Appl. Phys. B 85, 461–466 (2006)
DOI: 10.1007/s00340-006-2316-9

Applied Physics B
Lasers and Optics

D.S. BOMSE✉
D.J. KANE

An adaptive singular value decomposition (SVD) algorithm for analysis of wavelength modulation spectra

Southwest Sciences, Inc., 1570 Pacheco Street, Suite E-11, Santa Fe, NM 87505, USA

Received: 7 April 2006/Revised version: 5 May 2006
Published online: 24 June 2006 •

ABSTRACT Optical interference fringes due to unwanted etalons are often the limiting uncertainty in diode laser spectroscopic trace gas measurements. Temporal variations in the fringe spacings, phases, and amplitudes introduce systematic baseline changes that limit useful signal averaging times to ~ 1000 s, and constrain minimum detectable absorbances to between one and three orders of magnitude worse than the fundamental limiting noise sources (shot noise and/or detector thermal noise). We describe an adaptive numerical filtering method based on singular value decomposition (SVD) that shows, for the system studied, a fivefold reduction in baseline drift due to unwanted etalons over a one-week-measurement period. The adaptive algorithm is fast (< 1 ms per computation), robust, and uses linear methods. It is computationally equivalent to principal component analysis (PCA). The test system was acetylene detected using a near-infrared telecommunications laser operating at 6541.96 cm^{-1}. The gas detection limit was 20 ppb (1σ) over the one-week measurement.

PACS 07.05.Kf; 07.07.Df; 07.60.-j; 42.60.-v; 42.62.-b; 42.62.Fi

1 Introduction

Theoretical detection limits for diode laser absorbance measurements are in the 10^{-8} to 10^{-7} $Hz^{-1/2}$ range – defined by laser shot noise or detector thermal noise – yet practical diode laser spectrometers rarely perform better than 10^{-5} $Hz^{-1/2}$ short term and 10^{-6} long term (~ 1000 s) [1]. Unwanted optical interference fringes (etalons) cause this difference between theoretical and actual performance. To be pedantic, interference fringes are not noise. If they remained fixed, they would introduce at worst a constant absorbance offset. But, etalons are interferometric and their amplitudes, periods, and phases are all sensitive to temperature, mechanical stress, and vibration. The result is short and long term drift in measurements of small absorbances.

One might expect etalons to be removed easily by digital filtering methods. A number of factors, however, complicate etalon recognition and elimination:

– Fringe periods are often close to absorption linewidths. Most atmospheric pressure broadened absorption lines are ~ 0.2 cm^{-1} FWHM. This width also matches a fringe free spectral range due to an etalon that is 2.5 cm long and corresponds to the distances between optics within many optical trains.
– Diode laser scan ranges include only a few fringe periods. This is particularly true of distributed feedback (DFB), near-infrared lasers that current tune by ≤ 1 cm^{-1}. It is difficult to determine the frequency and phase of a sine wave precisely when using only a few wave periods.
– Spectra often include multiple, interfering etalons. Each etalon is characterized by different temperature and vibration sensitivity.
– Most diode lasers do not scan linearly with changing current. As a result, etalons do not appear sinusoidal.

Given these constraints, most researches have emphasized the mechanical or electronic etalon reduction methods and implemented frequent recalibrations using zero gas. Good optical design (using focusing mirrors instead of lenses, antireflection coatings on transmissive optics, wedged optics, and wide apertures) keeps interference fringe amplitudes in the low-10^{-4} equivalent absorbance range. Active reduction techniques including mechanically dithering the optics [2], oscillating Brewster plates [3], and electronic wavelength smearing [4], and numerical filtering methods [5] typically reduce the time-averaged fringes to the 10^{-6} to 10^{-5} range. Allan variance measurements usually show signal-to-noise ratios improving for averaging up to ~ 1000 s [5, 6], then declining as long-term etalon drift becomes important. For example, Werle et al. provide one of the most detailed descriptions of wavelength modulation spectral data analysis and emphasize minimizing laser frequency errors to avoid response errors [5]. This maximizes the useable signal averaging time between measurements of baselines using zero gas, but no attempt is made to track etalon fringe changes during sample measurements. Tittel [7] and Silver [8] have successfully used post-analysis Kalman filtering to remove high frequency noise, but the approach is not well suited to slow baseline drifts.

We now describe an adaptive filtering algorithm that effectively discriminates actual WMS signals from unwanted time-varying etalon fringes. The approach, based on singular value decomposition (SVD) [9], is computationally straightforward, fast, and robust. The technique is demonstrated using

✉ Fax: 505-988-9230, E-mail: dbomse@swsciences.com

wavelength modulation spectroscopy (WMS) [1, 6, 10–13], near-infrared diode laser measurements of acetylene at atmospheric pressure. We emphasize the characterization and elimination of slow baseline drifts having timescales of hours and days and show that it is possible to reduce baseline drift without requiring frequent zero gas measurements.

2 Theory

The standard multivariate approach for extracting a gas concentration from a wavelength modulation spectrum is best expressed using a matrix formalism:

$$\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}\,, \tag{1}$$

where $\boldsymbol{b}$ is a vector comprising the N points of an experimentally acquired spectrum with one point per wavelength step, $\boldsymbol{A}$ is an $N \times M$ array ($N > M$) of reference waveforms that are the expected components of each experimentally acquired spectrum, and $\boldsymbol{x}$ is a vector of M unknown coefficients representing the amount of each of the reference waveforms (from $\boldsymbol{A}$) present in spectrum $\boldsymbol{b}$. One column of $\boldsymbol{A}$ is a low noise WMS spectrum of the target compound and the corresponding element of $\boldsymbol{x}$ is the gas concentration. The other columns of $\boldsymbol{A}$ are baseline components (dc offset, slope, and curvature). Additional spectra can be included for simultaneous measurement of multiple absorption peaks, and derivative spectra can be used to correct for laser wavelength drift or small changes in spectral line width.

There is a one-to-one correspondence between the elements of $\boldsymbol{x}$ and the amount of each reference spectrum that comprise the experimental spectrum, $\boldsymbol{b}$. Usually, for example, the first element of $\boldsymbol{x}$ is proportional to the intensity of the low noise WMS spectrum and, therefore, proportional to the target gas concentration [6].

A matrix solution to (1) is

$$\boldsymbol{x} = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T\boldsymbol{b}\,. \tag{2}$$

The computation is fast using conventional methods such as LU decomposition, Gauss–Jordan elimination, or singular value decomposition. In fact, the system of linear equations represented by the matrix notation need be solved only once for systems where the absorption line width remains relatively constant (*i.e.*, constant pressure and only small temperature changes). This approach assumes that each experimentally acquired spectrum is well represented by a linear combination of the columns of matrix $\boldsymbol{A}$, meaning that the components need to be known prior to starting the gas measurements. Etalons usually are not specifically included because they vary over time. This lack of accuracy in modeling the spectrum creates errors in gas concentration measurements as etalons may add to or subtract from the actual WMS $2f$ spectrum. Broad etalons, however, appear as baseline curvature while high frequency etalons act like high frequency noise and have little effect on the gas concentration calculation. It is the etalons in between – having free spectral ranges comparable to the absorption profile – that are troublesome, and are the subject of the adaptive algorithm described below.

Our adaptive approach is shown schematically in Fig. 1. The first matrix $\boldsymbol{A}$ is constructed at the start of the experiment by using known, low noise WM $2f$ and $3f$ spectra of the target compound for the first two columns. The $3f$ spectrum is included to correct for small wavelength drifts in the laser output [14]. The remaining columns of $\boldsymbol{A}$ are baseline (zero gas) wavelength modulation $2f$ spectra. These baselines are experimentally acquired and include whatever etalons are present at the time the spectra are acquired. The initial matrix $\boldsymbol{A}$ is used to analyze the first unknown spectrum $\boldsymbol{b}$. When the calculated $2f$ portion of the unknown spectrum $\boldsymbol{b}$ is subtracted from $\boldsymbol{b}$, the difference is the baseline portion of $\boldsymbol{b}$, and is used in the construction of the next matrix $\boldsymbol{A}$. The baseline spectra columns of $\boldsymbol{A}$ are treated like a first-in first-out queue, and $\boldsymbol{A}$ is continually updated to track baseline (i.e., etalon) changes. This provides the data analysis algorithm with the most recent and complete description of the wavelength modulation spectra. The computation is better able to distinguish between contributions due to the actual WMS $2f$ absorption and the baseline.

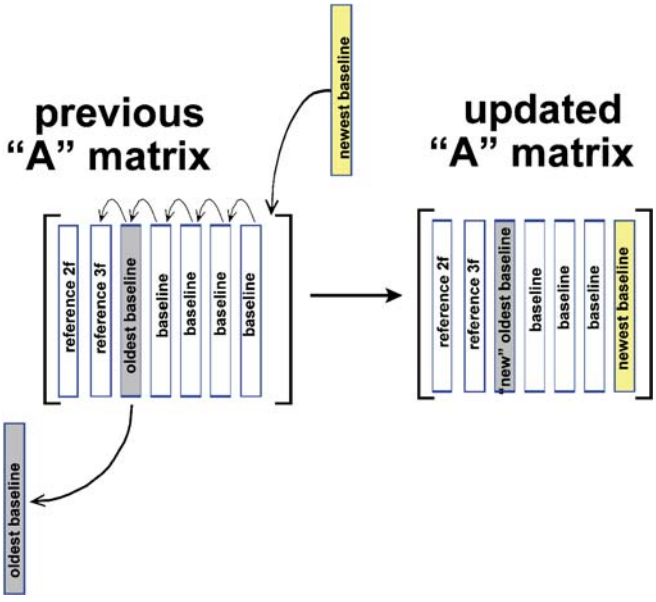

FIGURE 1 Schematic representation of the adaptive algorithm showing construction of, and revision to the rectangular matrix $\boldsymbol{A}$ that comprises the known, "low noise" $2f$ and $3f$ wavelength modulation spectra of the target gas plus baseline (zero gas) wavelength modulation $2f$ spectra. The baseline spectra are treated like a first-in-first-out queue, with the newest spectrum being added to the end while the oldest is dropped from the $\boldsymbol{A}$

The potential problem with this adaptive approach is that matrix $\boldsymbol{A}$ may become singular. Consider the case where the baseline is stable and the resulting columns of $\boldsymbol{A}$ are nearly identical. Successful implementation of our adaptive algorithm requires a method for solving (1) that can tolerate singularities: singular value decomposition (SVD) [9].

Singular value decomposition solves (1) by first factoring matrix $\boldsymbol{A}$ as

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{w}\boldsymbol{V}^T\,, \tag{3}$$

where $\boldsymbol{U}$ and $\boldsymbol{V}^T$ are both orthonormal, and $\boldsymbol{w}$ is diagonal. The coefficient vector $\boldsymbol{x}$ is found using

$$\boldsymbol{x} = \boldsymbol{V}\boldsymbol{w}^{-1}\boldsymbol{U}^T\boldsymbol{b}\,. \tag{4}$$

Matrix $\boldsymbol{U}$ is the basis set (eigenvectors) of $\boldsymbol{A}$, but not all of the basis vectors contain useful information. In our experience,

wavelength modulation spectra including etalons are well described by three or four basis functions. The rest are noise and are eliminated from the back-substitution step, (4), by setting the corresponding elements of $\boldsymbol{w}^{-1}$ (the eigenvalues of $\boldsymbol{U}$) to 0 instead of w_{ii}^{-1}. For this study, the useful w_{ii} were found to be $\geq 1\%$ of w_{00}.

The columns of $\boldsymbol{V}$ form the principal axes vectors of the multi-dimensional ellipsoids of constant fitting error, χ^2. The χ^2 value is defined in the usual way as

$$\chi^2 = |\boldsymbol{AB} - \boldsymbol{X}|^2 . \tag{5}$$

The boundaries of the ellipsoids are given [9] by

$$\Delta\chi^2 = w_0^2(\boldsymbol{V}_{(0)}\delta\boldsymbol{x})^2 + w_1^2(\boldsymbol{V}_{(1)}\delta\boldsymbol{x})^2 + \ldots + w_{M-1}^2(\boldsymbol{V}_{(M-1)}\delta\boldsymbol{x})^2 \tag{6}$$

where $\delta\boldsymbol{x}$ represents an incremental change in $\boldsymbol{x}$ and the w_i are the SVD eigenvalues (the diagonal elements of $\boldsymbol{w}$). Press et al. [6] show that it is possible to select a value of $\Delta\chi^2$ to match any target confidence level for a given number of degrees of freedom (size of M). This defines the SVD uncertainty for each eigenvector for the targeted overall confidence level.

Note that SVD is computationally equivalent [15] to principal component analysis (PCA), which is well known in chemometrics [16]. Our approach is similar to chemometric spectroscopic filtering methods that attempt to extract contributions due to target compounds from spectra that also contain absorbances due to multiple background species [17]. The key difference is that, in our case, the background features can not be well characterized prior to data acquisition. Instead, the time-varying etalon fringes are computed and removed in real time.

3 Experimental

Figure 2 is a schematic diagram of the optical and electronic experimental design. The laser was a 20 mW, NTT Electronics type NLK1556STG-BX distributed feedback (DFB) telecommunications diode laser having a nominal wavelength of 1529.6 nm and mounted in a butterfly package that also contained a thermoelectric heater/cooler, 10 kΩ thermistor, monitor photodiode, 30 dB optical isolator, and single mode output fiber equipped with an angle-polished, physical contact connector. An ILX Lightwave LDX-3722 diode laser controller provided laser temperature control and the laser drive current. The data acquisition system, based on a National Instruments E-series board, generated a symmetric triangle ramp (200 points total with a 2 kHz clock) that was combined using a differential amplifier with a sine wave generated by a lock-in amplifier (Stanford Research model 830), and then applied to the analog current control input of the diode laser controller. Typical operating conditions were a current sweep range from 80 to 140 mA plus a 25 kHz modulation at 15 mA peak-to-peak.

The laser output wavelength was confirmed using a Hewlett Packard model 86120C wavelength meter. Wavelength scanning rates and linearity were determined with a 0.22 cm^{-1} solid etalon.

The laser output was collimated using a fiber-coupled GRIN lens and directed through a non-polarizing, 50/50,

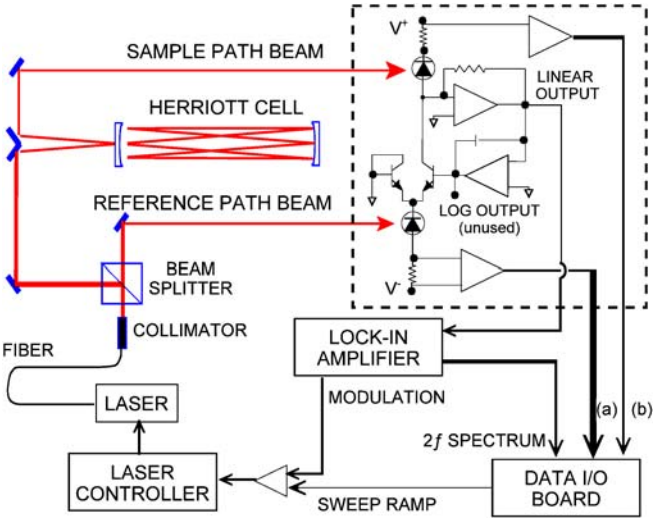

FIGURE 2 Schematic diagram of the optical and electronic apparatus. The Herriott cell used a pair of 25-mm-diameter, 10-cm focal length, protected silver spherical mirrors separated by 18.4 cm to provide 48 passes. The Haller–Hobbs noise-canceler circuit is highlighted by the *dotted line box*; for clarity we have omitted a cascode transistor used to reduced the sample path detector source capacitance, a 1 nF capacitor in parallel with the transistor that was added to prevent high frequency oscillation, and a unity gain buffer amplifier on the noise-canceler "LINEAR OUTPUT". The sample and reference path optical intensities were characterized by measuring the respective voltage drops across 1 Ω load resistors placed in series with the corresponding photodiodes. A pair of AD524 pin-programmable instrumentation amplifiers configured for 100× gain amplified the voltage drops due to the photocurrents, and the resulting output voltages were digitized by the National Instruments I/O board. Inputs (a) and (b) are the reference path and sample path intensities, respectively

5 mm cube beam splitter having exterior surfaces antireflection coated for 1550 nm. One portion of the beam was directed into a Herriott-style multiple pass optical cell [18] set for 48 passes at an 18.4 cm mirror separation. Light exiting the cell was focused by a 20-cm-focal length mirror onto an InGaAs photodiode (Fermionics 1000 W) having a 1 mm diameter active area. The photodiode connected to the sample side of a Haller–Hobbs noise canceller circuit [19, 20]; the circuit is highlighted in Fig. 2 by the dotted line box. The other portion of the beam was focused onto a second Fermionics model1000 W photodiode that was in the reference leg portion of the noise canceller. Our home-built noise canceller was identical to the original Haller–Hobbs circuit except that the photodiodes were biased at 3 V, a 1 μF ceramic capacitor was placed across the collector and emitter of the cascode transistor to prevent high frequency (several MHz) oscillation, and the photocurrents through each detector were measured by recording the voltage drops across 1 Ω sense resistors placed in series with the photodiode outputs. The cascode transistor and parallel capacitor are not included in the schematic diagram. The load resistors and their associated differential amplifiers are included; the voltage proportional to the reference path laser intensity is labeled (a) in Fig. 2 while the corresponding sample path signal is labeled (b).

Reflective losses within the Herriott cell reduced the sample leg laser intensity by almost exactly 50%. As a result the relative optical powers at the two noise canceller photodiodes were 1 : 2 for the sample and reference legs, respectively, which is optimal for circuit stability.

The linear (or high bandwidth) output from the noise canceller feedback loop was buffered by a unity gain inverting amplifier before being applied to the input of the lock-in amplifier. Wavelength modulation spectra were acquired by configuring the lock-in for second harmonic ($2f$) detection, and using the National Instruments board to record the analog output as the laser current was swept. Ten spectra were averaged to give an overall 1 s measurement time and the resulting spectra saved for later analysis.

A gas-tight cylindrical cover (2.8 cm i.d. ×20 cm long) over the Herriott cell kept the gas sample size to only 120 ml, allowing rapid exchange at flow rates $\sim 1\,\ell\,s^{-1}$. Acetylene mixtures were prepared by diluting a commercial, 101 ppm standard of acetylene in nitrogen with additional nitrogen or air using electronic flow controllers (Sierra Instruments) that were calibrated prior to use.

4 Results and discussion

Acetylene shows a strong peak at $6541.96\,cm^{-1}$ having an atmospheric pressure, room temperature absorption cross section of $2.15 \times 10^{-20}\,cm^2$, and has a significantly weaker peak $0.75\,cm^{-1}$ to lower frequency. These line positions agree well with Guelachvili and Rao [21]. The cross section was measured by direct absorption on the 101 ppm acetylene mixture over the 8.83 m Herriott cell optical path and is accurate to within ±20%.

Figure 3 shows the range of SVD eigenvalues (elements of $\boldsymbol{w}$) for all measurements. The best results required picking a cutoff value of $w_{min} \geq 0.01 w_{00}$. Larger values excluded all but one SVD eigenfunction and eliminated the benefit of the adaptive approach; smaller values just added noise as indicated by the lack of variation between corresponding eigenfunctions. Four basis functions – the $2f$ and $3f$ acetylene spectra plus two baseline spectra – are the minimum useful dataset size. Increasing to five by including one additional baseline spectrum provided a small improvement, but no further improvement occurred with larger data sets. The $2f$ and $3f$ acetylene spectra were acquired using the 101 ppm standard without additional dilution. The peak absorbance, $\alpha = 0.047$, is over 4000 times larger than the WMS noise level; therefore, the $2f$ and $3f$ spectra qualify as low noise reference spectra.

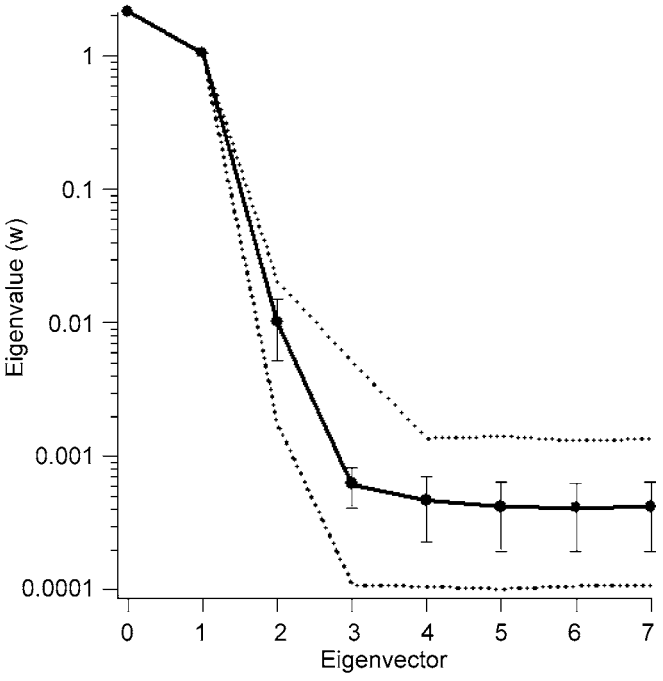

FIGURE 3 The range of SVD eigenvalues obtained for all experimental measurements using the adaptive algorithm. For this figure, eight spectra were included in each $\boldsymbol{A}$ matrix; the "low noise" $2f$ and $3f$ wavelength modulation spectra plus six sequential baseline spectra. The *solid line* is the mean value for each eigenvalue, the *error bars* show the standard deviations, and *dotted lines* are the maximum and minimum values

Five basis functions were used for all of the data analyses presented here. Each matrix $\boldsymbol{A}$ had dimensions 200×5. SVD computation times on a 3 GHz Pentium 4 desktop computer were 550 and 870 μs for 200-point spectra using 5 and 8 basis functions, respectively,

Acetylene concentrations were calculated as a linear combination of the first two eigenvalues:

$$[\text{acetylene (ppm)}] = 101.x_0 + 353.6x_1\,. \tag{7}$$

The coefficients were found by fitting a calibration data set that included 2.5 h of zero gas measurements as well as eight step changes in acetylene concentrations between 0 and 4 ppm (absorbances up to 1.8×10^{-3}). The calibration response was linear with $R^2 > 0.999$. Etalon fringe contributions are segregated into the third and higher eigenfunctions. SVD fit uncertainties, (6), were always significantly smaller than the point-to-point variations in concentration results. While this validated the SVD fit, it did not provide a useful concentration measurement drift.

Changes with time in the SVD eigenfunctions are shown in Fig. 4. Each panel includes the first four eigenfunctions. The three panels correspond to the start of the experiment and at 83 h (5000 min) and 166 h (10 000 min). In all cases, the first two eigenfunctions are dominated by linear combinations of the $2f$ and $3f$ reference spectra plus invariant portions of the system etalons, and remain unchanged. The third function shows the changes in the etalon fringes. The fourth is noise as indicated by the lack of clear structure and the predominance of high frequency signal variation. For clarity, the spectra in Fig. 4 show only the first 100 points in each spectrum. The laser wavelength was swept using a triangle ramp and the second half of each spectrum is a near mirror image of the points in the figure.

Key results are shown in Figs. 5 and 6. These data are taken from a week-long, continuous measurement of acetylene concentrations in gas flowing through the Herriott cell, and compare the results of the standard least squares analysis using a static reference matrix, $\boldsymbol{A}$, and the adaptive computational method obtained by continual revisions to matrix $\boldsymbol{A}$ plus singular value decomposition. Singular value decomposition computations used the adaptive algorithm, Fig. 1, with five basis functions. For both the static and adaptive methods, all columns of matrix $\boldsymbol{A}$ were normalized prior to computation. Figure 5 also shows measurement of the laboratory air temperature measured in the vicinity of the Herriott cell.

Most of the time, the sample contained 99.999% N_2 (zero gas) to best demonstrate baseline drift caused by unwanted

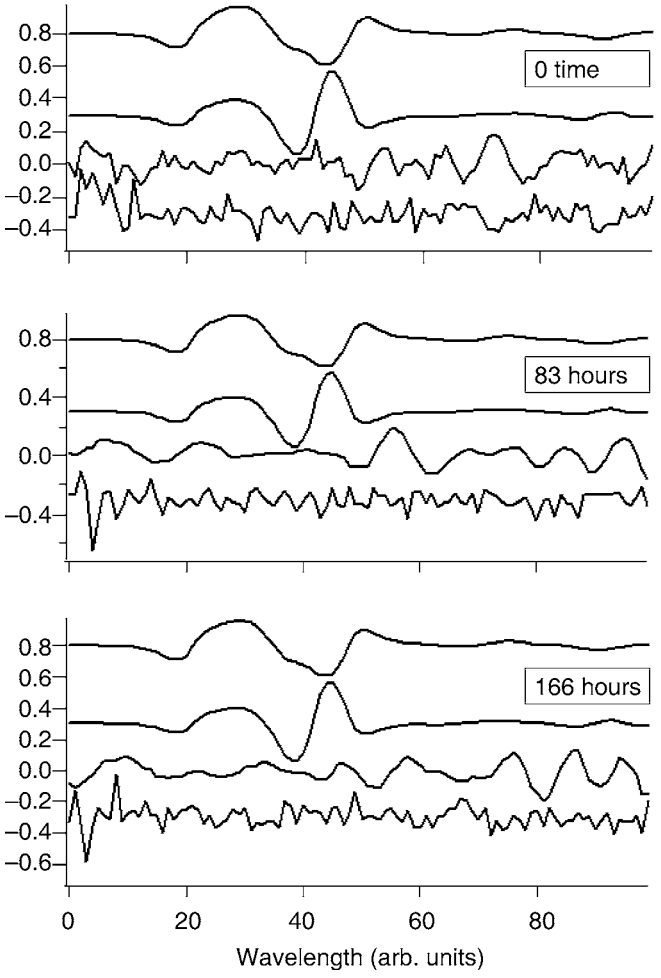

FIGURE 4 Variations in the first four SVD eigenfunctions at three different times during a one week long set of measurements. The first two eigenfunctions are dominated by the "low noise" $2f$ and $3f$ acetylene wavelength modulation spectrum, and do not change from panel to panel. The third function in each panel is dominated by etalon variations. The last function in each is almost exclusively due to noise

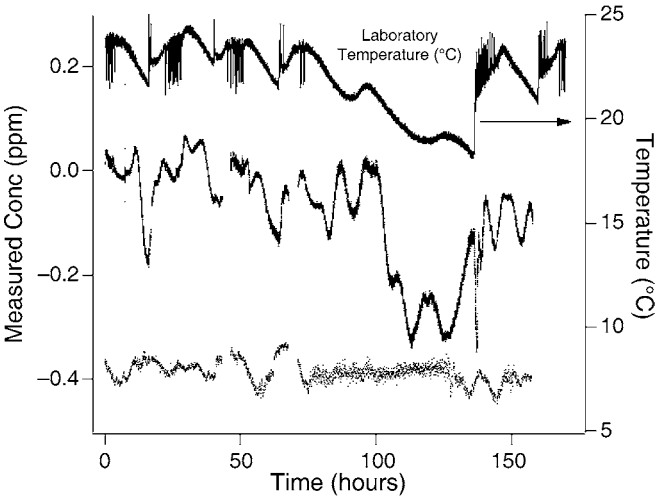

FIGURE 5 The zero gas measurement portions of a week long experiment comparing the static spectral analysis method (*center trace*) with the adaptive algorithm (*bottom trace*) The static algorithm has a mean value of 14 ppb and a standard deviation of 102 ppb ($\alpha = 4.5 \times 10^{-5}$). The adaptive algorithm data set has been offset by -400 ppb; it has a mean of -37 ppb and a standard deviation of 21 ppb ($\alpha = 9.3 \times 10^{-6}$). Gaps in the acetylene concentration data are due to periods of calibration and extended measurements of low concentration acetylene mixtures. The *upper trace* is the measured laboratory air temperature

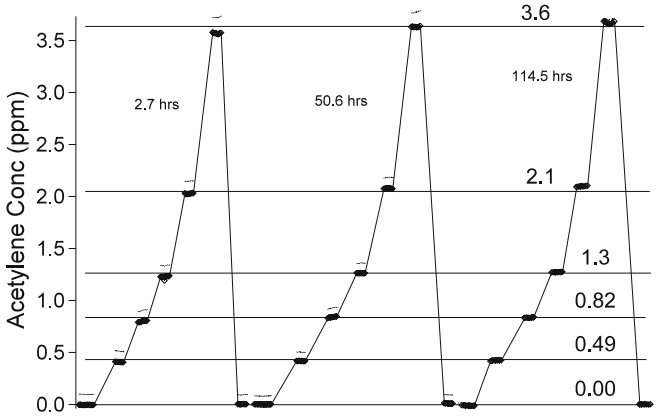

FIGURE 6 Comparison of responses to step changes in acetylene concentrations determined using the static (*black circle*) and adaptive (*open diamond*) algorithms. The adaptive algorithm data sections are connected with straight lines. Five such challenges were performed during the week-long experiment, and three sets of results are shown here

etalons. The experiment included six sets of step-response challenges at 2.6, 7.2, 16.8, 26.2, 50.6 and 114 h, and an extended period of measurement of 0.39 ppm acetylene ($\alpha = 1.8 \times 10^{-4}$) beginning at 48 h. Data were acquired at 1 s intervals, but only every 30th point is plotted. The zero gas results, Fig. 5, show a five-fold reduction in long term drift using the SVD approach compared with the traditional, static least squares fitting method. The static matrix $\boldsymbol{A}$ had 200×5 elements comprising the 101 ppm acetylene $2f$ and $3f$ spectra, plus offset, slope and curvature baseline components. Over the full week the standard deviations in the zero gas measurements were 102 ppb ($\alpha = 4.5 \times 10^{-5}$) and 21 ppb ($\alpha = 9.3 \times 10^{-6}$) for the static and adaptive methods, respectively. The adaptive algorithm shows some increase in high-frequency noise (e.g., the region between 85 and 135 h); but this is more than offset by the reduction in drift.

There is some correlation between the temperature changes and baseline drift obtained using the static algorithm. The region between 75 and 125 h is most intriguing. This covers a weekend when the laboratory heating/cooling system was turned off. The temperature changes are slow. Some diurnal variation is apparent, but the temperature data are free of sudden spikes due to the cooling system. The adaptive algorithm performs best here showing a standard deviation of 10 ppb ($\alpha = 4.4 \times 10^{-6}$) while, conversely, the static algorithm is at its worst with a large amount of baseline drift between 100 and 112 h. The slow temperature change implies a slow rate of etalon change, and the adaptive algorithm can track these baseline changes well.

The SVD adaptive algorithm does not alter measurement linearity. The six step challenge data sets gave nearly identical results. Once again, SVD performs better than the conventional data analysis method with most of the variations due to baseline fluctuations from etalons. Figure 6 shows three of the step challenges. The straight lines connect data analyzed using the adaptive algorithm while the solid dots (no connecting lines) show the static results. The average concentration

error from SVD and conventional analysis was 30 and 73 ppb, respectively, for all five step challenges. In all cases, including the step challenges not displayed, both methods showed excellent and repeatable linearity. The differences are entirely due to offsets arising from baseline changes (i.e., etalons). The agreement between the two methods seen in the 114.5 h data is due to coincidence. The etalon drift has, for this measurement period, returned the baseline to its zero value. The other data sets, however, show offsets ~ 100 ppb (absorbance differences of 4.5×10^{-5}) characteristic of systematic etalon-induced errors.

5 Conclusions

The singular value decomposition method shows a significant improvement in WMS measurement precision, and is fast and reliable. It can eliminate the need for regular zero gas baseline measurements and, thereby, reduce instrumentation cost and complexity. The SVD computation is linear; this avoids the need for a starting "guess" solution as is typical in iterative, non-linear algorithms [22]. Starting parameters obtained from an initial calibration and expressed in (7) remained effective throughout the week-long experiment. The SVD approach should remain well-behaved as long as the measurement period is short compared to the rate of etalon change which is, for the most part, dependent on the thermal stability of the optical train.

The results obtained here, a long-term (one week) minimum detectable absorbance of 9.3×10^{-6}, were obtained without any active etalon suppression methods. We note that a commercial instrument, Delta-F Corporation's DF-700 series diode-laser-based, ultra-trace moisture analyzers do include active etalon suppression and extensive signal averaging (~ 1000 s per measurement), and show at least a factor of two reduction in drift when implementing the adaptive algorithm.

The adaptive algorithm, Fig. 1, applies only to generation of reference matrix $\boldsymbol{A}$ implying that (1) can be solved using the least squares formalism, (2), instead of the SVD method. The real power of the SVD approach, however, comes from the ranking of importance of the eigenfunctions defined by the magnitudes of the eigenvalues in $\boldsymbol{w}$, the ease of removing the unwanted basis functions from the back-substitution calculation, and from SVD's robustness in handling (nearly) singular matrices. While the work presented used wavelength modulation spectra, the SVD approach is generally applicable to any form of high resolution spectroscopy in which time-varying optical interference fringes create a significant source of measurement uncertainty.

The biggest drawback to the SVD approach is that the basis functions in $\boldsymbol{U}$ are not always physically meaningful, and the gas concentration will be a linear combination of two or more such contributions.

ACKNOWLEDGEMENTS This work was supported by NASA SBIR Phase II contract number NAS9-02020 and US Department of Energy grant DE-FG02-05ER84302.

REFERENCES

1 J.A. Silver, Appl. Opt. **31**, 707 (1992)
2 J.A. Silver, A.C. Stanton, Appl. Opt. **27**, 1914 (1988)
3 C.R. Webster, J. Opt. Soc. Am. B **2**, 1464 (1985)
4 D.T. Cassidy, J. Reid, Appl. Phys. B **29**, 279 (1982)
5 P.W. Werle, P. Mazzinghi, F. D'Amato, M. De Rosa, K. Maurer, F. Slemr, Spectrochim. Acta A **60**, 1685 (2004)
6 D.C. Hovde, J.T. Hodges, G.E. Scace, J.A. Silver, Appl. Opt. **40**, 829 (2001)
7 D.P. Leleux, R. Claps, W. Chen, F.K. Tittel, T.L. Harman, Appl. Phys B **74**, 85 (2002)
8 J.A. Silver, W.R. Wood, Proc. SPIE **4817**, 82 (2002)
9 W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge, New York, 1986), Chapts. 2.9 and 14.3
10 D.S. Bomse, A.C. Stanton, J.A. Silver, Appl. Opt. **31**, 718 (1982)
11 D.B. Oh, M.E. Paige, D.S. Bomse, Appl. Opt. **37**, 2499 (1998)
12 J. Reid, D. Labrie, Appl. Phys. B **26**, 203 (1981)
13 T. Fernholz, H. Teichert, V. Ebert, Appl. Phys B **75**, 229 (2002)
14 D.C. Hovde, Filtering to Measure Gas Concentrations from Spectral Features, US Patent 6,615,142; issued Sep. 2, 2003. The full patent is available in PDF format free of charge at Southwest Sciences' web site: http://www.swsciences.com/technology/patents.html
15 M.E. Wall, A. Rechtsteiner, L.M. Rocha, *Singular Value Decomposition and Principal Component Analysis*: In: *A Practical Approach to Microarray Data Analysis*, ed. by D.P. Berrar, W. Dubitzky, M. Granzow (Kluwer, Norwell, MA, 2003)
16 M.A. Sharaf, D.L. Illman, B.R. Kowalski, *Chemometrics* (Wiley-Interscience, New York, 1986), Chapts. 5 and 6
17 A. Hayden, E. Niple, B. Boyce, Appl. Opt. **35**, 2802 (1996)
18 D.R. Herriott, J.J. Schulte Jr., Appl. Opt. **4**, 883 (1965)
19 P.C.D. Hobbs, Opt. Photon. News **2**, 17 (1991)
20 K.L. Haller, P.C.D. Hobbs, Proc. SPIE **1435**, 298 (1991)
21 G. Buelachvili, K.N. Rao, *Handbook of Infrared Standards II* (Academic, Orlando, 1993)
22 D.G. Lancaster, D. Richter, R.F. Curl, F.K. Tittel, Appl. Phys. B **67**, 339 (1998)