

Performance Characteristics of a Silicon Photodiode (SiPD) Based Instrument for Fast Functional Optical Tomography

Christoph H. Schmitz, Mario Löcker, Joseph Lasker, Andreas H. Hielscher, and Randall L. Barbour

SUNY Downstate Medical Center, Brooklyn, NY 11203, U.S.A.

ABSTRACT

We describe the design rationale, performance features, and operating characteristics of a newly constructed CW-NIR tomographic imaging system that is capable of continuous, real-time imaging of large tissue structures. Results from phantom and clinical studies are presented and discussed.

Keywords: Dynamic imaging, Optical tomography, Instrumentation, Mammography, Optical encephalography

1. INTRODUCTION

We have introduced the concept that optical tomography, applied in a continuous, fast imaging mode, can serve as a new tool to study vascular reactivity in large tissue structures.¹ This idea is based on the observation that because hemoglobin is nearly always confined to the vascular space, and is a principal contrast agent in tissue in the near-infrared region, temporal variations in optical measurements can be interpreted as arising from the natural vascular rhythms. While information obtained from discrete or time-averaged imaging studies may prove useful in certain applications, we nevertheless recognize that the collection and analysis of essentially continuous measures can provide for an entirely new class of information about tissue function, *i.e.*, the spatiotemporal dynamics of vascular reactivity in large tissues.

Control of the vascular response has been studied for many years, and much is known (see accompanying reports by Graber *et al.*² and Landis *et al.*³, for additional discussion). For instance, it is well appreciated that tissue perfusion is not uniform in either space or time. This spatiotemporal heterogeneity arises from the combined influence that local metabolic, humoral and central autonomic control mechanisms have on vascular smooth muscle. These effects vary among the different structural elements of the vascular tree, but are principally targeted to the arterial side of the microvasculature (*i.e.*, arterioles).⁴ The influence of these effectors varies from moment to moment, and serves to establish a perfusion state that is richly varying yet tightly controlled. Derangements in the control mechanisms can exacerbate disease processes and in certain instances may even be the primary cause. While this relationship is widely appreciated, notably lacking is a suitable noninvasive technology that allows for characterization of these processes in large tissue structures. It is our view that the method of dynamic optical tomography is uniquely positioned to provide this information. Further, such measures might also provide valuable insights regarding intended or unintended influences that pharmaceutical agents can have on the vascular response. Similarly, we believe that this approach can be extended to other application domains where knowledge of the dynamics of highly turbid media would be useful. One area of interest is in the mixing of powders, which is a process important in the chemical and pharmaceutical industries. For instance, it is known that mixing for longer times does not necessarily produce better results and can in fact lead to the re-segregation of components. Information obtained by monitoring the dynamics of the moving solids and voids in a cross-sectional view might lead to devising protocols that in turn lead to more reproducible pharmaceutical

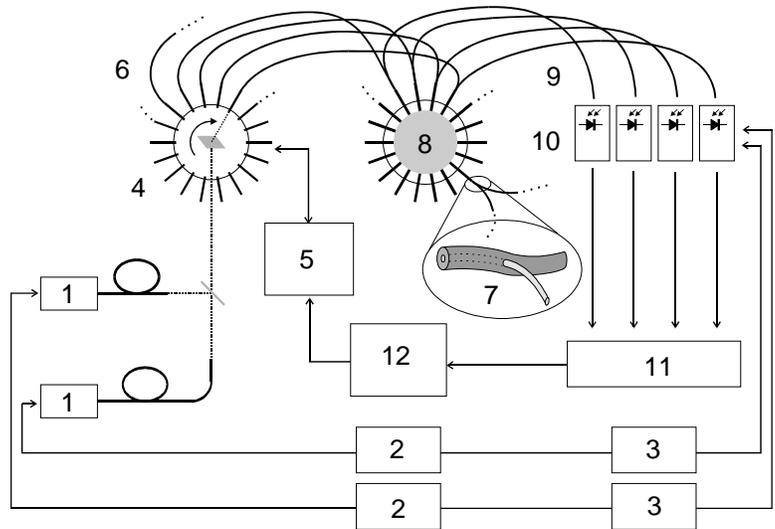


Figure 1: Block diagram of the instrument setup. 1) laser diodes, 2) laser drivers, 3) phase shifters, 4) D-MUX, 5) motor controller, 6) source branch, 7) bifurcated fiber bundle, 8) target, 9) detector branch, 10) detection, 11) data acquisition, 12) PC.

formulations.

To study dynamic behavior, an appropriately design data capture system must be available. In this report we discuss strategies we have adopted to accomplish this.

2. DESIGN RATIONALE

Our goal in system design has been to engineer a unit that provides for automated instrument step-up and fast data collection while retaining a modular design that can be scaled to meet specific applications. The system we adopted is based on a recently described design that provides for fast parallel data capture with a large dynamic range and employs DC illumination.¹ Our design approach also includes development of integrated system software to provide for (1) instrument setup and control, (2) image reconstruction and display and (3) interactive analysis of the time-series image data. While outside of the scope of this description, we have recently introduced the capability of real-time 2D and 3D image reconstruction using methods based on the descriptions given in an accompanying report by Pei *et al.*⁵ Evidence of these capabilities was presented in a series of live demonstrations given at the 4250th SPIE conference.

Figure 1 shows a block diagram of the instrument. The beams of two (or more) laser diodes (1) operating in the 800-nm range are collinearly overlaid and coupled into one of multiple source fiber bundles (6). Fast switching between source fiber bundles is made possible by means of an optical demultiplexer (D-MUX, 4). This unit is under microprocessor (5) control and employs a brushless DC servomotor to provide for fast, precise positioning. Each source fiber bundle (1-mm dia.) forms one branch of a bifurcated fiber bundle (7) and joins the other branch (3-mm dia., 9), which is used for light detection on the target surface (8). Each of the detector fiber bundles terminates at one silicon photodiode of a multi-channel detector module (10). The detector channels incorporate analog signal conditioning hardware such as adjustable gain stages, lock-in amplification, and sample-and-hold circuits in order to improve signal quality and for timing purposes. The output voltages of the detector channels are measured by means of a data acquisition board (11) and stored on a PC (12). For the purpose of lock-in detection, the laser diodes are current-modulated in the 5–10 kHz range by the laser drivers (2) and digital phase shifters (3) are used to optimize the phase angle between the measured signal and the reference. The mean optical output power of the laser diodes is in the range of 100 mW; the optical power incident on the target is about 30 mW.

2.1. Light Source

Fig. 2 shows a detailed view of the D-MUX. The unit houses a brushless DC servomotor that moves a gold-surfaced mirror, mounted 45° to its shaft, in a start-stop fashion. It is essential that the mirror make a complete stop for a short time (~10 ms) in order to minimize degradation of system precision due to variations in light intensity during the detection process. The chosen motor (Pacific Scientific model R23GENA-RS-NS-01), and accompanying microprocessor control unit (SC902AN-001-01) allows for flexible implementation of customized motion control protocols. This unit is capable of performing ~50 14°-start-stop motions per second without noticeable overshooting or ringing. The D-MUX unit shown houses as many as 25 source fiber bundles that are available for target illumination. The software-programmable motor controller allows absolute indexing of motor position and easy implementation of complex motion protocols. This allows for user selection of the number, timing and order of source fibers used for illumination. Because the maximum switching speed



Figure 2: Photograph of the optical D-MUX. A) complete view of the unit (height ~15") with optical fibers attached to it. B) Detailed view of the rotating mirror with the in-coupling optics removed.

between two source positions has an upper limit, increasing the number of sources will reduce the image-framing rate. For example, with 17 source positions and 50 Hz switching rate, ~3 full tomographic data sets per second can be acquired. As described in a recent report,¹ we have chosen to implement a time-multiplexing scheme instead of a previously described frequency-multiplexing scheme⁶ in

order to achieve a large dynamic range. This capability is critical for full-view tomographic measurements.

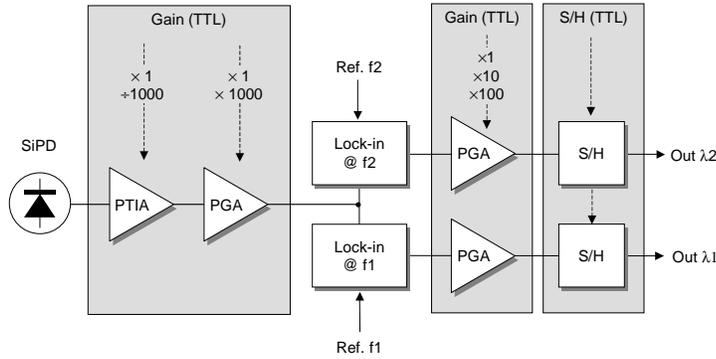


Figure 3: Block diagram of detector channel.

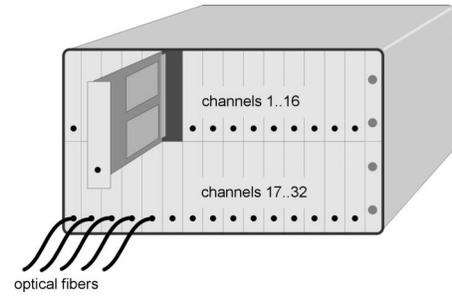


Figure 4: Schematic of the detector module.

2.2. Detection

Fig. 3 shows a block diagram of one detector channel. We have chosen to use a silicon photodiode (SiPD) as the photosensor because of its excellent linearity (better than $1:10^{10}$), high sensitivity ($6.5 \times 10^{-15} \text{ WHz}^{-1/2}$), and ease of operation. The signal amplification scheme and timing sequence used has been recently described in a report by Schmitz *et al.*¹ Motivating our approach was the need to achieve fast detection over a large dynamic range, coupled with fast source switching. This is achieved by synchronizing adjustment of the detector sensitivity for all channels with source movement, thereby achieving on-the-fly adaptive gain control.

Briefly, the signal amplification scheme used converts the photocurrent generated by the incident light to a voltage by means of a programmable transimpedance amplifier (PGIA) whose transimpedance value (gain), and hence photosensitivity, can be changed by a factor of one thousand. Subsequently, a programmable gain amplifier (PGA) can be set either to a gain of unity or to an additional gain factor of one thousand. Following appropriate amplification, the signal is then fed to one or more lock-in amplifiers (LIA, two are shown) to remove DC offsets, signals from ambient light, and other noise. By modulating the lasers at distinct frequencies, different wavelengths of light, simultaneously measured, can be distinguished.

Sample-and-hold circuits (S/H) are used to permit parallel data capture and to facilitate system timing. A second PGA located prior to each S/H circuit allows for adjusting the level of the demodulated (DC) signal, thereby improving noise immunity when transmitting analog signals to the data acquisition board. This also serves to effectively increase the usable dynamic range of the analog-to-digital conversion.

Fig. 4 shows a schematic of the multi-channel detector module. The unit is capable of measuring up to 32 channels in parallel. Each detector board can be equipped with either one or two lock-in amplifier modules, each carrying two LIAs. This allows for the simultaneous measurement of as many as four different light sources. Each detector channel is connected to a back plane, which provides supply voltages, control signals, and read-out lines. Depending on the number of modulation frequencies used, up to two phase shifter cards can be inserted, each containing two digital phase shifters. The detector module also contains connections to interface analog and digital ports from other hardware used in the instrument (DAQ board, power supply, motor controller). It also contains a system timing board that serves to store predetermined gain settings and establishes two-way communication with the motor controller unit.

2.3. System Timing

Precise instrument timing is crucial for the applications we seek to pursue because many different system tasks have to be scheduled in a way that allows for maximum image framing rates without sacrificing measurement precision. Fig. 5 shows schematically how the different system tasks proceed in time as the sources are switched. Tasks that have to be scheduled are: servomotor positioning, parallel multichannel light detection, and data capture. We have chosen to use the motor controller as the system clock to control these tasks. This unit contains several freely programmable digital I/O lines that allow synchronized communication of the motion protocol with external hardware.

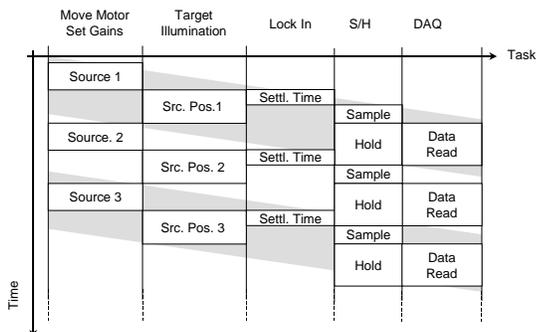


Figure 5: System timing diagram.

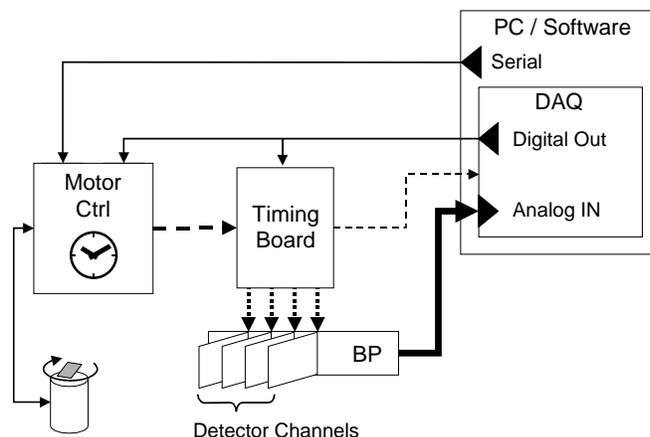


Figure 6: Hardware communication.

Optical signal detection is performed for all detectors in parallel to maximize the achievable frame rate. The detection process for the individual detector channel is composed of three sub-tasks: setting the detector to the appropriate gain, allowing the lock-in amplifier to settle, and triggering the S/H circuit. Adjusting the gain setting for each detector with respect to each source position is crucial because this provides the dynamic measurement range that is necessary to allow for high-quality signals when probing large tissue structures. Switching of these is achieved in a time well below 1 ms. After a detector has been set to its measurement range and the mirror has stopped at its new position, the actual measuring process, *i.e.* the settling of the lock-in amplifiers, takes place. The choice of a settling time (or detection bandwidth) for a lock-in amplifier is always a tradeoff between instrument response (ideally immediate, *i.e.*, fast-settling LIA) and the signal-to-noise ratio (SNR, which should be large, *i.e.*, slowly settling LIA). The consequences of this tradeoff for our application have been discussed in detail¹, and have led to a LIA design with a settling time of 6–7 ms. After this amount of time has elapsed, the S/H circuit is triggered, thereby freezing the signal level. This is followed by allowing the motor to proceed to the next position, while the detector channels are simultaneously read out by the data acquisition board. Depending on the number of detector channels/wavelengths used, this can require up to ~1 ms to complete. The entire cycle is repeated according to a pre-selected number of acquisition time points at a rate on the order of 50 Hz. This speed can be increased further to approximately 75 Hz by optimizing the motor performance.

The corresponding hardware implementation of the above timing scheme is shown as a block diagram in Fig. 6. The program determining the motion protocol is downloaded from the PC to the motor controller, which then waits idly until one of its control inputs is addressed by one of the digital I/O lines provided by the multi-function data acquisition (DAQ) board. These I/O lines are also used to write the gain settings for each source-detector combination to the timing board, where they are stored in a volatile memory prior to measurement. When executing the motion program, the motor controller moves the mirror to a source position, allows the appropriate time for the LIAs to settle (~7 ms), and then triggers the S/H circuit and data acquisition.

3. INSTRUMENT FUNCTIONALITY

3.1. Instrument Setup

The opto-mechanical components of the instrument have proven very stable: once optically aligned, the devices require no elaborate setup procedure. The instrument is operated intuitively by interacting with a graphical user interface that has been created using LabVIEWTM 5.0 software.

A measurement is prepared by a few minor setup procedures that can be performed on an initially appearing setup screen. Before performing a measurement, DC offsets are determined. Although lock-in detection is used, the detector channels show some residual DC offsets that can drift slowly and that vary for each channel and gain setting. By turning off the light source and performing a baseline measurement, these values can be determined and used for measurement correction. In the current version of the instrument, turning off the light source and initiating the offset measurement must be done manually; afterwards all measured data are corrected automatically. The offset correction can be skipped if desired. The user interface also allows adjustment of the phase shift for the reference signals. Because digital phase shifters are very stable, this procedure need be performed only once.

An important and potentially tedious task, should it be implemented using manual methods, is determination of the appropriate gain settings for all source–detector combinations. Which gain settings are appropriate depends on the particulars of the target. To facilitate this procedure we have implemented an automated routine that accomplishes this with a simple push of a button. We have tested the fidelity of this scheme and found that it can correctly and reproducibly identify the appropriate gain settings even on relatively dense media such as living tissue, and bring the signal levels for each detector to within a desirable measurement range (typically 0.3–5.0 V). The interface also allows the choosing of gain settings manually by moving virtual sliders that determine the detector gains for each source position. This functionality can be used to achieve specific gain settings or to check and overwrite the results from the automated setup. Gain settings can be stored to a file and reloaded as desired. Finally, the setup screen allows the specification of the file path and name under which the data will be stored.

3.2. Performing a Measurement

Data can be either acquired as raw data or normalized to an initial period over which the readings are averaged. The length of this period can be specified in number of time points. When starting a measurement, the raw data for all source–detector pairs are displayed in real time, forming a time–evolving 2D map that color codes deviations from the initial period. We have found this feature to be extremely helpful in providing real–time feedback to the operator, and we use it to assist in optimizing provocation protocols for clinical studies. This feature is also quite useful for identifying system problems such as inappropriate gain settings or degraded source–detector performance. The plot axes can be adjusted “on the fly” during the measurement to allow for improved visualization of the measurement data.

3.3. Instrument Configuration and Operation Modes

Depending on the needs of a particular application, it may be desirable to adjust the number of source positions used. Adding or omitting source positions is straightforward and is implemented using the motion–control software. In general, when adding a new source position, one has to determine the exact motor position at which the mirror must stop in order to maximize coupling efficiency and stability. We accomplish using a utility program that changes the mirror’s stop positions by a small amount every 20 measurement cycles, while acquiring data from a static phantom. By comparing the measured signal amplitude for each source as a function of motor position, optimal motor positions for each source can be identified. An example of this calibration is shown in Fig. 7.

A case of special interest is the use of a single source position. Because the mirror is not moved, only the last three tasks shown in Fig. 5 need be executed. As the data acquisition rate is primarily determined by settling time of the LIA, the theoretical upper limit for data acquisition rates in this mode is on the order of 150 Hz. In practice, this number is reduced due to additional time lags in the timing sequence. Using this approach we can routinely achieve parallel data readout rates continuously at a speed of more than 70 Hz. We are currently exploring this fast measurement mode for the ability to detect previously reported⁷ neural activated changes in the optical scattering properties of brain tissue.

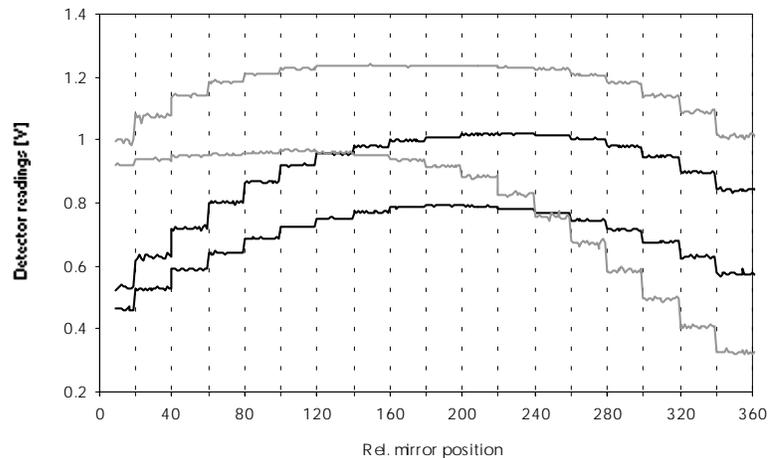


Figure 7: Calibration measurement for optimizing in–coupling efficiency and stability. Exemplary curves for four source positions are shown. The curve flanks show degrading in–coupling performance indicated by decreased signal amplitude and increased signal fluctuations.

4. INSTRUMENT PERFORMANCE

4.1. System Specification

Table 1 summarizes some of the more important performance characteristics of the instrument.

The detector sensitivity is defined by the noise equivalent power (NEP) as determined from dark noise measurements. The dynamic range of the instrument is defined by the ratio between the NEP and the signal saturation limit. Precision and long-term stability studies have been conducted and are discussed in detail below.

Because of the large dynamic range of measurement, optical cross-talk between detector channels can be important. This can occur as a result of light leakage within the D-MUX assembly, at the measuring head, and/or within the detector module itself. To test for leakage within the D-MUX, we compared the measured intensities of light exiting a target vessel when light from the source fiber is directed to the vessel or not, with the D-MUX illuminated by the laser source all the while. Leakage within the D-MUX can result in light reaching the target through other fiber bundles housed in the unit. Should the degree of leakage be significant, then corresponding appreciable intensities will be recorded even when the source fiber is not illuminating the target. We have observed that the cross-talk value is < 1 part in 10^4 for any source-detector configuration.

Cross-talk within the detector module can occur from light leakages occurring at the fiber-photodetector interfaces. This was quantified by comparing signal intensities when light from the target is directed to a detector channel and when it is not, with light received from the co-located source fiber (the largest signal) directed to one detector channel all the while. Comparison of the measured values showed that detector channel cross-talk was < 1 part in 10^6 for all source-detector pairs.

Fig. 8 shows the result of a linearity measurement. System linearity was determined by adjusting the source light intensity using calibrated neutral density filters and recording the detector readings over a wide dynamic range. Plotted are the averaged values for the readings from the source-co-located detectors, versus the nominal OD value.

Fig. 9 shows the results of system precision measurements, as a function of the detector gain. The columns indicate the mean value of the coefficient of variation (CV) of the signals acquired on a homogeneous, static phantom for data collected over a 2.6 min period (500 continuous measurement points) for different gain settings (1=lowest gain, 9=highest gain); the error bars indicate the standard deviation of the CVs. For gain settings 1-6 the fluctuations in detector readings are remarkably low, with CVs well below 1.0%. In this region the detector noise is very low, and source noise contributions (laser noise, in-coupling deficiencies) dominate. At gain 7 the CV value increases to 2.5% (corresponding to a SNR of 40). At high gain values, the measurement precision degrades exponentially, revealing increasing signal variability. This is due to increased

Table 1: Performance characteristics of the instrument

Parameter	Value
Detector Sensitivity	10 pW
Dynamic range	$1:10^9$ (180 dB)
Precision	CV $< 1\%$, all channels
Long-term stability	$< 1\%$ over 30 min
Cross-talk (D-MUX)	$< 1:10^4$
Cross-talk (detector channels)	$< 1:10^6$
Data acquisition rate	3.3 Hz @ 17 sources

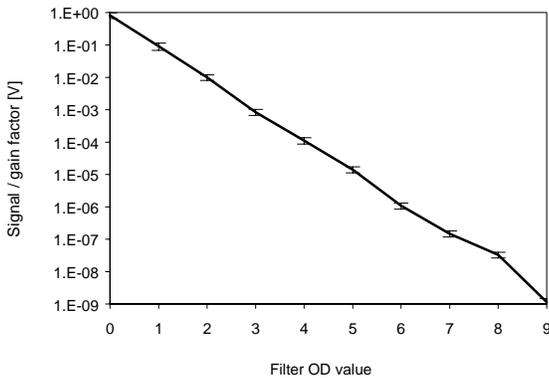


Figure 8: Evaluation of linearity performance and dynamic range of the instrument.

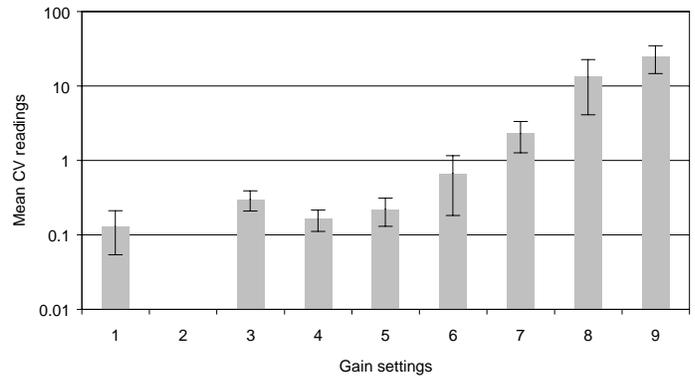


Figure 9: Precision study: Signal variability over 500 time points as a function of detector gain.

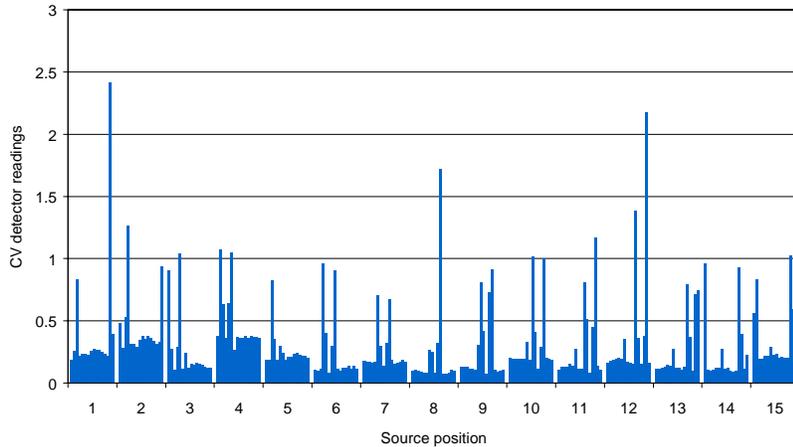


Figure 10: Long-term stability study (5000 time points). Each group of bars represent a source position. Individual bars indicate the stability of a particular detector at that source position.

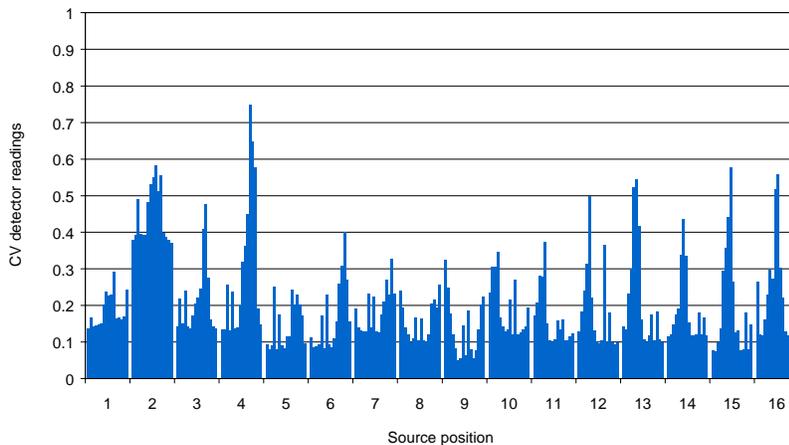


Figure 11: Short-term stability study (300 time points).

detector noise at higher gain settings, which limits overall measurement precision.

Fig. 10 shows the result of a long-term stability study that we conducted over a period of nearly 30 minutes (5,000 measurement points). The measurements were taken on a homogenous, static phantom having optical properties that required use of only the lowest 6 gain settings. Therefore, fluctuations in the signal can be attributed mainly to changes in in-coupling efficiency, laser power, and drifts in the detection electronics, rather than to detector noise. Results are presented in a bar graph format with the bars organized into groups according to the different source positions, with each bar representing the CV of one detector channel. This representation allows for easy identification of detector channels or source positions having notably good or bad performance. For example, the CVs for source no. 4 are systematically larger than for most other sources. Upon careful inspection, detector channels 11 and 15 can be shown to have slightly worse noise performance than all other channels. The majority of source-detector combinations show a long-term stability of better than a 1.0% CV, with most having values $<0.5\%$.

The short-term precision of the instrument is revealed in Fig. 11, where the results of a 360-data-point phantom measurements are plotted in a bar graph similar to that in Fig. 10. Using the same gain settings as in the long-term study, most CVs are now below 0.5%.

5. RESULTS

5.1. Breast Measurements

We have developed a folding hemispherical structure, intended as a measuring head for the female breast, that adapts to the instrument¹. We started initial studies on healthy volunteers to evaluate system performance such as the achievable signal levels and the quality of the fiber-tissue interface. We have further tested the ability to detect vascular responses to different stimuli in order to explore various examination protocols. Breast measurements were performed at a laser wavelength of 785 nm. Fig. 12 shows the time course of detector values for a selected source position. In this study, 1,200 consecutive time points were acquired at a frame rate of 3 Hz (6.8 min). Following a rest period of 250 time points, the volunteer was asked to place the hand contralateral to the breast under investigation into a container of ice water for the next 150 time points. This was followed by a recovery period of 200 acquisitions. Following this the subject was ask to initiate a series of slow, deep breathing exercises for 250 time points. This was then followed by a breath hold, which lasted for 100 time points. The last 250 points show a second recovery period.

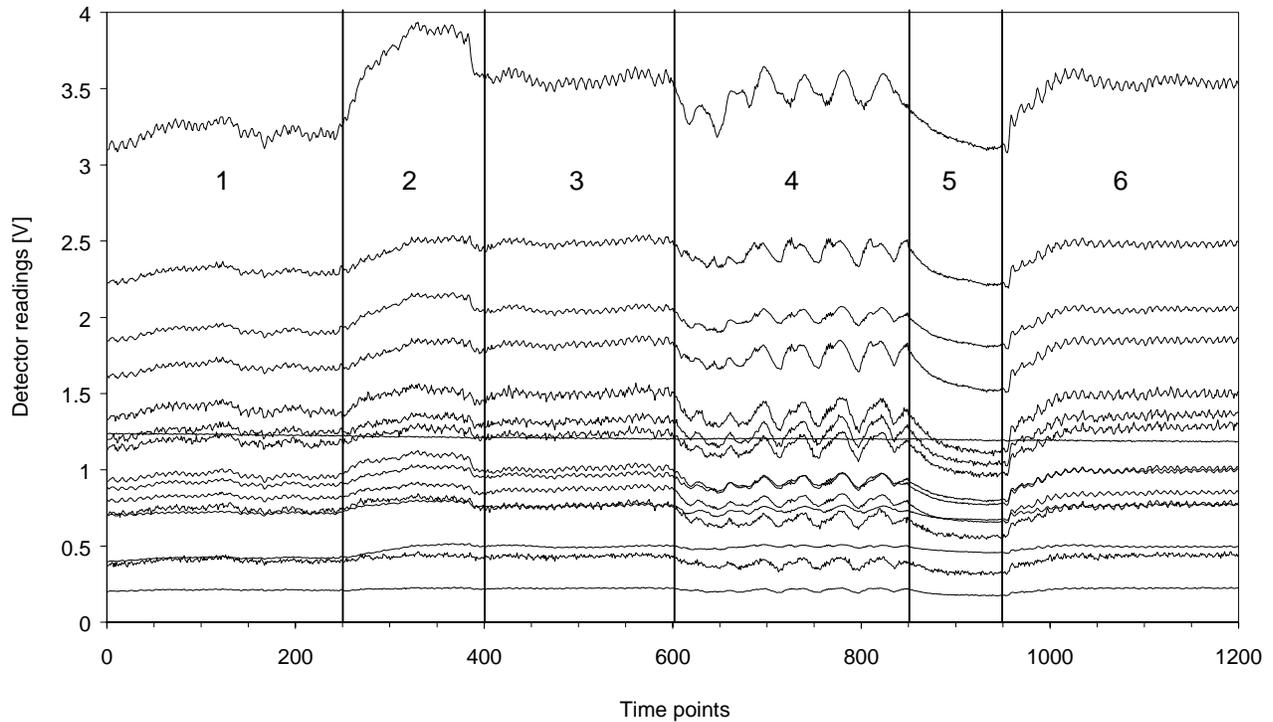


Figure 12: Detector readings for one exemplary source position obtained from an *in vivo* breast measurement on a healthy volunteer. Largest source–detector separation was ~ 8 cm. The response to various exercises is clearly revealed: 1) rest, 2) cold shock, 3) recovery, 4) deep breathing, 5) breath hold, 6) recovery.

The signals obtained are of high quality and reveal a host of physiological dynamics. Throughout most of the measurement, a prominent frequency at ~ 0.3 Hz is observed, which corresponds to the normal respiratory rate of the subject. During the deep–breathing study the respiratory signal changes in amplitude and frequency according to the breathing activity. The breath hold abolishes the respiratory rhythm and reveals an expected time–dependent vasocongestion occurring in response to increased venous resistance. Close inspection of the data shows a cardiac frequency, especially during the deep–breathing and breath–hold periods. The data show also a low–frequency modulation in the ~ 0.1 Hz range, at which vasomotion is known to occur.

5.2. Real Time Dynamic Optical Tomographic Imaging

Using the instrument in conjunction with image reconstruction methods based on those described in an accompanying report,⁵ we were able to demonstrate 2D and 3D optical dynamic imaging in real time. For this purpose, we set up an Ethernet connection between the DAQ computer and an SGI Origin 200 UNIX workstation and created a data–transfer protocol that establishes the transfer of individual data sets for each time point. For each of these data sets the reconstruction

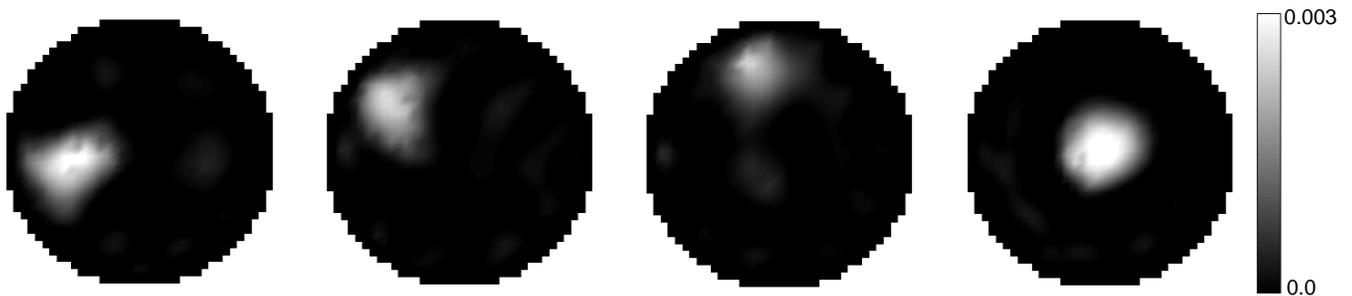


Figure 13: Four maps of changes in μ_a of the real–time reconstructed image time series of an $80\text{-}\mu\text{m}$ –diameter wire moving in a 1%–Intralipid background. Maps show frames 113, 118, 123, and 133 of the image series.

algorithm is launched, and the results are displayed in a color-encoded map. The time required to reconstruct a single image is ~ 30 ms for 2D and ~ 300 ms for 3D (number of unknowns $\approx 5,000$). Fig. 13 shows the time evolution of the reconstructed images obtained from a 2D phantom study. In this experiment, an $80\text{-}\mu\text{m}$ -diameter titanium wire was stirred slowly in a 1% IntralipidTM suspension contained in a cylindrical DelrinTM cup of 70-mm inner diameter (76-mm outer dia.), causing changes in detector readings on the order of 1%. This result demonstrates not only the computational speed of the algorithm but also its robustness against measurement noise in the presence of weak signals. The identified spatial location of the wire was accurate over the entire cross-sectional area of the phantom, although the dimensions of the wire are strongly overestimated by the reconstruction method.

6. DISCUSSION

In this report we describe the design rationale and performance features of a newly developed automated optical tomographic system that employs CW illumination. The system is specifically designed to acquire continuous measurements at high framing rates and with high dynamic range. The system is modular in design and employs scalable hardware. System control is accomplished using a virtual control panel operated within a LabVIEW environment. This allows for automated system setup and user selection of a broad range of measurement parameters. The instrument also features a fully integrated real-time image reconstruction package, together with software for interactive time-series image analysis (not shown). In addition, a number of measuring heads have been developed for specific applications, as previously reported.¹ The modular design allows for housing of the system within various formats. A photograph of the system housed within a portable instrument cart is shown in Fig. 14. System performance studies reveal excellent precision and stability and considerable sensitivity to various simple homeostatic vascular challenges. Studies examining the utility of acquired data for quantifying and characterizing vascular reactivity in various tissues (*e.g.*, breast, head, limbs) is underway.



Figure 14: Photograph of the dynamic optical tomographic imager developed in our lab.

7. ACKNOWLEDGEMENTS

This research was supported by NIH grants RO1-CA 66184, R01 AR46255-01, and the NYC Council Speaker's Fund for Biomedical Research: Towards the Science of Patient Care.

REFERENCES

1. C. H. Schmitz, H. L. Graber, H. Luo, I. Arif, J. Hira, Y. Pei, A. Bluestone, S. Zhong, R. Andronica, I. Soller, N. Ramirez, S.-L. S. Barbour, and R. L. Barbour, "Instrumentation and calibration protocol for imaging dynamic features in dense-scattering media by optical tomography," *Applied Optics*, **39**, 6466-6485 (2000).
2. H. L. Graber, Y. Pei, and R. L. Barbour, "Imaging of spatiotemporal coincident states by dynamic optical tomography," Proc. SPIE 4250 (2001, accompanying report).
3. G. Landis, S. Blattman, T. Panetta, C. H. Schmitz, H. L. Graber, Y. Pei, and R. L. Barbour, "Clinical applications of dynamic optical tomography in vascular disease," Proc. SPIE 4250 (2001, accompanying report).
4. J. Ross, Jr., "Structure-Function Relations in the Peripheral Circulation," Chap. 6 in *Best and Taylor's Physiological Basis of Medical Practice*, 11th Ed., J. B. West (Ed.), Williams & Wilkins, Baltimore (1985).
5. Y. Pei, H. L. Graber, and R.L. Barbour, "Normalized-constraint method for minimizing inter-parameter cross-talk in reconstructed images of spatially heterogeneous scattering and absorption coefficients," Proc. SPIE 4250 (2001, accompanying report).
6. A. M. Siegel, J. J. A. Marota, and D. A. Boas, "Design and evaluation of a continuous-wave diffuse optical tomography system," *Optics Express* **4**, 287-298 (1999)
7. A. Villringer and B. Chance, "Non-invasive optical spectroscopy and imaging of human brain function," *Trends in Neuroscience* **20**, 435-442 (1997).