CHINESE JOURNAL OF PHYSICS VOL. 49 , NO. 1 FEBRUARY 2011

Mapping Caustic Positions: Efficient Optimization of Elastic Constant Values in Non-cubic Crystals

Madeleine Msall[1,*] and Timothy L. Head[2]

[1]*Department of Physics and Astronomy, Bowdoin College, Brunswick, ME USA*
[2]*Department of Physics, Abilene Christian University, Abilene, TX USA*

(Received April 12, 2010)

Phonon images can be easily simulated if the elastic constants are known. The inverse problem of predicting the elastic constants on the basis of known experimental caustic positions or group velocities is less tractable. We explore methods of determining the elastic constants for crystals of low (non-cubic) symmetry with image matching techniques.

PACS numbers: 62.20.de; 66.70.-f; 62.30+d; 81.40.Jj

I. INTRODUCTION

Literature values for the elastic constants are often extracted from ultrasound measurements of the phase speed along high symmetry directions. These experiments involve plane acoustic waves that travel in the wave vector direction. Phonon imaging techniques give access to a wide range of information about propagation along low-symmetry directions, but do not involve plane wave excitation. The localized excitation in imaging experiments excites nonequilibrium phonons that carry energy along the group velocity direction, which is not parallel to the wave vector direction except in cases of high symmetry. Phonon images map the sharp boundaries between high- and low-flux regions, called *caustics*, which are directly related to folds in the acoustic wave surface. The caustic positions reflect the underlying crystalline symmetry and are very sensitive to small changes in the elastic constants. Thus, the map of caustic positions, know as a *phonon image*, is a sensitive indicator of the elastic constant values.

Fig. 1a shows a phonon image of [010] $CaWO_4$, made using standard techniques [1, 2]. In these measurements, phonons are generated by pulsed photoexcitation of a copper film deposited on a surface parallel to the image plane. The sample is held at a low temperature ($\sim$2 K) so that many phonons travel ballistically through the sample. Nonequilibrium phonon pulses are detected at the opposite surface with a superconducting aluminum bolometer. Raster scanning of the laser position allows us to measure changes in the phonon flux through the sample as a function of the propagation angle. The integrated flux for each raster position is shown as a brightness level on the phonon image.

Simulated images are generated in 2 steps: 1) Phonon group velocity is calculated

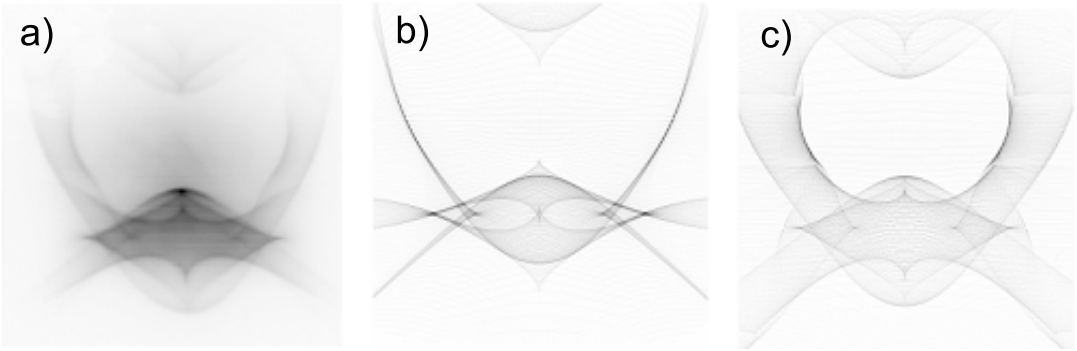

FIG. 1: [010] face of CaWO4 a) Experimental phonon image; b) simulated image using elastic constants of Gluyas et al.; c) simulated image using elastic constants of Farley et al. [5, 6].

for an isotropic distribution [9] of phonon wavevectors by using elasticity theory with a particular model set of elastic constants; 2) Phonon trajectories for these wavevectors are then calculated and their intersection with the detection surface mapped [4]. The resultant image displays the flux arriving at the detection surface for a fixed excitation point, but this is symmetrically equivalent to the experimental configuration of a scanned excitation point and fixed detector. Figs. 1b and 1c show the dramatic effect of changing the values of the elastic constants on the simulated images.

In an experimental phonon image, the sharpness of the flux pattern is influenced by the measurement parameters (e.g., source and detector sizes) and the number and type of scattering processes the phonons undergo enroute to the detector. The human eye is extremely good at identifying transition regions, even when broadened by scattering, and can quickly decide whether the caustic positions in a given experimental and simulated image are well-matched. The widths and shapes of particular features are quickly processed to form a judgement on the quality of the match. An initial response to the data in Fig. 1 is that neither simulation based on previous literature values of the low-temperature elastic constants completely reproduces the experimental image features but that more of the overall structure is reproduced using the Farley values.

In contrast, the differences between the simulated and experimental group velocities, shown in Fig. 2, are smaller and change less dramatically when the elastic constants are altered. Fig. 2 includes the uncertainty of the experimental measurement due to the finite temporal and spatial resolution of the source and detectors, but there is also some systematic difficulty in identifying the ballistic arrival time because of the pulse broadening from scattering [2]. It is possible to attempt to optimize the agreement between experimental

[9] For our well-bonded and moderately rough surface, the acoustic power transmitted from the metal film to the $CaWO_4$ should have a Lamberts law cosine dependence. Since our matching technique depends upon the identification of sharp features, the slow variation due to this cosine correction can be neglected. In cases of large acoustic mismatch and a loosely bonded film resulting in experimentally observable critical cone channeling, a modification of the isotropic assumption will be necessary. [3]

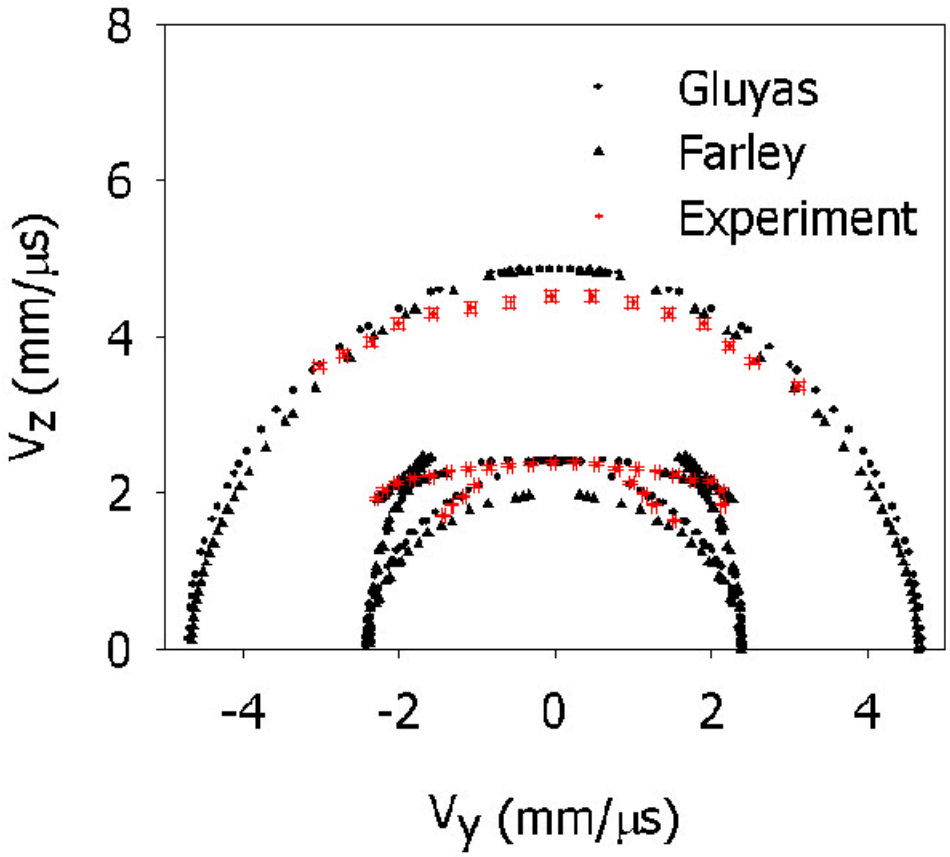

FIG. 2: Values of CaWO4 group velocity along the line of mirror symmetry in Fig. 1 ($v_x = 0$).

and theoretical group velocity values, but only in regions away from caustics where there is a relative lack of sensitivity to small changes [7]. This, in addition to the systematic question of how to choose the ballistic arrival time, makes group velocity comparisons less attractive.

Instead, we seek an automated mechanism for judging the quality of the match between experimental and simulated images. When studying a crystal like $CaWO_4$, which has tetragonal symmetry and 7 independent elastic constants, optimizing the image match is impossible without some automated comparison. The problem of image matching is well known in computer science, with very active fields of research in computer vision, feature recognition, and image identification. In most cases, the match process must take into account geometrical distortion and limited visibility. Both of these issues may be present in our experimental data, but we assume that the primary challenge will be from the latter as the contrast in experimental phonon images is often markedly less than in simulations.

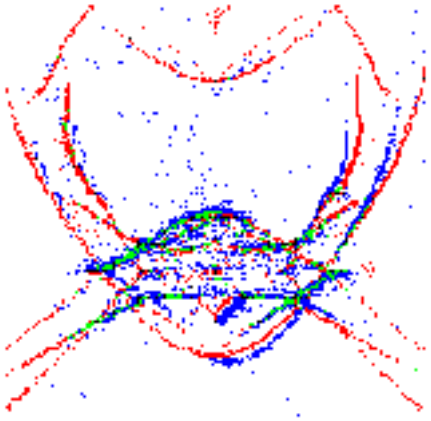

FIG. 3: An example of the binary map comparison of the experimental data (Fig. 1a) and a simulated image (from the first set of elastic constants in Table I). The blue points are points which pass the threshold test for the experimental data only; the red points for the simulated data only; and the green points indicate image points where both experimental and simulated data pass the threshold tests.

The issue of geometrical distortion should be minimized by careful calibration of the experimental scanning distances and thorough knowledge of the orientation of the image plane.

II. AUTOMATIC IDENTIFICATION OF CAUSTIC POSITIONS

Direct visual identification of caustic positions takes advantage of the fast pattern recognition of the human brain, but it is subjective and imprecise. Computational variations of the visual comparison method extract numerical values of the caustic positions in an experimental image and compare those positions to predicted caustic positions obtained, e.g., from calculations of the directions where the Jacobian vanishes or from simulated images. The challenges are then to develop efficient root finding routines for the Jacobian or to develop robust methods for the extraction of caustic positions from experimental and theoretical images with finite resolution and to develop a metric for judging relative positions.

The first step in identifying caustic positions from images is to identify regions of rapid change from high to low flux. This is done by calculating the local derivative of each image point from the average of the differences between the flux value at that point and its nearest neighbors. A binary map (Fig. 3) is then created where all points on the image at which the local derivative exceeds a threshold value are marked with a 1 and all other points with a 0. This operation is analogous to edge or ridge detection in image analysis. The selection of an appropriate threshold determines the scale of features that will be included in the analysis. Alternately, if Jacobian root finding is used to predict the caustic directions, a binary map, which marks the image position for those directions along

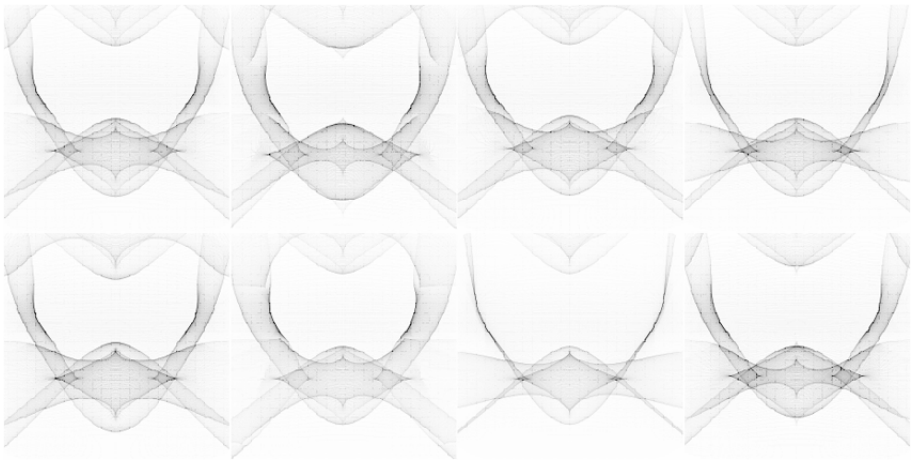

FIG. 4: Series of CaWO4 simulations with individual elastic constants increased by 10%. The first image in the upper right has a fixed constant set. Subsequent images have a single 10% change in the value of a single constant compared to the first image. Working from left to right and top to bottom, the changed constants are $C_{11}, C_{12}, C_{13}, C_{16}, C_{33}, C_{44}$, and C_{66}.

which the value of the Jacobian is below a minimum value, can be created.

The simplest method for quantifying the match between experimental and simulated binary maps is to maximize the number of positions at which both the experimental and the simulated maps have the same value. However, in order to take into account small offsets in the 2 images, it may be useful to check the neighboring points as well. In this case, one can calculate the distance to the closest point in the comparison image that has a corresponding value to the experimental image and try to minimize this "match distance" over the entire image. Fig. 3 illustrates the overlap of an experimental and simulated image. The level of noise in the binary map of the experimental data is determined by the choice of threshold for the local derivative. This value needs to be low in order to preserve some of the lower contrast experimental features, leading to some spurious points in the noncaustic regions with high scattered flux. The binary map of the simulated data, on the other hand, can be produced with a higher threshold and tracks a larger number of caustic features.

III. SIMULATIONS

The computation time for a single phonon image depends upon the density of the grid of k-space points used in the simulation. The images in Fig. 1, for example, are based on a look-up table of group velocities with 720,000 points and are produced in under 30 seconds on our Dell computers [8]. A lower resolution look-up table with 180,000 points produces an image with all the caustic features but shows a distinct line pattern from the simulation grid. This line pattern can be removed by choosing an appropriate threshold level in the binary map, and so it is effective and efficient to use the lower resolution look-up

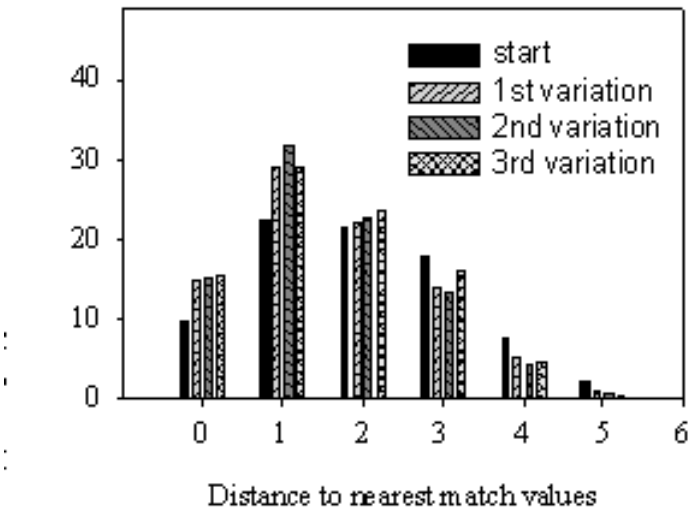

FIG. 5: Histogram of distance to nearest match point as constants are refined.

table when comparing multiple simulations to the experiment.

Many of the caustic features are determined by multiple elastic constants in concert so that the appearance of a feature can be affected by small changes of more than one constant. In our $CaWO_4$ example, the width of the curved side caustics can be changed dramatically by an increase of any one of 4 separate elastic constants, as seen in Fig. 4. Thus, it is not sufficient to minimize the match distance between experimental and simulated images by varying a single constant and then proceeding sequentially to the next. The entire 7-dimensional space must be checked for global minimization. Our first attempts looked at 4 variations of each of the 7 constants, requiring the computation of the match distance for $4^7 = 16{,}384$ images. Using the lower resolution simulated images and minimal output of the image files, this set of comparisons can be completed in $\sim$10 h. The match calculation has linear run time, and so an increase to 5 variations for each constant would increase the run time by 1.25^7 ($\sim$5) times. If, however, we can constrain the search by fixing even one of the constants, we can improve the run time by a factor equal to the number of images desired.

For $CaWO_4$, the transverse group velocity along (010) depends only upon the value of C_{44} and the material density ($v = [C_{44}/\rho]^{1/2}$). We choose a value of $C_{44} = 35.0$ GPa in good agreement with our measured group velocity and previous literature values [2, 5, 6]. Once C_{44} is fixed, we optimize the image match based on variations in the remaining 6 constants. We find that we can use either minimization of the match distance or maximization of the number of exactly overlapping points with equivalent results. We have used minimization of the match distance in the results presented here, but note that in larger searches, the

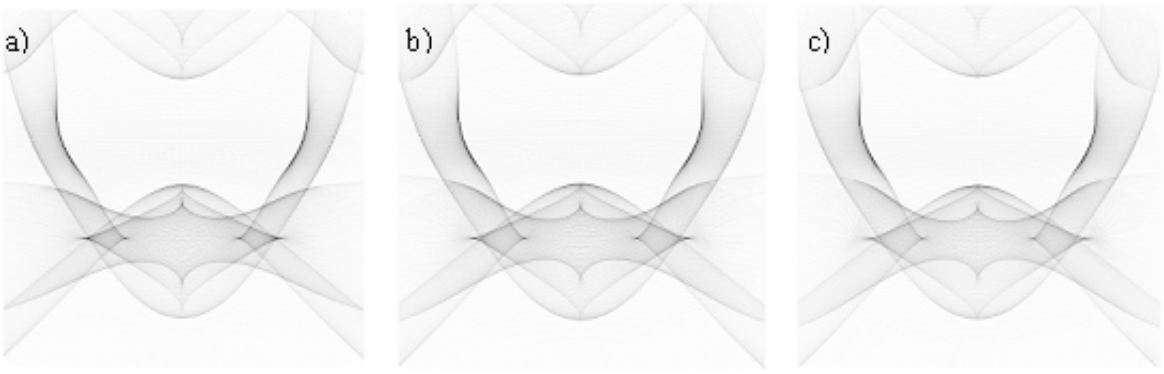

FIG. 6: Progressively refined images to match the experimental image of Fig. 1a. a) 5–10% variation, b) 2–4% variation, and c) 1–2% variation.

calculation of overlap is slightly more efficient.

IV. DISCUSSION AND CONCLUSIONS

An important factor in any minimization problem is the question of whether the solution converges on a single global minimum. Given the dependence of single image features on a constellation of constants, there are many local minima encountered in the search. Thus, the choices of initial constant values and the size of variations are critical to the search outcome. In the case of $CaWO_4$, initial constant values were derived by optimizing the match by varying each of the single constants (beginning with their average literature values) in turn. After this starting set was obtained, the first multidimensional search allowed a variation of $\pm 5\%$ and $\pm 10\%$ in each constant. The best matching values were then used to seed a search that allowed a variation of $\pm 2\%$ and $\pm 4\%$ in each constant; the results then seeded a search that allowed a variation of $\pm 1\%$ and $\pm 2\%$ in each constant. Table I shows the evolution of the elastic constants throughout this progressive process. Fig. 5 shows the increase in the fraction of closely matched points at each step.

The steady improvement in the average match distance as the magnitude of the variations is decreased suggests that the program can locate a minimum. Fig. 6 shows the (relatively small) changes in the phonon image as the match is refined. At present, however, the direct search requires a good choice of initial elastic constants. A poor choice of the initial constants can cause the search procedure to favor matches to a limited subset of features at the expense of the overall pattern. The situation could be greatly improved by the development of heuristic or stochastic methods that avoid trapping in local minima. Such a simulation set would effectively solve the problem of extracting reliable elastic constant values directly from a single well-calibrated phonon image.

TABLE I: Sequence of refinement of elastic constant values.

Variation	C11	C12	C13	C16	C33	C44	C66
Single constants	151.7	65.1	45.0	18.8	130.0	35.0	40.0
6constants	143.9	58.6	47.3	19.7	136.5	35.0	40.0
±5–10%	(-5%)	(-5%)	(+5%)	(+5%)	(+5%)	(fixed)	(none)
6 constants	146.8	57.4	45.4	20.1	136.5	35.0	40.8
±2–4%	(+2%)	(-2%)	(-4%)	(+2%)	(none)	(fixed)	(+2%)
6 constants	148.3	57.4	45.4	19.9	137.9	35.0	40.4
±1–2%	(+1%)	(none)	(none)	(-1%)	(+1%)	(fixed)	(-1%)

References

* Electronic address: mmsall@bowdoin.edu

[1] M.E. Msall, T.L. Head, and J.S. Jumper, submitted to Solid State Communications (2010).
[2] T.L. Head, M. E. Msall and D. S. Jumper, submitted to Chinese Journal of Physics (2010).
[3] A.G. Every, G.L. Koos and J.P.Wolfe, Phys. Rev. B **29**, 2190 (1984).
[4] G.A. Northrop, Comp. Phys. Comm. **28**, 103 (1982).
[5] J. M. Farley, and G. A. Saunders, J. Phys. C **5**, 3021 (1972).
[6] M. Gluyas, F. D. Hughes, and B. W. James, J. Phys. D **6**, 2025 (1973).
[7] A.G. Every and W. Sachse, Phys. Rev. B 42, 8196 (1990).
[8] The computers have 3.0 GHz Pentium D processors and 2 GB of RAM running CentOS 5.4 using a Linux Kernel of 2.6.18.