CARS Spectral Fitting with Multiple Resonant Species Using Sparse Libraries

Andrew. D. Cutler[*] and Gaetano Magnotti[†]

The George Washington University, Mechanical and Aerospace Engineering,
1 Old Oyster Point Road, Suite 200, Newport News, VA, 23602

Abstract

The dual pump CARS technique is often used in the study of turbulent flames. Fast and accurate algorithms are needed for fitting dual-pump CARS spectra for temperature and multiple chemical species. This paper describes the development of such an algorithm. The algorithm employs sparse libraries, whose size grows much more slowly with number of species than a conventional library. The method was demonstrated by fitting synthetic "experimental" spectra containing 4 resonant species (N_2, O_2, H_2 and CO_2), both with noise and without it, and by fitting experimental spectra from a H_2-air flame produced by a Hencken burner. In both studies, weighted least squares fitting of signal, as opposed to least squares fitting signal or square-root signal, was shown to produce the least random error and minimize bias error in the fitted parameters.

I.   Introduction

Coherent anti-Stokes Raman spectroscopy (CARS) is a non-linear spectroscopic technique in which three laser beams, two pump beams and a Stokes beam, are focused and crossed at their focal point, and a fourth signal laser beam is generated at the intersection through a four-wave mixing process. Resonances associated with Raman active molecular rotational-vibrational transitions strengthen the CARS signal in a manner dependent on gas composition and temperature. The wavelength of excitation or "Raman shift" corresponds to the difference in energy between a pump beam and a Stokes beam (longer wavelength) photon. In the broadband CARS technique one of the laser beams, the Stokes beam, is spectrally broad whereas the other pump beams are spectrally narrow, and this enables multiple resonances to be excited. The signal beam is spectrally broad and carries with it the spectral signature of the gases present (over some range of wavelengths). This signal beam is dispersed by a spectrometer and spectra are compared to analytical models to obtain information on composition and temperature. In the dual-pump CARS technique, originally developed by Robert Lucht and coworkers,[1,2] the two pump beams are at different wavelengths and two different ranges of Raman shift may be accessed.

In recent years, efforts have been made to utilize the CARS technique to obtain experimental data for development and validation of computational models of supersonic combustion. Data sets have been acquired in a hydrogen-fueled supersonic combustor using the CARS technique[3] to measure temperature, and the dual-pump CARS technique has been used to measure temperature and composition;[4] in Ref. 3 the resonant species was N_2 and in Ref. 4 they were N_2, O_2, and H_2. Both mean flow and turbulence statistics (variances and covariances) were derived from the data, although the uncertainty in the latter was high due both to instrument error and to the small number of measurements on which to base the statistics. More recently, data have been acquired in a hydrogen-combustion heated supersonic coaxial jet with the intent to obtain turbulence statistics of lower uncertainty.[5] These efforts required analysis of large numbers of experimental spectra containing up to three different species.

Analysis of dual-pump CARS experimental spectra is typically performed with the Sandia CARSFT code as modified by Lucht[1,2,6] and O'Byrne.[4] The CARSFT code is able to accurately compute spectra for a range of molecular species and, in the fitting mode, to read an experimental spectrum and fit it to a theoretical spectrum by repeatedly computing spectra and minimizing the residual between the theoretical and experimental spectra. This fitting process becomes very slow when the number of molecular species increases, due to a large spectral range, many energy level transitions, and a large number of iterations. Additionally, when fitting multiple variables, CARSFT has problems in finding a global minimum, sometimes returning a result dependent upon the initial guess. At NASA Langley, with the exception of Tedder et al.,[5] CARS data has usually been fitted using libraries of spectra.[3,4] Libraries are computed ahead of time using CARSFT. Theoretical spectra are then interpolated from the library rather than being calculated from scratch. However, conventional library size scales to the power of the number of parameters varied (temperature plus resonant chemical species) and tends to become uneconomically

---

[*] Professor. Associate Fellow AIAA
[†] Graduate student. Member AIAA

large for more than three fitted species. Recently, dual-pump broadband CARS techniques have been developed that employ broadband dye lasers of larger spectral width and simultaneous detection of up to 6 resonant species was demonstrated;[7] such spectra could not be fit using conventional library methods and are difficult if not impossible to fit using CARSFT in its fitting mode.

Recently, we have developed fitting algorithms that use sparse libraries (as opposed to the "fully-packed" libraries used previously) to enable fitting of a larger number of resonant chemical species and rapid fitting of large experimental data bases acquired in combustion turbulence studies. The objective of the current work is to describe these techniques and their application to fitting dual-pump CARS spectra containing N_2, O_2, H_2 and CO_2.

II.   Method

The fitting method employs a library of theoretical spectra calculated separately, using CARSFT or some other similar code. Theoretical spectra are interpolated from the library at each iteration of the fitting algorithm. The algorithm finds the parameters of the fit, which are the temperature, mole fractions of resonant species, and a horizontal wavelength shift that minimize a residual function between the experimental and theoretical spectra.

A Library Structure

Previously, libraries used in fitting CARS spectra were fully populated on a regular grid of temperature and species mole fractions, with size (number of spectra) that scaled with the power of the number of resonant chemical species (e.g., Ref. 4). The new method takes advantage of the fact that the resonant spectra of the different species generally do not overlap very much. (Indeed, when doing dual-pump CARS the second pump frequency offset is typically chosen so as to minimize overlap.) The number of points in the sparse grid is effectively reduced compared to the fully populated grid, while interpolation from the sparse library is of comparable accuracy to interpolation from the fully packed library, except at wavelengths where resonances of multiple species overlap. The library is only sparse in the species concentration coordinate directions; that is to say, the sparse grid of species concentrations is fully repeated at each library temperature.
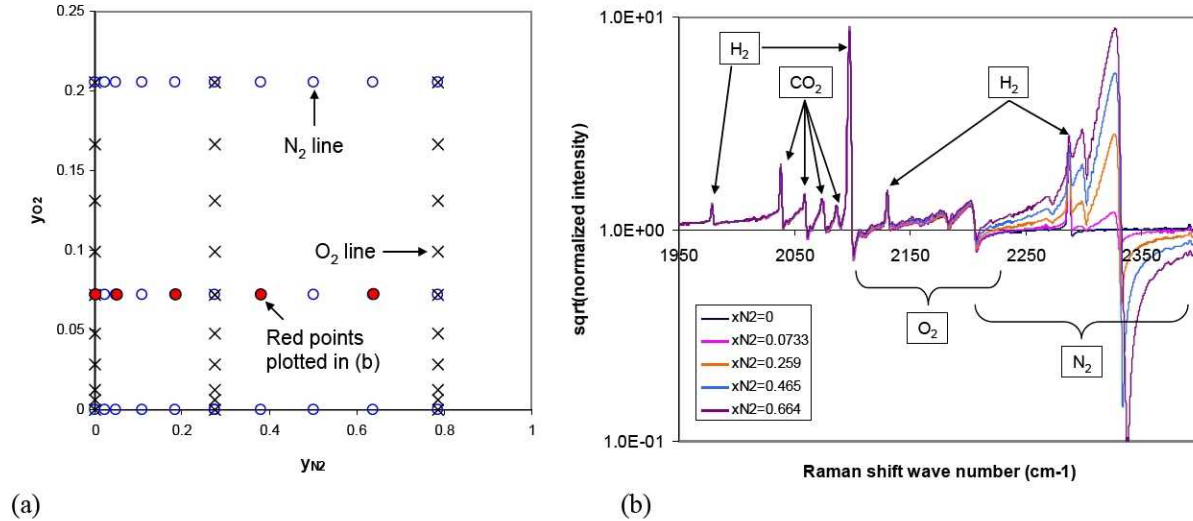

(a)   (b)

Figure 1. (a) Library grid points in weighted coordinates in the N_2 and O_2 plane (10 points and 3 nodes along each line); (b) typical spectra along an N_2 line of the library at compositions indicated by the red points in (a), normalized by non-resonant background susceptibility

The species concentration coordinates are chosen so that the variation of a given species affects primarily the portion of the CARS spectrum that is resonant for that species, leaving the rest of the spectrum relatively unchanged. Thus, the grid is orthogonal in CARS non-resonant susceptibility weighted fraction coordinates defined below:

$$y_n = x_n \chi_{NR,n} / \chi_{NR}$$

Here, x_n is the mole fraction of species n, N_d is the number of resonant species, species N_d+1 is the non-resonant species (thus, $\sum_{n=1}^{N_d+1} x_n = 1$), $\chi_{NR,n}$ is the CARS non-resonant susceptibility of species n, and χ_{NR} is the non-resonant susceptibility of the mixture as defined below:

$$\chi_{NR} = \sum_{n=1}^{Nd+1} x_n \chi_{NR,n}$$

The location of grid points in a plane of a typical sparse grid in which N_2 and O_2 vary, plotted in weighted fraction coordinates, is shown in Figure 1(a). For each species there is a line of points at the "nodes" of the other species; for the case shown in the figure there are three nodes for each species and 10 points along each line giving a total of 60 points. By comparison, in the comparable fully packed library each point is a node and there would be $10 \times 10 = 100$ points for the example of Figure 1(a). It is assumed that there always exists an additional species that has no resonances in the range of wavelengths considered, and that makes up the sum of the mole fractions of all the species to 1; in this example the non-resonant species is water vapor. Some points in the library define compositions that have negative amounts of the non-resonant species. Spectra at these physically unrealizable compositions can none-the-less be computed and may be required for interpolation to realizable compositions in instances when the composition contains only a small amount of the non-resonant species.

The library spectra signal intensities are normalized by the non-resonant susceptibility of the mixture, χ_{NR}, prior to interpolations. Figure 1(b) contains plots of the square root of the normalized signal intensity with N_2, O_2, H_2, and CO_2 resonant (see Section III for more details of this calculation) along a N_2 coordinate line in the library (the red circles) shown in Figure 1(a). In these spectra the shape varies only in the region where N_2 is resonant, which would not be the case if, say, the mole fraction of N_2 was varied at constant mole fraction of the other resonant species.

Table 1 Library size at each temperature given 10 points and 3 nodes along each line.

no. chem species	sparse library size	packed library size	ratio packed to sparse
1	10	10	1.00
2	60	100	1.67
3	270	1000	3.70
4	1080	10000	9.26
5	4050	100000	24.69

Library sizes (at each temperature level) for typical fully-packed libraries with 10 points per coordinate direction are compared to the sizes of a sparse library with 10 points per line and 3 nodes in Table 1. Computational times of fitting algorithms involving libraries scale roughly with library size, as does computer memory usage. The computational savings of using a sparse library become substantial for three chemical species or more, and make the fitting of up to at least 5 chemical species practical (which is impractical with a fully-packed library).
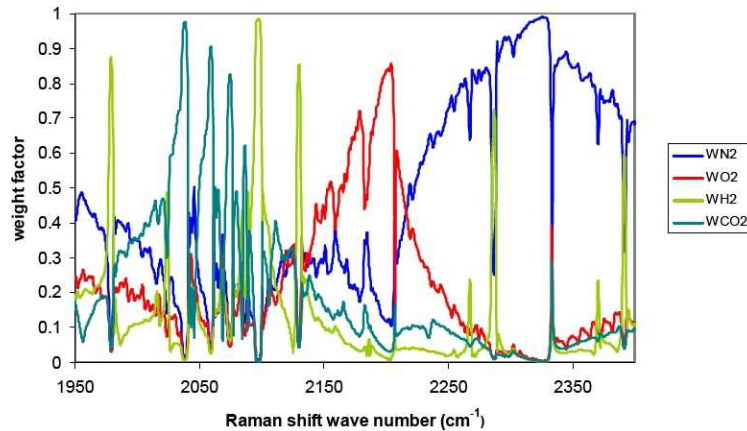
B   Interpolation from the Library

Figure 2. Weight factor distributions derived from the library and used in interpolation of spectra from the library

Before spectra can be interpolated from the sparse library, a weight function distribution is determined for each resonant species, W_n. This function is an estimate of the typical contribution to the intensity of the spectra of a gas mixture from that particular species. For each resonant species the average normalized spectrum is formed by

averaging (across all temperatures) the library spectra that have the maximum mole fraction of that species and zero of the other resonant species. The weight functions are the average spectra renormalized at each wavelength by the sum of the average spectra for all resonant species; thus the sum of the weight functions of all the species is 1. Typical weight function distributions computed for the library with N_2, O_2, H_2, and CO_2 resonant are shown in Figure 2. The origins of the various peaks and valleys may be identified by comparing to the typical spectra in Fig. 1(b), where the resonances in the various species can be identified. Thus, for example, the weight function for N_2 tends to 1 at wavelengths where N_2 is resonant, has some sharp dips where H_2 resonances overlap the N_2 resonances, tends to zero where other species are resonant and N_2 is not resonant, and where no species are resonant has the same value (~0.25) for all species.

The method for interpolating a spectrum at a particular library temperature to a desired composition has several steps. A complete spectrum (including all the resonances), T_n, is interpolated to the desired composition from all the library lines in a given species coordinate direction, one interpolated spectrum for each species. Thus, in Fig 1(a), one spectrum would be interpolated using all the lines where N_2 is varied (blue circles), a second spectrum would be interpolated using all the O_2 lines (black x's), etc., repeated for all the species. Each of these spectra will be similar, but the n^{th} spectrum will be more accurate in the vicinity of the resonances of the n^{th} species since it is interpolated from the lines in the n^{th} species direction (which have more closely spaced points).

The first step in computing the n^{th} spectrum described above is to interpolate along each line in the n^{th} species coordinate, using quadratic interpolation, a spectrum at the desired concentration of that species. This results in a regular "reduced" grid of spectra at the node points of the remaining species, $T_{n,\mathbf{i}}$, where $\mathbf{i}=i_1,i_2,\ldots i_j,\ldots i_{Nd}$ where $j=n$ is omitted; the reduced grid has N_d-1 dimensions and contains O^{Nd-1} spectra, where O is the number of nodes in each coordinate direction. Since the reduced grid is regular, the n^{th} spectrum is found using Lagrange polynomial interpolation as defined by the equation below:

$$T_n = \sum_{i_1=1}^{O}\sum_{i_2=1}^{O}\ldots\sum_{i_j=1}^{O}\ldots\sum_{i_{Nd}=1}^{O} T_{n,\mathbf{i}} \left(\prod_{\substack{ii=1\\ii\neq i_1}}^{O}\frac{y_1-y_{1,ii}}{y_{1,i_1}-y_{1,ii}}\right)\left(\prod_{\substack{ii=1\\ii\neq i_2}}^{O}\frac{y_2-y_{2,ii}}{y_{2,i_2}-y_{2,ii}}\right)\ldots\left(\prod_{\substack{ii=1\\ii\neq i_j}}^{O}\frac{y_j-y_{j,ii}}{y_{j,i_j}-y_{j,ii}}\right)\ldots\left(\prod_{\substack{ii=1\\ii\neq i}}^{O}\frac{y_{N_d}-y_{N_d,ii}}{y_{N_d,i_{Nd}}-y_{N_d,ii}}\right)$$

The summation and product of a sequence in the above equation is omitted for $j=n$; y_i is the composition being interpolated to and $y_{i,ii}$ are the compositions at the reduced grid points. Thus, for 4 resonant species and 3 nodes in each direction, as in the example of Fig. 1, the reduced grid is three-dimensional and there 27 terms in the summation; each summed term consists of the product of 27 terms that multiplies a theoretical spectrum value. The general formulation of this equation allows for efficient programming via "do" loops for arbitrarily large numbers of species.

The final interpolated spectrum is the sum over the species of these spectra, weighted by the weight function distribution defined below:

$$T = \sum_{n=1}^{Nd} W_n T_n$$

The resultant spectrum is accurate for all resonances.

C   Residual Functions

Experimental spectra are analyzed to determine composition and temperature by interpolating from the library of theoretical spectra and minimizing the residual functions defined below. S_k is the square root of the experimental signal intensity (in square-root counts), T_k is the square root of the normalized theoretical signal intensity, σ_k^2 is the variance of the experimental noise in the signal intensity (in counts), k is the detector pixel number (wavelength), N_p is the number of pixels, and υ is the number of degrees of freedom (N_p minus the number of fit parameters). X is a scaling factor that rescales the theoretical spectrum so that it best fits the experimental spectrum, and is recomputed at each evaluation of the residual. The equation used to calculate X is slightly different for each residual; they are derived by differentiating the definition of the residual with respect to X, setting the result to zero, and solving analytically for X.

Three residual functions are considered. The first residual is the mean square error of the signal intensity:

$$R_1 = \frac{1}{N_p-1}\sum_{k=1}^{N_p}(S_k^2 - X_1 T_k^2)^2$$

American Institute of Aeronautics and Astronautics

The resulting fit then is simply a least squares fit to the intensity. The second residual is the "reduced chi-square statistic":

$$R_2 = \chi^2 = \frac{1}{\upsilon} \sum_{k=1}^{N_p} \frac{(S_k^2 - X_2 T_k^2)^2}{\sigma_k^2}$$

The resulting fit is then a weighted least squares fit to the intensity. Assuming that the noise model is correct, a value of $\chi^2 \gg 1$ indicates that the fit is poor and $\chi^2 \sim 1$ indicates a good fit. If $\chi^2 < 1$ the model has too many free parameters and is fitting the noise. Snelling et al.[8] showed that both the random and bias errors in single-shot CARS measurements of temperature were dramatically reduced with weighted least-squares fitting compared to un-weighted. The third residual is the mean square error of the square root of the signal intensity and the resulting fit is simply a least squares fit to the square root of the intensity.

$$R_3 = \frac{1}{N_p - 1} \sum_{k=1}^{N_p} \left(S_k - X_3 T_k \right)^2$$

The CARS signal intensity is roughly proportional to the square of the mole fraction of resonant species. Least square fitting of the square root of the signal intensity emphasizes the lower concentration resonant species and the non-resonant regions of the spectrum compared to least squares fitting the signal intensity.

All previous single- and dual-pump CARS studies at NASA Langley (e.g., Refs. 3, 4, 5) used least squares fitting to square-root intensity. (CARSFT has the options to least squares fit square-root intensity or a weighted square-root intensity.) The effect on the fitted parameters of the choice of residual function when fitting noisy CARS spectra will be studied in this paper.

D   Fitting Algorithm

The fitting algorithm searches the library, interpolates spectra and minimizes the residual function. Fitted parameters are temperature, mole fraction of each resonant species, and a horizontal shift parameter that is applied to the data to correct small constant wavelength shifts in the data, Δ (cm^{-1}).

The algorithm includes a quick fitter to get an estimate of the solution. The quick fitter performs a direct search by computing the residual for every library spectrum and interpolating the composition and temperature that produces the minimum residual. This provides a good initial guess for the solution which is then refined by an iterative fitter. (This two-step process tends to avoid the possibility of the iterative fitter not converging to the global solution as it might if given an arbitrary initial guess.)

The iterative fitter employs the Levenberg-Marquardt algorithm (LMA), as implemented by Lahey Fortran, which minimizes the function defined below:

$$R = f(\mathbf{y}) = \sum_{k=1}^{N_p} \left\{ F_k(\mathbf{y}) \right\}^2$$

F is a function of the species mole fraction and pixel number, and can be identified by comparison with the definitions for the residuals. Rather than include temperature as a parameter of the fit, the algorithm is implemented to find the best fit at several fixed temperatures (including both library values and at temperatures intermediate between library values) and the minimum residual is computed at each temperature. The temperature that produces the overall minimum is interpolated using quadratic interpolation from the residuals at the fixed temperatures. The resulting algorithm was found to always converge to sensible parameter values that globally minimized the residuals.

The horizontal shift is found in a separate iterative cycle applied both as part of the quick fitter and as part of the LMA fitter stages. The equation used to calculate Δ at each iteration is slightly different for each residual and is derived from the equation for the residual by setting $T = T(\nu - \Delta)$, where ν is the wave number, differentiating this equation with respect to Δ, setting the result to zero, and then solving algebraically for Δ. The resulting equation contains $dT/d\Delta$ which is found by fitting a spline to $T(\nu)$ assuming $\Delta = 0$, which amounts to a local linear approximation to the slope. The equation is applied repeatedly until the accumulated shift ceases to change; the method typically converges within 2-3 cycles.

III. Numerical Experiments

Numerical studies were conducted to verify the accuracy of the new fitting algorithm. A set of spectra was generated numerically in which the concentrations of four resonant species (N_2, O_2, H_2, and CO_2) were varied at random and the spectra were fit to a sparse library using the algorithm. The Sandia code CARSFT (modified by Lucht[2] for dual-pump CARS and further modified by O'Byrne[4]) was used to calculate all the spectra. The spectral range of computation for this study was 1950 cm^{-1} to 2400 cm^{-1} and the second pump frequency offset was 650 cm^{-1}. The first pump line width was assumed to be negligible, the second was 0.07 cm^{-1}, and the instrument line width was 2.0 cm^{-1}. The Voigt line shape model was used for all species except H_2 where the Galatry model with narrowing parameter 0.047 was used.[9]

The library contained 9720 spectra: there were 9 temperatures from 800 K – 2400 K in 200 K increments, and 10 compositions for each species that were non-uniformly spaced with mole fractions in the range 0-0.21 for O_2, 0-0.79 for N_2, 0-0.5 for H_2, and 0-0.25 for CO_2. The library is sparse – a fully packed library with these dimensions would contain 90,000 spectra.

Two hundred synthetic "experimental" spectra were generated using CARSFT that were uniformly and randomly distributed in the same temperature and mole fraction ranges as the library, but constrained so that the sum of the resonant mole fractions did not exceed 1 and N_2 mole fraction was greater than 0.3 for temperatures greater than 1600 K; this last condition is arbitrary but justified based on the fact that high temperature is not possible without combustion of fuel and air, leaving as one of the products N_2. These spectra were rescaled by an arbitrary factor to be in the units of the detector analog-to-digital (A/D) converter or "counts" and noise was added to the spectra to simulate experimental noise.
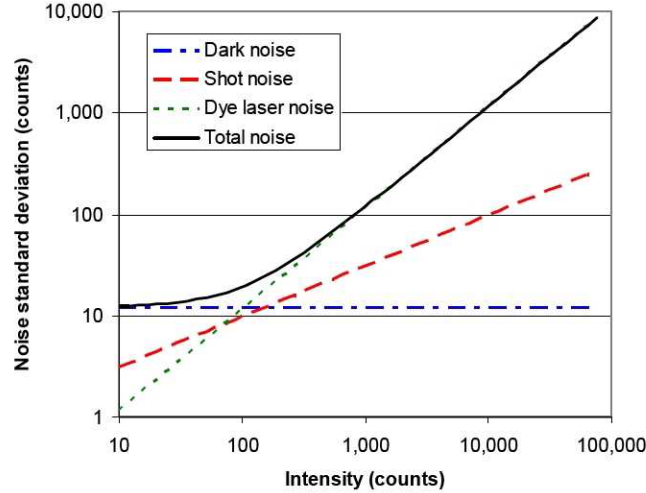

Figure 3. Model for standard deviation of signal noise.

The noise was assumed to be Gaussian distributed with variance dependent on experimental signal intensity S^2 according to a model that included the dark noise of the detector, photon shot noise, and broadband dye laser mode noise;[10] the variance (in counts) is given by the equation below:

$$\sigma_k^2 = (D/N_e)^2 + S_k^2 / N_e + (S_k^2 C_{bbdl})^2 \quad \textbf{Equation 1}$$

D is the dark noise (12.6 electrons), N_e is the number of electrons per count (1.05), and C_{bbdl} is the intensity of the broadband dye laser noise as a fraction of intensity (0.115). These model coefficients were initially believed to simulate an existing CARS system at NASA Langley (see Section IV). Dark noise and broadband dye laser mode noise are expected to be Gaussian distributed. Shot noise obeys Poisson statistics that tend to Gaussian as the signal level increases (signal counts > 10). Figure 3 shows the noise standard deviation plotted as a function of signal intensity, and also the individual contributions of dark noise, shot noise and dye laser noise. The ratio of total noise to signal does not go to zero but tends to a constant as the signal counts become very large; dye laser mode noise is dominant very much above 100 counts while the dark noise is dominant very much below this level.

The library took 75 hours to compute on a single processor of a PC, an average of 0.5 min/spectrum. Fitting the 200 experimental spectra to the library took about 150 minutes, 0.8 min/spectrum. If CARSFT in the fitting mode had been used, there would have been hundreds or thousands of spectrum evaluations to fit a single spectrum, taking many hours at best, if a converged fit could be found at all.

Two typical experimental spectra and fits to them are shown in Figure 4: these are based on the same theoretical spectrum that has been rescaled to two different signal levels. The low signal-to-noise (S/N) case has 25 detector counts at the non-resonant background (NRB) baseline level and the high S/N case has 400 counts at the NRB baseline level (note that the plot shows square root of signal level). Figure 4(a) shows the spectra for both cases and the fits to them using residual R_2. Figure 4(b) shows the low S/N spectrum, the synthetic experimental spectrum prior to addition of noise, and the fits for each of the residuals. Note how fits using R_1 and R_2 are almost identical to each other and close to the no-noise spectrum, while the fit using R_3 is low where the signal is low and the noise highest. The low fit is expected since the noise distribution of the signal at a given pixel has mean equal to the no-noise intensity, but the noise distribution of the square root of the signal is skewed towards zero, and has a mean less than the square root of the no-noise mean. Obviously the best fit is one that recovers the no-noise spectrum since it will have the correct values of the fit parameters, and R_1 and R_2 gives superior results in this case.
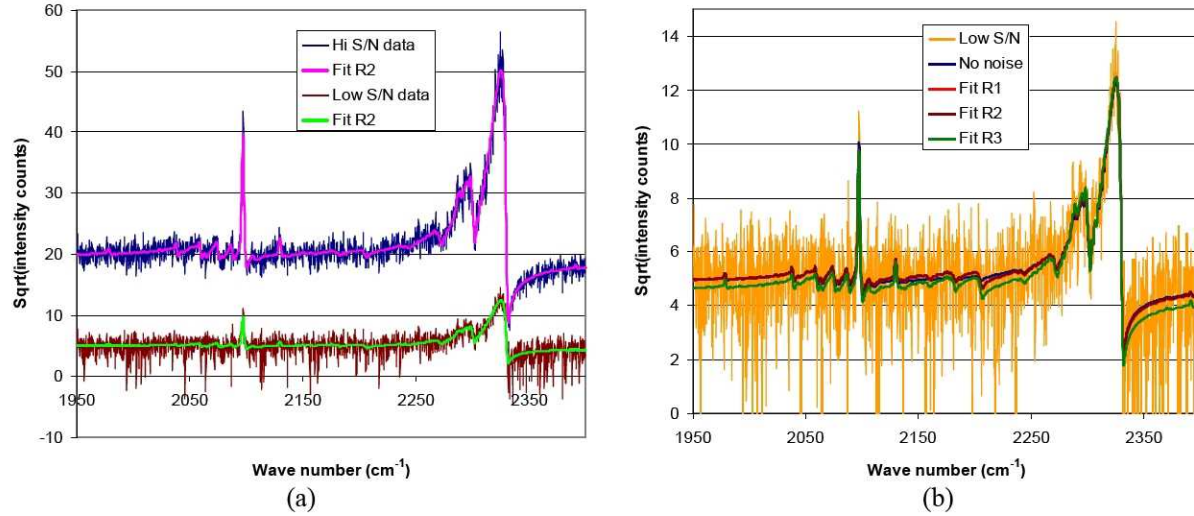

Figure 4. Fits to synthetic experimental spectra with noise: (a) low and high signal-to-noise ratio (residual R_2 minimized), (b) low signal-to-noise ratio with all three residuals fitted and prior to addition of noise.

Table 2. Mean error and standard deviation of error in fitted temperature and composition fitting 200 noisy spectra. (a) No noise (baseline), (b) R_1 minimized, (c) R_2 minimized, (d) R_3 minimized.

(a)

fit error	res	T (K)	mole fraction			
			N_2	O_2	H_2	CO_2
mean	R_1	1.3	0.0001	0.0001	-0.0003	0.0003
	R_3	1.2	0.0003	0.0002	-0.0001	0.0004
st dev	R_1	5.4	0.0012	0.0014	0.0014	0.0013
	R_3	5.1	0.0008	0.0010	0.0007	0.0009

(b)

results R_1	NRB (counts)	T (K)	mole fraction			
			N_2	O_2	H_2	CO_2
mean	400	2	0.0014	-0.0007	-0.0001	-0.0015
	25	5	0.0027	-0.0007	-0.0003	-0.0014
st dev	400	25	0.006	0.012	0.006	0.017
	25	35	0.008	0.017	0.008	0.020

(c)

results R_2	NRB (counts)	T (K)	mole fraction			
			N_2	O_2	H_2	CO_2
mean	400	-2	0.0007	-0.0001	0.0000	-0.0001
	25	3	-0.0004	-0.0003	-0.0016	-0.0004
st dev	400	13	0.003	0.003	0.003	0.003
	25	29	0.006	0.011	0.005	0.009

(d)

results R_3	NRB (counts)	T (K)	mole fraction			
			N_2	O_2	H_2	CO_2
mean	400	3	0.0010	0.0006	-0.0005	0.0005
	25	46	0.0239	0.0126	0.0046	0.0157
st dev	400	16	0.003	0.003	0.003	0.003
	25	54	0.011	0.021	0.009	0.021

Table 2 shows mean and standard deviation of the error (fit value minus true value) for the fitted parameters for 200 experimental spectra at different temperatures. Table 2(a) is a case with no noise (R_1 are R_3 are minimized). Tables 2(b,c,d) are cases with noise where R_1, R_2 are R_3 are minimized respectively. The standard deviation of the errors in the no noise case, attributed to errors by the fitting algorithm, are ~5 K and less than ~0.0014 in mole fractions, while the mean errors are negligible. These errors are small as compared to errors in fitting noisy spectra and thus the fitting algorithm introduces no significant errors. Comparing Tables 2(b,c,d) to each other it may be seen that residual R_2 produces consistently the best results, both in the mean and the standard deviation of the error in both temperature and composition. For example, for the low signal-to-noise case the mean and standard deviation of error in temperature are 3 K and 29 K respectively for residual R_2, but 46 K and 54 K for R_3. Results are similar for the mole fraction errors. Results for R_1 though worse than R_2 are better than for R_3.

IV.   Measurements in Hencken Burner

Experimental broadband dual-pump CARS data acquired in the H_2-air flat-flame produced by a Hencken burner were fitted using the algorithm. The Hencken burner flame has previously been carefully studied;[11] it was found to be very nearly adiabatic and the composition and temperature could be calculated using adiabatic, equilibrium assumptions. Burner flow rates were between 44 and 63 standard liters per minute while H_2-air equivalence ratios were varied in the range ~0.1-3.5.

CARS data were acquired using the portable system at NASA Langley[5] described most recently in the companion paper by Magnotti[12] (elliptical beam shaping was not used in the present experiment). For the present experiment laser beams pulsed at 20 Hz and ~10 ns pulse length, one at 532 nm (injection seeded, doubled Nd:YAG), one at 552.8 nm (narrowband dye laser with line width ~0.07 cm^{-1}), and one at 604.5 nm (broadband dye laser with full-width half-maximum line width of 9.6 nm), are focused and crossed in a planar BOXCARS phase matching configuration. The signal is separated by dichroic mirrors and focused at the entrance of a 1 meter spectrometer with a 2400 lines per inch grating. Spectra were normalized by a spectrum (averaged over ~200 pulses) acquired in argon (which is non-resonant) to remove the effect of the broadband dye laser line shape. Library spectra were generated using CARSFT; the Voigt lineshape model was used for N_2 and O_2 and the Galatry model was used for H_2 with narrowing parameter of 0.047.
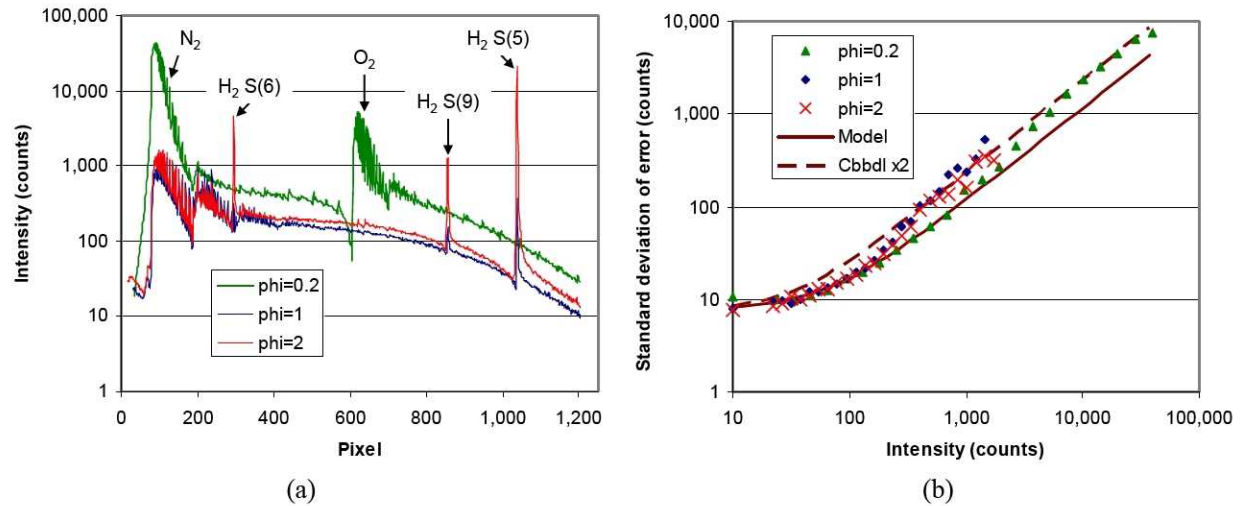

Figure 5. CARS Hencken burner data at three equivalence ratios: (a) averages of 200 spectra, (b) noise levels compared to the noise model.

The noise model of Equation 1 was employed in evaluating residual R_2. The parameters of the noise model were estimated from experimental spectra. From about 200 experimental spectra in the argon, which is non-resonant, the ratio of the standard deviation of the noise to the signal level (excluding the wings of the spectrum where the dark noise is significant) was found to be ~0.115; this is a good estimate of C_{bbdl}. Noise properties were also estimated from Hencken burner spectra at nominal equivalence ratios of 0.2, 1, and 2. The average of 200 spectra at each of these equivalence ratios is shown in Fig. 5(a); these spectra have not been normalized by the argon spectrum since noise depends on absolute not normalized signal level. For each equivalence ratio, the noise error in each pixel of

10 spectra was found by applying a small correction to make the total signal of each spectrum (summed over all the pixels) the same as for the average spectrum, to compensate for shot-to-shot random variations in the laser energy, and then subtracting the average spectra. The errors were binned according to the signal level of the pixel and the standard deviation of the error was found for each bin. The standard deviation of the error is plotted as a function of signal level in Fig. 5(b) at each equivalence ratio. Also shown is the noise model with $D=8$ electrons and plotted both with $C_{bbdl}=0.115$ (labeled "model") and $C_{bbdl}=0.23$ (labeled "Cbbdl x2"). The "model" line is thus a model of the noise in the non-resonant argon spectrum. The Hencken burner spectra data agree with the "model" line at lower signal levels, where the signal is generated by the non-resonant portion of the spectra, and then deviate to higher levels where the signal comes from resonant portions of the spectra. Noise levels seem to asymptote to the top model line where $C_{bbdl}\sim0.23$. The reason the broad-band dye laser noise is higher in the resonant parts of the spectra is that none of these spectra are fully resolved by the spectrometer; thus, the interaction between the spectrally unresolved structures of the dye laser spectrum and the unresolved lines of the resonant spectrum produces greater noise than interaction with the featureless non-resonant spectrum. For the purpose of computing R_3 an intermediate value is used: $C_{bbdl}\sim0.16$. Both theoretical and experimental spectra were multiplied by the reference argon spectrum to recover spectra in the original intensity units before computing R_2; spectra in this form were used both directly and indirectly in computing the standard deviation of the noise error from Equation 1.

Problems were encountered in accurately fitting the isolated H_2 rotational lines (the S(5), S(6), and S(9) shown in Fig. 5(a)). The apparent spectral location of these lines deviated a small amount from the location of the lines in the library spectra after fitting N_2 and O_2 (error is less than the line width), even after correcting for small nonlinearities of the spectrometer with respect to wave number. Also the line shape was not always well represented by the library spectra. Since these are isolated rotational lines, at the resolution of our instrument the line shape contains no useful information. This problem was solved by applying the average level in a narrow range around the line center to each pixel in that range; this averaging was done to both the theory spectrum and the experimental spectrum prior to each residual evaluation. The result is that the algorithm fits only the area under the spectrum in the specified wave number range. An additional problem encountered with the H_2 rotational lines was that the temperature that produced the best fit to these lines was different than the calculated flame temperature. Several researchers report errors measuring temperature using H_2 CARS that are associated with line-shape modeling (e.g., Ref. 11). If all three H_2 lines plus the N_2 and O_2 resonant regions are included in the fit then the fitted temperature is some compromise between the temperature that best fits the H_2 lines and the temperature that best fits the N_2 and O_2 resonant regions; this compromise temperature depends strongly on which residual is minimized. Since the spectral modeling of N_2 and O_2 has been widely validated for fitting for temperature (e.g., Ref. 11), and since S(5) is stronger than S(6) or S(9) over all our temperature range, only S(5) is included in the fitting; the other lines are omitted from the residual by excluding small wave number regions each side of each line center. The result is that H_2 has no effect on the fitted temperature but the algorithm fits for H_2 mole fraction. Another problem is that none of these H_2 rotational lines are significantly populated when room temperature is approached (between 300K and 500 K). At these temperatures the method becomes ill-posed for composition and small noise errors in the spectrum are fitted as large mole fractions of H_2. Tedder et al.[7] have proposed using the S(3) line which is strong at both room and flame temperatures.

Figure 6 (a-d) shows the mean and standard deviation of the fitted temperature and the species mole fractions from 200 spectra at several equivalence ratios plotted as a function of equivalence ratio. Results are shown for fits minimizing each type of residual. For the purposes of the plots, the standard deviations are multiplied by 10. Also shown is the calculated temperature and mole fraction distribution based on flow rates to the burner; uncertainty bands on the calculations are associated with uncertainties in the flow rates. The error between theoretical and mean measured temperature is typically less than about ±50 K and the error in mean mole fractions is less than about ±0.02. Mole fraction of O_2 is measured high, although the results of minimizing residual R_2 appear to be the best. Errors in N_2 and H_2 are similar for all three residuals.

The standard deviation in fitted temperatures and mole fractions are consistently lowest for residual R_2, followed by R_3 and then R_1. In some cases the trend is very pronounced; in the case of temperature or O_2 mole fraction the reduction from R_1 to R_2 is a factor of one half or less. This is an important result for the use of CARS to measure turbulent fluctuations in flames since instrument random error limits the minimum turbulent fluctuation level that can be resolved. These results are consistent with the results of the numerical experiments in Section III in that minimizing the chi-square residual, R_2, reduced random error (noise) and bias errors in some fitted parameters. The standard deviations for this residual are ~50K for temperature, ~0.005-0.01 for O_2, ~0.01-0.015 for N_2, and ~0.01-0.03 for H_2. The Hencken burner flame may have some real fluctuations in temperature and composition, so the random errors in this CARS method may be less than suggested by these results.
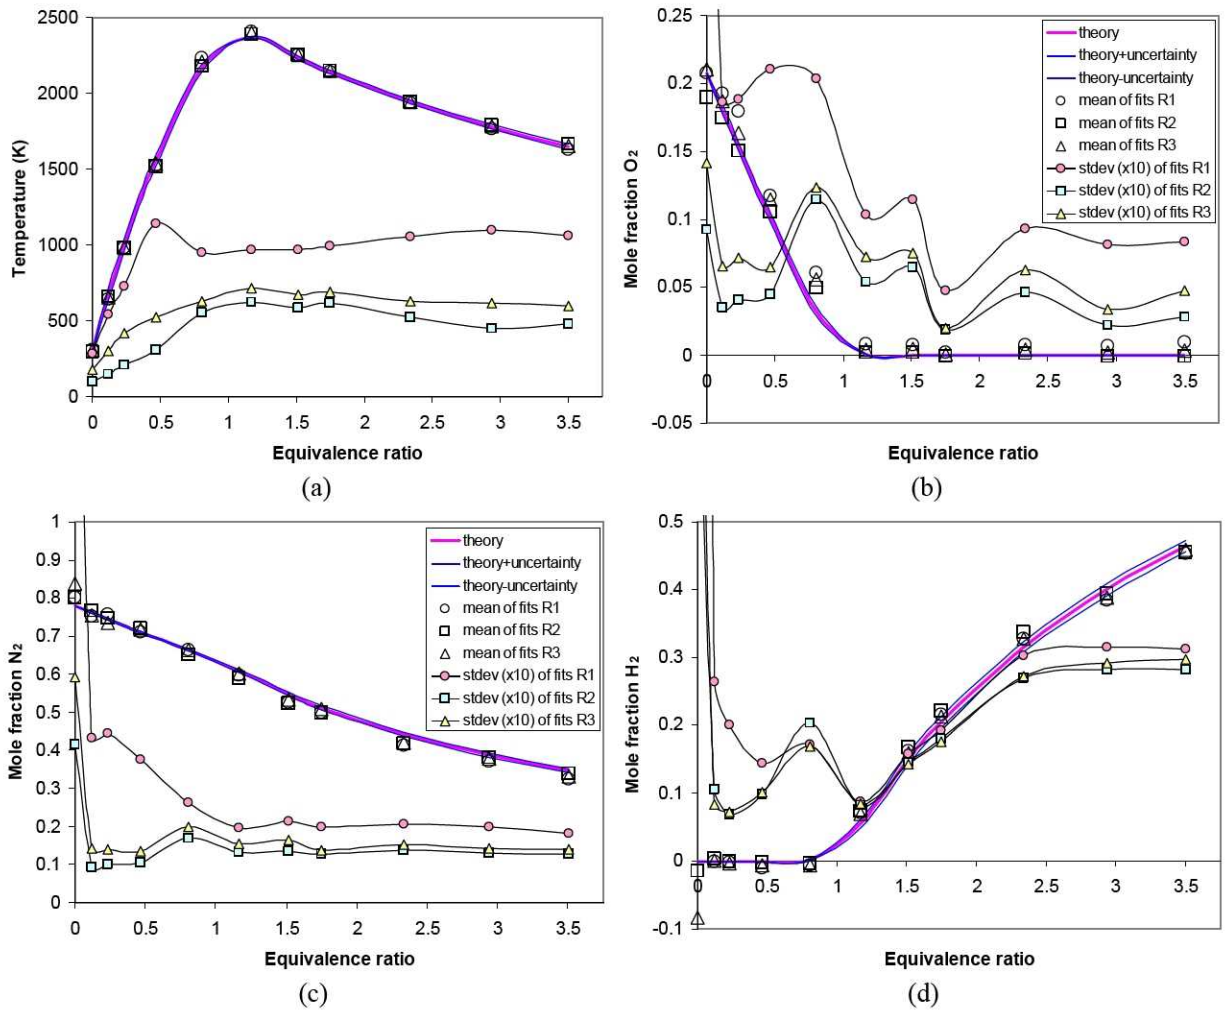
Figure 6. Fits to CARS Hencken burner data; mean and standard deviation of fitted parameters as a function of equivalence ratios. Standard deviation is multiplied by 10. (a) Temperature, (b) mole fraction O_2, (c) mole fraction N_2, (d) mole fraction H_2.

V.   Conclusions

This paper has described the development of an algorithm for fitting experimental CARS spectra containing multiple resonant species to theory. The algorithm employs a library of previously calculated theoretical spectra and thus is much faster than methods that compute new spectra at each iteration of the solver. Library spectra are assumed calculated by some other algorithm. The library structure is orthogonal in a non-resonant susceptibility weighted mole fraction coordinate and is sparse; its size thus grows much more slowly with number of species than a conventional fully-packed library. The method of interpolation from the library is discussed; several residual functions minimized by the fitter were considered, and the fitting method was described. Synthetic "experimental" dual-pump CARS spectra containing 4 resonant species (N_2, O_2, H_2 and CO_2), both with noise and without it, were fitted by the algorithm and the method was demonstrated to be both robust and accurate (there were negligible errors in fitting noise-free spectra). Experimental dual-pump CARS from a Hencken flat-flame burner reacting H_2 and air were also fitted. In both studies, weighted least square fitting of signal intensity, as opposed to least squares fitting signal intensity or square-root signal intensity, was shown to produce the least random error and the same or less bias error in the fitted parameters. Standard deviation of the error (random error) was reduced by factor of one half or more in some parameters by weighted least squares. Use of weighted least squares is thus particularly important in turbulence studies where minimum resolved turbulence levels are limited by random error. Some experimental issues associated with fitting isolated rotational lines of H_2 were also resolved.

VI. Acknowledgments

VII. References

[1] Lucht, R.P., "Three-laser coherent anti-Stokes Raman scattering measurements of two species," *Optics Letters*, Vol. 12, No. 2, February 1987, pp. 78-80.

[2] Hancock, R.D., Schauer, F.R., Lucht, R.P. and Farrow, R.L., "Dual-pump coherent anti-Stokes Raman scattering measurements of nitrogen and oxygen in a laminar jet diffusion flame," *Applied Optics*, Vol. 36, No. 15, 1997.

[3] Cutler, A.D., Danehy, P.M., Springer, R.R., O'Byrne, S., Capriotti, D.P., DeLoach, R., "Coherent Anti-Stokes Raman Spectroscopic Thermometry in a Supersonic Combustor," *AIAA J.*, Vol. 41, No. 12, Dec. 2003.

[4] O'Byrne, S., Danehy, P.M., Tedder, S.A., Cutler, A.D., "Dual-Pump Coherent Anti-Stokes Raman Scattering Measurements in a Supersonic Combustor," *AIAA Journal*, Vol. 45, No. 4, p. 922-933, April 2007

[5] Tedder, S.A., Danehy, P.M., Magnotti, G., Cutler, A.D., "CARS Temperature Measurements in a Combustion-Heated Mach 1.6 Jet", AIAA-2009-0524, 47th AIAA Aerospace Sciences Meeting, Orlando, FL, January, 2009

[6] Palmer, R.E, The CARSFT Computer Code for Calculating Coherent Anti-Stokes Raman Spectra: User and Programmer Information, Sandia report SAND89-8206, Feb 1989.

[7] Tedder, S.A., Wheeler, J. L., Danehy, P.M., "Width Increased Dual-Pump Enhanced CARS (WIDECARS)", Submitted to *Applied Optics* on 11/30/09.

[8] Snelling, D.R., Smallwood, G.J., Sawchuk, R.A., Parameswaran, T., "Precision of multiplex CARS temperatures using both single-mode and multi-mode pump lasers," *Applied Optics*, Vol. 26, No. 1, pp. 99-110, Jan 1987.

[9] Kojima J., Nguyen Q., "Quantitative analysis of spectral interference of spontaneous Raman scattering in high-pressure fuel rich H2-air combustion," *Journal of Quantitative Spectroscopy & Radiative Heat Transfer*, Vol. 94, 2005, pp 439-466, 2005.

[10] Greenhalgh, D.A., Whittley, S.T., "Mode Noise in CARS Spectroscopy," *Applied Optics*, Vol. 24, No. 6, pp. 907-913, March 1985.

[11] Hancock, R.D., Bertagnolli, K.E., Lucht, R.P., "Nitrogen and Hydrogen CARS Measurements in a Hydrogen/Air Flame Using a Near-Adiabatic Flat-Flame Burner," *Combustion and Flame* 109:323-331 (1997)

[12] Magnotti, G., Cutler, A.D., Danehy, P.M., "Beam Shaping for CARS Measurements in Turbulent Environments," AIAA-2010-1400, 48th AIAA Aerospace Sciences Meeting, Orlando, FL, January, 2010