 suppressions, the image is not only de-rotated but averaged around each feature, corresponding to a net total of 83,000 averages. The SNR is still not high enough with this degree of averaging; only after the travel times of the quiet datacube, obtained through the same averaging process, are subtracted, do the hyperbolae of the kernel become clearly visible. A similar averaging algorithm is used to produce the sound-speed kernel (Figure 3.12), albeit the wavefield was simulated in the presence of only one anomaly.

3.5.1 Source Kernels

Although the magnitudes and the hyperbolic features are somewhat different, kernels obtained from simulation and theory (Figure 3.5), possess a strong resemblance. It is

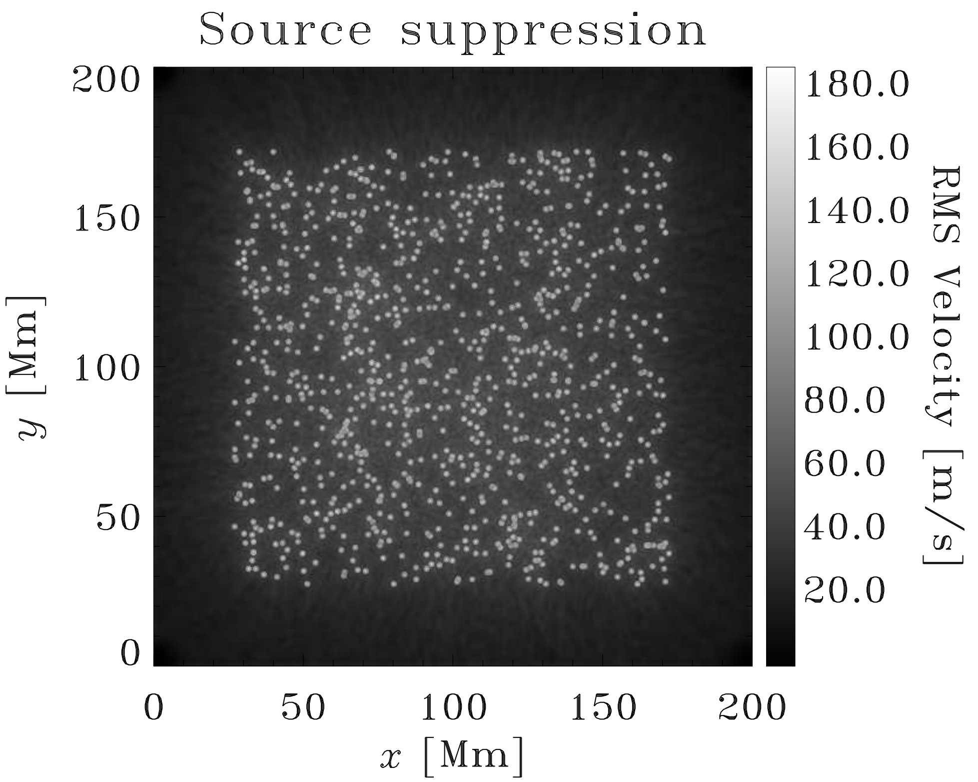

Figure 3.9 The time averaged RMS of the difference in velocities from a simulation with one thousand randomly located 1 Mm-sized source suppressions and its unperturbed counterpart. In the vicinity of a source suppression, acoustic velocities are altered; therefore, the hundreds of local maxima in this difference map show where the sources are diminished. The perturbations are approximately limited to a 150×150 Mm^2 interior square, allowing for travel-time shifts associated with all the anomalies to be computed.

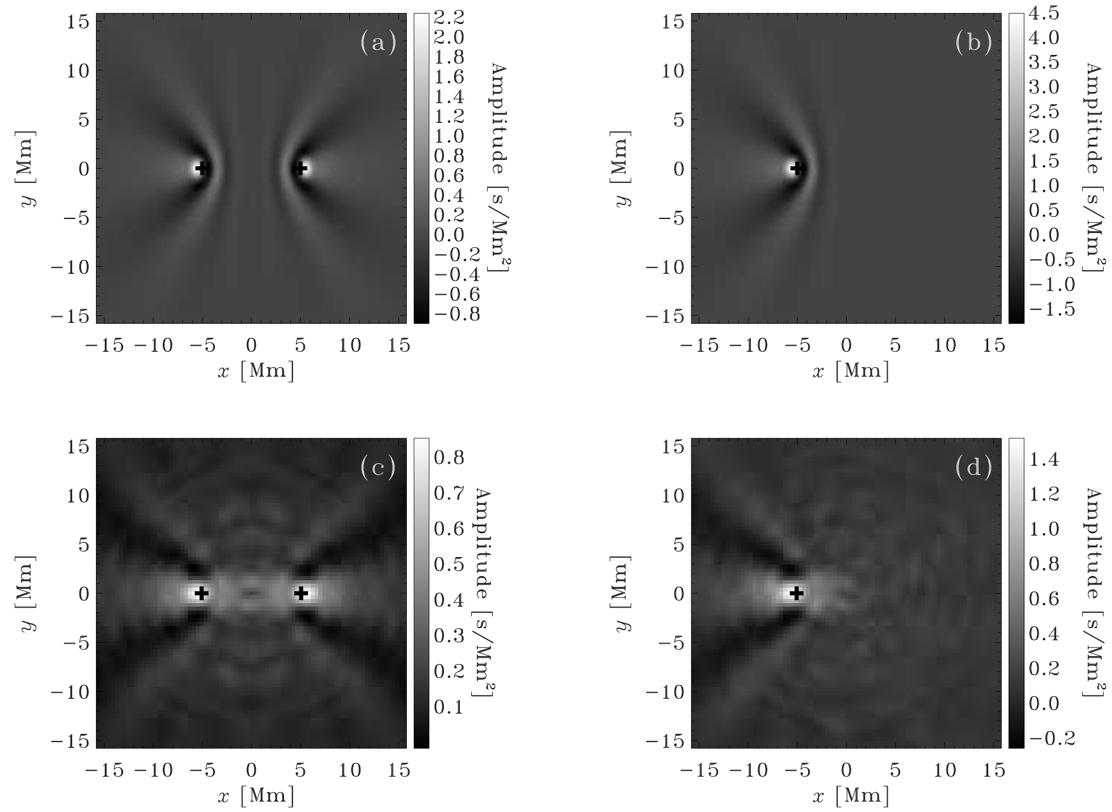

Figure 3.10 Source perturbations kernels for the f-mode. The panels (a) and (b) show the mean and left to right one-way kernels from Gizon & Birch (2002) respectively. Crosses show the measurement points. Upon systematically de-rotating and averaging the travel-time shifts associated with the thousand locally deactivated sources from Figure 3.9, subtracting the noise and scaling the travel times as described in appendix C.1, the kernels in panels (c) and (d) emerge. The lack of elliptical features in these kernels is perhaps explained by the independence of the sources from the wavefield, in the theory of Gizon & Birch (2002) and in the simulations. Panel (d) is especially compelling because according to theory, left to right one-way travel-times should only be shifted for those points which lie closer to the left measurement point. The power profile of the f-ridge used to recover the travel-time shifts is shown in Figure 3.1.

seen from Figure 3.1 that the f-mode power distribution is able to match the observations (and the excitation model of Gizon & Birch (2002)) quite well, so frequency effects are not significant contributors to the differences in the hyperbolic structures. One possible explanation could be the damping, included in theory but not in the simulations. The absence of solar-like damping has two important effects: (1) mode power distribution and line-widths differ from the solar counterparts and (2) waves are correlated across larger distances. The latter implies that travel-time shifts of waves remain coherent over longer length scales than in the Sun or the model of e.g., Gizon & Birch (2002). The other issue concerns the differing magnitudes of the kernels; perhaps the method described in appendix C.1 to scale the travel times is only approximate. Moreover, the theoretical kernel of Gizon & Birch (2002) has a spatial integral of zero, while the kernel obtained from the simulation has a positive bias of 0.5 seconds (calculated as the mean travel-time shift of a small off-center region).

The travel-time shifts observed in this case are due to reductions in the cross correlation amplitude because of local depressions in source strengths. The cross correlations are biased in the sense of decreased amplitude because all the sources have been suppressed, possibly leading to a non-zero value of the integral of the kernel. In an attempt to correct for this, we tried another case where approximately half the perturbed sources were suppressed and the other half amplified. To obtain a meaningful average kernel (shown in Figure 3.11), the sign of the travel times for the source suppressions was flipped while the source amplification counterparts were left untouched. In this case, the integral of the kernel is much smaller than the amplitude, indicating that the bias in the kernel of Figure 3.5(c) is most probably caused by the systematic sense in the perturbation (suppression).

3.5.2 Sound-Speed Kernels

The sound-speed kernel of Figure 3.12(a) was derived for the same measurement distance (12.4 Mm) as the one in Figure 3.12(b) (reproduced from the lower right panel of Figure 10 from Birch, Kosovichev & Duvall (2004)). Although there are many differences in the approaches, not the least being the background model and

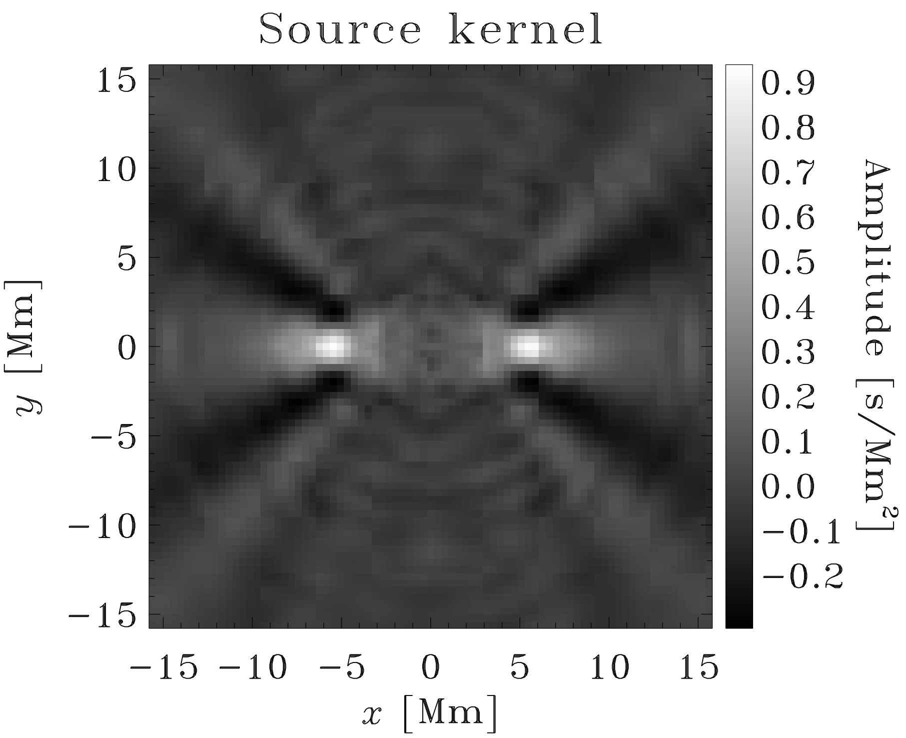

Figure 3.11 The source kernel obtained by the suppressing approximately half the sources and doubling the strengths of the rest. The integral of the kernel is an order of magnitude smaller than the amplitude of the kernel.

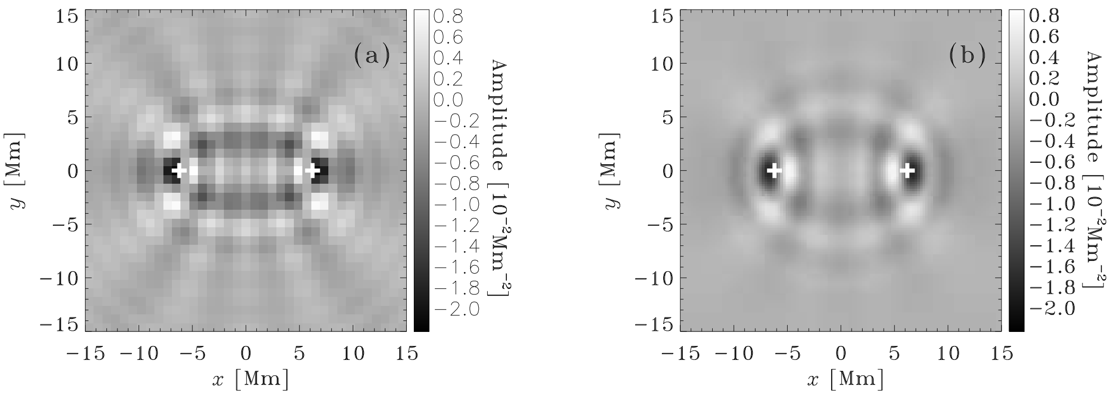

Figure 3.12 Panel (a): sound-speed kernel from 24 hours of simulated data (12 perturbed + 12 quiet). Panel (b): sound-speed kernel reproduced from the lower-right panel of Figure (10) of Birch, Kosovichev & Duvall (2004). The crosses mark the measurement points. The kernels have been multiplied by the sound speed at the photospheric level in both cases; units are in $10^{-2}\mathrm{Mm}^{-2}$. The agreement is striking. The circular features in panel (a) are generated by the repeated de-rotation of travel-time shifts and hence are mainly noise. Both hyperbolic and elliptic features are visible in this kernel. The power profile of the modes utilized to construct the kernel in panel (a) is shown in Figure 3.1(b).

damping rates, the agreement is excellent. This result also illustrates the ability of the method of noise subtraction to remove the noise without affecting the signal itself. Although not shown here, we have noticed that as the central frequency of the wavepacket used to construct the kernel increases, the ellipses become even more 'elliptical', an indicator of ray-like behavior. In general, there is a definite dependence of the shape of the kernel on the filter parameters.

3.6 Conclusions

A numerical method to compute the 3D wavefield in a solar-like medium was discussed and implemented. The concept of noise subtraction, a technique whereby the realization noise can be significantly reduced was introduced. This method is quite useful in reducing computational cost by making it possible to achieve large SNR even with short simulations. Results from the simulations were validated using comparisons between the travel times obtained (1) from the surface focusing method of time-distance helioseismology, (2) through the application of approximate Born sensitivity kernels, and (3) ray theory. The results agree rather well, showing the validity of these independent approaches for this particular situation.

We ran a simulation with a thousand randomly located suppressed sources which subsequently was analyzed by the feature method of Duvall, Gizon & Birch (2006) to obtain travel-time shifts associated with the average diminished source. The result closely resembles the source perturbation kernel of Gizon & Birch (2002). A sound-speed kernel for parameters similar to a case considered by Birch, Kosovichev & Duvall (2004) was derived. The agreement between Born theory and simulations in conjunction with time-distance helioseismology is impressive. It lends greater confidence in the technique of direct measurements of kernels from data (Duvall, Gizon & Birch, 2006) and the method of noise subtraction. It is interesting to note that the structure of the kernels is quite sensitive to the relevant frequency bandpass; high-frequency wave kernels have pronounced elliptic features displaying ray-like behavior. The source and sound-speed kernels shown here contain aspects of multiple scattering and in general, are not bound by the assumptions of the Born approximation. We can also derive sensitivity kernels for various types of perturbations like anomalies in density, pressure, Γ_1 etc., that are relatively difficult to compute in the Born limit.

The price we currently pay for the ability to perform these simulations is reduced realism. Damping, not accounted for in our computations, is an important contributor, affecting travel times (e.g., Woodard, 1997) and in general changing the distribution of modal power and the linewidths. A further approximation in our simulations is the use of an altered background state; the standard solar model S makes it difficult to perform linear simulations due to the inherent convective instability of the near-surface layers. It is important for sophisticated forward models to be able to incorporate damping and an accurate solar model.

The effects of diffractive or wavefront healing, a hotly contested phenomenon in geophysics, are currently being investigated in the context of helioseismology through these simulations. We will be able to place bounds on detectability, the accuracy of inversions and study vector effects like flows. It will also be very exciting to disentangle magnetic field effects and determine if we have a reasonable understanding of the interior structure of sunspots.