

4: The design and evolution of CW DOT instrumentation

Accurate temporal measurement of cerebral hemodynamic activity requires DOT instrumentation capable of acquiring noninvasive real-time imagery with minimal temporal distortion. Since off-the-shelf DOT instruments are not commercially available, most researchers either design and build their own equipment or collaborate with other laboratories which already have DOT instruments. However the ambiguities introduced by using borrowed instrumentation with nebulous performance specifications and little, if any, formal documentation leads most serious researchers in the field of photon migration imaging to pursue the design of their own equipment.

But before any instrumentation can be built, the system requirements must first be understood so that the proper design trades can be made. Therefore Chapter 4 begins by introducing the spectral, temporal, and optical flux ranges over which DOT measurements are performed. Optical sources and detectors are then discussed with respect to their suitability for use in DOT instrumentation.

Tomographic optical measurements, by their nature, require some form of source encoding to disambiguate the various optical sources. Thus the need for source encoding is explained and a number of popular source encoding techniques are described. An attempt at improving methods of source encoding led to the development of a new optical encoding technique, Pulse-TDM, which allows for advanced capabilities like individualized gain control while providing very high dynamic range and minimal temporal skew in the optical imagery. The concept of individualized gain control is then introduced and is validated through an in-vivo measurement.

The design evolution of a variety of DOT instruments is then presented, ranging from CW1, the first tomographic system we constructed, to CW4, the flexible high-speed frequency encoded DOT system used to perform the temporal hemodynamic response measurements discussed later in this dissertation. Lessons learned from the evaluation and use of each system were applied towards the design of future systems. The results of in-vivo performance evaluations and laboratory measurements used to characterize some of these instruments are also included in Chapter 4.

4.1 Spectral, temporal, and flux ranges for DOT

Understanding the limits over which an instrument must operate is critical to developing a practical design. The design requirements for CW DOT instrumentation are unique in requiring both the multiplexing of many optical sources and the detection of optical flux over a very wide dynamic range. In partial compensation for this, the near-infrared spectral range permits the use of efficient optical sources and sensitive detectors, and the relatively low temporal bandwidth enables a low noise floor.

Spectral range

The spectral range of DOT measurements generally extends from 600nm to 1000nm. Optical absorption from Hb and, to a lesser extent, HbO₂ limits measurements below ~600nm. A combination of increasing water absorption, poor detector responsivity, and low spectral contrast limits the utility of measurements beyond 1000nm.

Temporal range

DOT measurements can span the frequency range from 10mHz up to 20Hz or greater. Hemodynamic measurements typically range from 0.1Hz to 10Hz. The upper bound is determined by the vascular impulse response and the lower bound is determined either by the duration of the interstimulus

interval, or by temporal factors unique to each experimental paradigm. Unfortunately the hemodynamic signals of interest occupy a frequency spectrum which includes other biogenic signals, as shown in Table 4.1.

Table 4.1. The temporal ranges of most biogenic signals of interest to DOT. The typical hemodynamic modulation depth is around 1%. Although actual signal strengths vary greatly as a function of probe location and body position, pulsatile modulation is often the strongest signal, followed by the respiration signal. Mayer wave modulation can range from undetectable levels to as much as a few percent. Measurements of intrinsic scattering changes (“fast” signals) detect small changes in membrane and cytosolic index caused by neural activity, and require the highest temporal response and sensitivity. These measurements can extend from 10Hz to as much as 500Hz or more, however the signal is so weak as to be virtually undetectable.

Measured Parameter	Response Timescale	Frequency Band
Heart rate	~1 sec period	1-3 Hz
Respiration	~5 sec period	0.2-0.3 Hz
Mayer waves	~10 sec period	0.08-0.12 Hz
“Fast” signal	~0.050 sec	20+ Hz
Hemodynamic response	~0.2-10 sec	0.1-5 Hz

Optical flux range

DOT measurements cover a very wide dynamic range – from the shot noise limit at the low end all the way to microwatts of detected power at the high end, as shown in Table 4.2. This places great demands on the detection circuitry of DOT instrumentation designed for use with human subjects.

4.2 Sources and detectors for DOT and NIRS

4.2.1 Optical Sources

The two most common optical sources used for DOT are laser diodes and light emitting diodes (LEDs). Most are fabricated using a specially-doped layered structure consisting of either a ternary or a quaternary alloy with the empirical formula: $\text{In}_{1-x}\text{Ga}_x\text{As}_{1-y}\text{P}$. CW laser diodes are the most common optical sources in DOT instruments today. Modern multiple quantum-well devices are efficient, linear, powerful, compact, provide a very high flux density, and have a narrow spectral linewidth. They also have excellent temporal response and are easy to modulate, but they must never be overdriven or else permanent damage to the cleaved exit facets may result.

LEDs are inexpensive, efficient, and rugged, but as DOT sources they do have their limitations. They are available in a wide range of wavelengths since indirect bandgap materials can be used, however their spectral bandwidth is wide enough (~ 20nm to 40nm FWHM) to complicate blood oxygenation calculations at the wavelengths most suitable for DOT. Blood volume measurements at ~800nm are more forgiving than multispectral oxygenation measurements, since the relative absorption changes for Hb and HbO₂ complement each other across the isosbestic point, so

measurements with an 805nm LED should be about as accurate as an 805nm laser diode for determining blood volume.

Laser diodes

CW laser diodes are available in two basic varieties: single mode and multimode. The gain-per-unit length within the semiconductor material is extremely high, enabling the creation of completely monolithic lasing structures. The end mirrors are often the cleaved facets of the die itself, since the index mismatch between InGaAs and air reflects about 30% of the light back into the cavity, which is ideal for the output coupler of a semiconductor laser [76, 77]. If electrical-to-optical efficiency is important, the back facet can be coated for high reflectivity to reduce the backside loss, however there must be some light leakage for the monitor photodiode to provide a feedback signal to the drive circuitry if closed-loop operation is desired.

In order to achieve the areal current densities required for lasing at high efficiency (and to prevent the material from thermally degrading), the dimensions of the active area must be confined to a narrow region within the material. Vertical confinement of the light is usually achieved through index-guiding, using a dual heterostructure (two-junction) design. This generates a sandwich-like structure with a thin high-index gain layer. The high index acts to confine the flux within a very thin (<1um) vertical waveguide region, in the same manner as light is confined within the core of a graded-index fiber. Single mode lasers achieve horizontal confinement through a similar index-guiding mechanism and multimode lasers are gain-guided, as depicted in Figure 4.1, however the actual output spectrum is a function of both output power and temperature, as shown in Figure 4.2 [77]. Since multimode lasers provide better temporal and spectral stability, they are preferred for DOT applications.




Figure 4.1. Vertical and horizontal confinement mechanisms in laser diodes. All modern laser diodes use a double heterostructure design for vertical confinement and employ either gain- or index-guiding for horizontal confinement. Gain-guided lasers can provide single longitudinal

mode operation, but thermally induced mode-hopping can cause significant changes in output power. Index guided lasers produce a multimode output with a broader spectrum than single mode lasers. However the spectral bandwidth is only 2-4 nm, and they don't mode-hop, so they provide better power stability than single mode lasers [78].

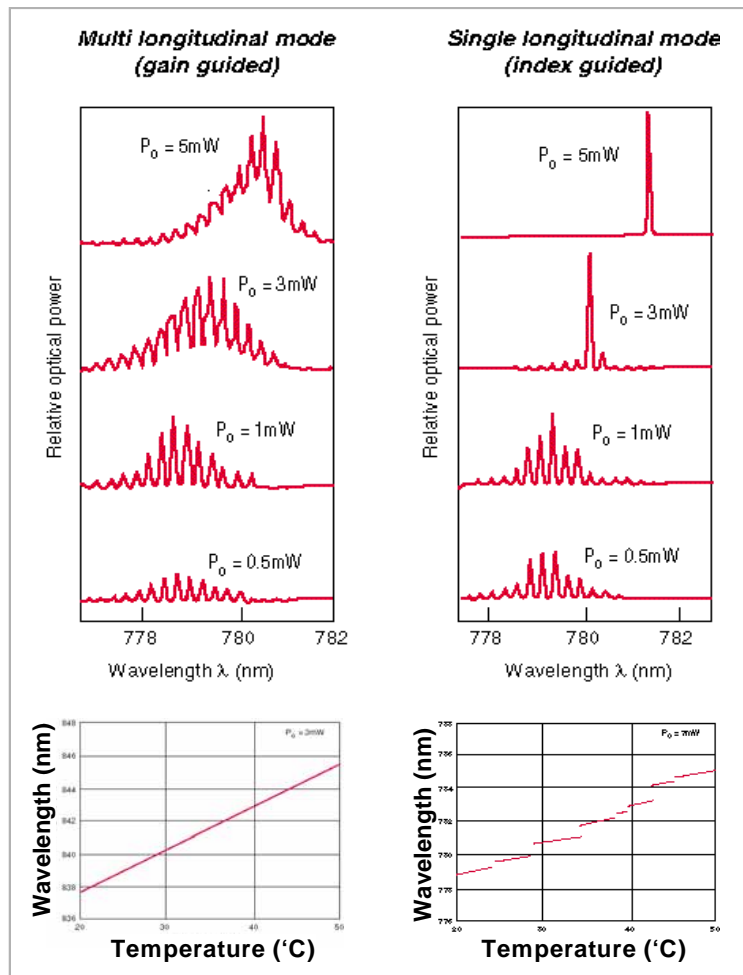


Figure 4.2. The spectral properties of multimode and singlemode laser diodes. Multimode laser diodes simultaneously emit across multiple longitudinal (Fabry-Perot cavity) modes over the full output power range. As a consequence, their output power remains stable and their spectrum shifts smoothly with temperature. Singlemode lasers initially lase across multiple modes, but revert to singlemode operation at high power levels. As a result, the output spectrum exhibits discrete hops in both center wavelength and output power as a function of junction temperature. Since these hops are difficult to predict in advance, single longitudinal mode operation of conventional index-guided lasers requires a stable drive current and active temperature stabilization – both of which are difficult to achieve with DOT source encoding.

Optode coupling issues with laser diodes

Since the emitting areas of both types of laser diodes are only a few μm across, the beam is spatially coherent and fiber coupling is relatively straightforward. Beam quality issues such as ellipticity (due to diffraction from the small emitting aperture), far-field uniformity and astigmatism are rarely an issue for DOT applications, since optodes usually contain large source fibers with core diameters exceeding $200\mu\text{m}$. The outer surface of the AR-coated window on standard 5.6mm and 9mm hermetic laser diode packages is about $300\mu\text{m}$ from the exit facet of the laser and the elliptical beam has diverged to less than $200\mu\text{m}$ in diameter at this point, so larger core fibers and fiber bundles

can simply be butt-coupled directly to the window of the laser diode package. An example of a simple fiber coupling fixture is shown in Figure 4.3. As long as no collimating optics are used, the risk of inducing lasing instability through cavity-coupling effects (i.e. the inadvertent creation of a resonant Fabry-Perot cavity between the exit facet and the fiber endface) remains slim. Smaller single-strand fibers may require a convex lens to reimage the beam onto the fiber endface, however the risk of cavity-coupling is quite high with this configuration. The fiber endface, or the laser itself, can be intentionally defocused or tilted slightly to prevent reimaged retroreflections off of the fiber endface from destabilizing the laser cavity.




Figure 4.3. Drawing of a machined laser diode mounting fixture designed to couple laser diodes in standard 9mm packages to large diameter optical fibers equipped with SMA 905 connectors. The fixture was designed to provide 100um of clearance between the window surface and the fiber endface.

Optode coupling issues with LEDs

The emitting area of LEDs is quite large and diffuse, so fiber-coupling is inefficient unless large area fiberoptic bundles are used. The fiber-coupling techniques popular in the telecommunication industry are geared towards launching a few mW at 1330nm or 1550nm into single low-loss fibers and are not commonly available for LEDs at wavelengths below 800nm. The best and simplest way to employ LEDs for DOT measurements is to place them directly on the optode assemblies. This minimizes optical coupling losses and can simplify the hardware design, however the presence of electronic components within the optode assembly may preclude simultaneous DOT/fMRI measurements with such instruments.

Driving laser diodes and LEDs

The electro-optical characteristics of a typical multimode CW laser diode are shown in Figure 4.4. CW laser diodes have forward voltage drops below 2V, sub-nanosecond risetimes, and a dynamic junction resistance of only a few Ohms, so modulating laser diodes and LEDs at up to VHF frequencies with low voltage standing wave ratios is relatively straightforward. With careful bias control, modulation depths approaching 100% are achievable. Since laser diodes are extremely susceptible to damage from excessive output power, all laser bias supplies should be equipped with overvoltage protection. This can range from circuits as simple as a large Zener diode, Transzorb[®], or ZNR-type surge arrester from V+ to ground preceded by a fast-acting fuse, to the more complex commercial SCR crowbar circuits with integral fault alarms. Although these circuits will not provide nanosecond response times, it is quite likely that, except for a direct mechanical short within the supply, the rate of voltage rise due to a component failure in a standard series-pass linear power supply

will never exceed the response time of either of these circuits, so they should provide adequate protection under normal circumstances.




Figure 4.4. The electro-optical properties of a typical multimode gain-guided laser diode. Lasing will only occur at drive currents greater than the threshold current, above which the coherent optical output is nearly linear with increasing drive current. The main disadvantage of laser diodes is their frailty. The surface of the output coupler facet is very fragile, thus the maximum output power ratings of CW laser diodes should never be exceeded or else permanent damage to the exit facet will result.

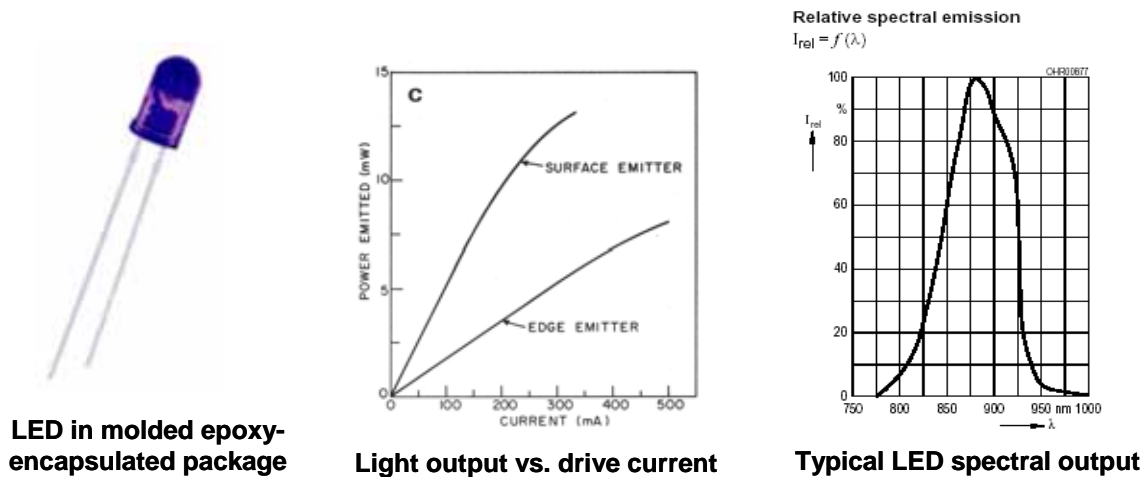


Figure 4.5. The electro-optical properties of a typical IR LED. Emission begins at very low drive currents and is reasonably linear over the full output range. LEDs are far more rugged and their output power is more stable than laser diodes. Since their degradation mechanisms are thermally mediated, they can be driven at extremely high peak currents so long as the duty cycle is derated accordingly. The two main disadvantages of LEDs for DOT applications are broad spectral linewidth and poor coupling efficiency into optical fibers and fiber bundles.

LEDs are usually slower and offer a much broader linear operating range than do laser diodes, however modulation rates greater than 100MHz have been achieved with some devices [77]. LEDs are

more electrically and thermally robust than laser diodes, so electrical biasing tolerances can be much wider. The electro-optical characteristics of a typical LED are shown in Figure 4.5.

Circuits for driving both LEDs and laser diodes can be quite similar, with the main distinction being whether the modulation is linear, as in RF applications, or consists of fixed-current digital drive pulses. A trade can be made between driver efficiency and electrical noise generation. A class-A drive circuit consumes more power than a simple on-off-keyed circuit, but it generates very little EMI, since the absolute change in current consumption during modulation is minimized. Examples of both circuits are shown in Figure 4.6. For analog operation, both devices should be biased at the 50% flux point to maximize dynamic range. For digital or DC pulsed operation, lasers should be biased just below threshold when in the “off” state to reduce thermal transients, thus minimizing the spectral shift.




Figure 4.6. Examples of an analog (left) and a class-A digital pulse drive circuit (right) suitable for open-loop operation of both laser diodes and LEDs. The value of R_{shunt} is chosen to reduce the operating junction voltage by about 500mV, which effectively detours all of the bias current around the diode. Both laser diodes and LEDs have a relatively constant $\sim 1.5V$ to $2V$ forward drop, and are easy to drive with analog or digital signals so long as the bias currents are properly set. Laser diodes require sufficient bias to overcome internal losses before any lasing can occur. Above this threshold, the optical output is very linear with increasing current. LEDs are rugged and can tolerate significant electrical stresses which would destroy laser diodes in microseconds, so their biasing is less critical.

4.2.2 Optical Detectors

The three most common optical detectors used in DOT instrumentation include silicon PIN photodiodes, silicon avalanche photodiodes (APDs), and photomultiplier tubes (PMTs). A spectral comparison between a silicon PIN detector and a photoemissive PMT detector is shown in Figure 4.7, and datasheets for each of the devices mentioned below are included in the Appendix.

PIN photodiodes are very inexpensive ($\$$), very linear, and have a very large dynamic range. Although their quantum efficiency is quite high, they have no internal gain, so they are not very sensitive. PIN photodiodes are best suited for small optode separations, where penetration depth and sensitivity are not a problem. Silicon photodiodes have a relatively high quantum efficiency over the 600nm to 900nm spectral band, which simplifies the radiometry. Since the junction capacitance per-unit active area is large, photodiodes are rarely used at the VHF frequencies employed by “RF”

instruments. Photodiodes are best in CW instruments used for small animal and neonate measurements, where optode separations are small and signal levels are relatively large. The Burr-Brown OPT101 detector/preamplifier module is a popular silicon detector. It provides a higher noise floor than a comparable photodiode equipped with a discrete preamplifier, but it is very easy to use and the electrical shielding requirement is minimal [79].

APDs and APD modules (APD + power supply and preamp) are more sensitive than PIN photodiodes and have a much higher gain-bandwidth product, but they have less dynamic range and are more expensive (\$500 +). Like photodiodes, silicon APDs have high quantum efficiency over the 600nm to 900nm band. APDs have sufficient gain to be practical in RF instruments, however they are more commonly used in CW instruments designed for use on adult humans, where larger optode separations and denser pericortical structures are the norm.

The Hamamatsu C5460 Series APD modules were used in a number of DOT instruments. Each module contains an on-board temperature-compensated high voltage bias supply and transimpedance preamplifier [80].

PMTs and PMT modules are comparable in sensitivity and cost to APDs below ~800nm and they can provide far more gain, however they are very sensitive to supply voltage and stray magnetic fields (an issue when MRI magnets are nearby), they have a strong spectral gain dependence, and their linearity is poor. Advantages of PMTs include a very high gain-bandwidth product and the ability to perform RF (i.e. nonlinear) mixing within the PMT, which is a valuable plus for RF DOT systems. The Hamamatsu H5783-20 is a compact PMT module with a photocathode spectral response out to ~850nm [81]. The Hamamatsu H6573 Modulated PMT Module is a bit larger and has poorer NIR spectral response, but it has an external modulation input which can be used for RF mixing and downconversion. This makes it popular for heterodyne and homodyne detection.

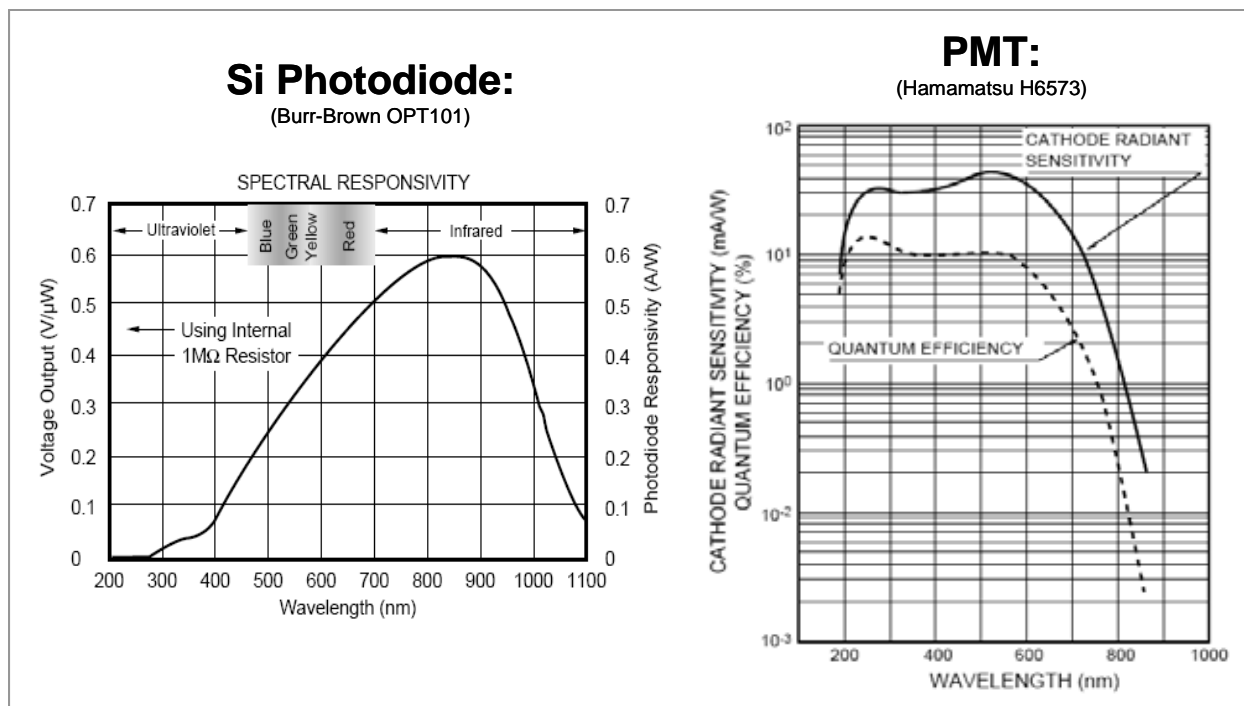


Figure 4.7. A comparison of the spectral characteristics of silicon detectors and photoemissive detectors. The quantum efficiency of silicon detectors (not shown) is relatively high and flat.

This creates the uniformly sloping responsivity vs. wavelength shown above, since the number of photons per Watt is directly proportional to wavelength. APDs are similar to photodiodes, although their spectral response is somewhat narrower since the photon absorbing region in APDs is thinner than in PIN photodiodes. The quantum efficiency of the photoemissive surface of the PMT shown begins at around 8% at 600nm and decreases to only 0.025% at 850nm. The large gain-bandwidth product provided by the electron multiplier within the PMT significantly offsets this disadvantage. Newer GaAs photocathode materials are being developed which offer much higher quantum efficiency below 900nm, however their lifetime is limited.

4.3 Encoding Techniques

4.3.1 Source Encoding

Most cortical DOT measurements employ an optical probe assembly consisting of an array of source and detector optodes distributed over the surface of the scalp, with the optode locations chosen to cover the cortical regions of interest. In order to measure all of the optode pairs simultaneously, some form of source encoding is required. Encoding provides a means by which the light from each source can be uniquely identified by the detection circuitry, so that every optical signal which arrives at the detectors can be traced back to its respective source optode.

Domains in which source identity can be encoded include time, frequency, and phase. Certain domains may not be suitable because they contain the experimental information. For example, wavelength-division multiplexing, which operates in the spectral domain, is a form of encoding used to increase the data capacity of fiberoptic networks. Most tissue imaging studies, however, rely on the spectral features of chromophores to provide vital information, so wavelength-division multiplexing is impractical for DOT.

With DOT, the number of *source* optodes can directly affect both the demodulation complexity and the amount of source power that each detector receives. However the number of *detector* optodes has no effect on demodulator design or source power. This situation is analogous to broadcast radio, in which the receiver design depends only on the type of transmission (AM, FM, digital) and the frequency allocation of each station, but is not affected by the number nor the locations of the radio receivers actually in use.

Given a fixed constraint on source power and assuming that the detectors receive no stray light, it can be argued that both time and frequency encoding can provide similar shot noise-limited sensitivity. This sensitivity decreases as a function of the square-root of both the frame rate and the number of source optodes. Unlike sensitivity however, dynamic range does depend on the encoding scheme. With shot noise-limited sensitivity, time-encoding offers a greater dynamic range than frequency-encoding, since there is no self-generated background flux with time-encoding.

4.3.2 Dynamic range and signal span

The dynamic range of a DOT measurement is greatly affected by optode geometry. As the optode separation increases, the detected flux decreases more or less exponentially with distance. For an optode array with a spacing ratio of 5:1, the signal span (the span in detected flux caused by changes in optode spacing) will vary with the dimensions of the array, from as low as 8dB for a small optode array used on mice to as high as 85dB or more for a larger array used on humans, as shown in Table 4.2. This means that the system dynamic range will need to be very large to perform multi-optode DOT for brain function on adult subjects over coverage areas which may exceed 100cm². Similar measurements performed on smaller creatures (rats, mice, etc.) with the same optode geometry, scaled down to cover a far smaller (~1cm²) area, would generate a smaller signal span, and thus would require a lower dynamic range.

Table 4.2. The effect of optode spacing on the optical attenuation range. The signal span is a function of both the maximum optode spacing and the optode spacing ratio.

Creature	Optode spacing ratio = 5:1 (spacing)	Signal Span (dB)	Optode spacing ratio = 10:1 (spacing)	Signal Span (dB)
rat	1cm - 2mm	8.5	1cm - 1mm	9.6
neonate	5cm - 1cm	42.6	5cm - 5mm	47.9
adult	10cm - 2cm	85.1	10cm - 1cm	95.7

Note that although the hemodynamically-mediated attenuation changes are inherently unpredictable, the signal span attenuation is quite predictable and repeatable within reason (give or take a few dB). If a single-gain instrument were designed to perform all of these measurements, it would require an instantaneous dynamic range of 120dB or more, which is extremely difficult to achieve in practice. A better approach would be to choose an instantaneous dynamic range of 60dB, but to include an additional switched-gain capability to shift this 60dB window to where it is needed the most. This is far more practical from an engineering standpoint, since it preserves the instantaneous dynamic range for detecting hemodynamic fluctuations, and will be discussed further below.

4.3.3 Theoretical performance limits

As a consequence of the relatively large (~1.5eV) photon energies in the near-IR spectral band, the ultimate limit in sensitivity results from the shot noise generated by the incoming photons themselves. The upper bound for optical power density delivered to human test subjects is dictated by federal regulations. Thus, it is important that the instrument and probe assembly be made as photon-efficient as possible.

Since the arrival rate of individual photons is completely random, there will always be some uncertainty in the photocurrent produced by even an ideal (100% efficient, noise-free) quantum detector. This uncertainty is called “shot” noise, and its magnitude is proportional to the square root of the photocurrent. So in an ideal shot noise-limited measurement, the SNR grows as the square root of the photon flux (measured in photons, not in watts, since these are quantum detectors). Since shot noise is a fundamental property of signal detection in a quantized universe, obtaining shot noise-limited performance is the best that any equipment can achieve. Once this limit has been reached, no amount of electronic enhancement at the detector end can further reduce the power spectral density below the shot noise floor. What can be done, however, is to increase the number of photons which reach the detector during a measurement. There are two ways to accomplish this: boost source power and increase photon collection efficiency. This will increase the SNR, but only by the square root of the flux increase. This leads to an important fundamental concept:

- The maximum achievable SNR of any optical detection system is determined by the number of source photons detected (i.e. the photocharge collected) during each measurement.

Electrical noise, stray light, tissue absorption, and optical coupling losses only serve to reduce this value even further.

4.3.4 Technical issues

There is much interest in using DOT to obtain real-time three-dimensional imagery of cerebral hemodynamics within the human brain. In order to achieve this, a number of important physiological, optical, and electrical issues must be addressed. Some of these include:

The wide range of probe geometries encountered in research environments

Test subjects can range from mice to men, with probe dimensions and optode spacings sized accordingly. For a DOT measurement system, this represents a very large dynamic range. The detected power between optode pairs can span from around 9dB in rats to greater than 90dB in humans. This presents one of the toughest challenges for DOT instrument design, since the entire flux range must be measured within a single frame time or data will be lost.

The sensitivity limits imposed by the NIR spectral window and the quantum nature of light

As a consequence of the relatively large photon energies in the near-IR spectral band, the ultimate limit to sensitivity results from the shot noise generated by the incoming photons themselves. The upper bound for optical power density delivered to human test subjects is dictated by federal safety regulations. Thus, it is important that the instrument and probe assembly be made as photon-efficient as possible, so as not to further degrade sensitivity.

Temporal skew

This term represents the time lag between the first and last optical measurements collected within a single “frame” of data. With DOT this can introduce significant blood volume and oxygenation errors, since the optical measurements vary both spectrally and spatially with time. Temporal skew can be minimized by operating at high frame rates or by employing specific encoding methods, some of which are discussed below.

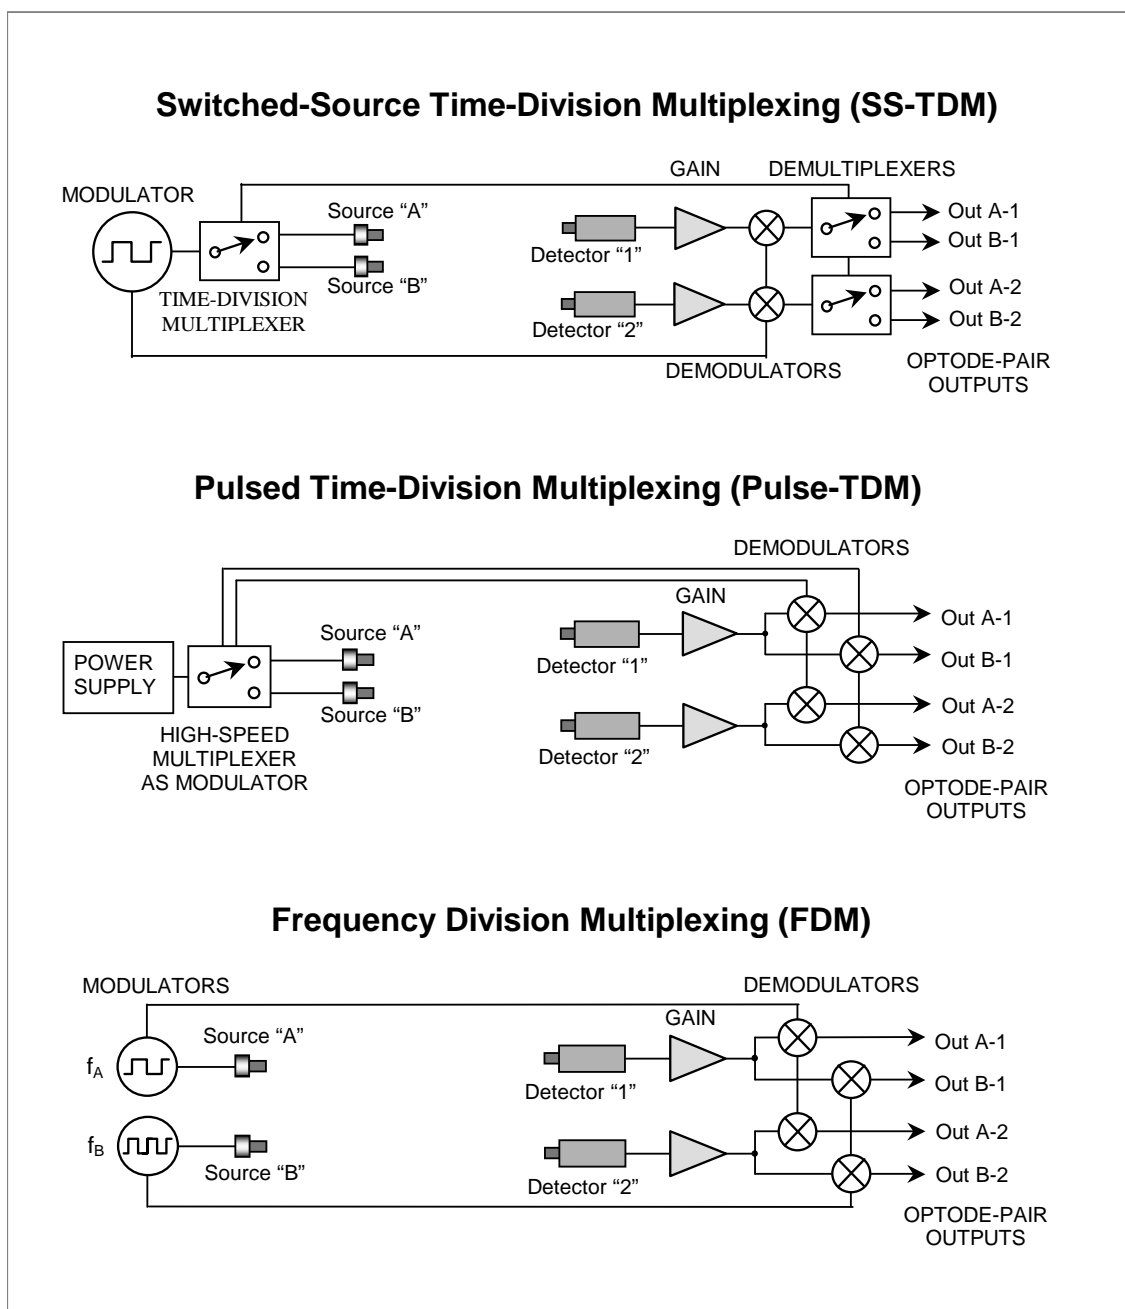


Figure 4.8. Block diagrams of three DOT systems based upon the three encoding schemes discussed in the text. Demodulation can either be performed using electronic circuitry, or the preamplified detector signals can be digitized and later demodulated using digital signal processing. For simplicity, each system here is depicted with only two sources and two detectors.

What follows is a brief discussion of the relative merits of three different encoding schemes for a hypothetical probe geometry consisting of "M" source optodes and "N" detector optodes. Block diagrams depicting the implementation of the three schemes are shown in Figure 4.8. With all three encoding schemes, demodulation can either be performed with analog circuitry, or the raw detector signals can be digitized, stored to memory, and later demodulated in the digital domain.

Switched-source time-division multiplexing (SS-TDM)

All M sources are modulated at the same carrier frequency and are cycled through consecutively, with each operating for $1/M$ of the frame time. Each detector is followed by its own synchronous demodulator, which recovers the signal from all M sources in turn.

Features of SS-TDM:

Maximum dynamic range: Since only one source operates at any one time, the self-generated background level is zero, providing the greatest possible dynamic range.

Very low crosstalk: Sequential source operation minimizes the potential for interchannel crosstalk.

High temporal skew: The time delay between the first and last source measured during each frame creates approximately 1 frame time of temporal skew in the recorded data.

Simple Construction: This is the easiest of the three systems to design and build. The detector/preamplifier must operate over a very large dynamic range, although temporal performance is not as critical as with the other encoding schemes.

Pulse-modulated time-division multiplexing (Pulse-TDM)

Each of M sources is pulsed individually and consecutively at a rapid (\sim kHz) rate, with a duty cycle inversely related to M. Each of N detector sites demodulates all M sources using M time-gated synchronous demodulators.

Features of Pulse-TDM:

Maximum dynamic range: Since only one source operates at any one time, the self-generated background level is zero, providing the greatest possible dynamic range.

Low crosstalk: Some crosstalk can be introduced through inadequate settling time, although it can be easily reduced in trade for temporal response.

Minimal temporal skew: Since the cycle time through all M sources is around 1ms, temporal skew is negligible from a DOT standpoint.

Complex construction: Very circuit- or computation-intensive (a total of $M \times N$ demodulators are required) but straightforward, from a design standpoint. For Pulse-TDM, the detector/preamplifier must perform well in the time-domain, providing a clean step response with minimal ringing and rapid settling over a very large dynamic range.

Frequency-division multiplexing (FDM).

Each of the M sources is simultaneously modulated at an anharmonically-related audio frequency (to minimize crosstalk from intermodulation distortion). Each of N detectors demodulates the N sources coherently through synchronous detection, or incoherently through envelope detection.

Features of FDM:

Lower dynamic range: Since all sources operate simultaneously, the self-generated background level is large, creating a higher noise floor and thus lower dynamic range than the TDM encoding schemes.

Moderate crosstalk potential: Since the modulated envelopes of the optical signals are essentially incoherent, they can periodically combine to produce large peak photocurrents in the detector, increasing the potential for intermodulation distortion.

No temporal skew: Since all sources operate simultaneously, the temporal skew is essentially zero.

Complex construction: Very circuit-intensive with on-board demodulation (a total of $M \times N$ demodulators are required) or computation-intensive if digital demodulation is employed. For FDM,

the detector/preamplifier must also perform well in the frequency-domain, with a large, flat passband and low intermodulation distortion at high signal levels.

Source encoding design trades

PARAMETER:	SS-TDM	Pulse-TDM	FDM
Detector Bandwidth:	LOW	HIGH	HIGH
Crosstalk:	VERY LOW	LOW	HIGH
Temporal skew:	HIGH	NONE	NONE
Dynamic Range:	HIGHEST	HIGHEST	LOW
Individualized Gain Control:	EASY	EASY	HARD
Circuit complexity:	SIMPLE	COMPLEX	COMPLEX
Avg. Delivered Power:	LOWEST	LOWEST	HIGH
Source count limit:	NO LIMIT	MANY	MANY
(Traded for):	(time)	(time)	(bandwidth)

4.3.5 Using Individualized Gain Control to improve performance

Dynamic range limitations with DOT instrumentation

The dynamic range of a DOT measurement is strongly influenced by the optode probe geometry. As the source-detector optode separation increases, the detected flux decreases more or less exponentially with distance. For a rodent-scale probe geometry of ~1 x 1cm and a 5:1 optode spacing ratio, the detected flux would vary by about 10dB between the smallest and the largest optode spacing. For a 2cm geometry, this attenuation ratio grows to 21dB. For an adult probe geometry of 10cm, the attenuation ratio can exceed 100dB. Contrast this static attenuation with the actual information of interest to brain researchers for instance - the signal modulation created by the cerebral hemodynamics - which typically spans less than 1dB. Thus, the dynamic range of most DOT measurements is completely dominated by the static optical attenuation ratio resulting from both the probe geometry and the bulk absorption and scattering properties of tissue.

The concept of individualized gain control

The main consequence of using a DOT system with insufficient dynamic range is the loss of hemodynamic information at larger optode spacings, which limits tomographic resolution. The desire to improve tomographic performance motivated the development of individualized gain control (IGC) as a means of improving the capability of DOT instrumentation.

The static optical attenuation between each of the optode pairs is relatively stable over time and is of little experimental value. If it can be quantified and then removed from within the signal span of the measurement, the dynamic range can be preserved for the more important hemodynamic signals of interest. One means of accomplishing this is to dynamically adjust the gain of the preamplifiers to compensate for the unique static attenuation between each optode pair. For maximum benefit, the attenuation between every optode pair should be compensated individually. Although such “individualized” gain control requires additional circuitry, it should enable a DOT system to more closely approach the theoretical limits of performance, simultaneously capturing both local and distant optode information to provide higher quality tomographic reconstructions. The circuitry required to perform these gain adjustments will depend on the encoding scheme. Regardless of whether demodulation is performed in hardware or in software, IGC must be implemented prior to demodulation to be effective.

With SS-TDM, individualized control can be achieved with a single adjustable gain stage following each of the N detectors. Since SS-TDM uses only one demodulator per detector, these gains must be dynamically-switched in synchrony with the sources. However, since the source cycle rate is relatively slow, dynamic gain-switching should be easy to implement.

Pulse-TDM employs one demodulator per optode pair, so N gain stages could be dynamically-switched following each detector and ahead of the demodulators as with SS-TDM, or the gains within each of the $M \cdot N$ demodulator channels could be set once and then remain fixed during the measurement, simplifying the design but requiring M times more gain circuitry.

FDM also employs one demodulator per optode pair, but in the frequency domain, dynamic gain-switching is not required. So $M \cdot N$ statically-programmed gain stages can be used, however each must be preceded by a narrowband filter for frequency selection.

The implementation of IGC with all three encoding techniques is shown in Figure 4.9.

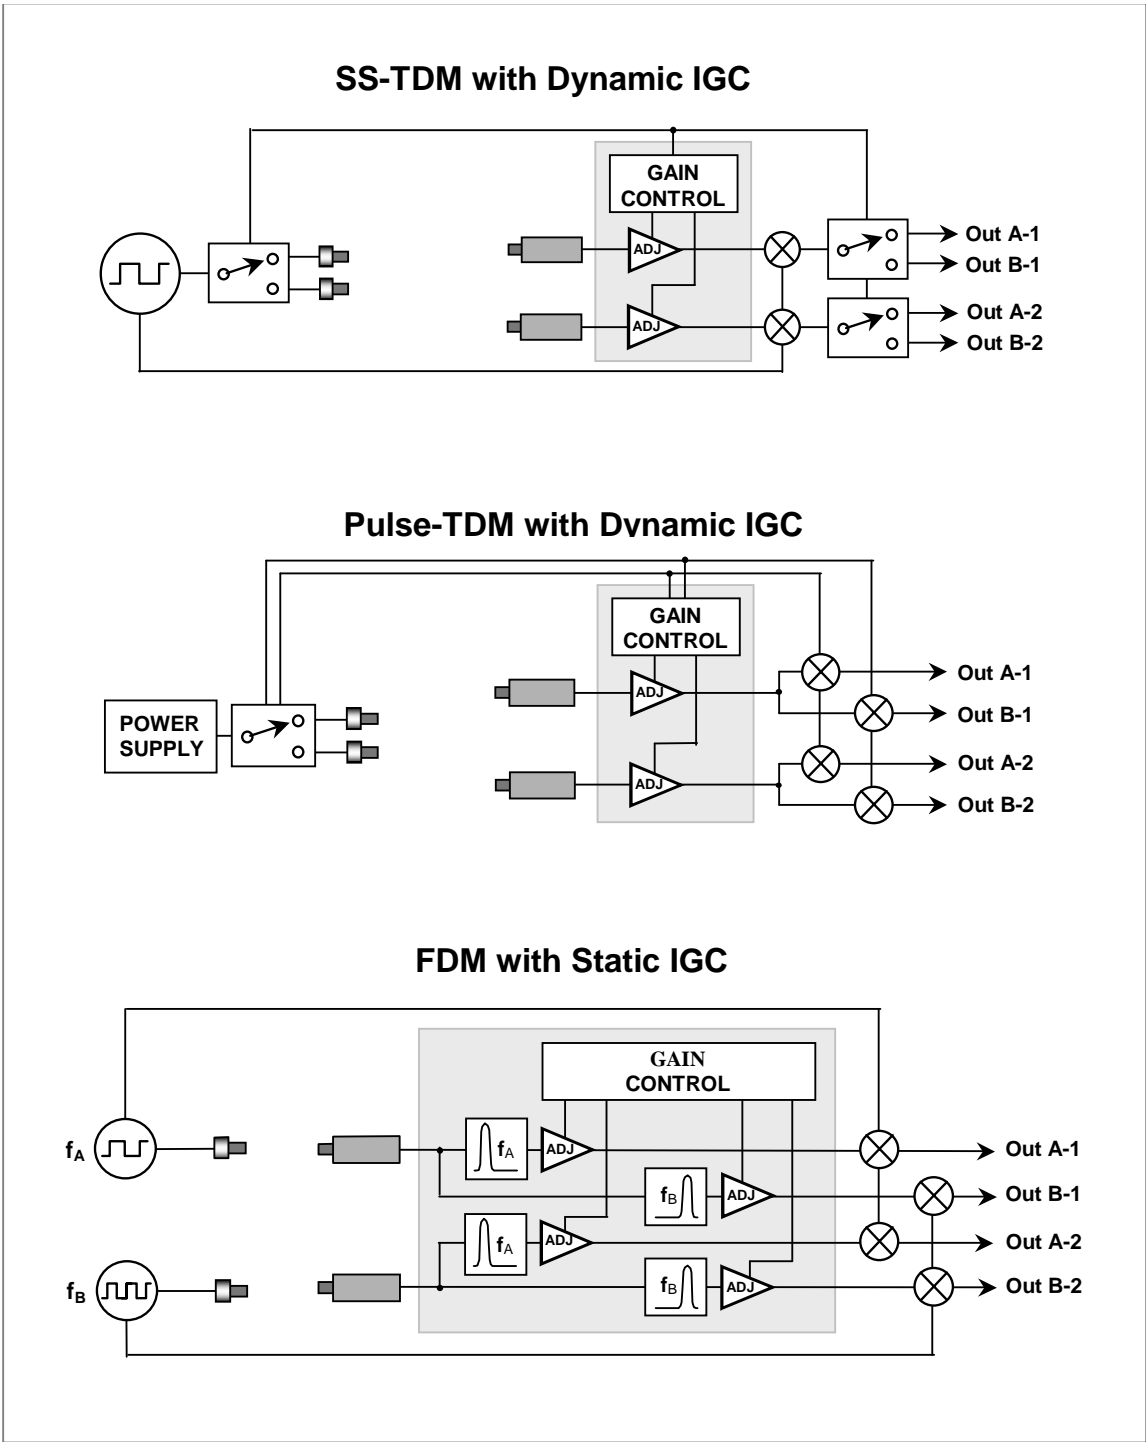


Figure 4.9. Examples of the same three DOT systems equipped with individualized gain control. Both time-encoding techniques require dynamic gain-switching. With frequency-encoding, the gain values can remain stable during the measurement, although pre-gain filtering is required.

Some issues with IGC

Although individualized gain control may improve system performance downstream of the detector, it cannot compensate for the fact that the detector and its preamplifier must operate over the entire dynamic range spanned by the optical signals. Therefore significant effort should be spent in optimizing the detector circuitry to perform well over this dynamic range. TDM systems require detectors with good temporal response and clean settling behavior. FDM systems require good frequency and phase performance, along with high linearity to minimize intermodulation distortion.

The increase in circuit complexity with IGC can be significant, although most of the new circuitry involves the replication of relatively simple amplifiers constructed from inexpensive components. To date, the dominant hardware expense with many CW DOT instruments has been the cost of both sources and detectors. Both APD and PMT detector modules currently cost between \$500 and \$1,000 apiece,⁹ while the additional component cost for an entire 16-source IGC amplifier/demodulator circuit would run about \$100 per-detector. So, aside from requisite increases in volume, weight, and power consumption, the addition of IGC would incur only a modest increase in overall cost.

In order to control and track all of the gain levels effectively, computer control is recommended. Dynamic gain-switching can be performed by programming an on-board memory or serial shift register with each of the $M \times N$ gain values. Once all the gains have been loaded, the computer interface can then be disabled to both reduce EMI and to free up the computer for data collection. During operation, the appropriate gain values are addressed and synchronously incremented by timing pulses generated within the instrument.

Proof-of-principle testing

A “proof-of-principle” prototype system was constructed and tested to evaluate the feasibility of individualized gain control with pulse-TDM encoding. Specifically, this work was performed to answer the following questions:

- Is IGC practical to implement in a CW DOT system?
- If so, then how will the addition of IGC affect the overall performance?
- Even if it improves electrical performance, will IGC actually improve the quality of physiological data?

The system, shown in Figure 4.10, contains four laser diode sources and one silicon PIN detector, which is sufficient to simulate four optode pairs for a variety of optode spacings and detector flux levels. The master timing generator consists of a binary ripple counter driven by a quartz crystal oscillator. It governs both the modulation and demodulation of all four sources, and also synchronizes the IGC gain state selection to the appropriate optode pairs. An 8-to-1 analog multiplexer served as a 4-channel synchronous demodulator. Synchronous detection, which is easy to implement with pulse-TDM encoding, improves noise and stray light immunity, and also reduces the electronic noise floor of the system.

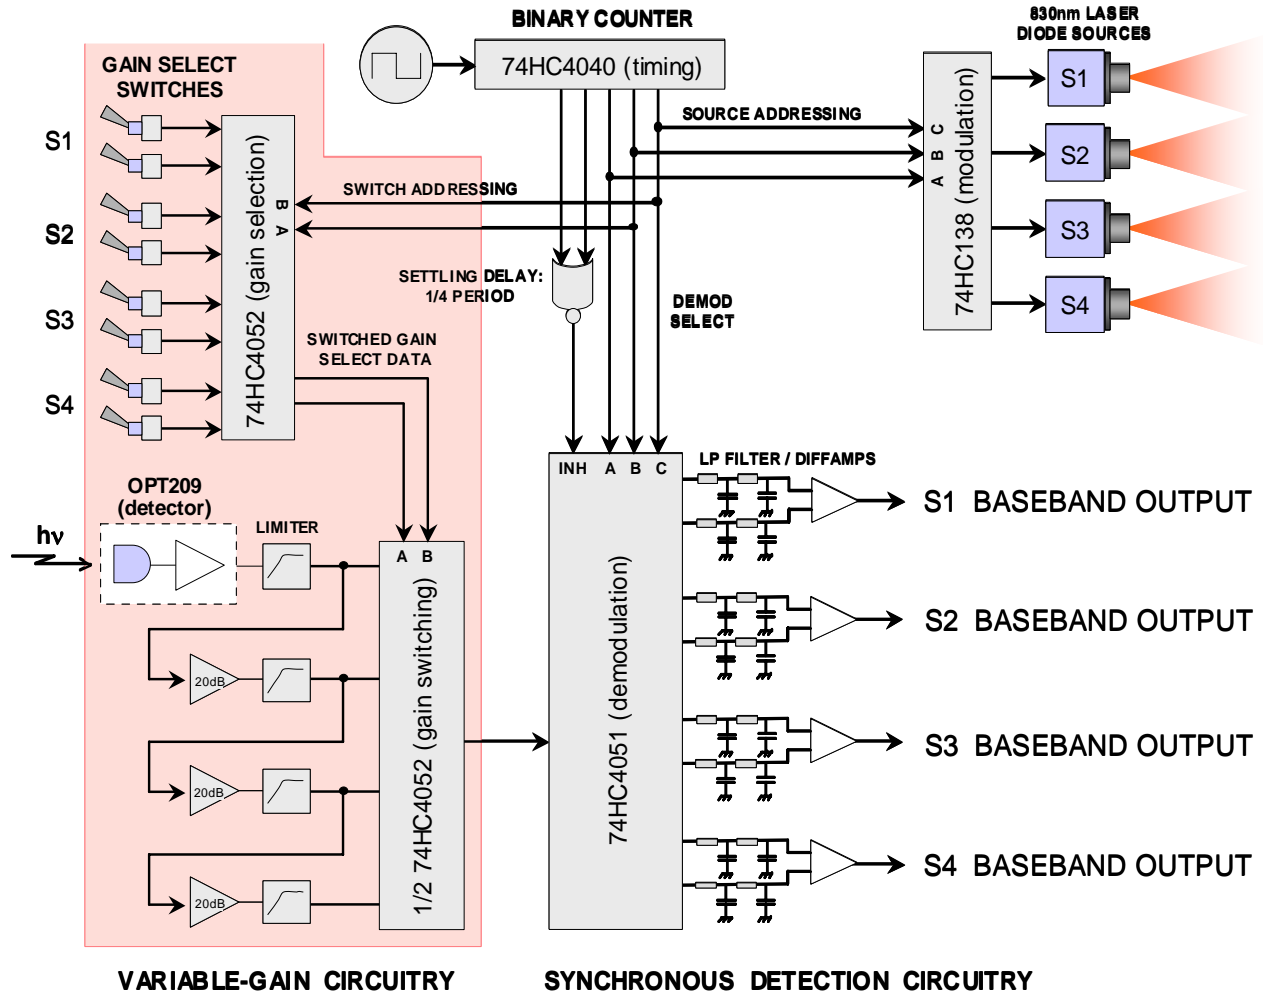


Figure 4.10. The time-encoded IGC Evaluation System, which includes four laser diode sources and one Burr-Brown OPT209 amplified silicon detector. It employs pulse-TDM encoding, and provides four dynamically-switched gain states of 1, 11, 121, and 1331, which amplify the AC-coupled detector output. The gain states for each of the four source-detector pairs are selected manually using toggle switches. A timing diagram for one complete 4-source cycle is shown in Figure 4.11.

One of the virtues of time-encoding should be the ability to minimize the signal corrupting effects of capacitive crosstalk, since electrical transients can settle prior to sampling. Settling time is even more important with IGC, due to the additional gain-switching involved. To test this hypothesis, the entire system was constructed on a solderless breadboard to emulate the worst-case environment from an electrical crosstalk perspective.

The timing of this system is based upon a repeating 8-step cycle, as shown in Figure 4.11.




Figure 4.11. This timing diagram for the IGC Evaluation System shows the sequencing of the four laser sources (*top*) and the resulting detector and IGC gain amplifier outputs (*middle*). IGC gain values are chosen to maintain a similarly large output swing for all four sources, which is why the weaker source signals appear noisier. Gain-switching occurs simultaneously with the transition to the next source in the sequence, creating the glitches visible in the “IGC AMP OUTPUT” waveform. Two repetitively acquired voltage samples, representing the “signal” and “reference” levels for each of the four laser pulses (not shown), are lowpass filtered and then differenced to perform the demodulation for each optode pair.

Results

Crosstalk

The most likely source of crosstalk for a time-encoded system is during transient recovery; when a weak optical signal is measured immediately following an extremely strong signal. This situation was simulated by strongly coupling one source (Source #3) to the detector and mechanically modulating its intensity so that the optical flux drove the demodulator output from 0V to 4V at an IGC gain of unity. The outputs from the three other sources were carefully concealed so that no optical signals were detectable, even at the highest IGC amplifier gain of 1331. The variation in voltage at these other three outputs was recorded for each of the four gain states. The crosstalk component was then scaled to correct for the value of the IGC gains used during the measurement, and the results are shown in Table 4.3.

Table 4.3. The interchannel crosstalk vs. channel number and gain value. The large spread in crosstalk values is probably due to capacitive coupling between specific circuit elements, thus the minimum value is more representative of achievable performance. The full-scale signal was received on channel #4 (not shown), which operated at a gain of 1 to simulate the worst-case condition of maximum signal flux.

Interchannel Crosstalk (dB)				
Gain	Channel 1	Channel 2	Channel 3	Minimum
1	-103.2	-109.8	-127.1	-127.1
11	-118.2	-98.3	-93.0	-118.2
121	-127.5	-108.9	-106.7	-127.5
1331	-126.1	-106.2	-112.4	-126.1

Dynamic Range

The dynamic range was defined as the ratio of the RMS noise floor to the maximum output swing of 5V, as recorded by the data acquisition card. The measured ratio was then scaled to account for the IGC gain span. The effective bandwidth of these measurements was from 0.02Hz to 0.5Hz, which was chosen to represent the noise level in a hemodynamic measurement following either strong temporal filtering, or block addition with a bandwidth extending out to 5Hz and following about one hundred trials. These results are shown in Table 4.4.

Table 4.4. The dynamic range vs. gain span. Dynamic range is defined as the ratio of the RMS noise floor to the maximum output swing of 5V, which is then scaled for the gain span. Voltage data were collected at a rate of 50 samples/second, and 100 samples were averaged to provide an effective noise bandwidth of ~0.5Hz for each measurement. An IGC gain span of 1331:1 increased the dynamic range by more than 40dB. The measurement noise, which dominated at the higher IGC gains, could be reduced by increasing the number of samples averaged, or by extending the duration of the measurement (through block-averaging), which reduces noise while preserving the temporal response. Both would increase the dynamic range beyond the values shown.

Dynamic Range		
Gain (V/V)	DNR (dB)	Enhancement (Δ DNR/ Δ Gain)
1 (fixed)	94.7	
1 to 11	115.3	97%
1 to 121	128.4	41%
1 to 1331	137.0	24%

Linearity

Linearity was measured relative to an EG&G UV-100 silicon detector coupled to a variable-gain transimpedance amplifier. The UV-100 detector was operated in the photoconductive mode in order to minimize noise and improve sensitivity, since synchronous detection was not used in the reference channel to maximize linearity. Both detectors were fiber-coupled, and the source light was directed through a diffuser plate to minimize speckle artifacts. The gain of the measured channel was varied as needed, and the gains for the three other channels were kept at unity in order to accentuate the transient effects caused by gain-switching. At the lowest gain setting, some nonlinearity was observed, but speckle noise dominated for the three higher gain measurements. These data are shown in Table 4.5.

The speckle noise was produced by interference between specularly reflected light with optical pathlength differences less than the coherence length of the laser diode sources. Normally tissue scattering destroys the spatial coherence of the laser light by introducing random pathlength differences which are much greater than the coherence length of the laser diode. However for these measurements, the overall scattering was insufficient, so enough spatial coherence was preserved to render the fibers slightly “microphonic” (i.e. sensitive to vibration or mechanical deflection).

Table 4.5. The optical-to-electrical linearity vs. gain state over the full 0-5V output range. Nonlinearity was at or below 0.1% for all four gain states.

Nonlinearity	
Gain	0 - 5V swing
1	~0.1%
11	<0.1%
121	<0.1%
1331	<0.1%

Prefrontal cortical response to Valsalva Maneuver

To evaluate the utility of IGC with physiological measurements, this system was used to monitor the hemodynamic response of the prefrontal cortex to Valsalva maneuvers. The detector and three 830nm sources were optically coupled to the forehead using flexible silica fiberoptic bundles. The source-detector spacings were 16mm, 32mm, and 48mm, and the radius of curvature of the forehead was approximately 80mm. The optode assembly used for these measurements is shown in Figure 4.12 and the results are shown in Figure 4.13.

The subject breathed normally during the first 30 seconds of the measurement to establish a baseline. At T=30 seconds, the subject attempted to forcibly exhale through sealed lips and nose for a period of about ten seconds, raising both intrathoracic and intra-abdominal pressure. At T=40 seconds, the subject resumed normal breathing. This was repeated five times in order to collect five datasets – once for each static gain state, plus one with IGC enabled.



Figure 4.12. A photo of the optode assembly used to perform the Valsalva measurements. Separate 3mm diameter silica fiberoptic bundles were used to couple light from three 830nm laser diode sources to the forehead, and back to the silicon detector (thick bundle on far right). The fourth source was not used for these measurements. The assembly was placed on the forehead as shown above and held in place with an elastic band. All four optodes were spaced by approximately 16mm, and the radius of curvature was about 80mm.

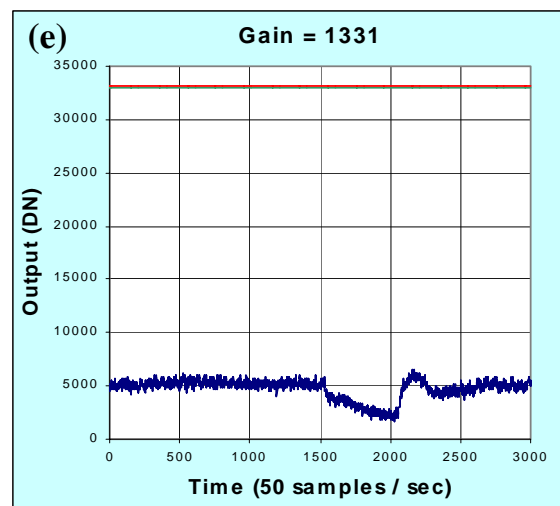

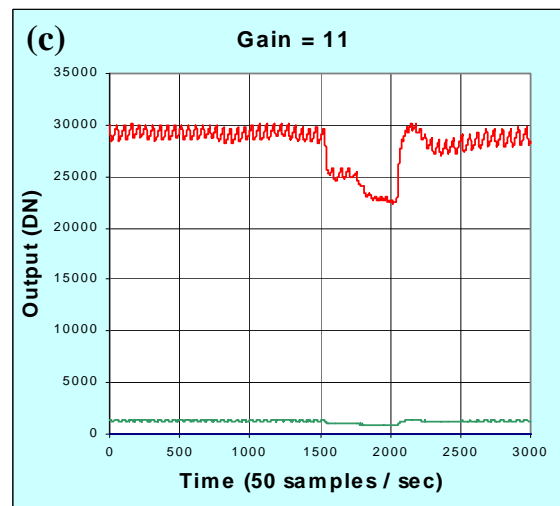
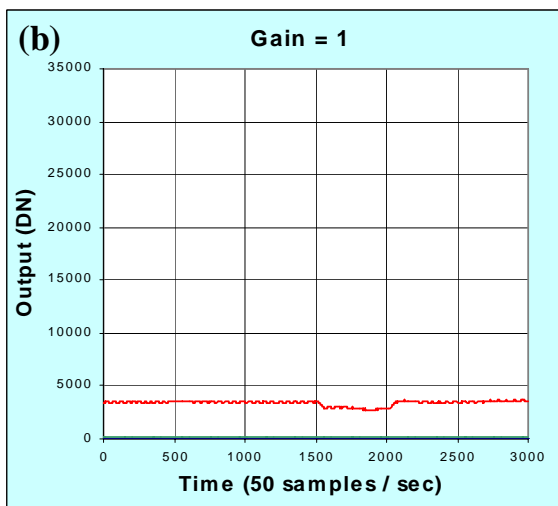
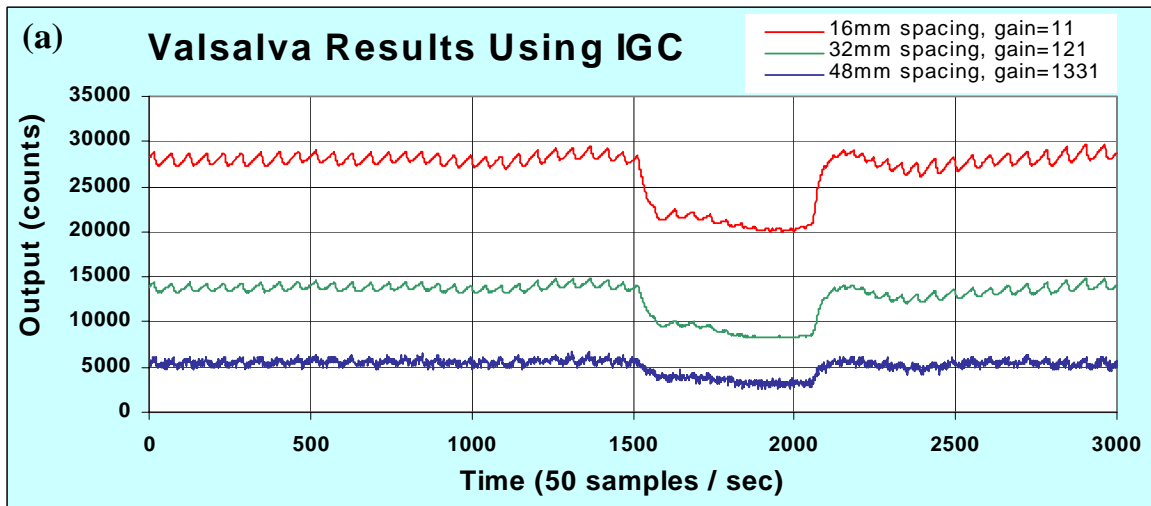


Figure 4.13. Optical measurement of the hemodynamic response to Valsalva maneuvers. The resolution limit of these plots represents the finite dynamic range of the A/D converter: 32,768 digital numbers (DN). With IGC, the detector gain was dynamically adjusted, providing high-quality simultaneous data for all three source-detector separations as shown in (a). The pulsatile intensity modulation (a result of the ventricular contractions of the heart) is clearly

visible in all three waveforms. (b) thru (e) show the waveforms resulting from fixed-gain operation. At low gain, the closest source provides an adequate signal, but the two farther sources are corrupted by quantization noise. At higher gains, the farther sources are detectable, but the output levels of the nearer sources exceed the dynamic range of the A/D converter, and all of their data is lost.

Discussion

Three important performance metrics for the detector channels in DOT instrumentation are crosstalk, dynamic range, and linearity. Signal crosstalk from nearby sources can corrupt measurements at large optode spacings, where the optical flux is weakest. Dynamic range is critical to achieving good depth resolution, since it allows the simultaneous use of a wide range of optode spacings. Linearity is important for maintaining an accurate measure of oxyhemoglobin and deoxyhemoglobin concentration changes over the full range in detected flux, especially for small optode spacings, where signals may be strong enough to exceed the linear range of the detectors.

Crosstalk

Even low levels of interchannel crosstalk can introduce serious errors in DOT measurements, because the signals reaching a detector from its nearest-neighbor sources can be thousands of times stronger than those arriving from more distant sources. Although techniques such as block-averaging can improve the SNR of weak signals, there is no simple means of removing or negating the effects of crosstalk. Thus the most prudent approach is to minimize crosstalk as much as possible through good instrument design and encoding practices. In this case, the concern was whether the use of dynamic gain-switching would reduce channel isolation, thus leading to greater interchannel crosstalk than with fixed-gain operation. These measurements revealed that dynamic gain-switching did not substantially increase the interchannel crosstalk above those same gain levels used with fixed-gain operation. This result highlights one of the advantages of encoding in the time-domain.

Since these crosstalk measurements were performed over a period of many seconds, they represent DC, or static crosstalk values. Although dynamic crosstalk is likely to exist, it was not considered an issue with DOT for two reasons: 1) The modulation depth created by cerebral blood flow changes are typically below 5% of the steady-state optical signal. Thus, the dynamic crosstalk would have to exceed the static crosstalk by a factor of 20 in order to become significant, and this is extremely unlikely. 2) The major temporal components of most hemodynamic measurements are typically below 20Hz, which should result in a minimal increase in crosstalk over static levels.

Dynamic Range

The 16-bit data acquisition card used to evaluate the performance of this system provided about 15 bits of resolution in unipolar mode, with a collective noise and uncertainty of around 2 counts RMS. This limited the dynamic range at low IGC gain settings, since the output-referred noise floor of the instrument was well below the quantization noise floor of the data acquisition card. This explains the 20.6dB increase in dynamic range achieved by raising the IGC gain span from 1 to 11, which is nearly a 1-to-1 enhancement. The enhancement drops with further increases in IGC gain, as detector and preamplifier noise components begin to dominate.

Although boosting the gain by a factor of 1300 to achieve a meager 42dB increase in dynamic range may at first seem futile from the standpoint of diminishing returns, it serves an important function. A large IGC gain will raise the value of the instrument's output-referred noise well above the input-referred noise floor of the data acquisition card, so block averaging (a correlation technique commonly used to improve the SNR in biophysics experiments) can reduce both the shot and electronic noise components well below their initially acquired values. As long as the instrument noise is sufficiently greater than the data acquisition noise floor, the effective dynamic range can be

increased well above the 42dB value quoted here, albeit at the added cost of a significantly extended measurement time.

Linearity

Output nonlinearity at high flux levels can be caused by “soft” compression within the photodetector. Both photomultiplier tubes and solid-state detectors exhibit a gradual, progressive reduction in quantum efficiency at high flux levels. In PMTs, compression can be caused by current spreading resistance in the photocathode, space-charge effects close to the anode, and voltage drops along the dynode bias ladder caused by high current draw. This may be compounded by a reduction in quantum efficiency over time with cumulative exposure (“aging”) in certain PMTs.

Solid-state detectors exhibit nonlinearity at lower flux levels in the zero-bias, or “photovoltaic” mode of operation. This is due to insufficient removal of the photogenerated carriers from the junction region. External biasing both increases the extent of the depletion region within the detector and improves carrier removal from the depletion region. Silicon detectors operating in the reverse-biased or “photoconductive” mode can be linear over as many as nine orders of magnitude in flux, and often exhibit a wider spectral response as well. Unfortunately, reverse-biasing solid-state detectors increases their dark current, which increases the noise. The thermally generated component of the dark current, which dominates for silicon detectors, can be reduced through cooling, but the added cost and bulk of detector cooling circuitry may be prohibitive for many applications.

Linearity is somewhat less critical for DOT than crosstalk or dynamic range, since unlike near-infrared spectroscopy, optical tomography is primarily based upon differential absorption measurements. Amplitude nonlinearity in a DOT system would mostly affect the background optical measurements, and should only minimally affect the hemodynamic signals because their modulation depth is so small. Since linearity is both monotonic and often stable over time, the linearity of each detector in time-encoded systems can be quantified and corrected for in software using lookup tables. Frequency-domain artifacts introduced by nonlinearity, such as intermodulation distortion, are more difficult to eliminate, but these would not occur in time-encoded systems.

This system uses a Burr-Brown OPT209, which is a monolithic device containing a zero-biased silicon detector coupled to a transimpedance amplifier with an internal $1\text{M}\Omega$ low-noise metal film feedback resistor. Although the noise-equivalent power of this device is quite poor when compared to a similar detector/preamplifier using discrete feedback components, it represents the lower end of the price/performance range for commercially available detectors. The linearity of the OPT209 is specified to be within 0.01% over the full 0 to $10\mu\text{W}$ optical flux range of this system. Although prior measurements indicated significant nonlinearity at low gain levels, these were in fact traced to a nonlinearity in the reference detector, which was transiently saturated by the high peak power of the optical pulses. Recent measurements yielded nonlinearities below 0.1% for all four gain states within the full 0V to 5V output range.

Prefrontal cortical response to Valsalva Maneuver

The motivation for increasing dynamic range was to enable simultaneous optical measurements at multiple source-detector spacings. This was tested with an in-vivo measurement: monitoring the cerebral hemodynamic response to the Valsalva maneuver. With IGC enabled, all three sources provided high quality data with clearly visible cardiac pulsatile modulation, even though the span in detected flux range was greater than 6,000:1. A fixed-gain instrument, which was emulated here by selecting identical gain values for all four sources, lacked the dynamic range to effectively acquire more than two neighboring sources simultaneously. At a gain of 1, only the closest source, which generated the strongest optical return, provided an adequate signal. The signals from the more distant sources were too weak, and were buried in noise. For gain values of both 11 and 121, two neighboring

sources were detectable, with the other source signal either buried in noise or saturated at +5V. At a gain of 1331, only the farthest source was not saturated, and it was the only useable signal detected.

Summary

A significant improvement in dynamic range was achieved with essentially no penalty in crosstalk or linearity by adding individualized gain control to a single channel prototype for a time-encoded DOT system. This was confirmed through both electrical and in-vivo measurements. Although this prototype was intentionally constructed with little concern for grounding and shielding, the channel-to-channel crosstalk was still minimal. By increasing the dynamic range of DOT equipment, IGC techniques may improve the fidelity of optical brain imaging in the future.

4.4 Developing the DOT instruments

The development of hardware is an evolutionary process. Although revolutions (creative ideas, innovations) sometimes occur, the majority of design improvements occur slowly. Prototypes are built, lessons are learned, the design is changed, and performance improves. This ongoing cycle of improvement is sometimes referred to as “spiral development” – homing in on the performance goals by iteratively improving the equipment design in multiple steps. The DOT instrumentation in our laboratory evolved in a similar fashion. This section discusses the evolution of CW DOT instrumentation over a period of many years – from the construction of the first CW imager to the development of the highly flexible, frequency encoded systems in use today.

4.4.1 The Prototype DOT Imager: CW1

CW1 was conceived as an inexpensive and portable continuous-wave DOT system. It was constructed on a shoestring budget, using the technology available and affordable at the time. CW1 contains 18 laser diode sources (9 at 780nm and 9 at 830nm) and 16 silicon detectors, and can acquire 288 independent measurements in less than 4 seconds. These data can then be processed using a variety of imaging algorithms. The general issues involved in the design of diffuse imaging equipment are first discussed and then the instrument is described, along with the technical issues that influenced its design. Some of the technical challenges involved in performing DOT over large optode areas are also discussed.

Introduction

Diffuse Optical Tomography offers the capability to simultaneously quantify the tissue concentration of both oxyhemoglobin (HbO₂) and deoxyhemoglobin (Hb) [82-84]. Two or more near-infrared sources, with wavelengths specifically chosen to straddle the isosbestic point of the oxy/deoxyhemoglobin absorption spectrum, illuminate the tissue at various locations. The flux distribution at the tissue surface thus contains both spectral and spatial information about subsurface absorbers.

In this paper, we discuss the engineering issues involved in the design of diffuse imaging equipment [85-87]. We then describe the design and evaluation of a continuous-wave (CW) DOT system that can image brain function in animals and humans with 4-second temporal resolution.

Background

Ballistic photon imaging is usually performed in a shot noise-limited, photon-starved environment due to the low percentage of photons that actually penetrate the tissue unscattered. Thus the photon-counting detectors require only a modest dynamic range. Diffuse imaging, on the other hand, employs many sources and detectors distributed over a large region of illuminated tissue, so each detector may operate over an effective photon path length ratio in excess of 10:1. With total optode source-detector geometries spanning less than 1cm, as in rat brain studies, this effect is negligible, and the hardware

can easily accommodate the 60-80dB of dynamic range required for rather good quality imagery. On large subjects such as neonates or adult humans, however, optode dimensions of between 4 and 8cm are required. Under these conditions, that same photon path length ratio of 10:1 now translates into a flux attenuation range on the order of $10^6:1$! In order to obtain all the information available, each detector must be capable of operating over a 120dB dynamic range, which is rather large from an electronic perspective. Although many techniques have been developed to solve this problem, all of them require a large capital investment in dollar cost and an increase in both measurement time and hardware complexity.

In order to collect enough photons to reach 120dB, the source power must increase and the detection noise must decrease. The source power will ultimately be limited in human patients by regulatory issues or, in experimental settings, by the vasodynamic effects of tissue heating, or in rare cases, thermal necrosis. The detection noise floor can be reduced by improving the detectors and circuitry, but only to a point. Once the electrical noise contribution is reduced below the noise introduced by the fundamental uncertainty in the photon arrival rate, the system has become photon noise-limited or “background-limited”. The only way to further reduce the noise level is to improve the statistics by increasing the total number of photons collected during each detector sample – which, for a fixed photon arrival rate, means a longer dwell time. This noise reduction is proportional to the square root of the total number of photons collected, and hence to the square root of the total photon collection time. Thus, a fundamental tradeoff exists between measurement time and dynamic range.

If sufficient power and measurement time is available and a signal can be detected, the next challenge begins: what to do with these signals and how to keep them clean. Direct digitization of a 120dB signal would require a 22 bit A/D converter, although logarithmic compression may reduce this to a far more practical 16 bits. Now that the signals can be detected and digitized, the entire system must be capable of preserving their quality. This means that detector channel-to-channel crosstalk, multiplexer settling, and feedthru should be around -120 dB. This can best be achieved by performing the A/D conversion directly following detection. Electromagnetic interference, power, and ground isolation all must approach the quantization limit – about half a DN (Digital Number) or so. All of these requirements can be met, but it would require a carefully constructed system based upon a well-engineered design.

Many other, perhaps better, solutions exist - such as using two or more lower dynamic range fixed-gain detectors at each optode location, or perhaps using gain-switching to reduce the detector gain by known increments when a certain signal voltage threshold is exceeded. All of these still require the same level of attention to signal quality and integrity as discussed above.

Diffuse imaging techniques can be grouped into two general categories: “scalar” and “vector” techniques:

Scalar techniques measure the optical flux exiting the tissue. The system we describe below is a scalar system. There is additional useful information available in the form of the average photon path length. Some light travels deeply into the tissue and returns to the surface, most of the light passes through the classic banana-shaped region, and a small amount travels the shortest path between the source and the detector. A localized absorption at a specific depth will attenuate only the small number of photons which pass through that region, leaving the rest unaffected [88, 89]. Since the average photon path length, and thus the transit time, varies significantly with depth, the magnitude of the photon flux versus the average photon transit time provides a measure of the relative absorption vs. tissue depth – a valuable piece of information. (Although most in-vivo imaging - especially brain imaging - requires light to travel through various tissue types: skin, bone, cerebrospinal fluid, etc., and the resulting optical heterogeneity complicates image reconstruction, these basic concepts still apply).

Vector techniques measure both the magnitude and the average propagation delay, either directly, or in the form of amplitude and phase-shifts relative to the modulation frequency of the light source.

Two common vector approaches in use today are the *time-domain* approach [90] and the *frequency-domain* approach [86, 91, 92]. A time-domain system employs picosecond-wide optical pulses, time-gated photon-counting detectors, time-to-amplitude converters, and the like. Typical frequency-domain systems use a radiofrequency (RF) modulated light source, photomultiplier tubes or fast photodiodes feeding tuned RF amplifiers, and an RF in-phase/quadrature (I/Q) phase detector followed by postdetection filters.

Although time-domain systems are flexible and can detect both ballistic and diffusely scattered photons, they are expensive and, due to their wideband nature, require significant averaging in the digital domain to improve the signal-to-noise ratio (SNR). Frequency-domain systems use simpler, lower-cost components and provide greater SNR. The simplicity and cost savings result from the widespread commercial availability of excellent RF circuitry at modest cost, specifically around common satellite and radar receiver intermediate (IF) frequencies such as 70MHz, 140MHz, 200MHz, etc. The greater SNR stems from the fact that, unlike the wideband nature of time-domain measurements, frequency-domain measurements are usually performed at single RF frequencies, so they occupy a much narrower bandwidth. In the ideal case, the detector electronics need only occupy an RF bandwidth commensurate with about twice the signal acquisition rate (to capture both sidebands), much as the coherent detection system described below. Although this may be impractical due to component instabilities, bandwidths in the Hz to low kHz range are quite practical.

Regardless of the techniques employed, all DOT instrumentation should be designed with the following parameters in mind: Large optode separations lead to significant optical attenuation, so multiple optode spacings will require a **large dynamic range**. The system should be as **linear** as possible over this dynamic range in order to keep the measurements both accurate and precise. **Stray light rejection** is important for systems that must operate outside of the sheltered confines of the laboratory, especially in clinical settings, which often contain large amounts of “optical pollution”. Both optical and electrical **crosstalk** should be reduced to levels commensurate with the system dynamic range, if possible. Good **long-term stability** means fewer and less frequent calibration cycles, and good **temporal response** is important for functional imaging, however it comes at the cost of an increase in the noise floor, which reduces dynamic range.

Designing the CWI DOT system

Our goal was to develop a prototype portable diffuse imager which can be used as a research tool for both characterizing tissue optical properties and guiding us toward future hardware and software design improvements. For simplicity, we decided on a magnitude-only system. Our optical sources were low-power commercially available laser diodes. Although more powerful light-emitting diodes were available, coupling that light into the various core sizes of both plastic and glass fiber would have been very difficult and far less optically efficient. Also, the spectral bandwidth of an LED is far broader than that of a laser diode – wide enough to affect the accuracy of the blood volume and oxygenation calculations. Our detector options ranged from multi-anode PMTs to discrete commercial-grade photodiodes with external preamplifiers. We compromised by choosing a monolithic photodiode/preamplifier IC housed in a clear 8-pin DIP package (OPT209 from Burr-Brown, Tucson, AZ). This offered both the convenience and low-cost of a solid-state detector, combined with the electrical and optical isolation of an integrated preamplifier. Although the parasitic capacitances of the monolithic preamp were about twice that of a well-designed discrete circuit, the simplicity and the relaxed electrical shielding requirements for the monolithic preamp swayed our decision.

We wanted the prototype system to provide a useable (i.e. linear) dynamic range of at least 80dB. Dynamic range and linearity are closely paired, since each must be defined with reference to the other. For example, a statement of a given dynamic range value is, in itself, meaningless unless the criteria

for determining the limits of that range are expressly stated. In our case, our goal was 80dB of dynamic range with less than a 1% deviation from a least-squares-fit line. Although this linearity value was somewhat arbitrarily selected, it gave us a good starting point around which to base our design. Since valuable information can be gleaned by linearizing even strongly saturated signals, we decided to “bracket” the signal swing within the dynamic range of our A/D converter.

In order to meet the 80dB objective, crosstalk, drift, settling time, and feedthru errors had to match this goal. Optical crosstalk was minimized by concealing each detector package in opaque heatshrink tubing and providing sufficient separation to further attenuate any stray reflections within the metal housing. Each detector fiber was sheathed in opaque tubing as well, which also served to protect the fragile cladding of the PMMA (acrylic) fibers from abrasion. Electrical crosstalk occurring at the front end was not significant since the only high impedance node in the preamp was physically removed from the DIP package and the modulation frequency was in the kilohertz range. Ground loops were minimized by using an electrically isolated power supply, a battery-powered portable laptop computer, and a single-point earth connection to safety ground. Crosstalk through power lines was minimized with on-board regulation and the use of separate power supply decoupling filters at each opamp.

To reduce settling time errors, we planned for a 5τ dwell time prior to conversion. For a post-detection time constant of 40ms, this gave us a minimum delay of 200ms per source. Faster acquisition rates could easily be achieved, albeit at the cost of higher crosstalk. Feedthru among multiple channels within the analog multiplexer was minimized by switching only demodulated (baseband) signals and by reducing the DC currents through the switches by placing a buffer directly at the multiplexer output. This also provided better linearity, since switch impedance varies with common-mode voltage.

We chose to use a rather fast 133 kilosample-per-second, 16 bit A/D converter (Analog Devices AD7884) for a couple of reasons. The high sample rate allowed us to rapidly scan through the 16 demodulated outputs during high data rate measurements, while still permitting us to oversample and then average in the digital-domain during slower measurements. The parallel data output design also made it easy to interface with all of the digital I/O boards used in our lab. A summary of our performance goals for the prototype DOT imager is shown in Table 4.6.

Table 4.6. Performance goals for the prototype DOT imager.

PARAMETER	GOAL
DYNAMIC RANGE	10,000:1 (80dB)
NONLINEARITY	<1% over the 80dB dynamic range
SETTLING TIME	<300ms to 0.1%
CROSSTALK	<0.01%
DIGITAL RESOLUTION	16 bits
SOURCE CHANNELS	9 at 780nm and 9 at 830nm
SOURCE OPTICAL POWER	~5mW
DETECTORS	16 Si photodiode/preamplifiers (OPT209)
MODULATION TECHNIQUE	single-phase squarewave AM with coherent detection
POSTDETECTION BANDWIDTH	10 to 20Hz
STRAY LIGHT REJECTION	<1% error under normal illumination levels
PACKAGING ISSUES	must be portable, compact, and extremely rugged
POWER REQUIREMENTS	120VAC +/-10%, 50-60Hz
PATIENT SAFETY	leakage current <1 μ A, case-to-gnd impedance <1 Ω

Since we planned to perform some measurements in a clinical setting, the unit had to operate within a normal hospital environment. This placed a severe stray-light rejection requirement on our design. We needed to detect picowatts of source signal under an ambient optical background in the microwatt range (as seen by the detector). We solved this problem by using synchronous detection.

Synchronous detection, is a technique which offers substantial advantages when detecting small signals buried in random, uncorrelated noise (with respect to the modulation frequency). In order for coherent detection to work, the source must be intensity-modulated in some known fashion, usually in the form of a 50% duty cycle squarewave or sinusoidal waveform, which we create by modulating the bias current through the laser diode sources.

Each photodetector preamp output is first highpass filtered to remove low frequency signals from stable interference sources like sunlight and 1/f (flicker) noise generated by the electronics. The medium frequency components produced by fluorescent lamps at 120Hz and its harmonics are somewhat attenuated as well. This filtered preamp output is then fed into a double-balanced mixer. The mixer, which is gated by the same modulator that controls the laser diode source intensity, synchronously rectifies the weak modulated source signal, which appears as a small DC voltage at the mixer output. All other spurious modulated optical signals (including those produced by line-powered lamps, computer terminals, multiplexed LED displays, etc.) which are not phase-coherent to the source will exit the mixer in the form of frequency-shifted AC signals.

A lowpass filter placed at the output of the mixer strongly attenuates these incoherent signals, leaving only the small DC voltage proportional to the magnitude of the source energy detected. The time constant of this lowpass filter controls the post-detection bandwidth, which significantly reduces the noise level, but also limits the rate at which the desired signal intensity changes can be detected. It is important to set this bandwidth as low as practical to minimize noise, since the noise is proportional to the square root of the post-detection bandwidth.

The modulation frequency, which is typically between 100Hz and 100kHz, can be adjusted slightly to shift the detection passband away from narrowband interference sources that happen to fall near harmonics of the modulation frequency. These interference sources can sometimes generate intermodulation products (undesired signals created in the demodulation process, equal in frequency to the difference between the interference frequency and the modulation frequency or one of its harmonics) within the postdetection bandwidth, which are sometimes difficult to remove otherwise.

We used a nominal modulation of 2kHz, which was chosen to shift the detection band out to a frequency region where the flicker noise was low, yet to where the detector/preamps were still capable of responding to the third harmonic at 6kHz. Although third harmonic response was not critical for operation, it served to improve the “average-to-peak” voltage ratio of the received signal, and so provided a little more dynamic range at the high end. The post-detection time constant was set at 40ms, and a dwell time of ~200ms was chosen to provide sufficient settling for each detector channel. This time constant was chosen to both provide significant noise reduction and yet remain fast enough to perform basic functional imaging studies.

Both source and detector selection was under direct, real-time computer control. This both simplified the digital design and offered us the flexibility to operate as few or as many sources and detectors as were needed. A photo of the prototype imager in operation is shown in Figure 4.14 and a functional block diagram is shown in Figure 4.15 below.




Figure 4.14. A photo showing CW1 in operation. The optode assembly is immersed in a fluid scattering phantom with a ~1cm absorber suspended in the fluid by means of a thin wire supported by the alligator clipstand. The DOT image of the absorber is visible on the computer display, which correctly reveals its location at the lower right edge of the optode array.

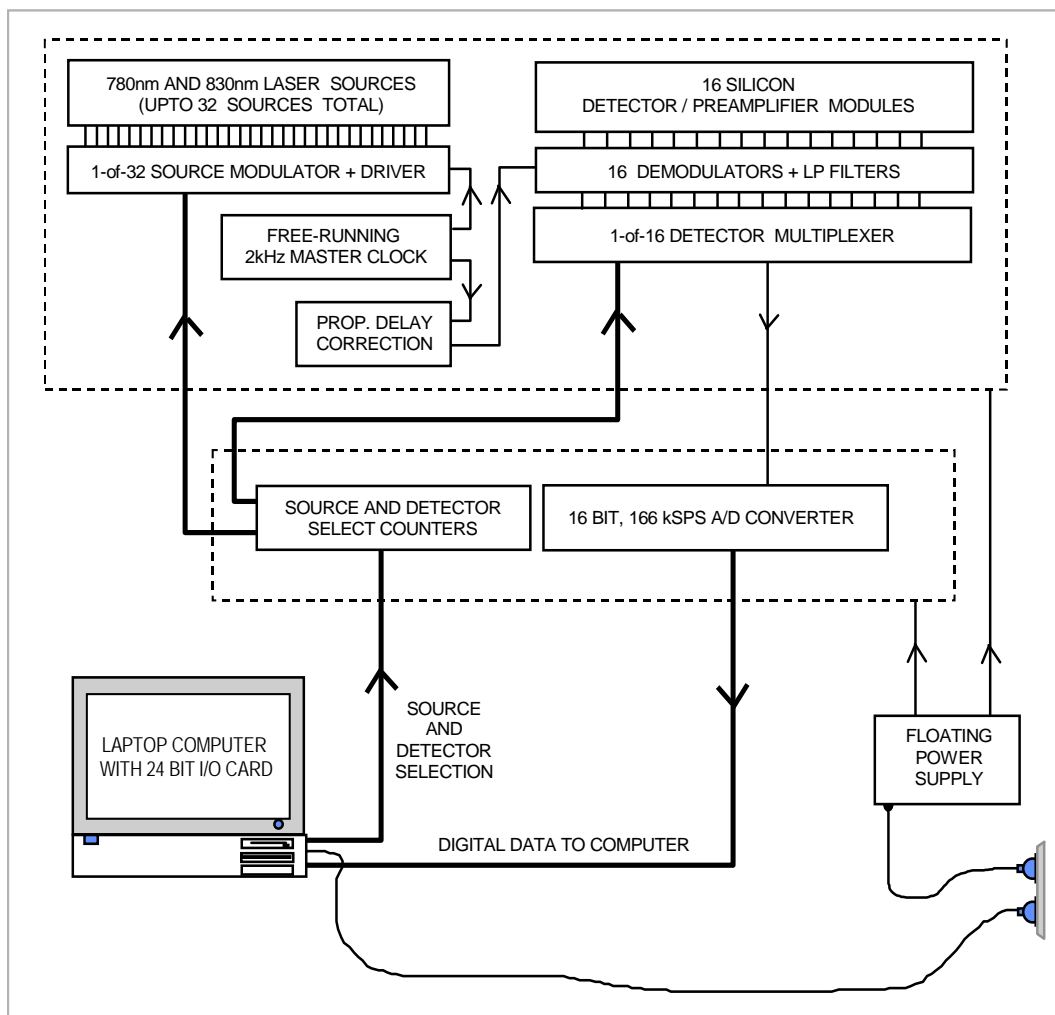


Figure 4.15. A block diagram of the CW1 prototype DOT imager, which contains nine 780nm sources and nine 830nm sources. The analog outputs of all 16 detector channels, each consisting of an OPT209 detector followed by a synchronous demodulator, are digitized on-board and then read directly to computer memory using a commercial PCMCIA data acquisition card. All functions, including source and detector selection, were computer-controlled.

DOT system performance evaluation

Our measurements revealed the overall system performance to be slightly better than expected. Crosstalk and stray light rejection were not the performance limiters, as had been feared. The lasers provided enough power to handle an optode geometry of 4cm x 6cm (for piglet studies), yet the detectors could handle a 1cm x 1.5cm geometry (for rat studies) without saturating. This meant that from a detector sensitivity standpoint, we guessed well. A summary of our test results is shown in Table 4.7.

Table 4.7. Results of the measurements performed on the prototype DOT imager.

PARAMETER	MEASURED VALUE
NOISE EQUIVALENT POWER	<40pW RMS (measured with 32 samples per dwell)
DYNAMIC RANGE	~45,000:1 (92dB) @ 0.75% nonlinearity ~25,000:1 (88dB) @ 0.05% nonlinearity
LONG-TERM STABILITY	+/-1% of reading in 30 minutes (half-scale output)
INTERCHANNEL CROSSTALK	<1:20,000 (-86dB)
STRAY LIGHT REJECTION	~1 DN signal change from darkness to normal ambient using cool-white fluorescent and incandescent lamps
TEMPORAL RESPONSE	~20Hz
POWER DELIVERED TO TISSUE	2mW @ 780nm, 8mW @ 830nm

The system noise floor was obtained by calculating the standard deviation of a number of readings taken at the lowest resolvable source intensity. The noise equivalent power was then calculated for an SNR of unity, using the power-per-unit-flux value obtained from the linearity measurement.

Dynamic range and linearity are related: the dynamic range can only be defined with reference to a specified linearity limit. This is why two dynamic range/linearity values were presented in Table 4.7. The incident beam from an optical source was sampled with a “monitor” fiber that led to a calibrated power meter. The linearity was then measured by comparing this value to the DOT instrument reading over a wide range of flux levels. The flux was varied using neutral density filters. A diffuser was used to reduce the spatial coherence of the source, which prevented modal noise in the fibers from interfering with the measurement.

Long-term stability was measured by noting the total drift between initial and final readings for a fixed probe geometry on a static phantom.

Interchannel crosstalk was measured as a single detector channel was alternately driven between the noise floor and full-scale output using the diffuser/ND filter technique described earlier. The largest level change among the other fifteen channels was recorded.

Stray light rejection was measured by operating the system with a static phantom in a dark room and then turning on both the fluorescent lights and a computer display located about a meter away. The largest change among the sixteen detector channels was recorded. The detected power in the visible band was ~1.5uW. The temporal response was defined by the bandwidth of the single-pole postdetection filters, which were set to 20Hz. This was not directly measured.

The power delivery was measured at the end of ~1 meter of 1mm dia. acrylic fiber using a calibrated power meter. No index-matching with the detector was attempted. This will underestimate the actual power by ~4% for in-vivo measurements due to the improved coupling through index-matching from perspiration trapped at the fiber/skin interface.

Future systems under development

We are currently involved in the design of a number of systems ranging from general multipurpose imagers to specific purpose-built units. Based on the lessons learned during the development of the DOT imager described above, we have designed, and are currently building, an autonomous diffuse imager for optical mammography studies. This system will collect and store complete frames of data within an on-board memory module. Both the detector bandwidth and the sample rate will be

adjustable to optimize the acquisition speed. We plan to double our acquisition rate by operating two lasers simultaneously and resolving the two signals using phase diversity. We are also developing a stroke monitoring system for humans, using four high power laser diode sources and discrete detectors, held close to the scalp to minimize fiber coupling and attenuation losses. Although all four lasers operate at the same wavelength, each source-detector pair will be isolated through a combination of phase diversity and tissue attenuation.

Phase Diversity

The prototype DOT imager discussed above can only energize one source at a time, which limits our acquisition rate. A more efficient technique would be to exploit the phase diversity afforded by coherent detection by modulating each laser wavelength at the same 2kHz frequency but in phase quadrature with each other. Double-balanced mixers are insensitive to coherent signals that arrive exactly 90° out of phase with the demodulation clock. When a quadrature signal passes through the mixer, the DC level of the resulting signal averages out to zero, just as with uncorrelated noise. The mixer, fed by a detector signal containing both in-phase and quadrature components (generated by the two laser sources), will demodulate the in-phase signal only, and will completely ignore the quadrature component. The same holds true for the second mixer, fed with a “quadrature” demod clock. Double-pole postdetection filters are used to attenuate the strong second harmonic component produced by the quadrature source.

There are some limitations to this technique, however. Interchannel isolation is critically dependent on accurate and stable phase control throughout the system. The good news is twofold: Since the propagation delay through most components varies little with voltage or temperature, long term stability is expected to be excellent. A four-channel prototype system has maintained an interchannel isolation of more than 50dB for many months. Furthermore, since the *difference* between μ_a at 780nm and at 830nm is relatively small, the amplitude difference between the two signals is only around 10:1 (20dB), so an interchannel isolation of 40dB is more than adequate for our needs.

Two disadvantages with this technique, however, are a reduction in dynamic range and a slight increase in the noise floor when using background-limited detectors (from the additional flux generated by the quadrature source). An example of such a system is shown by the schematic in Figure 4.16.

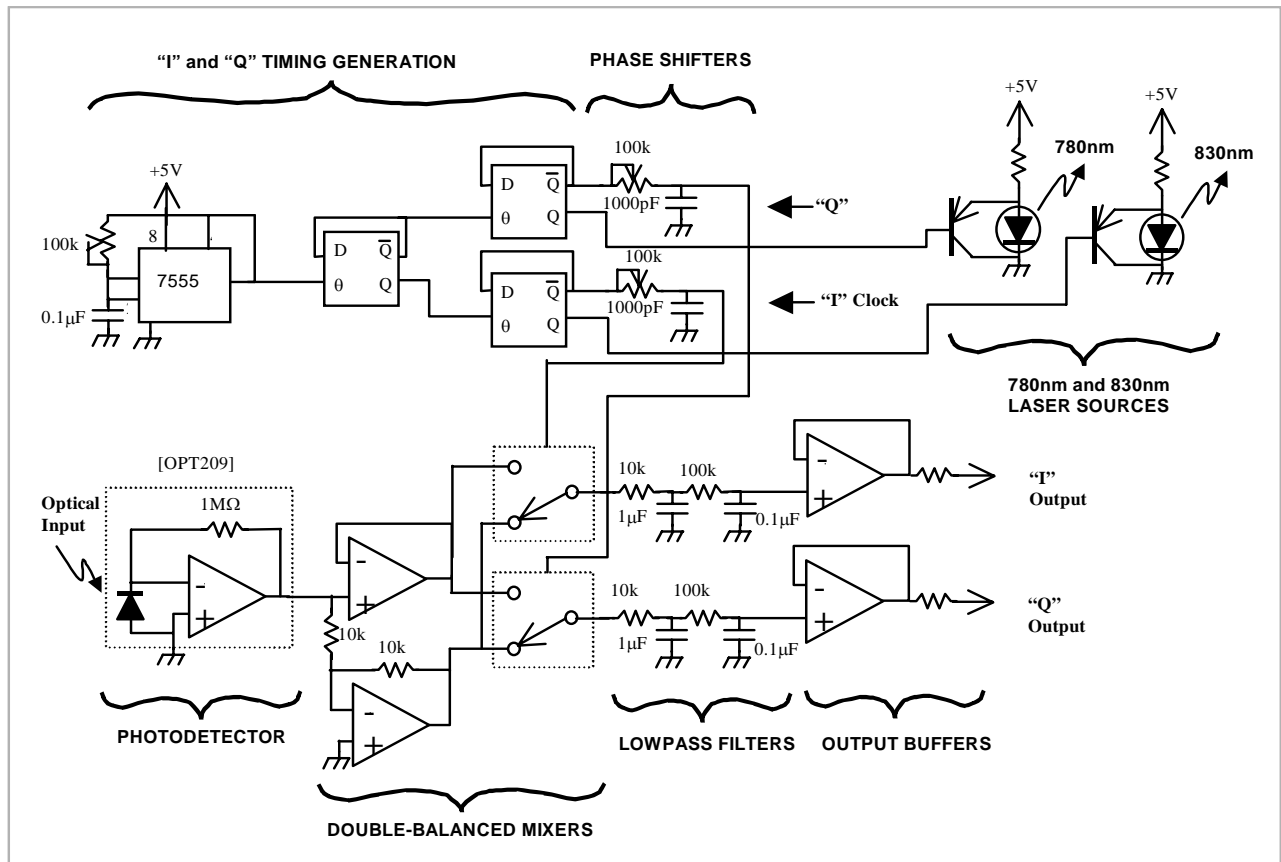


Figure 4.16. The schematic diagram of a single detector phase diversity prototype system. A four-channel version was later developed to evaluate phase diversity in more complex two-wavelength systems.

Summary

We discussed a number of issues that influence the design and development of diffuse imaging instrumentation. We then described the construction and testing of a prototype time-encoded CW DOT imager, which was later used to acquire DOT images of rat brain function following forepaw stimulation. Future time-encoded prototype DOT systems (the Animal series) will explore the use of phase diversity and gain-switching to further improve sensitivity and expand dynamic range.

4.5 The “Animal” Instrument series and their progeny

4.5.1 Animal I

Animal I was originally conceived of as a prototyping testbed for exploring the capabilities of two new multiplexing techniques, a phase-encoding technique called phase-division multiplexing and a time-encoding technique called pulse-TDM, which were being developed to eliminate the temporal skew problem inherent to the standard “switched-source” time-division multiplexing used in the CW1 imager. The results were promising. Phase-division multiplexing worked well for performing simultaneous dual-wavelength measurements. It allowed a doubling of throughput without increasing the channel bandwidth. When properly adjusted, the interchannel crosstalk was less than 0.1% (-60dB). Its main disadvantage was that only two sources could be used per modulation frequency in any one region of the head. Other disadvantages were a 6dB reduction in dynamic range and the need for stable phase alignment. For its initial application (single-point blood volume fraction and Hb/dHb

ratio measurement), these disadvantages did not pose a problem. Only two sources were required for these initial measurements. Dynamic range was adjusted manually, with the gain controlled through resistor substitution. Phase alignment was easy to adjust and remained stable for months.

The phase-encoding technique validated by Animal I was later used in the design of the Optical Fetal Monitor, an instrument developed to monitor fetal blood oxygenation.

4.5.2 The Optical Fetal Monitor

The Optical Fetal Monitor (OFM) was developed in an attempt to quantify fetal blood oxygenation noninvasively, through a simple transabdominal measurement. It measures the amount of light absorbed by the fetus at two different near-infrared wavelengths. The fetal modulation signal can then be uniquely resolved from the maternal signal by its frequency content.

The optical flux reaching each detector optode will be modulated by both maternal and fetal cardiac activity as the elastic arterial walls expand slightly and then relax following each ventricular contraction. This appears as a real-time modulation of blood volume within the tissue. Any light passing through tissue perfused by these vessels will exit with a small amplitude modulation in the shape of an inverted sawtooth. A sharp drop in optical signal during ventricular contraction (systole) is followed by a slower signal rise as the blood passively exits the capillaries and drains into the venous compartment on its way back to the heart (diastole). This is illustrated in Figure 4.17.

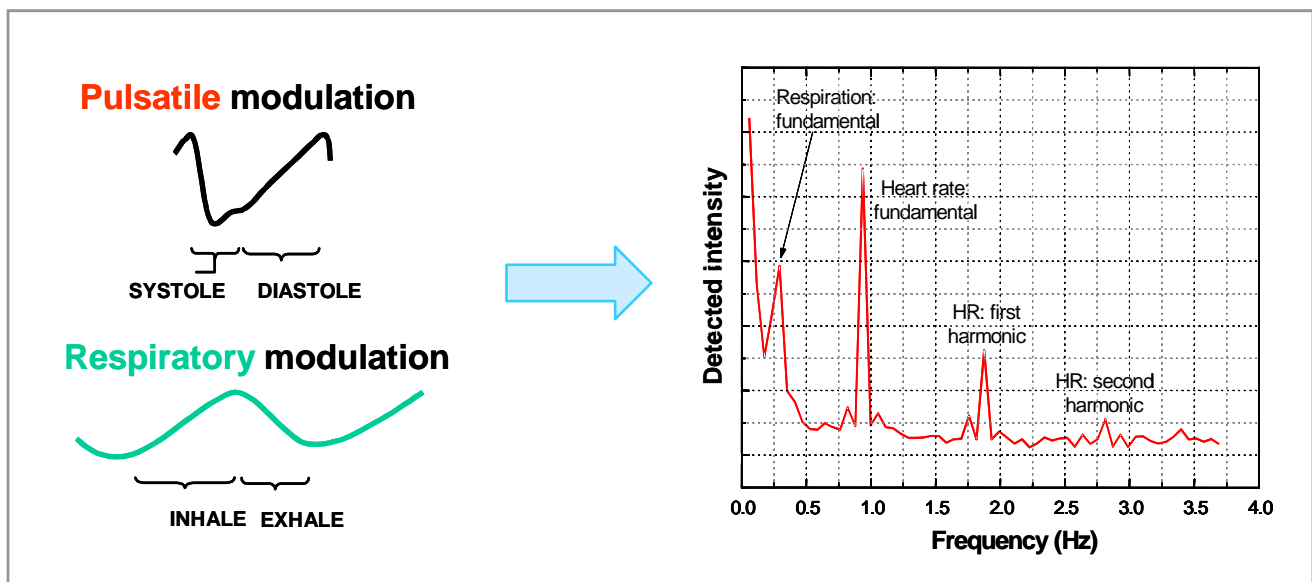


Figure 4.17. Pulsatile and respiratory pressure changes within the vascular system amplitude modulate the reflected optical signal. These optical modulation waveforms are distinctive, and contain sufficient temporal information to allow both the maternal and fetal heart rate to be simultaneously measured.

Since the mother and fetus have separate hearts and thus distinct heart rates, the blood passing through the fetus will be “frequency tagged” by the fetal heart rate. A Fourier transform can then be performed on the detected signals to discriminate the fetal signal from the maternal signal, since the mother’s heart rate, and hence her contribution to the amplitude modulation, should be far slower. By performing measurements at two or more optical wavelengths, it should be possible to measure the oxygen saturation of both the maternal and fetal blood over time using standard pulse oxygenation calculations. Although absolute calibration is probably not possible due to optical shunting errors, a

decrease in fetal oxygenation, which may indicate inadequate placental or fetal perfusion, should be detectable.

The prototype system consisted of a single wavelength light source (an 808nm fiber-coupled laser diode) and four separate optical detectors (Burr-Brown OPT209). Synchronous demodulation was used to both reduce extraneous interference from ambient light sources and to improve the signal to noise ratio of our measurements. The optical source was electrically modulated at a frequency of around 2kHz or so. This frequency could be varied slightly to avoid aliasing with harmonics from AC-powered light sources. The detectors simply detect any in-band optical flux they see, which can include sunlight, electric lighting, or even the light from computer CRT displays. Somewhere amongst all of that lies our weak modulated optical signal.

We retrieve that signal by first passing the detector output through a high pass filter (to remove any slowly varying components produced by sunlight or incandescent lamps) and then we feed the output into a synchronous detector. This circuit is similar to an RF mixer: it electrically multiplies the filtered detector signal with a “carrier” signal using analog switches. This carrier signal is the same 2kHz signal used to gate the source on and off. The only difference is that a slight time delay is added to the carrier to make up for the time it takes for the optical signals to propagate through the detector circuitry. This way the gated light from the optical source and the carrier signal arrive at the synchronous detector at exactly the same time. This is important because the synchronous detector will act as a rectifier for the gated optical signal – and only that signal. All other stray light signals (which are all uncorrelated to the gated source) will exit the synchronous detector as frequency-shifted AC signals. But the small signal component from the gated source (the same one which hopefully passed through both the mother and the fetus) will be “synchronously” rectified, and will appear as a small DC voltage.

Since this synchronous process is linear, the DC level tracks the amount of gated light reaching the detector (yes, the demodulator *multiplies* two signals, but think of it as multiplying the gated source signal by the number “1”). The output of the synchronous detector then passes through a lowpass filter to eliminate all the other frequency-shifted AC signals which we don’t want.

The bandwidth of this filter is important, since it sets the frequency response of our measurement and it thus determines the noise floor. A wide bandwidth would allow us to measure lots of heartbeat harmonics, but at the price of lower sensitivity due to a higher noise floor. A narrow bandwidth would yield much better sensitivity, but might not be fast enough to catch the second or third heartbeat harmonics, which are important in discriminating the fetal signal from the maternal signal.

A drawing of the prototype OFM unit is shown in Figure 4.18. All functions are controlled by the computer. The 808nm laser source module generates the optical signal which reaches the optode assembly through a single 1mm diameter silica fiber. Its output is modulated by an oscillator located in the control module. The four detectors, each spaced one inch (25mm) apart on the optode assembly, couple to the patient through short lengths of 3mm diameter PMMA optical waveguide. The detector outputs travel back through the umbilical cable to the control module, where each output goes to its own amplifier. Since the detector located farthest from the source will receive the least light, its signal receives the most gain (1000x). The more proximal detectors receive proportionally less gain (100x, 10x, 1x). This is necessary, since the amount of signal reaching the distal detector can be ten thousand times weaker than that reaching the most proximal one.

Scaling the gain is an easy way of conserving our dynamic range for more important things, like fetal heartbeat signals. Since each patient is likely to have a unique distribution of adipose tissue and an unknown vascular structure, the actual detector signal levels will vary from mother to mother, and our dynamic range had to allow for this.

The four synchronously rectified and filtered DC outputs then pass from the control module to the signal processing computer. They are digitized at a constant rate of 1000Hz and stored in memory.

While this is going on, the computer grabs a portion of the data and uses that time sequence to generate a Fourier transform, which is then displayed. This presents the data, sampled in the time-domain, as a 2-D plot in the frequency-domain. Signal power components modulated at different frequencies now appear distinct from each other. In this way we can measure the modulation created by both the mother's heartbeat and the fetal heartbeat as discrete quantities – which is important for blood oxygenation measurements.




Figure 4.18. A diagram of the Optical Fetal Monitor, showing the modular format. It consists of five modules: the laser source module, the control module, the optode assembly, the signal processing computer, and the system power supply. Optical power from the laser source module is fed to the optode assembly in contact with the mother's abdomen. The weak optical signals reflected and scattered back to the skin surface (curved red lines) are detected by four OPT209 detectors, which amplify the photodiode current and buffer the resulting signal. These four signals then travel through shielded cables to the control module, where they are synchronously rectified and then filtered. The filtered signals are then sampled by a PCMCIA ADC card in the computer and stored in memory for later analysis. Approximately 98% of the light traveling from the source to the most distal detector will pass through the fetus en-route.

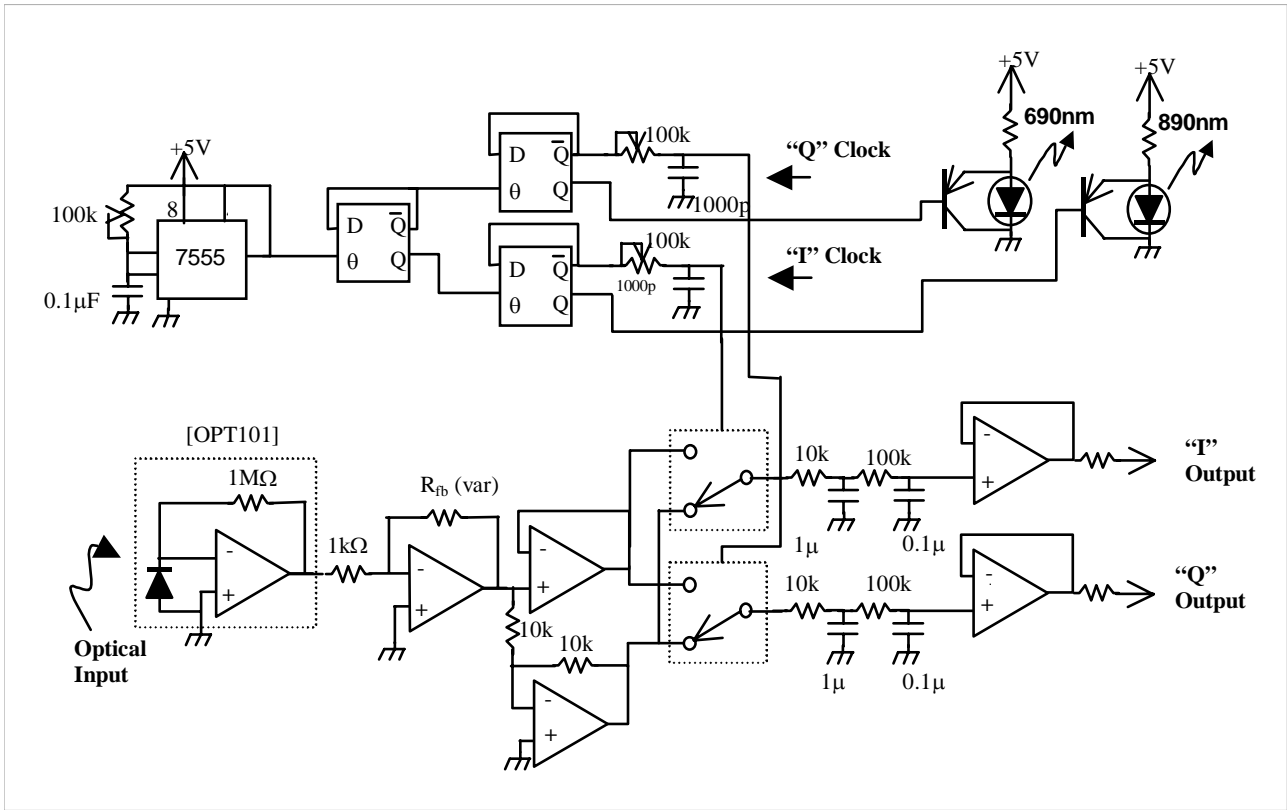


Figure 4.19. A schematic diagram of the more advanced two-color version. The lower portion of the circuit, containing the detector, gain stages, and demodulators, was replicated four times – once for each detector channel. The initial one-color instrument contained one 808nm laser diode source and only contained a single-channel demodulator. The improved instrument contained two laser diode sources and employed phase-encoding. This permitted both optical signals to be demodulated separately, despite the use of a single photodetector and gain path for each channel.

Table 4.8. Results of a Monte-Carlo calculation which quantified the fraction of light reaching the fetus vs. optode spacing. At the smallest optode spacing, only 5% of the light actually reached the fetus. This fetal-crossing light only spends a small portion of time scattering within fetal tissue, so maternal modulation would dominate, and little if any fetal signal would be detected. At the largest optode spacing, nearly all of the light reached the fetus and spent a large fraction of time within fetal tissue, however some maternal modulation is unavoidable, since the light must transit the abdominal wall in order to reach the fetus. This “optical shunting” complicates the quantification of fetal blood oxygenation.

SOURCE-DETECTOR SEPARATION	25mm	50mm	75mm
TOTAL SIGNAL DETECTED	1580mV	820mV	36mV
FRACTION PASSING THROUGH FETUS	5%	65%	98%

Unfortunately the two-color version was never completed, so no further measurements were collected with this instrument.

4.5.3 Animal II

Pulse-TDM encoding offered even more promise than phase-encoding as an expandable scheme to ultimately replace Switched-Source TDM in area measurements for DOT. In order to evaluate Pulse-TDM, Animal I was then upgraded to perform simultaneous measurements at three wavelengths, and was renamed “Animal II” to signify this change. It employed Pulse-TDM encoding at 685nm, 780nm and 830nm, with a ~1ms risetime (DC-350Hz bandwidth) and a temporal skew of less than 200us. It was specifically designed to search for intrinsic optical scattering events (“fast signals”) in the whisker barrel cortex of rats following whisker stimulation. The source optode consisted of three 220um fibers arranged in a triad and secured with cyanoacrylate adhesive, and the detector optode consisted of a single 400um silica fiber. Both were secured inside the bore of a polystyrene pipette to provide both strain-relief and a means of mechanical support during placement. Measurements were conducted over a long period of time and the instrument performed well, however no event-related intrinsic scattering changes were detected.

4.5.4 The NIRS3 four-wavelength system

Since Animal II demonstrated that Pulse-TDM encoding was a viable technique, a nonimaging point measurement system suitable for performing Near-Infrared Spectroscopy (NIRS) was designed around this technique. The NIRS3 instrument is a time-encoded CW instrument designed for research involving photon migration in both animals and humans, and contains four pairs of optical sources and eight identical optical detectors. The detected signals are demodulated internally and fed in analog form to a panel-mounted connector which is compatible with the PCM-DAS16/16 data acquisition board and interface cable, manufactured by Computer Boards, Inc., Middleboro, MA.

Features and Specifications:

Laser diode sources

Designation:	A	E	B	F	C	G	D	H
Wavelength (nm):	680	680	780	780	830	830	980	980
Manuf. Part No.:	HL6738MG		DL7140-201		HL8325G		L9801E3P1	
Peak Power (mW):	10	10	10	10	10	10	10	10
Avg. Power:	1.25	1.25	1.25	1.25	1.25	1.25	1.25	1.25
Modulation Rate:	Jumper-selectable							
Duty Cycle:	12.5%							
Gateability:	By source, software-controlled							
Gating Depth:	Greater than 103dB							
Optical Interface:	SMA902 female (ref. Figure 4.3)							

Detectors

Sensor:	Silicon APD
Type:	Hamamatsu C5460-01 avalanche photodiode module [80]
Preamp:	DC-coupled inverting TIA, gain = -1E7 Volt /Amp
Active Area:	3.0mm dia.
Spectral Response:	400nm to 1000nm
Responsivity:	1.5E8 Volts/Watt at 800nm
Bandwidth:	DC to 100kHz
NEP (800nm):	20fW/Hz ^{1/2}
Output noise:	1mVrms
Max. flux:	60nW

Detector Channels

Gain Scaling -

Gain Setting:	“1”	“3”	“10”	“30”	“100”	“300”
Actual value:	1.01	3.33	10.00	30.39	100.0	297.25
Source Designation:	1, 3, 5, 7			2, 4, 6, 8		
Output Voltage Range:	-4.75 to +0.25VDC			0 to +5.0VDC		
Temporal response:	~3ms risetime, ~100Hz bandwidth					
Nonlinearity:	Less than 0.3% from a linear LSF line over a 5V span					
Static Crosstalk -						
Intrachannel:	Maximum 0.022%					
Interchannel:	Maximum 0.16% (worst-case conditions)					
Dynamic Crosstalk:	Less than 0.01%					

System

Typical Noise Performance (std. dev., 100Hz bandwidth) -

Gain Setting:	1	3	10	30	100	300
Actual value (mV):	1.25	1.25	1.34	2.00	5.21	15.42
Dynamic Range:	4000	4000	3731	2500	960	324
RMS noise vs. detector channel gain vs. number of averages / bandwidth:	(see plot)					
System Stability vs. Time:	(see table)					

Safety

Power:	110 to 130VAC, 60Hz
Fuse:	3AG (AGC) type, 4 Amp Fast-Blow
Leakage:	<1uA from AC lines to safety ground
Ground Impedance:	<1 Ohm from case to ground pin on power plug

Operating Conditions

Altitude:	0 to 15,000ft
Humidity:	0 to 95% RH, noncondensing
Temperature -	
Operating:	55°F to 90°F
Storage:	30°F to 120°F
Compliance:	Y2k / Y3k compliant

Circuit Design Approach

The purpose of this instrument is to measure the flux which reaches each detector from up to four different source wavelengths simultaneously. This is achieved using a technique called pulsed time-domain multiplexing, or Pulse-TDM for short, in which each of the four sources is time-interleaved and pulsed at a duty cycle of 12.5%. Thus the detectors see only the flux from one source wavelength at a time. The demodulation circuitry decommutates these pulses and synchronously rectifies them, yielding four analog outputs whose amplitude corresponds to the amount of flux being detected from each of the four source wavelengths. This is shown pictorially in Figure 4.20.

Each of the eight detectors will generate these four outputs. This means that a total of 8 x 4, or 32 independent analog channels must be made available to the PC Card for data acquisition. Since it was

decided to use one 16-input ADC card, the choice was made to multiplex the data into two banks of 16 analog channels each. In order to eliminate ambiguity as to which bank of data is which, one bank will be ground-referenced and the other bank will be referenced to -4.75V . This voltage was chosen to provide a 250mV safety margin above the -5V lower dynamic range limit of the ADC, to handle some voltage drift with time. Synchronization between the ADC sample rate and the analog multiplexer is achieved by dividing down the ADC convert pulse and using this clock to control the multiplexing process.

A block diagram of the system is shown in Figure 4.22 and the timing diagram is shown in Figure 4.21. Timing is derived from a crystal clock, and drives both the laser pulsing circuitry and the demodulation circuitry. The optical pulses generated by the lasers are detected by the silicon APD detectors. The electrical signals pass through a variable-gain amplifier and then on to the demodulator.

The clock circuitry was designed to be extremely flexible so that the timing could be varied as needed. Both the modulation clock and the settling delay period are jumper-selectable.

The laser driver circuitry uses a class A design, similar to that shown in Figure 4.6, to minimize radiated and conducted EMI at the cost of increased power dissipation. Since this unit was designed to operate in a climate-controlled environment from 120VAC line power, the additional 10 Watts of power consumption was not considered to be a problem.

Each laser can be gated independently. This is achieved by sending an 8-bit serial data word to the unit, using DIO#1 and DIO#2 as normal and inverted clocks, and DIO#3 for the data. Each bit controls one laser, and is loaded as D1 . . . D8. The ordering is shown in Table 4.9 below:

Table 4.9. The laser diode control register settings.

Register Bit Order:	D1	D2	D3	D4	D5	D6	D7	D8
Panel Label:	G	C	H	D	E	A	F	B
Wavelength:	830nm	830nm	980nm	980nm	680nm	680nm	780nm	780nm
Software Label:	5	6	7	8	1	2	3	4

The variable gain amplifier (VGA) consists of two identical series-connected amplifier stages. Analog switches are used to select various feedback resistor combinations to close the loop around the two opamps. Two separate gain stages were chosen to provide a wide (1000 to 1) gain selection range while still allowing the use of standard operational amplifiers with modest gain-bandwidth products. This both reduced cost and minimized the risk of oscillation from excess loop gain at low settings. The additional propagation delay introduced by the second stage is minimal in comparison to the temporal response of the APD, and so it can be safely ignored. Discrete resistor programmed gain (as opposed to logarithmically-variable gain) was selected because it provides very stable and reproducible gain values, each of which can be changed relatively easily in hardware, if desired. Gain stability and repeatability are important for certain optical measurements.

The gain of all eight preamps is controlled by a 32-bit dataword sent in serial fashion, beginning with D1 and ending with D32. DIO#7 and DIO#6 are clock and clock_bar, and DIO#5 is data. Each nybble controls the gain of one preamp. The preamp control register settings are shown in Table 4.10 below:

Table 4.10. Variable-gain preamplifier control register settings.

Register Bit Order: D1-D4				Channel 1 Gain	Detector Channel	Data Bits
D1	D2	D3	D4			
0	0	0	0	1	1	D1 - D4
0	0	0	1	3.3	2	D5 - D8
0	0	1	0	10	3	D9 - D12
0	0	1	1	33	4	D13 - D16
1	0	1	0	100	5	D17 - D20
1	0	1	1	330	6	D21 - D24
1	1	1	1	1000	7	D25 - D28
					8	D29 - D32

The gain values were selected to cover as much dynamic range as practical, given the saturation limits of the APD detector modules. The lowest gain setting was set at 1V/V to exploit the maximum 10Vpp output swing available from the APD module. The highest gain setting was chosen to ensure that the noise floor would not be ADC quantization-limited after coaddition of enough samples to reduce the effective noise bandwidth to ~0.1Hz (this represented the combination of temporal averaging down to 3Hz and coaddition of 1000 time-locked recordings). For the APD detectors, the RMS noise over a 100kHz bandwidth was previously measured at 1mV. So compressing this to 0.1Hz yields about 1.2uV of detector noise (flicker noise is being neglected here because it will be reduced both by the avalanche gain and by the modulation process.) The ADC operates over the +/-5V input range, which yields a voltage of 153uV per digital count. It was assumed that the ADC quantization noise floor would be around 1 count (which is about triple the theoretical limit) due to differential nonlinearity and conversion noise. So the noise of the amplified signal must be at least three times this value, about 450uV RMS, to avoid being quantization noise limited. This indicated that a gain of about 400V/V would be sufficient. A 3x gain step was judged to be adequate for DOT measurements, so a maximum gain value of 1000V/V was chosen to provide a convenient gain step sequence: 1, 3.3, 10, 33, etc. up to 1000V/V. Note that, since these gains are determined by discrete resistors, they can be changed to set the gain steps to other values, if needed.

To prevent external flux from debiasing the amplifiers, their inputs are AC-coupled. Since there is no obvious feedback to the user when the APD or the amplifiers are clipping or are saturated by excessive flux, a clipping detection circuit was designed. When the APD output is saturated, the red “DC” overload indicator illuminates. When the amplifier stage begins to clip, the red “AC” overload indicator illuminates.

The demodulator is designed to both demultiplex and demodulate the signals from all four laser wavelengths simultaneously. This is performed with a circuit which functions like a multichannel scaler: during the first of eight time intervals, the first laser wavelength is operating, and the detected signal is fed into the first “bin” (a two-pole RC filter) for averaging. During the second time interval all lasers are off, and this reference signal is averaged in the second “bin.” The third and fourth time intervals are for the second laser wavelength, and so on. The signal and reference bins for each source are then fed to a unity-gain differential amplifier stage which removes any common-mode interference and converts the signal to normal mode (i.e. ground-referenced) form. This demodulator design is an improvement over classic double-balanced demodulator designs using gain inversion, because it performs both demultiplexing and demodulation with one analog multiplexer.

The outputs from all 32 outputs (8 detector channels x 4 laser wavelengths) go to a 16 wide 2-to-1 analog multiplexer. The input selector for this mux comes from a five stage binary counter clocked by the “ADC convert” pacer clock output on the PC card.

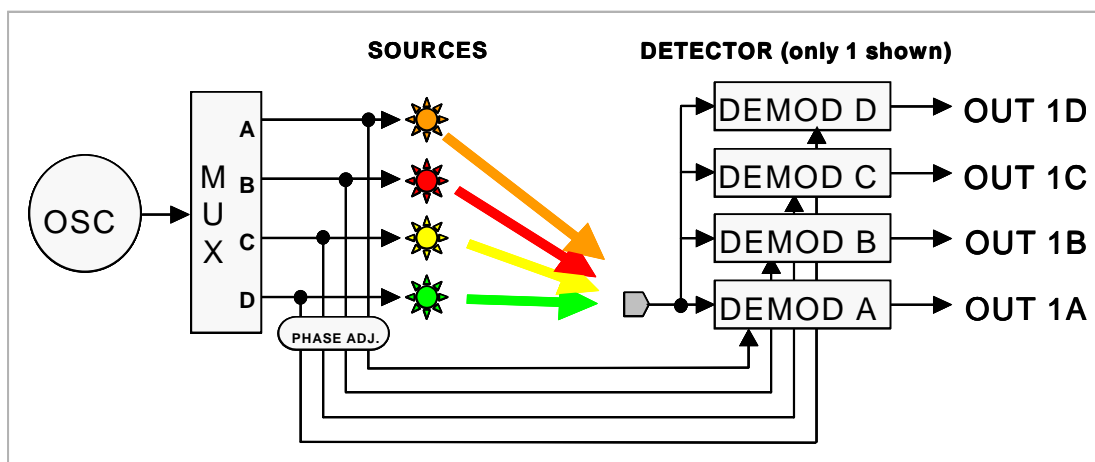


Figure 4.20. A simplified block diagram depicting the Pulse-TDM concept.




Figure 4.21. A timing diagram for the NIRS-3 system. Each of the four sources is pulsed at a 12.5% duty cycle in rapid succession. The detector signal from each of the four sources is then synchronously demodulated to generate four analog voltages, each of which corresponds to the instantaneous flux reaching the detector from that source.

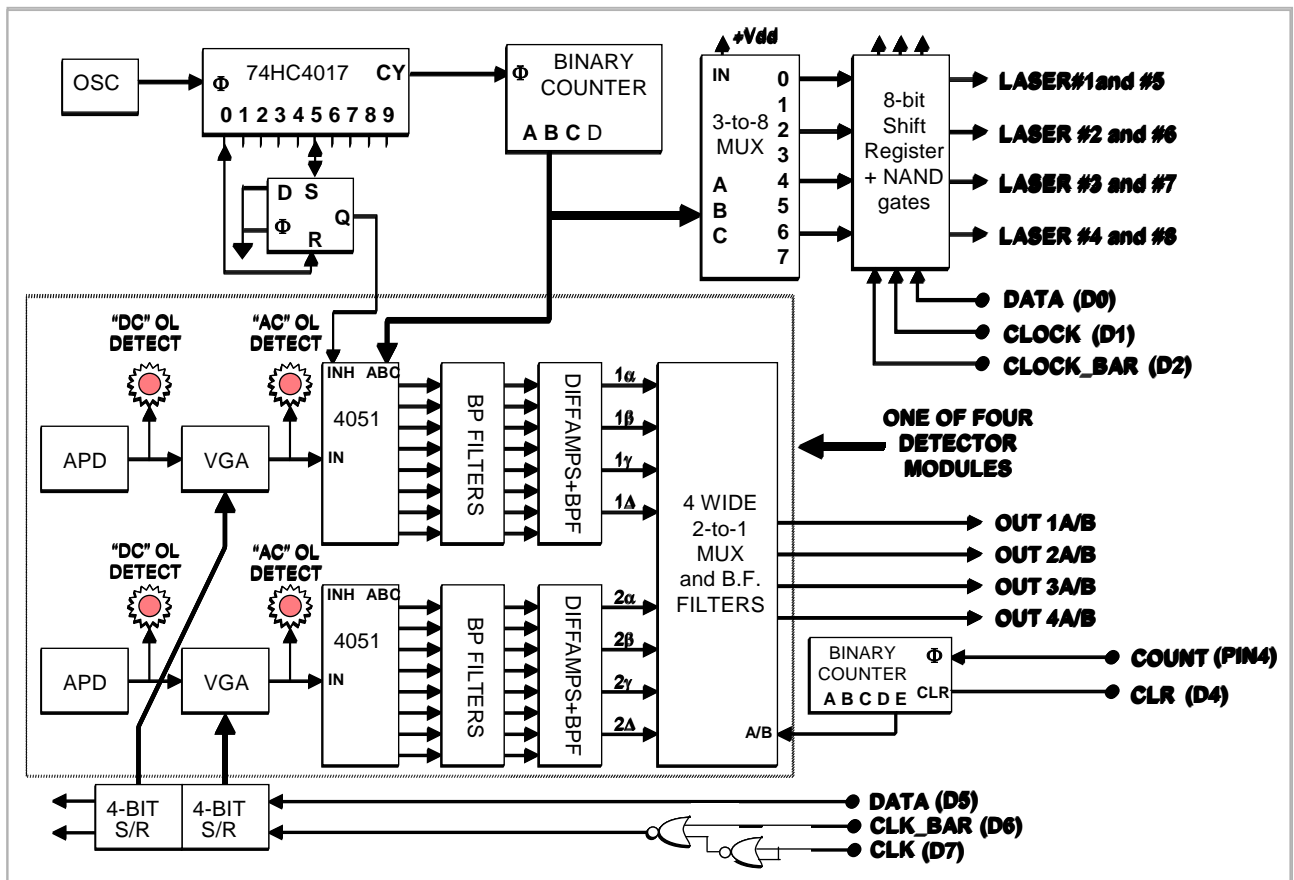



Figure 4.22. A block diagram of NIRS-3. A quartz crystal clock is divided down in frequency to drive both the laser pulsing circuitry and demodulation circuitry. The optical pulses from the lasers are converted to voltages by the silicon APD detector modules, which are then fed through variable-gain amplifiers and then demodulated to recover the slowly varying DC signals. The 32 analog signals are then digitized by an external PCMCIA-style ADC card. Both source selection and gain settings are computer-controlled through three digital input lines. Complementary clock lines protect against EMI-induced glitches during a measurement.

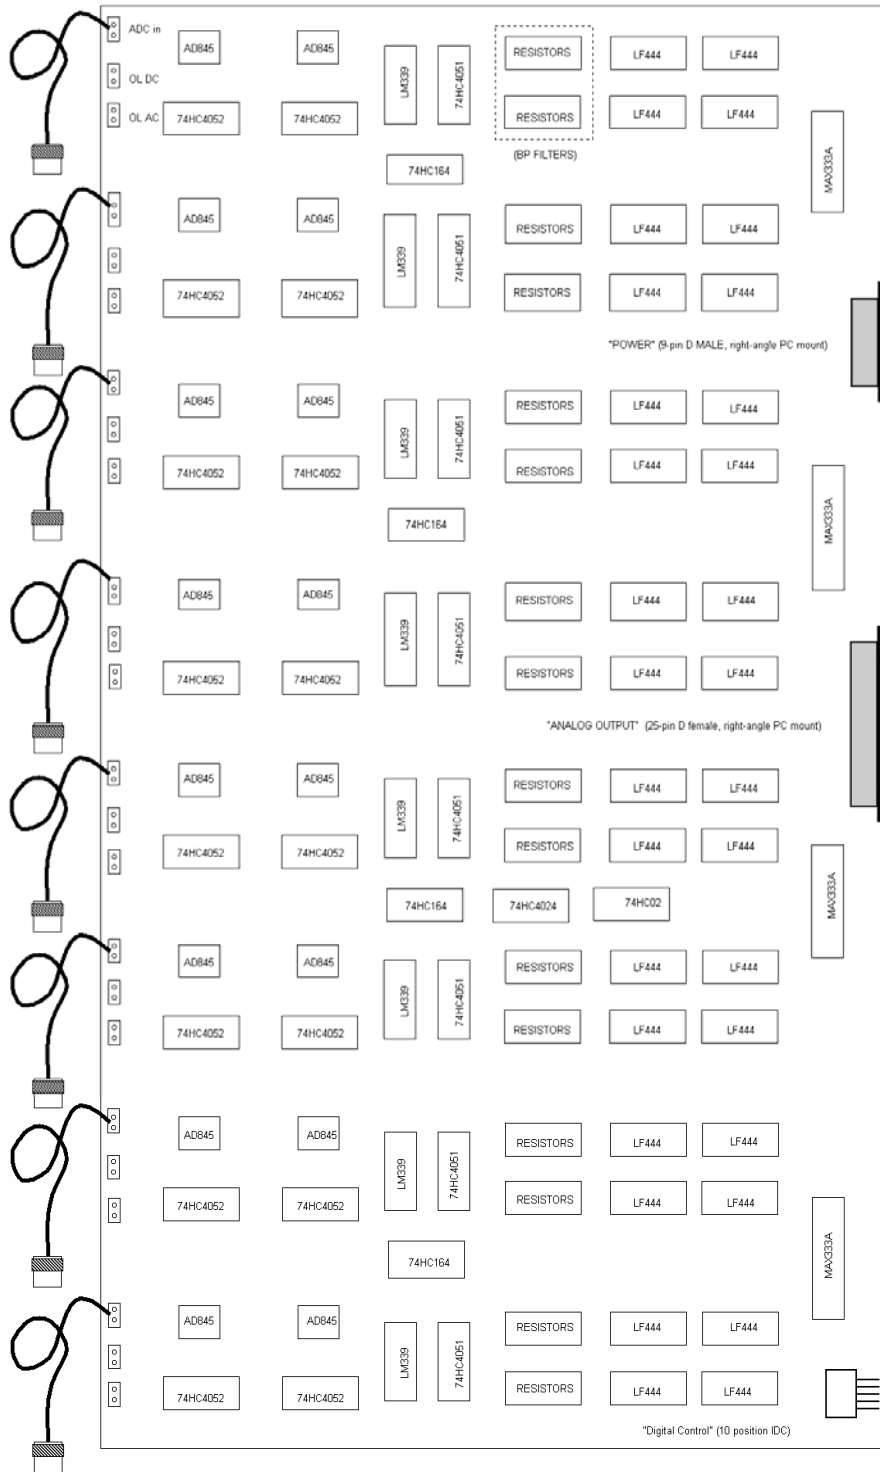
Mechanical Assembly Drawings



Analog Board Component Layout

Analog Board component layout details
Last Rev. 8/31/00

A. Siegel



POWER Conn. 9 pin D female

FUNCTION	PIN#
+12V	1
+12V RTN	2
GND	3
-12V RTN	4
-12V	5
+5V RTN	6
+5V	7
-5V	8
-5V RET	9

CABLE DETAILS:

ADC Conn.

FUNCTION	PIN#	GC
OUT1/2A	1	---
OUT1/2B	2	---
OUT1/2C	3	---
OUT1/2D	4	---
OUT3/4A	5	---
OUT3/4B	6	---
OUT3/4C	7	---
OUT3/4D	8	---
A GND	9	---
A GND	10	---
DATA	11	---
CLK_BAR	12	---
CLOCK	13	---
OUT5/6A	14	---
OUT5/6B	15	---
OUT5/6C	16	---
OUT5/6D	17	---
OUT7/8A	18	---
OUT7/8B	19	---
OUT7/8C	20	---
OUT7/8D	21	---
A GND	22	---
	23	---
CLEAR	24	---
ADC PACER	25	---

PCM-DAS16 Conn.

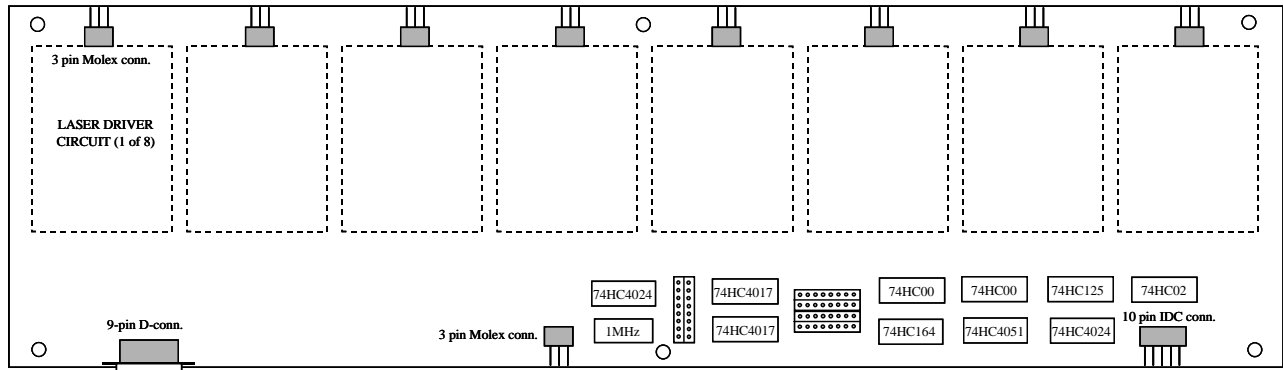
FUNCTION	PIN#
Channel 0	2
Channel 1	4
Channel 2	6
Channel 3	8
Channel 4	10
Channel 5	12
Channel 6	14
Channel 7	16
Analog Ground	1
Analog Ground	18
DIO#5	32
DIO#6	33
DIO#7	34
Channel 8	3
Channel 9	5
Channel 10	7
Channel 11	9
Channel 12	11
Channel 13	13
Channel 14	15
Channel 15	17
DIO#4	31
Counter 3 Out	47

DIGITAL CONTROL Conn.

FUNCTION	PIN#	Connects to . . .
A GND	1	Float for now
"A"	2	"A"
"B"	3	"B"
"C"	4	"C"
A GND	5	Float for now
"INH"	6	INH"
Float for now	7	Float for now
Float for now	8	Float for now
Float for now	9	Float for now
Float for now	10	Float for now

Digital Board Component Layout

LASER conn.	
FUNCTION	PIN#
Drive Pin	1
Monitor PD	2
Case GND	3




POWER Conn.	
FUNCTION	PIN#
+12V	1
+12V RTN	2
GND	3
--12V RTN	4
--12V	5
+5V RTN	6
+5V	7
--5V	8
--5V RET	9

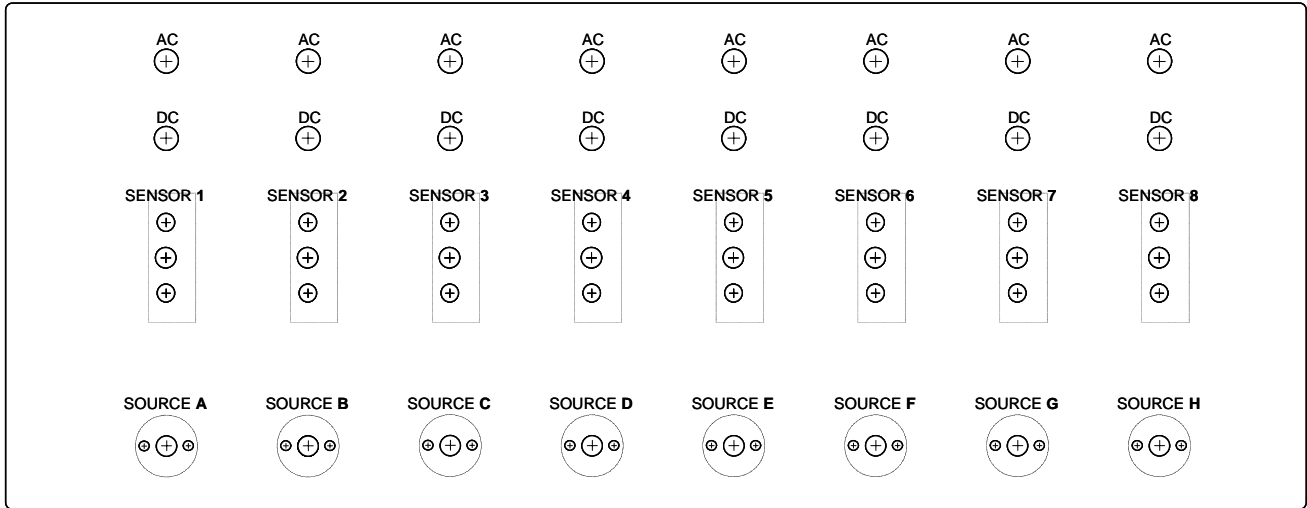
ADC Conn.			PCM-DAS16 Conn.	
FUNCTION	PIN#	GOES TO . . .	FUNCTION	PIN#
CLOCK	1	-----	DIO 1	28
CLOCK_BAR	2	-----	DIO 2	29
DATA	3	-----	DIO 3	30

DIGITAL CONTROL Conn.		
FUNCTION	PIN#	Connects to .
ANALOG GND	1	Float for now
"A"	2	"A"
"B"	3	"B"
"C"	4	"C"
ANALOG GND	5	Float for now
"INH"	6	"INH"
Float for now	7	Float for now
Float for now	8	Float for now
Float for now	9	Float for now
Float for now	10	Float for now



Back Panel Sketch



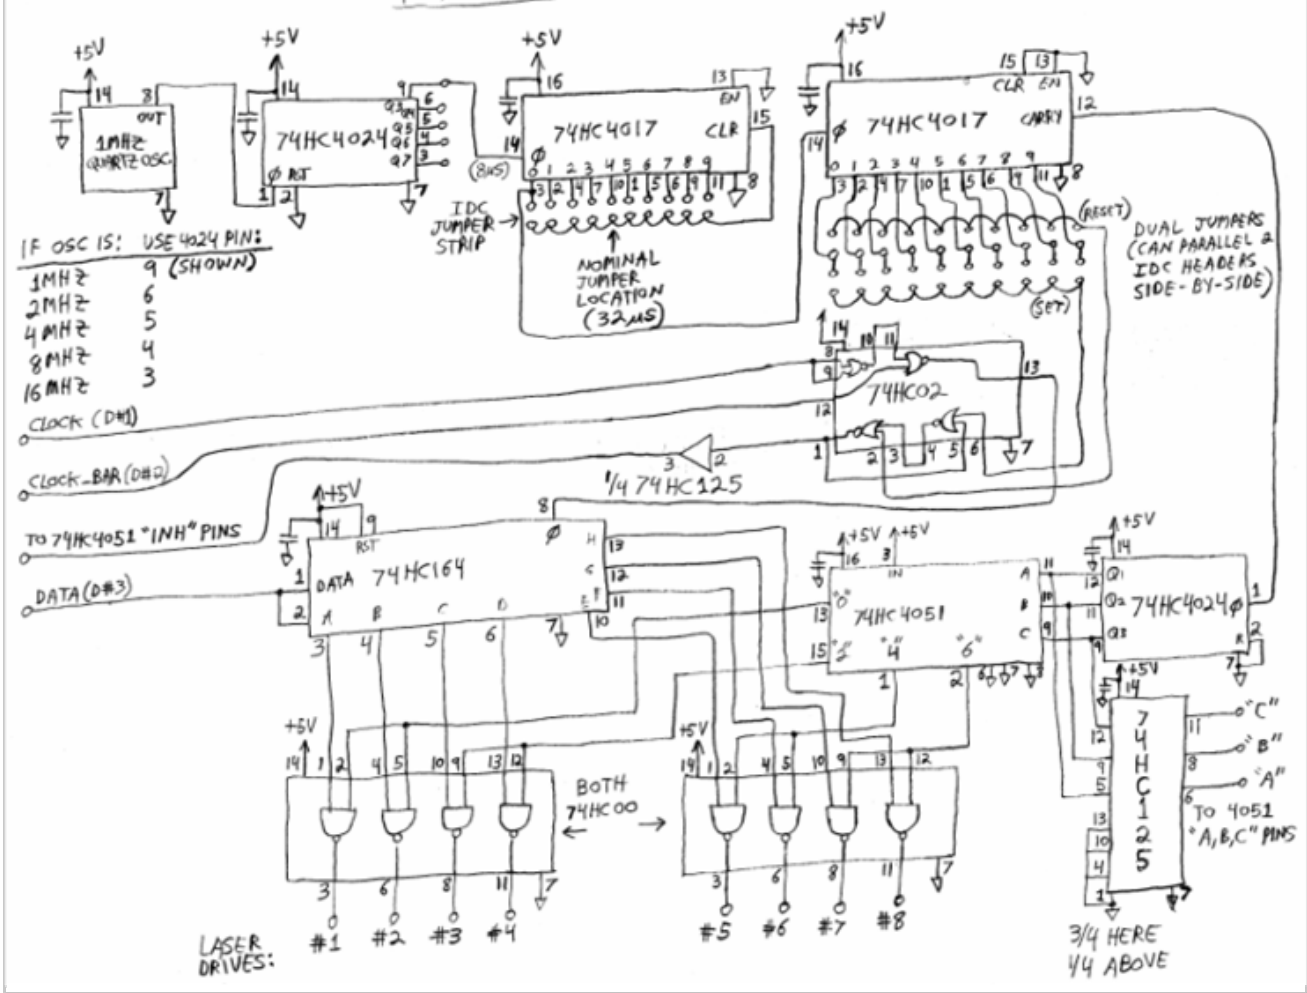
Front Panel sketch



Electrical Schematics





MASTER TIMING GENERATOR



Test results

Dynamic Performance:



Laser Source Gating Depth:

Source Wavelength	Gating Depth
680nm:	-138.0 dB
780nm:	-123.3 dB
830nm:	-124.1 dB
980nm:	-103.0 dB

Static Crosstalk:

680nm Crosstalk:

Detector #1 Only, Gain = 1V/V:
(Intrachannel Crosstalk) To Channel: 980nm: 830nm: 780nm: 680nm (main):
Crosstalk (%): -0.0147 -0.0226 -0.0116 4.32VDC

Detector #2 Only, Gain = 300V/V:
(Interchannel Crosstalk) To Channel: 980nm: 830nm: 780nm: 680nm:
Crosstalk (%): 0.0663 -0.1605 -0.1383 0.0932

830nm Crosstalk:

Detector #1 Only, Gain = 1V/V:
(Intrachannel Crosstalk) To Channel: 980nm: 830nm(main): 780nm: 680nm:
Crosstalk (%): -0.0161 4.27VDC 0.0039 -0.0119

Detector #2 Only, Gain = 300V/V:
(Interchannel Crosstalk) To Channel: 980nm: 830nm: 780nm: 680nm:
Crosstalk (%): -0.0140 0.0172 0.0300 0.0032

Gain Scaling:

Gain Setting:	GAIN = 1	GAIN = 3	GAIN = 10	GAIN = 30	GAIN = 100	GAIN = 300
Measurement:	1.01	3.33	10.00	30.39	100.00	297.25

Linearity







Goal: To measure the flux linearity of Detector #1 for all six gain states

Test Conditions: Detector #1, 830nm source "C", 200sps, 10 sec. All other sources off, other detectors at 10E9 V/A.



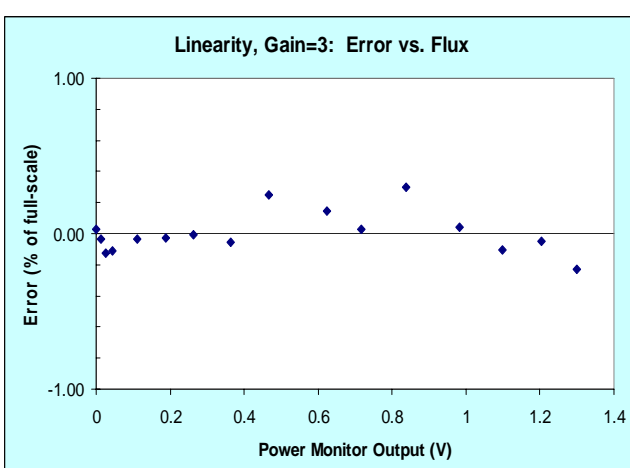



Measurements: For each flux level, measure at each gain stage, and with a power meter at 10E9 V/A.

Filename	Gain	Source	Meter	Signal(V)	Signal(DN)	Line	Diff	Error(%)
L_1_a	1	off	0	0.158	1032	0.174	-0.0165	-0.270
L_1_b	1	on	0.03	0.218	1429	0.214	0.0043	0.071
L_1_c	1	on	0.08	0.277	1816	0.280	-0.0026	-0.043
L_1_d	1	on	0.22	0.461	3023	0.465	-0.0035	-0.057
L_1_e	1	on	0.58	0.944	6189	0.941	0.0038	0.062
L_1_f	1	on	0.92	1.399	9169	1.390	0.0092	0.150
L_1_g	1	on	1.31	1.927	12628	1.905	0.0215	0.352
L_1_h	1	on	1.56	2.240	14682	2.236	0.0045	0.074
L_1_i	1	on	1.85	2.608	17090	2.619	-0.0115	-0.188
L_1_j	1	on	2.18	3.063	20074	3.055	0.0077	0.127
L_1_k	1	on	2.65	3.677	24099	3.677	0.0007	0.011
L_1_l	1	on	2.93	4.048	26527	4.047	0.0011	0.018
L_1_m	1	on	3.375	4.586	30057	4.635	-0.0484	-0.794
L_1_n	1	on	3.66	5.019	32895	5.011	0.0079	0.130
L_1_o	1	on	3.84	5.257	34451	5.249	0.0075	0.123
L_1_p	1	on	4.11	5.621	36840	5.606	0.0152	0.248


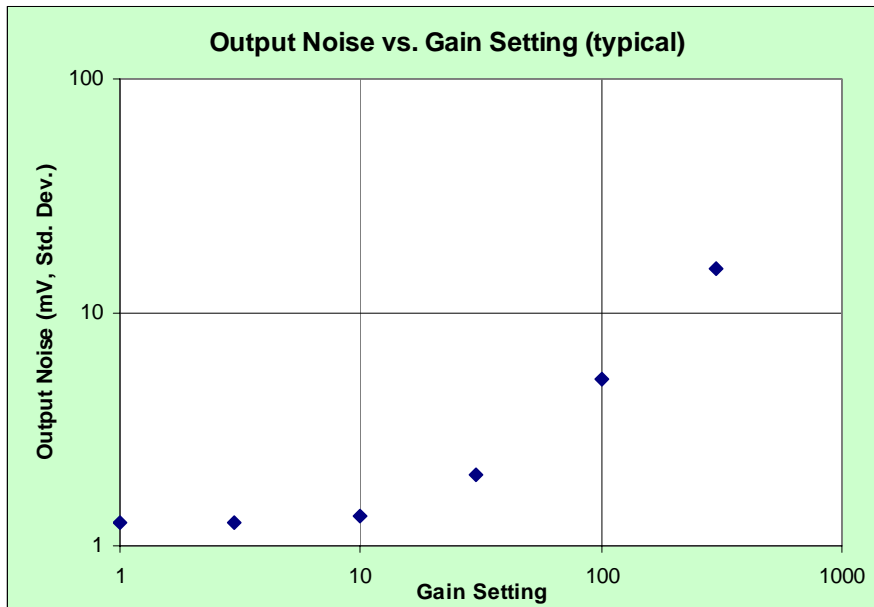
Linearity Plots:



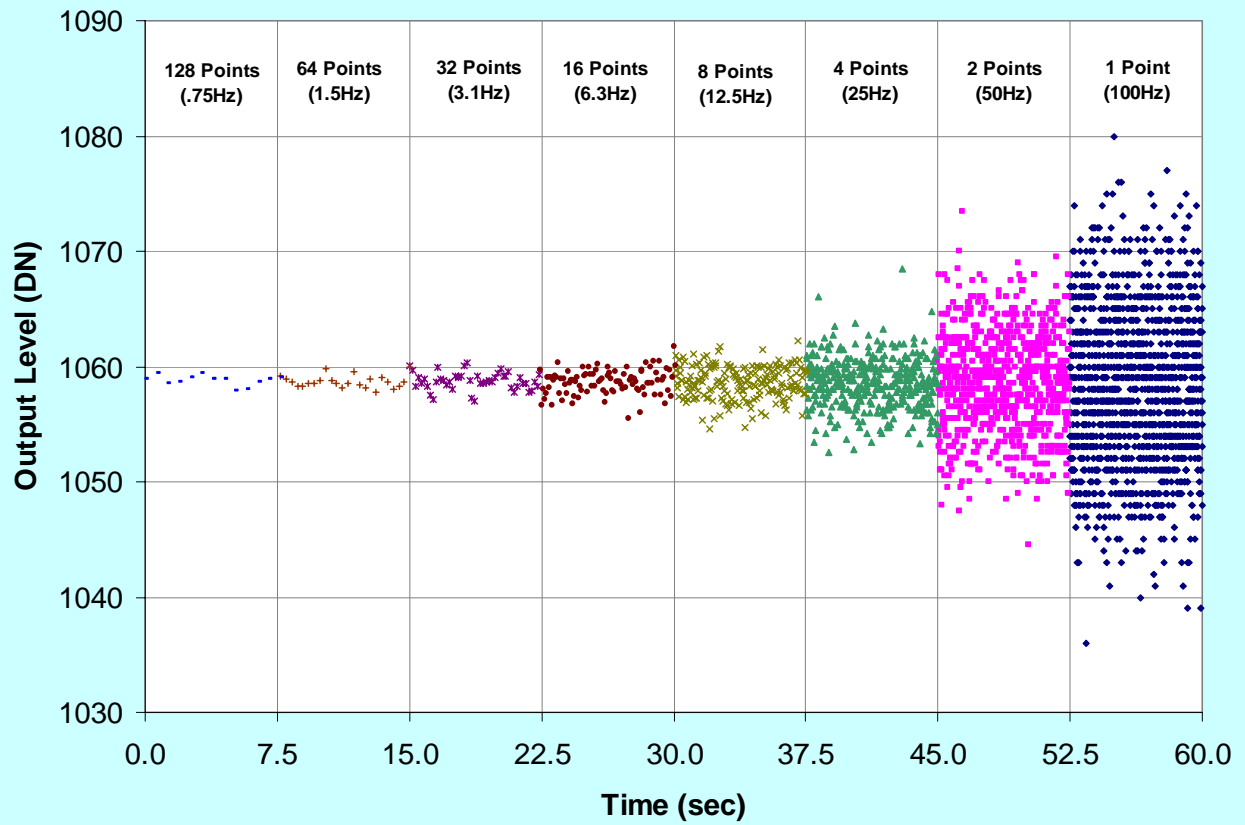
Linearity:



Noise Performance:



Noise Level vs. Number of Coadded Points: Gain = 3



Long-Term Drift: Over 10 minutes, all units in millivolts, all sources were off.

	Gain:	300	100	30	10	3	1
Channel #1	980nm	0.142	0.179	-0.084	0.318	0.046	0.285
	830nm	-0.153	0.079	-0.037	0.241	-0.049	0.211
	780nm	-0.090	0.157	-0.081	0.285	-0.108	0.224
	680nm	0.212	0.385	-0.055	0.304	-0.089	0.188
Channel #2	980nm	0.277	0.136	-0.068	0.224	-0.098	0.231
	830nm	-0.179	0.116	-0.099	0.276	-0.086	0.232
	780nm	-0.437	-0.051	-0.196	0.241	-0.058	0.220
	680nm	0.243	-0.058	-0.132	0.217	-0.005	0.231
Channel #3	980nm	-0.283	-0.332	-0.084	0.242	0.027	0.199
	830nm	0.700	-0.222	-0.054	0.280	0.044	0.279
	780nm	0.510	-0.140	-0.178	0.237	-0.033	0.246
	680nm	0.048	0.134	-0.163	0.203	-0.033	0.185
Channel #4	980nm	0.877	0.248	-0.048	0.224	-0.087	0.243
	830nm	0.118	0.017	-0.050	0.260	-0.125	0.217
	780nm	0.481	0.034	-0.073	0.318	-0.107	0.200
	680nm	0.092	0.148	0.007	0.279	-0.112	0.237
Channel #5	980nm	0.133	-0.140	-0.097	0.013	0.033	0.025
	830nm	-0.170	-0.235	0.052	0.042	0.008	-0.006
	780nm	-0.369	-0.269	0.034	0.003	0.007	0.010
	680nm	0.001	-0.294	0.012	-0.008	0.046	0.065
Channel #6	980nm	-0.165	0.428	-0.025	-0.003	0.065	0.014
	830nm	-0.304	-0.487	0.050	0.017	-0.024	0.029
	780nm	-0.062	0.205	0.070	-0.011	-0.023	-0.012
	680nm	-0.015	0.092	0.135	-0.024	-0.070	0.038
Channel #7	980nm	0.082	0.034	0.016	-0.015	-0.007	-0.054
	830nm	-0.550	-0.014	-0.012	0.020	0.004	-0.026
	780nm	-0.270	0.009	0.002	-0.013	-0.028	-0.020
	680nm	0.221	0.053	0.014	-0.059	0.043	0.023
Channel #8	980nm	0.928	0.300	-0.003	-0.020	0.039	0.024
	830nm	0.411	-0.210	0.046	0.034	0.042	0.030
	780nm	0.368	-0.079	0.075	0.093	-0.047	0.012
	680nm	0.096	-0.189	-0.024	-0.033	0.031	-0.006

4.5.5 Animal III

Both Animal II and NIRS3 showed that Pulse-TDM worked well. However, at around this same time, the CW1 imager was being replaced by a frequency encoded imager soon to be known as CW4. One of the problems soon discovered with frequency encoding was that some detector channels were being saturated by high flux levels from nearby sources, so that weaker signals from more distant sources were no longer detectable. This led to the idea of dynamic gain-switching: using variable gain amplifiers to scale the detector channel gains according to the amount of flux which they are exposed

to. Since Pulse-TDM is a time-resolved encoding technique, dynamic gain-switching seemed like a natural extension of the TDM concept. Animal II was then upgraded to include variable-gain circuitry, and the result was renamed Animal III.

Animal III was built upon the chassis of Animal II in order to determine whether variable-gain amplification can be used to enhance the dynamic range of a Pulse-TDM system. The goal was to demonstrate that dynamic gain-switching could provide a substantial improvement over “fixed-gain” designs for large-area (i.e. human cortical) DOT measurements. It showed that variable gain amplifiers (VGAs) constructed using switched-gain techniques work well together with Pulse-TDM multiplexing.

This upgrade involved a complete redesign of the electronics, and the conversion from the older 37-pin D connector to a new 50 pin IDC-style connector interface to the newer PCMCIA ADC cards. Animal III included independent variable-gain controls for each of the four sources, with four gain selections: 0dB, 20dB, 40dB, and 60dB. In principle, these four sources could represent either four different wavelengths, or four different locations, or any combination thereof.

A Variable Gain Amplifier (VGA) circuit capable of providing 60dB of gain control with greater than 100dB of channel isolation was desired. However since Pulse-TDM operates in the time domain, signal crosstalk and channel isolation are primarily governed by the preamplifier settling time, and can in principle approach the very low static crosstalk level of the circuitry. Thus the channel isolation should be extremely high (better than 120dB, the anticipated static crosstalk level of the VGA). Also, drift and DC offsets should not be an issue, since offset errors cancel out with differential measurements.

Preamp settling was improved by including input and output voltage limiting stages. These improved the saturation recovery of the amplifier output stages by preventing very strong optical signals from disturbing the DC operating point of the amplifiers, possibly corrupting the weaker signals that may follow. The AD829 opamp is well-suited to this application. It is a voltage-mode opamp with an enormous GBWP, yet is unity-gain stable. Its compensation pin connects directly ahead of the output stage, so diode clamping is easy to implement.

DC restoration was also performed between each gain stage in an effort to keep the signal levels within the dynamic range of the analog switches. Since light is detected as power (i.e. the square of the E-field), all signal excursions will be above the “zero flux” point - thus the output of a “quantum” detector is inherently DC-restored. Keyed clamping was not used due to concerns over saturated pulse broadening (i.e. there was no consistently “flat” region to clamp to). The time constants of the DC restoration circuitry were intentionally set quite long (about 100ms, as compared to the 1024us source cycle time) to be much slower than the ~12ms baseband temporal response to prevent “recovery crosstalk.” This slow a time constant will reduce transient (but not steady-state) ambient light rejection, which should not be a problem if jacketed fibers and well-designed optode assemblies are used.

A decision had to be made as to whether to employ static gain-switching or dynamic gain-switching. It turns out that, although counterintuitive, less circuitry is required to implement dynamic gain-switching. Static switching requires a 4-to-1 mux and set of demodulation switches for each of the four sources. Since space is limited on the Animal system breadboard, this could present a problem. If dynamic switching were used, then only one 4-to-1 mux and a single 8-to-1 analog switch would be required to perform both gain selection and demodulation. For this application switching speed is not a problem, since there is a generous (~62.5us) detector/preamp settling time delay built into the timing anyway. For these reasons, dynamic gain switching was chosen for the VGA.

The final VGA design consisted of four serially-connected gain stages with DC-restored inputs and voltage-clamped outputs, as shown in Figure 4.10. This output-switched VGA design was preferred over the stage-insertion approach - “switching in” higher gains stage by stage - for a number

of reasons. Switching only amongst the output nodes eliminated the risk of charge injection or electrical transients upsetting the charge-sensitive DC-restoration nodes at the amplifier inputs. Also, since amplifier outputs are servostabilized through negative feedback, settling is rapid. A gain-switching increment of 20dB was chosen to give a good compromise between low stage count and adequate gain resolution. This provided a gain range of 1000:1, or 60dB, to work with. Eight toggle switches, one pair for each source wavelength, were used to select the gains for each source in binary fashion, providing the four selectable gain states. Two 4-to-1 analog muxes (one 74HC4052) were incremented to read the settings of the eight switches. A 4-to-1 analog mux (1/2 74HC4052), controlled by the gain switch settings, then selected the appropriate tap points in the amplifier chain. The output of the 74HC4052 then fed into a 74HC4051, here used as a 1-to-8 analog demultiplexer for demodulation. A more detailed discussion of Animal III, along with the results of an in-vivo validation measurement, can be found in Section 4.3.5.

4.5.6 Gertrude

After CW1 was in use, it soon became apparent that the Burr-Brown photodiode detectors had poor sensitivity. This limited optode spacings, which subsequently limited the optical penetration depth, thus reducing the utility of CW1 in humans. Other detector options included avalanche photodiodes and photomultiplier tubes, however these were both physically large and very expensive. Since DOT instruments typically employed 16 or more detectors, this cost difference could amount to thousands of dollars per instrument. One possible solution to this was to increase the source power delivered to the scalp. CW1 provided about 4mW of optical power per source. If this source power could be safely increased by a factor of 100, then an SNR comparable to that achievable with APD detectors could be obtained at a far lower cost with simple photodiodes. Thus, before future instruments were constructed, it was important to determine whether adequate optode separations could be achieved by increasing source power alone. The system constructed to evaluate this possibility was named Gertrude.

Gertrude was an instrument constructed in an early effort to determine the upper range of source-detector separations achievable on an adult human subject using only amplified silicon photodiode detectors. Since the Burr-Brown OPT209 detectors available at this time were not very sensitive, high source powers were required to achieve measurements at even modest (~4cm) optode spacings. For the measurement shown in Figure 4.23 the 808nm source power, estimated to be around 500mW, was delivered through a 1mm silica core fiber and terminated with a 5mm length of 3mm diameter PMMA fiber which was placed in direct contact with the scalp. This short length of PMMA fiber served as a simple power diverger, which spread the beam waist to around 2mm diameter, thus reducing the optical power density at the tissue surface by about a factor of four. Despite this, the 40 minute exposure resulted in a raised welt which took about a week to resolve. As a result, it was decided that increasing the source powers above ~4mW was unwise, for both medical reasons (ouch!) and legal reasons. It was soon discovered that the legally mandated upper power limit for cutaneous exposure to near-IR light was in the low mW range. Thus, if the SNR of DOT measurements was to be improved, it must be done through a reduction in “N,” not through an increase in “S.”



Figure 4.23. An image showing Gertrude in use. The gray die-cast boxes enclosed the high-power 808nm laser source and an electrolytic strain gauge respirometer which was used to monitor ventilation. The three OPT209 detectors were located on the optode assembly and were optically coupled to the scalp through light pipes constructed from 1cm lengths of 3mm diameter PMMA fiber.

4.6 The frequency-encoded DOT imager: CW4

CW4 was designed to be an improved version of CW1. It contains laser diode sources like CW1, but the silicon photodiodes were replaced with avalanche photodiode modules for improved sensitivity. This sensitivity improvement may then be traded for better temporal response.

CW4 would employ frequency-encoding, and demodulation would be performed in the digital domain instead of inside the instrument as with CW1. The reasoning behind this approach was twofold:

Shifting the demodulation process into the digital domain would afford far more flexibility in how the data was processed. Changes in postdetection bandwidth (i.e. varying the temporal response) could be made after the data was collected – in fact the same data could be processed with digital filtering algorithms at multiple bandwidths.

The detector electronics could be greatly simplified. Only optical detection and gain adjustment would be implemented in hardware. The analog signals from the imager would be digitized by an external ADC card in a computer located nearby.

Frequency encoding provided very low temporal skew, which was vital for obtaining accurate spatio-temporal hemodynamic measurements. But this came at the cost of a limited dynamic range, since the detected flux from nearby sources could easily saturate the detector if the gain values were set high enough to measure the flux from more distant sources. This limitation was only an issue when

measuring human subjects, and frequency-encoding worked quite well for the small optode geometry used for the rodent measurements discussed in Chapter 6.

Other disadvantages with this instrument included a lack of real-time imagery and the need for extremely large (i.e. GB-size) raw datafiles. The processing delays complicated the measurements, since the cycle time between measurements and imagery was tens of minutes, hampering efforts at probe adjustments or the detection of defective optodes. Future upgrades will include a real-time imaging capability, which will greatly simplify detector gain selection and probe alignment.

Characterizing the CW4 Imager

The CW4 system contains nine 780nm and nine 820nm laser diode sources, and sixteen silicon APD detectors. The raw detector signals are amplified and buffered in the chassis and then digitized in the host computer and saved directly to disk. The raw data can then be post-processed using various custom DSP algorithms employing FFT techniques to implement multiple bandpass filters and envelope detectors in the digital domain. This allows design trades formerly fixed in hardware to be varied at will, simply by modifying the DSP software. The measurements summarized in Table 4.11 were processed using a simple FFT filter algorithm and thus do not represent the true channel isolation capability of the instrument. Other parameters more accurately reflect the instrument performance, although these may also improve once advanced correction and calibration routines are developed.

Table 4.11. A summary of the features and measured performance of the CW4 instrument. The spectral measurements were performed with an Ocean Optics Model S2000 with a 100um multimode fiber, using the 00lBase32 software: Tint = 1ms , 50 Averages, Boxcar = 0. The power measurements were performed with a Thorlabs Model S20MM plus a Thorlabs Model ND10A 10% transmission (ND1) metal film neutral density filter. The wavelength settings were 781nm and 820nm.

LASER SOURCES	
780nm sources:	Sanyo DL7140
830nm sources:	Hitachi HL8325G
Spectral Range:	778-782nm, 820-822nm
Nominal Output Power:	2.5-3.5mW
Modulation:	Sinusoidal AM
Frequency Range:	4000Hz to 8000Hz
Frequency Spacing:	200Hz
Frequency Stability:	<0.1Hz (estimated)
Optical Interface:	SMA 905 Female Receptacle
Optical Power Drift:	<0.3% over 30 minutes
APD DETECTOR CHANNELS	
Detectors:	Hamamatsu C5460-01
Gain Uniformity (worst-case):	30%
NEP (G=12, BW=3Hz):	0.05pW
Instantaneous Dynamic Range:	83dB @ G=3, 40dB @ G=60
Nonlinearity:	4% @ 80dB
Nearest-Neighbor Crosstalk:	-48dB

More detailed information is provided in the tables and figures below:

CW4: Laser Wavelength, Output Power vs. Source Number

Source Number	Spec λ	Laser Type	Modulation Freq. (Hz)	Measured with 600um core SiO2 fiber		
				Measured λ (modes)	Meas. FWHM	Measured Power
1	830nm	Hitachi HL8325G	5787	819.8, 820.7, 821.2, 822.0	0.8nm/mode	1.10mW - 2.53mW
2	830nm	Hitachi HL8325G	6188	819.8, 820.7, 821.2, 822.0	0.8nm/mode	0.40mW - 2.64mW
3	830nm	Hitachi HL8325G	6010	820.0	0.8nm/mode	1.17mW - 3.77mW
4	830nm	Hitachi HL8325G	6614	819.8, 820.7, 821.2, 822.0	0.8nm/mode	1.76mW - 2.64mW
5	830nm	Hitachi HL8325G	6410	819.8, 820.7, 821.2, 822.0	0.8nm/mode	1.40mW - 2.70mW
6	830nm	Hitachi HL8325G	6983	819.8, 820.7	0.8 - 2.5nm	1.86mW - 2.77mW
7	830nm	Hitachi HL8325G	6794	820.7, 821.2, 822.0	0.8 - 2.5nm	1.42mW - 2.76mW
8	830nm	Hitachi HL8325G	7397	820.0, 820.3, 821.0, 821.3, 822.0	2.5nm	1.78mW - 3.35mW
9	830nm	Hitachi HL8325G	7184	820.7, 821.7, 822.3	0.8 - 2.5nm	1.44mW - 2.95mW
10	785nm	Sanyo DL7140	4006	777.2nm, 778.7nm	2.6nm	0.28mW - 3.57mW
11	785nm	Sanyo DL7140	4401	~781.8	1 - 1.5nm	2.18mW - 3.67mW
12	785nm	Sanyo DL7140	4196	780.8, 781.5	~2nm	0.33mW - 3.66mW
13	785nm	Sanyo DL7140	4808	~781.2	1 - 1.5nm	1.80mW - 3.09mW
14	785nm	Sanyo DL7140	4596	~781.3	1 - 1.5nm	0.90mW - 4.10mW
15	785nm	Sanyo DL7140	5208	~781.2	1.5nm	0.55mW - 3.59mW
16	785nm	Sanyo DL7140	5000	782.1, 782.8	1 - 1.5nm	0.18mW - 3.80mW
17	785nm	Sanyo DL7140	5605	~781.4	2nm	0.18mW - 3.12mW
18	785nm	Sanyo DL7140	5388	~781.5	1.5nm	0.60mW - 3.45mW

CW4: SOURCE OUTPUT POWER STABILITY vs. TIME

t = 0	t = 5 min.	t = 10 min.	t = 15 min.	t = 20 min.	t = 25 min.	t = 30 min.
3.05mW	3.05mW	3.04mW	3.04mW	3.05mW	3.04mW	3.04mW

CW4: SOURCE OUTPUT POWER vs. SYSTEM STATUS


Source	This source on only	All 9 820nm sources on	All 9 780 nm sources on	All 18 sources on
#1 (820nm)	3.00mW	2.83mW	2.79mW	2.78mW
#12 (780nm)	2.73mW	2.33mW	2.37mW	2.31mW

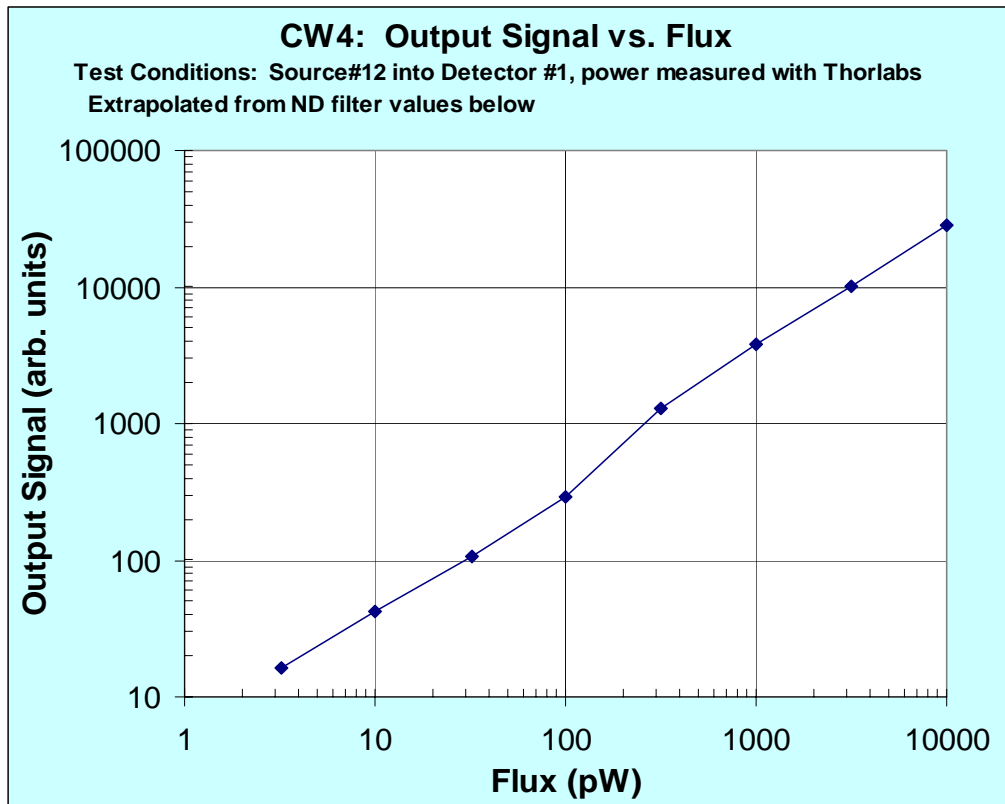
CW4: SYSTEM STABILITY

Time (min)	Detector #1		Change
	Signal	Noise(pp)	(%)
0	30690	70	0.00
5	30770	65	0.26
10	30900	70	0.68
15	30830	75	0.46
20	30835	40	0.47
25	30800	45	0.36

CW4: DETECTOR GAIN UNIFORMITY

Detector Number	Output Level (first meas.)	Output Level (second meas.)
1	0.88	0.88
2	0.77	0.76
3	0.61	0.61
4	0.79	0.78
5	0.70	0.70
6	0.78	0.78
7	0.77	0.83
8	0.78	0.78
9	0.68	0.68
10	0.73	0.72
11	0.76	0.75
12	0.66	0.66
13	0.72	0.71
14	0.78	0.76
15	0.75	0.74
16	0.71	0.72





Source Performance

The 780nm sources were all within ~2nm of their nominal value, however the 830nm lasers ranged from 820nm to 822nm – a significant change from their nominal value. This was an important discovery, since it affected our hemoglobin concentration estimates.

The nominal output powers were all below 4mW, a value which is in accordance with current safety guidelines. The power stability was quite good, varying by only 0.3% over a span of 30 minutes. Frequency stability was estimated to be within 0.1Hz at 8kHz, which was expected since all of the modulation frequencies were derived from a commercial quartz crystal oscillator module with a nominal thermal frequency stability of around 10PPM /°C, and the instrument was warmed up for about an hour prior to these drift measurements.

Detector Performance

The noise performance was only about a factor of two worse than the specified noise floor of the APD alone at high post-amp gain (typical NEP = 0.02pW/Hz⁻¹, = 0.032pW @ 3Hz). This is excellent, however it was measured under ideal conditions with only a single optode pair. Actual DOT measurements may employ lower gains due to the wide range in detected flux.

The instantaneous dynamic range of 83dB was quite high at a relatively low post-amp gain of 3. Given the limited capabilities of the 14 bit ADC card, it is likely that the dynamic range was ADC-limited for this measurement. As expected, the dynamic range drops to around 40dB at a post-amp gain of 60, which is representative of the gains employed during actual DOT measurements with multiple optode spacings. This is commensurate with the observed crosstalk levels of around -48dB, so it appears that human DOT measurements will likely be limited by the system noise floor and not by crosstalk unless significant block averaging is performed.

The nonlinearity of 4% was larger than expected, and is probably due to the inherent nonlinearity of the variable-gain amplifiers used in the post-amp gain stages. These components were designed for medical ultrasound instruments, which operate over extremely wide dynamic ranges as do radar receivers, and can thus tolerate modest levels of nonlinearity. The inherent nonlinearity of both APD module and the ADC card are specified to be well below this level. This may present some problems for DOT measurements in two ways. First, the nonlinearity directly affects the amplitude of the detected signals in a manner which cannot be easily deconvolved from the measurement, since the detected optical signals are all amplitude-modulated. Second, since multiple frequencies are present in the detector channel simultaneously, intermodulation will occur. As long as the magnitude of the in-band intermodulation components are sufficiently below the system noise floor, they can be ignored. The out-of-band IM components are rejected by virtue of keeping all of the modulation frequencies within a one octave frequency span.

Summary

In order to perform accurate quantitative DOT measurements, suitable instrumentation must be designed and constructed. Chapter 4 was written to address the many issues which influence the design and development of CW DOT instrumentation. I begin by covering the parametric signal ranges for cerebral DOT measurements. Then the performance features and limitations of optical sources and detectors suitable for DOT are discussed. The motivation behind source encoding is explained, and the development and validation of new encoding techniques such as Pulse-TDM are presented. The design evolution of a series of instruments is then presented as a conceptual chronology, beginning with the first DOT imager ever constructed in our lab and ending with a frequency encoded system still in use today.