

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Algorithm for Rapid Tomography of Gas Concentrations

Permalink

<https://escholarship.org/uc/item/3q86200w>

Author

Price, P.N.

Publication Date

2008-12-05

Peer reviewed



ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY

Algorithm for Rapid Tomography of Gas Concentrations

P.N. Price, M.L. Fischer,
A.J. Gadgil, and R.G. Sextro

**Environmental Energy
Technologies Division**

June 2000

Submitted to
Atmospheric Environment

LOAN COPY
Circulates
For 4 weeks
7th Street Warehouse
Lawrence Berkeley National Laboratory

Algorithm for Rapid Tomography of Gas Concentrations

P.N. Price, M.L. Fischer, A.J. Gadgil, and R.G. Sextro

Environmental Energy Technologies Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, California 94720

June 2000

Algorithm for rapid tomography of gas concentrations.

Phillip N. Price*, Marc L. Fischer, Ashok J. Gadgil, and Richard G. Sextro
Indoor Environment Department, Lawrence Berkeley National Laboratory
Berkeley CA 94720

June 28, 2000

Abstract

We present a new computed tomography method, the low third derivative (LTD) method, that is particularly suited for reconstructing the spatial distribution of gas concentrations from path-integral data for a small number of optical paths. The method finds a spatial distribution of gas concentrations that (1) has path integrals that agree with measured path integrals, and (2) has a low third spatial derivative in each direction, at every point. The trade-off between (1) and (2) is controlled by an adjustable parameter, which can be set based on analysis of the path-integral data. The method produces a set of linear equations, which can be solved with a single matrix multiplication if the constraint that all concentrations must be positive is ignored; the method is therefore extremely rapid. Analysis of experimental data from thousands of concentration distributions shows that the method works nearly as well as Smooth Basis Function Minimization (the best method previously available), yet is 100 times faster.

Keywords: computed tomography, concentration mapping, airflow, pollutant dispersion

1 Introduction

For over a decade, researchers have investigated the use of computed tomography (CT) as a way of mapping tracer gas or pollutant distributions in air (Todd and Leith 1990, Yost et al. 1994, Drescher et al. 1996). Path-integrated concentrations of the gas of interest are measured—often, though not always, with a Fourier-transform infrared spectrometer (FTIR)—typically along a few dozen optical paths. A computer algorithm is used to solve the inverse problem of determining a spatial distribution of gas concentrations that could have produced the observed set of path integrals.

Due to the time required to orient the FTIR, and the time required to measure an individual optical path in order to obtain a sufficiently high signal-to-noise ratio, FTIR experiments have usually taken several minutes to measure all of the optical paths, and some have taken over an hour. Given these speed limitations in collecting the data, there has not been a strong need for CT algorithms that work quickly. In particular, Smooth Basis Function Minimization (SBFM), the method that has so far been most successful for CT of gas concentrations in air, is computationally intensive, typically requiring several minutes (on a Pentium-class 300 MHz personal computer running compiled Mathematica code) to generate a reconstruction.

A recent set of experiments, fully described by Fischer et al. (2000), used a new instrument that carries out a complete measurement cycle of 28 optical paths in only 7 s,

*Address correspondence to this author.

during which each path is sampled for about 150 ms. In each of these experiments data were collected for more than 30 minutes, thus generating data from over 300 measurement cycles. Complete CT reconstruction of all of the measurement cycles, using SBFM with our current computational machinery, takes over 7 hours. The delay between data collection and reconstruction precludes real-time monitoring of experimental conditions, which would be useful for trouble-shooting and for ensuring that the desired experimental conditions have been attained. Also, practical industrial applications such as monitoring chemical plants will require rapid CT.

In this paper, we introduce a reconstruction algorithm that uses techniques from Bayesian modeling. Our innovations are (1) including prior information in a way that tends to remove pixel-to-pixel oscillations in the reconstructions (which are usually artifactual) while still allowing large concentration gradients (which are common in reality), and (2) a mathematical statement of the CT problem that can often be solved in a single step, thus avoiding the need for iterative convergence.

2 Methods

Several computational methods have been suggested for reconstruction of gas concentrations based on path integral data. Most methods divide a planar area of interest into pixels and attempt to assign a gas concentration to each pixel in such a way that the predicted path integrals match the observed integrals, while also satisfying other criteria related to smoothness of the spatial distribution.

The problem of pixel-based CT is to predict the vector of pixel concentrations γ (where γ_j denotes the concentration in pixel j) from \mathbf{y} , the vector of measurements y_i of path-integrated concentrations along rays indexed with i . Let \mathbf{S} be the “system matrix,” so that $S_{i,j}$ is the path length of path i through pixel j . Predicted path integrals $\hat{\mathbf{y}}$ can be calculated from predicted pixel concentrations $\hat{\gamma}$ with

$$\hat{\mathbf{y}} = \mathbf{S}\hat{\gamma}. \quad (1)$$

With a sufficiently large number of rays, and if some conditions on the spatial arrangement of rays are met, then the CT problem can be solved by finding the set of pixel concentrations that best fit the ray integrals, subject to the constraint that $\gamma_j \geq 0$; that is, by finding the vector $\hat{\gamma}$ that minimizes

$$\phi(\gamma) = \sum_i w_i (y_i - [\mathbf{S}\gamma]_i)^2, \quad (2)$$

where w_i is the statistical weight given to measurement i , as we discuss below.

Unfortunately, minimizing the mean-squared difference between measured and predicted path integrals does not lead to a unique solution: in order to attain a spatial resolution that is high enough to be useful, the number of pixels will be much larger than the number of measured path integrals. For instance, the experiments performed by Fischer et al. use 28 optical paths in a plane, for a chamber of 65 m², so if each pixel represents 1m² the system is underdetermined by more than a factor of two. There is therefore no unique set of γ values that minimizes $\phi(\gamma)$, even under the constraint that all of the concentrations must be nonnegative.

Although many sets of pixel values can lead to the path integrals being fit about equally well, most are non-physical or at least highly improbable, for instance involving concentrations that alternate between high and low values for adjacent pixels. This problem has previously been noted and discussed (Todd and Ramachandran 1994, Drescher et al. 1996,

Park et al. 2000) for the Algebraic Reconstruction Technique (ART) and related methods (Herman et al. 1973), all of which generate reconstructions that are far too “noisy,” in the sense of having too much concentration variability on small spatial scales.

Drescher et al. (1996) developed SBFM to constrain the CT reconstruction to physically plausible solutions. SBFM writes the predicted concentration as a superposition of a small number of smooth basis functions, whose parameters are estimated so as to minimize the weighted or unweighted mean-squared difference between predicted and measured path integrals. Previous work has used 2-dimensional Gaussian distributions, for which the parameters are position, amplitude, width in each direction, and the angle between the x -axis and the major axis of the Gaussian (Drescher et al. 1996 and 1997, Price 1999, Hashmonay et al. 1999). High predicted concentration gradients can be attained with small values for Gaussian widths, but oscillatory solutions are avoided because the number of local maxima cannot be more than the number of Gaussian basis functions, which is usually restricted to 5 or less. (Drescher et al. 1996, suggested a method of determining the best number of Gaussians to use, but in practice we have had success by fixing the number at 4 or 5). SBFM works quite well, but is unfortunately computationally intensive, as the search for the best-fit set of parameters must contend with many local minima of the goodness-of-fit function. Current computational methods for SBFM rely on simulated annealing (Metropolis et al., 1953; see Press et al., 1986) and a single CT reconstruction takes 100–200 seconds on a 300 MHz computer running compiled Mathematica code. Although improvements in computer speed, re-coding in a more efficient computer language, or algorithmic improvements that eliminate the need for simulated annealing may eventually allow rapid SBFM reconstructions, that prospect is still in the future. Moreover, extending SBFM to three dimensions will greatly increase the computational burden. We don’t expect rapid 3-D SBFM to be attainable in the near future.

2.1 Including prior information in computed tomography

The fundamental problem of CT is to select the desired reconstruction (or reconstructions) from among the large number of unlikely or unphysical reconstructions that generate similar path integrals. Since the path integrals alone cannot determine the desired reconstruction, some additional information must be included. One approach, Bayesian modeling, uses so-called “informative prior information”: knowledge that is used to fit a model but that is external to the data at hand. One can think of informative prior information as addressing the question “before seeing any data, what can I say about the spatial distribution of concentrations?” Highly certain beliefs can be given high statistical weight, while less certain information can be given lower weight (and thus is easily overcome if contradictory data are available).

Although none of the previously suggested gas concentration CT methods were derived from a Bayesian approach, they do implicitly include prior information. For example, SBFM with Gaussian basis functions assumes that the gas concentration distribution has a small number of local maxima, falls away smoothly from these maxima, and indeed has a specific mathematical form.

Park et al. (2000) recently applied Penalized Weighted Least-Squares (PWLS), a technique developed by Sauer and Bauman (1993) and Fessler (1994), to data collected by Drescher et al. (1996). Using a “penalty function” suggested by Fessler in the context of medical imaging, they searched for reconstructions with small concentration differences between adjacent pixels. This penalty assumes that concentration gradients should be small,

with the strength of this assumption being controlled by a parameter β . Reconstructions using this method had systematic errors, such as substantially underestimating the peak concentrations.

Other methods, such as ART and its relatives, are harder to analyze in terms of prior information, but they nevertheless fit some implicit model of what the spatial distribution of the gas should look like.

Ideally, prior information would comprise a complete statistical description of gas concentration distributions. Currently no such description exists, and in fact it is difficult to picture what mathematical form such a description would take. In principle, almost any two-dimensional map *could* represent actual concentrations; for example, one could construct an experimental chamber with laminar upward airflow, and release gas into it from a grid of release locations. Any distribution of gas concentrations in a plane could be attained by this method. Thus, given the underdetermined nature of CT reconstructions using current experimental technology, there is no method that will work perfectly for all realizable concentration distributions. Instead, the goal is to find a method that works for the types of gas distributions that are likely to be encountered in practice.

Examination of gas concentration distributions in a plane as directly measured in experiments (Drescher et al. 1996, Drescher et al. 1997, Fischer et al. 2000), calculated with computational fluid dynamics (Gadgil et al. 2000), and simulated with dye distributions in a scale-model water tank (Gadgil et al. 2000), reveals several features:

1. concentrations can have very large spatial gradients;
2. where very large gradients occur they are usually near sources of gas, whereas areas far from all sources generally have lower gradients, even if they have fairly high gas concentrations;
3. even with a single source, concentrations can have several local maxima in a plane.

Any CT method must be able to accommodate at least the features mentioned above. SBFM fits this description, but at great computational cost. Is there an alternative?

We propose a CT method that seeks reconstructed concentration distributions in which the third spatial derivative of concentration is near zero in each direction, at every location. We refer to this approach as the “low third derivative” (LTD) method. If the third derivative is zero, then the second derivative is constant, so the concentration itself is a quadratic function of position. Setting the third derivative *exactly* to zero everywhere won’t allow good fits, as that would force a single global quadratic form to fit the whole concentration distribution, whereas the intent is instead to generate solutions that are *locally* quadratic, at least approximately.

2.2 LTD algorithm and computation

The notation γ_j for the pixel concentrations hides the spatial relationship of the pixels, so we introduce an alternative notation. The plane is gridded into pixels, with n_r rows and n_c columns. Pixels are numbered from 1 to $n_{\text{pix}} \equiv n_r \times n_c$, with pixel 1 in the upper left corner of a map, and pixel n_{pix} in the lower right. Pixels can be specified by row k and column l , with the concentration in a pixel written as $\gamma(k, l) \equiv \gamma_{n_c(k-1)+l}$. Note that k and l denote row and column, not position along the x and y axes. Each pixel corresponds to an area in the x - y plane, but we use the notation that is standard for matrices, not for algebraic geometry.

The concentration difference between adjacent pixels approximates the first spatial derivative of concentration at the midpoint of the pixels. In order to convert this derivative into a physical units, it must be divided by the distance between pixel centers; for now, we remain in pixel units, so this derivative represents the concentration change per pixel, not (for example) per meter. If the pixels are square (as we recommend, and as is the case in the data analyses discussed below) then this distinction is merely a matter of scaling

The difference between first derivatives is a measure of the second derivative, and the difference between second derivatives is a measure of the third derivative. For example, the third derivative in the l -direction at the juncture between pixels (k, l) and $(k, l + 1)$ is given by

$$\frac{d^3\gamma(k, l)}{dl^3} = \gamma(k, l + 2) - 3\gamma(k, l + 1) + 3\gamma(k, l) - \gamma(k, l - 1). \quad (3)$$

Any prior information that can be expressed as the expected value of a linear combination of pixel values can be included in a linear model through matrix augmentation, a standard technique of Bayesian regression (e.g. see Gelman et al., 1995). In essence, we are adding more terms to the sum in Eq. 2; these terms penalize $d^3\gamma/dl^3$ terms in proportion to their distance from zero. This is implemented as follows. Create a new matrix \mathbf{M} by appending rows to the system matrix \mathbf{S} ; a new data vector \mathbf{y}' by appending elements to \mathbf{y} ; and assign statistical weight w for each of the new rows. Each row appended to the system constitutes a mathematical statement; for example, the statement that $d^3\gamma(k, l)/dl^3 = 0$ is implemented by constructing a row r of the \mathbf{M} matrix that has 1, -3 , 3 , and 1 in columns corresponding to pixels $(k, l + 2)$ through $(k, l - 1)$, and setting the r th element of the \mathbf{y}' vector to zero; the statistical weight for this statement is controlled by w_r . We refer to the weight w_r as the “prior weight,” and to equations such as Eq. 3 as “prior equations.”

The third-derivative prior equations defined above cannot be used for pixels near the walls, in the direction perpendicular to the wall (because, for example, pixels $(k, l + 1)$ and $(k, l + 2)$ would be outside the boundary of the room for $l > n_c - 2$). Using third-derivative prior equations alone would thus leave the pixel values at the edges of the room relatively unconstrained, possibly allowing a lot of variation in concentration among those pixels, which would probably be non-physical. Several remedies are available; the one we selected is, for the pixels near the walls, to include prior equations so that the *second* derivative perpendicular to the walls is small, and assign this prior equation twice the statistical weight as is given to the third-derivative prior equations.

Including prior information for every pixel makes the matrix system overdetermined, as required for a least-squares solution: given some measured path integrals, plus prior information about the third derivative in each direction for every pixel, there are now more equations (rows of the \mathbf{M} matrix) than unknowns (pixel concentrations γ_i). Finding the pixel values $\hat{\gamma}$ that minimize $\phi(\gamma)$, the weighted squared difference between \mathbf{y}' and $\hat{\mathbf{y}}' \equiv \mathbf{M}\hat{\gamma}$, is just the problem of weighted least-squares regression. If we ignore the constraint that all $\gamma_j \geq 0$, the solution can be found analytically (see Gelman et al. 2000, for example): define

$$\mathbf{H} \equiv (\mathbf{M}^T \mathbf{W} \mathbf{M})^{-1} \mathbf{M}^T \mathbf{W}, \quad (4)$$

where \mathbf{W} is a diagonal matrix whose diagonal elements are w_i . (Non-diagonal weight matrices can be used to model covariance between pixels—the \mathbf{W} matrix is the inverse of the variance-covariance matrix—but we ignore that issue here). The value of $\phi(\gamma)$ is minimized by letting

$$\hat{\gamma} = \mathbf{H} \mathbf{y}'. \quad (5)$$

For most reconstructions, all of the predicted pixel concentrations are positive or only slightly negative (in which case we simply set them to zero).

Solving a linear least-squares problem by direct calculation of \mathbf{H} , as suggested by Equation 5, is ordinarily not recommended because much more computationally efficient methods are available—matrix inversion is a slow procedure for a large matrix. Although usually not recommended, in our case it is actually highly efficient to directly compute \mathbf{H} since the calculation needs to be done only once for a given arrangement of optical paths and set of prior equations and weights; then a matrix multiplication is all that is required in order to perform a reconstruction.

A statistical weight w_i must be assigned to each equation—that is, to each row of \mathbf{M} . Short optical paths tend to have low path integrals (e.g. in a completely mixed room, path integrals are proportional to path length), so weighting all optical paths equally would give shorter paths less influence on the reconstruction. Indeed, a (hypothetical) point sample would fail to influence the reconstruction at all! To counteract this effect, we assign weights inversely proportional to path length, scaled so that the longest path has weight 1. For our current experimental set-up this is not a major issue, since the longest paths are only a few times longer than the shortest. Setting all of the weights of the path-integral equations equal instead does not alter the solutions substantially in any of the cases we have examined.

We assign a statistical weight ω to each of the third-derivative prior equations (using the same ω for every pixel). For pixels near the walls, we constrain the second derivative perpendicular to the wall, with weight equal to 2ω . The prior weight for the third-derivative prior equations, ω , is then the only adjustable parameter.

The experiments of Fischer et al. (2000) included short-path measurements (which are nearly point measurements) throughout the plane. In principle, one could select the weight ω in order to maximize the agreement with the point measurements, but such reliance on point measurement would largely defeat the purpose of CT. Fortunately, alternatives exist. In particular, we select the highest weight for the prior equations that still allows a good fit to the measured long-path integrals. In fact, because the reconstructions vary only slowly with the prior weights, precise selection of the weights is not necessary. Instead, we calculate and store \mathbf{H} for three or four widely varying prior weights, and use the one that produces reconstructions that give suitable agreement with the long-path integrals.

2.3 Performing reconstructions from a complete experiment

To perform reconstructions for an entire experiment, consisting of many measurements of every optical path, the approach is as follows. First, calculate and store \mathbf{H} matrices for three or four widely varying prior weights (e.g. ω values of 1, 10, 100, 500). Then, for each time step:

1. Construct the data vector \mathbf{y}' , by appending the prior values (a vector of zeroes, in the present application) to the long-path measurements \mathbf{y} .
2. Calculate predicted pixel concentrations $\hat{\boldsymbol{\gamma}} = \mathbf{H}\mathbf{y}'$;
3. Handle negative predicted concentrations, either by setting them to zero or by using the predictions as the initial guess for a constrained optimization;
4. Calculate predicted long-path measurements $\hat{\mathbf{y}} = \mathbf{S}\hat{\boldsymbol{\gamma}}$;

5. Calculate the agreement between predicted and measured long-path integrals, $R_{\text{long}}^2 = \text{Correlation}(\mathbf{y}, \hat{\mathbf{y}})^2$;
 - (a) If $R_{\text{long}}^2 < 0.92$, choose a \mathbf{H} matrix that was pre-calculated with lower prior weights, and return to step 2;
 - (b) if $R_{\text{long}}^2 > 0.96$, choose a \mathbf{H} matrix with higher prior weights, and return to step 2;
6. Accept the reconstruction $\hat{\gamma}$. Read in the data for the next time step, and go to step 1.

2.4 Including other prior information

Matrix augmentation allows the use of prior information other than (or in addition to) the second- and third-derivative priors suggested above. For example, the functional equivalent of PWLS with the penalty function suggested in Fessler (1994), and applied to gas concentration CT by Park et al. (2000), can be attained by including prior information that the first derivatives should be small. This is implemented by augmenting the system matrix with one row for each pixel, and the data vector with one entry for each pixel, so that each new row corresponds to a statement such as $\gamma(k, l + 1) - \gamma(k, l) = 0$. The statistical weights for these rows play the role of Fessler’s β parameter.

Prior information on individual pixel concentrations is trivial to add; such information might come from a CT reconstruction based on earlier data, or from a computational fluid dynamics solution.

3 Results and Discussion

To investigate the performance of the LTD algorithm, we analyzed data from the experiments described in Fischer et al. (2000). In these experiments, a tracer gas (methane) was released from a square-meter area source near the floor of a $7\text{m} \times 9\text{m} \times 11\text{m}$ room. During every 7-second interval, the path-integrated methane concentration was optically measured along each of 28 “long-path” rays that cross the room in a plane about 2m above the floor. In addition to the long-path measurements, the concentration was determined along each of 28 0.5-meter “short-path” rays, using telescopes and receiving optics suspended from cables in the interior of the room. The short-path measurements are used to examine the performance of the CT reconstructions, as follows: (1) perform a reconstruction using only the long-path measurements, (2) calculate the predicted short-path concentration, based on the reconstruction, and (3) quantify the agreement between predicted and measured short-path concentration.

Figure 1 shows a map of the long-path rays. In the figure, each ray is plotted with a width proportional to its concentration measurement (the path integral divided by the path length) during a 7-second measurement cycle during which each path integral was measured once. This figure illustrates the input data to the CT algorithm.

Figure 2 shows CT reconstructions, along with short-path measurements, for four consecutive 7-second measurement cycles in experiment 3. Methane was released just above the floor near $x = 4\text{m}$, $y = 2\text{m}$; a persistent eddy in the experimental chamber carries the gas counterclockwise after release, so concentrations in the measurement plane tend to be highest along the wall at $x = 7\text{m}$.

No single goodness-of-fit parameter provides an adequate summary of fit between predictions and measurements. For simplicity, here we discuss R_{short}^2 , the coefficient of determination between the short-path predictions and measurements, since this addresses one

major question of interest: is the actual gas concentration low where the reconstruction says it is low, and high where the reconstruction says it is high? The R_{short}^2 values for the sequence of reconstructions shown in Figure 2 vary over a wide range: they are 0.59, 0.45, 0.62, and 0.77. For the experiment as a whole, the median R_{short}^2 value is 0.68.

The short-path data sometimes show rapid temporal variability. For example, the measured concentration for the short-path sensor near $x = 6\text{m}$, $y = 1\text{m}$ changes by more than a factor of 2.5 from the second to the third measurement cycle in Figure 2. Gas movement during the measurement cycle can affect both the reconstruction itself and the amount of agreement between the reconstruction and the short-path measurements. Each optical path is sampled for only 150 ms during each cycle, so the measurement along a path does not represent the average concentration over the 7-second interval that it takes to measure all of the path integrals. If a high-concentration wisp of gas passes through the measurement path during the short interval during which the path is measured, the path integral will be higher than the 7-second average. The temporal variability can thus cause a discrepancy between the short-path measurements and their predicted values based on the reconstruction, even if the reconstruction does accurately reproduce the time-average gas distribution.

Table 1 summarizes the performance of the method in reconstructing several experiments performed by Fischer et al. (2000). For each experiment, we performed reconstructions beginning with the first appearance of substantial gas concentrations in the measurement plane, and ending 30 minutes later. To help assess the significance of temporal variability, we performed two series of reconstructions. In the first series, each reconstruction was based on data collected over a 7-second interval, and thus uses one measurement of each optical path. In the second series, each reconstruction was based on a moving average of data collected over a 21-second interval, in 7-second steps (so that each path is measured three times, and the average of the three is used for the reconstruction). Reconstructions were compared to short-path data collected during the 7-second or 21-second time interval, respectively. Averaging over three measurement cycles removes some of the rapid temporal variation in the gas concentrations that can reduce the agreement between the CT reconstructions and the short-path measurements.

As Table 1 shows, averaging over three measurement cycles substantially improves the worst reconstructions (those with the lowest 5 to 10% of R_{short}^2 values), while leaving R_{short}^2 for most of the rest essentially unaffected. This suggests that temporal variation is one of the reasons the poorest reconstructions are as poor as they are, but such variability is not a major factor in causing discrepancies between predicted and actual concentration distributions for most of the reconstructions.

Somewhat remarkably, ignoring the constraint that all γ_i must be non-negative almost never causes a problem: in practice, the matrix solution to the unconstrained problem rarely predicts any pixel concentrations that are substantially negative (e.g. with magnitude greater than 0.1 times that of the predicted peak). As discussed above, when negative pixel predictions do occur, we set them to zero. We have also solved the constrained system when the unconstrained solutions generated negative predictions, but the resulting solutions took more computer time and were no better, in terms of agreement with the short-path measurements, than were obtained from the unconstrained solutions with negatives set to zero.

To compare the LTD method to SBFM, we calculated SBFM reconstructions for the same experiments whose LTD reconstructions are summarized in Table 1. The SBFM reconstructions are consistently slightly better: the SBFM R_{short}^2 values exceed those from the LTD method for every experiment and every quantile except for the 0.05 and 0.10 quantiles of Experiment 1. SBFM's superiority (by this measure) is remarkably similar for all quantiles

and all experiments: the SBFM R_{short}^2 value is higher by about 0.05 ± 0.02 . For example, whereas the median R_{short}^2 values when the LTD pixel method is applied to the four experiments in Table 1 are 0.64, 0.73, 0.68 and 0.70, respectively, the corresponding SBFM values are higher by 0.08, 0.06, 0.04 and 0.03. A similar pattern is present in the other quantiles.

The similarity between SBFM and the LTD method goes beyond the similar values of R_{short}^2 for each quantile; in fact, both methods perform similarly for each individual reconstruction—when one method produces relatively poor reconstructions, the other tends to do so as well. In terms of reconstruction accuracy, SBFM remains the gold standard for CT of gas concentrations by a narrow margin. However, the enormous speed advantage of the LTD method, which is more than 100 times faster than SBFM, makes the LTD method appealing for most applications in which large amounts of data require analysis or when real-time reconstructions are desired.

The LTD method produces CT reconstructions in which path integrals agree with measurements and third spatial derivatives of the gas concentration are low. The method works well for data we have analyzed so far: measured average concentrations over an array of short optical paths are in good agreement with predicted concentrations from CT reconstructions. Modifications might improve the method. For example, currently the same prior weight is used for every pixel, in each direction. Allowing the weight to vary with position could probably create better reconstructions, but for this approach to be useful it must be possible to determine the best spatial variation of weights by analysis of the path integral data alone, without reference to point-sample data. It is not obvious how that can be done. One possibility is to produce a reconstruction with uniform weights, calculate the third spatial derivatives of the reconstruction, and check the extent to which the derivatives are consistent with the prior equations. If there are areas over which most of the derivatives are of small (large) magnitude, better fits may be possible by decreasing (increasing) the prior weights for pixels in those areas. Alternatively, spatial correlation in derivative values could be incorporated by using a non-diagonal weight matrix W ; again, this will be a useful approach only if the spatial covariance of derivatives can be estimated without reference to point-sample data.

By combining hardware allowing measuring all 28 optical paths in an experimental chamber within about 7 seconds (see Fischer et al., 2000), with the LTD algorithm described here for performing very rapid CT reconstructions from the resulting data, we have taken CT mapping of air pollutant concentrations out of the realm of proof-of-principle experiments and into the world of actual application. We now use these tools at Lawrence Berkeley National Laboratory to investigate airflows and gas dispersion in a large experimental chamber.

4 Acknowledgements

We thank Mathias J. Craig and Emily E. Wood for assistance with the experiments, and Carrie Schwalbe for assistance with experimental set-up, data analysis, and graphics preparation. This work was supported by the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

- [1] Drescher, AC, Gadgil, AJ, Price, PN, and Nazaroff, WW. Novel approaches for tomographic reconstruction of gas concentration and distributions in air: use of smooth

- basis functions and simulated annealing. *Atmospheric Environment*. 30:929–940 (1996).
- [2] Drescher, AC, Park, DY, Yost, MG, Gadgil, AJ, Levine, SP, and Nazaroff, WW. Stationary and time dependent tracer-gas concentrations measured by OP-FTIR remote sensing and SBFM computed tomography. *Atmospheric Environment*. 31:727–740 (1997).
 - [3] Fessler, JA. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Transactions on Medical Imaging*, 13:290–300 (1994).
 - [4] Fischer, ML, Price, PN, Thatcher, TL, Schwalbe, CA, Craig, MJ, Wood, EE, Sextro, RG, and Gadgil, AJ. Rapid Measurements and Mapping of Tracer Gas Concentrations in a Large Indoor Space. Submitted to *Atmospheric Environment* (2000). Also available from Indoor Environment Department, Lawrence Berkeley National Laboratory, as report LBNL-45542.
 - [5] Gadgil, AJ, Finlayson, EU, Fischer, ML, Price, PN, Thatcher, TL, Craig, M, Hong, KH, Houseman, J, Schwalbe, CA, Wilson, D, Wood, JE, and Sextro, RG. Pollutant transport and dispersion in large indoor spaces: a status report for the large space effort of the Interiors Project. Lawrence Berkeley National Laboratory report, LBNL-44791 (2000).
 - [6] Gelman, A, Carlin, JB, Stern, HS, and Rubin, DB. *Bayesian Data Analysis*. New York: Chapman and Hall (1995).
 - [7] Hashmonay, RA, Yost, MG, Wu, CF. Computed tomography of air pollutants using radial scanning path-integrated optical remote sensing. *Atmospheric Environment* 33:267–274 (1999).
 - [8] Herman, GT, Lent, A, Rowland, SW. Art: Mathematics and applications. A report on the mathematical foundations and on the applicability to real data of the Algebraic Reconstruction Technique. *Journal of Theoretical Biology*. 42:1-33 (1973).
 - [9] Metropolis, N, Rosenbluth, A, Rosenbluth, M, Teller, A, and Teller, E. Equation of State Calculations by Fast Computing Machines. *Journal of Chemical Physics*, 21:1087-1092 (1953).
 - [10] Park, KY, Fessler, JA, Yost, MG, and Levine, SP. Tomographic reconstruction of tracer gas concentration profiles in a room with the use of a single OP-FTIR and two iterative algorithms: ART and PWLS. *Journal of the Air and Waste Management Association* 50:357–370 (2000).
 - [11] Press, WH, Flannery, BP, Teukolsky, SA, and Vetterling, WT. *Numerical Recipes: the art of scientific computing*. New York: Cambridge University Press (1986).
 - [12] Price, PN. Pollutant tomography using integrated concentration data from non-intersecting optical paths. *Atmospheric Environment* 33:275–280 (1999).
 - [13] Sauer, K, and Bauman, C. A local update strategy for iterative reconstruction from projections. *IEEE Transactions on Signal Processing* 39:534–548 (1993).
 - [14] Todd, LA, and Leith, D. Remote sensing and computed tomography in industrial hygiene. *American Industrial Hygiene Association Journal* 51: 1165–1178 (1990).

- [15] Todd, LA and Ramachandran, N. Evaluation of algorithms for tomographic reconstruction of chemical concentrations in indoor air. *American Industrial Hygiene Association Journal* 55:403–417 (1994).
- [16] Yost, MG, Gadgil, AJ, Drescher, AC, Zhou, Y, Simonds, MA, Levine, SP, Nazaroff, WW, and Saisan, PA. Imaging indoor tracer-gas concentrations with computed tomography: experimental results with a remote sensing FTIR system. *American Industrial Hygiene Association Journal* 55: 395-402 (1994).




Figure 1: Plan view of the experimental chamber of Fischer et al., showing long-path rays. Each ray's width is proportional to the average gas concentration along the ray (i.e. the ray integral divided by the path length) as measured during a 7-second interval of one experiment. Distances are in meters.

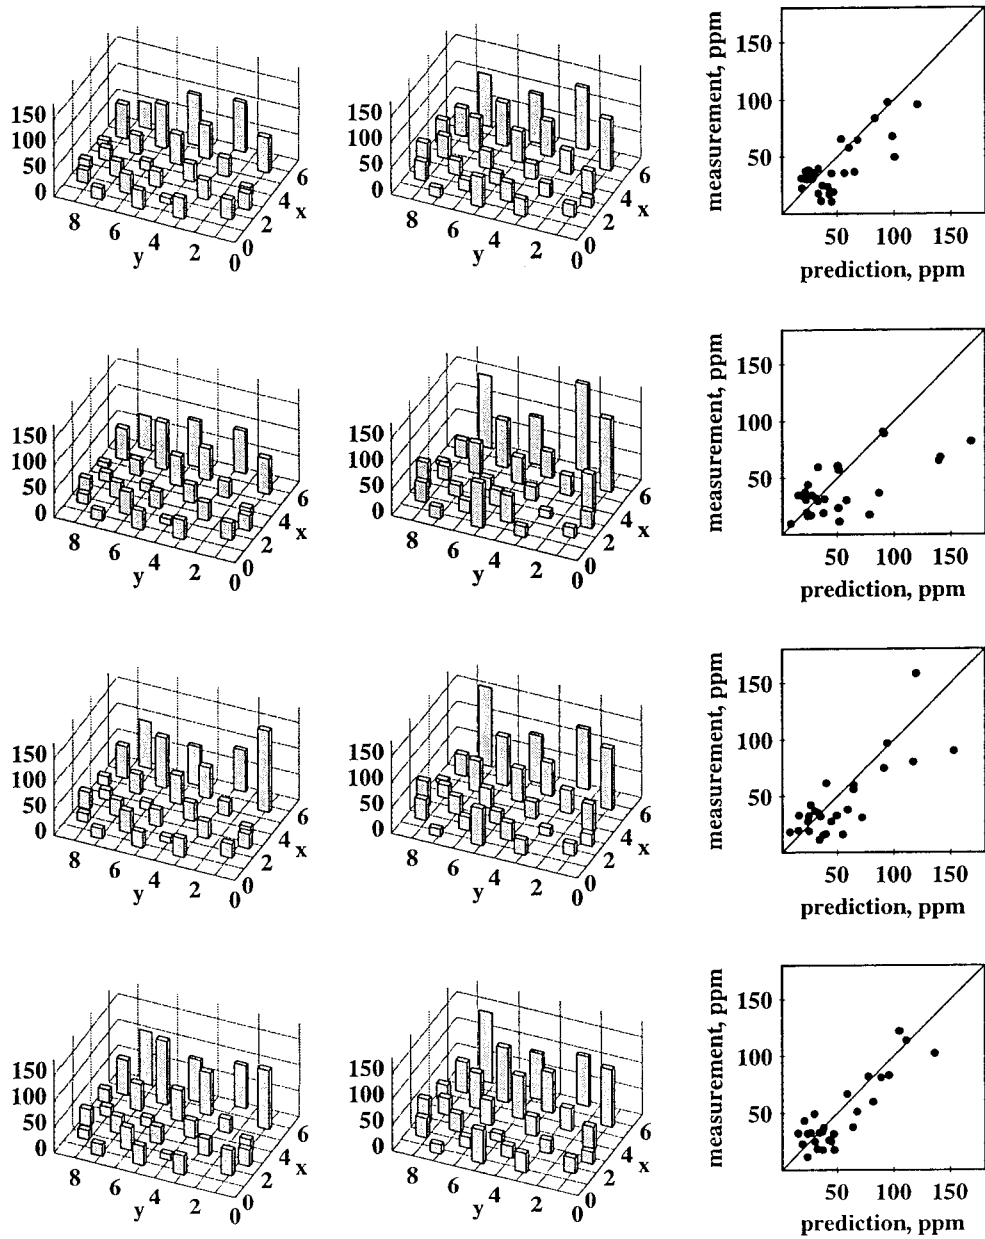


Figure 2: Measured (left column) and reconstructed (middle column) concentration at the location of each of the short-path sensors: the height of each bar shows the concentration (in ppm) at that sensor. Four consecutive 7-second time intervals are shown. The scatterplot in the right column plots measurement against prediction for each short-path sensor location. The diagonal line is the 1-to-1 line representing perfect agreement between prediction and measurement.

Table 1: Quantiles of R_{short}^2 for four of the experiments discussed by Fischer et al. (2000). For each experiment, reconstructions were performed from data in which each path integral was measured once (a 7-second interval), and from the average of three measurements (a 21-second interval). In each case, the reconstructions are based on data acquired for 30 minutes after the first appearance of tracer gas in the measurement plane.

Experiment number	Integration time (sec.)	Quantiles of R_{short}^2						
		0.05	0.10	0.25	0.50	0.75	0.90	0.95
1	7	0.34	0.46	0.57	0.64	0.73	0.79	0.82
	21	0.41	0.50	0.61	0.67	0.75	0.81	0.85
2	7	0.52	0.58	0.65	0.73	0.80	0.83	0.86
	21	0.58	0.62	0.68	0.76	0.81	0.85	0.87
3	7	0.32	0.43	0.57	0.68	0.74	0.79	0.82
	21	0.45	0.50	0.62	0.71	0.76	0.83	0.85
4	7	0.30	0.41	0.59	0.70	0.77	0.83	0.86
	21	0.37	0.49	0.63	0.73	0.80	0.85	0.87